

Paolo Spada¹ and
Tiago C. Peixoto²

The Limits of Representativeness in Citizens' Assemblies: A Critical Analysis of Democratic Minipublics

Abstract: *This article critically examines the widespread claim that minipublics, such as citizens assemblies, typically represent the broader population in democratic decision-making. Through systematic analysis, we identify four fundamental challenges to representativeness: small sample sizes, group effects that complicate output legitimacy, sampling biases from population lists, and low acceptance rates. We evaluate three common strategies used to justify small sample minipublics — stratification, supermajority voting, and second-best arguments — and demonstrate why these approaches fail to resolve the underlying representativeness problems. Rather than abandoning minipublics entirely, we propose three alternative ways forward: (1) scaling up and integrating multiple independent minipublics, (2) targeting specific inclusion failures rather than pursuing broad representation, and (3) leveraging non-domination claims instead of representativeness. While the first approach faces significant technical and cost challenges, the latter two offer more practical paths forward, particularly in addressing concrete democratic deficits in existing institutions. We conclude that minipublics remain valuable democratic innovations, but their legitimacy should be grounded in their ability to address specific inclusion failures and prevent domination by organized minorities rather than claims of broad population representation.*

Keywords: deliberative democracy; sortition; minipublics; collective intelligence; sampling strategies; representativity.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101037328

¹ Centre for Democratic Futures, University of Southampton, UK. p.spada@soton.ac.uk

² World Bank Group. tpeixoto@worldbank.org

1. Introduction

Robert Dahl introduced the idea of the ‘minipopulus’ in *After the Revolution? Authority in a Good Society* (1970): an assembly of citizens that could represent the views of the entire citizen body. The idea was not operationalized by Dahl, but his intuition focused on leveraging sampling methodology developed for survey analysis to include a broadly representative sample of the population in the assembly. In *Democracy and Its Critics* Dahl offers a more refined description:

Suppose an advanced democratic country were to create a ‘minipopulus’ consisting of perhaps a thousand citizens randomly selected out of the entire demos. Its task would be to deliberate, for a year perhaps, on an issue and then to announce its choices. The members of a minipopulus could ‘meet’ by telecommunications. One minipopulus could decide on the agenda of issues, while another might concern itself with a major issue. Thus, one minipopulus could exist for each major issue on the agenda... It could be attended — again by telecommunications — by an advisory committee of scholars and specialists and by an administrative staff. It could hold hearings, commission research, and engage in debate and discussion. (Dahl, 1989, p. 342)

In classical representative democracy, legitimacy is established through free and fair elections, where everyone has the right to vote. In town halls and similar forums, legitimacy is based on the principle that everyone can attend if they wish. However, participation in minipublics is not open to all; instead, their key justification is that they are representative.

For example, according to Curato, Farrell and Geissel (2021, p. 3), deliberative minipublics are defined as ‘carefully designed forums where a representative subset of the wider population come together to engage in open, inclusive, informed and consequential discussions on one or more issues’.

But how difficult is it in practice to achieve representativeness? Are most adopted minipublics representative of the population? If they are not, should we abandon them? Or are there other sources of legitimacy for minipublics that do not achieve representativeness?

In this paper we adopt a narrow procedural definition of minipublics as processes in which participants: (1) are randomly selected (often combined with some sort of stratification), (2) participate in informed deliberation on a specific topic, and (3) reach a public judgment and provide recommendations on that topic. Thus, in this text, ‘minipublics’ serves as a general term for a variety of practices such as consensus conferences, citizens’ juries, planning cells, and citizens’ assemblies. The focus on random selection of participants helps distinguish these processes from other democratic innovations designed to improve or deepen our existing democratic systems by promoting citizen deliberation and participation. The procedural nature of our definition allows us to evaluate under what conditions the claim of representativeness might be achieved.

Initially, the idea of a minipopulus remained a niche academic proposition, with rare pilots mostly by universities at the local level. Early academic literature was more cautious about the ability of minipublics to achieve statistical representa-

tiveness for the general population of a city or country (Goodin and Dryzek, 2006; Smith, 2009).

Over the past 15 years, we have observed an explosion of minipublics at multiple levels of governance. The OECD crowdsourced dataset showcases a significant increase in these democratic innovations starting around 2008 (see Figure 1). Most minipublics were adopted by high-income countries (Figure 2) and at the local level (see Figure 3), but some pilots have been adopted by low- and middle-income countries. Also, some of the most notable recent minipublics were implemented at the national and multinational levels. The dataset is incomplete, making these numbers a conservative estimate of the actual phenomenon, aptly termed a ‘deliberative wave’. Propelled by this wave and by high-profile national cases such as the citizens’ assemblies in Iceland (2009), Ireland (2012), and France (2019), minipublics have gained popularity beyond small circles of scholars and advocates. From CNN to the *New York Times*, and from the *Hindustan Times* (India) to *Folha de São Paulo* (Brazil), major newspapers around the world have covered stories on minipublics showcasing the global interest in this topic across multiple languages.

In the excitement generated by this unprecedented wave, some of the caution present in most of the early academic literature has been lost. Practitioners and academics who promote and study these democratic innovations often describe them as the best devices to represent the will of the people on a given topic, potentially overlooking the limitations highlighted in earlier research.




Figure 1. Source: Crowdsourced OECD dataset 2023, most likely underestimating the phenomenon.




Figure 2. Source: Crowdsourced OECD dataset 2023, most likely underestimating the phenomenon.




Figure 3. Source: Crowdsourced OECD dataset 2023, most likely underestimating the phenomenon.

Current proponents of minipublics put representativeness at the core of their definition and evangelical activities. In fact, it is one of their main selling points. For example, the OECD’s ‘Catching the Deliberative Wave’ Report (OECD, 2020), the United Nation’s ‘Democracy Beyond Elections Handbook’ (New-Democracy Foundation, 2018), Marcin Gerwin’s ‘Citizens’ Assemblies Guide’ (Gerwin, 2018), and Mosaic Lab’s ‘Facilitating Deliberation: A Practical Guide’ (White, Hunter and Greaves, 2022), all highlight that an advantage of citizens’

assemblies, compared to other mechanisms of participatory democracy, is their typical combination of random selection and stratification to form a body that is representative of the whole public. This general argument resonates with the media and the wider public. A recent illustration is an article in the *Guardian* (2023), which depicts citizens' assemblies as 'a group of people who are randomly selected *and* reflect the demographics of the population as a whole'.

It should be noted that claims of representativeness vary in their assertiveness. For instance, some may refer to citizens' assemblies as 'representative deliberative democracy', while others may use more cautious language, referring to assemblies' participants as being 'broadly representative' of the population (e.g. by gender, age, education, attitudes). Moreover, some practitioners, acknowledging the difficulty of representing the national population with the small samples typically recruited for minipublics, recommend the usage of super-majority decision rules as a mechanism to compensate for the error generated by lack of representativity; we will analyse the details of this strategy later.

This variation in terms used to describe representativeness and the strategies to strengthen the legitimacy of the minipublics' outputs should prompt an attentive observer to ask basic questions such as: 'Are existing practices of deliberative democracy representative?' 'If they are "broadly" representative, how representative are they?' 'Should we care about the representativity of the sample of participants *per se*?' 'What criteria, if any, are used to assess whether a deliberative democracy practice is more or less representative of the population?' 'Can their representativeness be improved, and if so, how?' These are basic questions that, surprisingly, have been given little attention in recent debates surrounding deliberative democracy. The purpose of this article is to bring attention to these basic questions and to provide initial answers and potential avenues for future research and practice.

This exercise is in our opinion particularly relevant because the popularity of minipublics has started to attract criticism. For example, the failed attempt in 2003 to introduce gender-equal language in Ireland's 1937 constitution was used by some commentators and academics to open a reflection on the use and misuse of the citizens' assembly.

Assemblies could stimulate productive debate but were not necessarily representative... *One government source said there was no appetite for further attempts at constitutional change before the next general election and that faith in assemblies had been eroded.* (Carroll, 2024)

Practitioners who have been at the forefront of developing minipublics over the years are beginning to criticize the lack of impact of current approaches, the way these tools have been used to manage protest movements rather than empower citizens, and highlight the 'danger of replicating old power asymmetries, in effect making people's sense of powerlessness worse' (e.g. Wilson and Mellier, 2023). These criticisms echo concerns from political theorists such as Cristina Lafont and Nadia Urbinati (2024), who have highlighted the risk of situating these democratic innovations within a new technopopulist agenda.

Many of these practitioners and academics are not calling for the abandonment of these valuable experimentations. Instead, they advocate a more critical reflection on these processes to improve their design and our understanding of their usage within existing institutional settings. They also caution against the positioning of minipublics as complete substitutes for other democratic institutions, and instead propose a more integrated approach to democracy that supports a large repertoire of methods, some based on sortition and some not, better exploring the trade-off that each method offers (Saward, 2021; Pogrebinschi, 2023).

In line with these calls for constructive experimentation, this article highlights the problem of uncritically emphasizing the ability of minipublics to represent the population and proposes some alternatives to better inform future efforts. Our work also connects to recent experimental evidence highlighting the potential danger of overclaiming representativeness in terms of legitimacy: even a minor decrease in the representativeness of a minipublic significantly reduces its legitimacy in the eyes of citizens (Germann, 2024).

In the next sections, we will focus on what we consider to be the four main challenges to the claims of representativeness of minipublics. First, we will examine the issue of sample size. Second, we will address group effects, and why they pose a challenge to representativeness. Third, we will consider sampling error, which is the error that occurs when observing a sample rather than the entire population. Fourth, we will examine the issue of non-response. Then we will discuss three strategies that have been used to address these problems and justify small sample minipublics: stratification, supermajority voting, and second-best arguments.

In conclusion we suggest that we should stop justifying minipublics based on their representativeness, given the issues with the current generation of processes, and we should leverage alternative strategies.

2. Challenge I: Small Samples Size

The median number of minipublics' participants in the OECD database is 40, with 75% of cases with a sample smaller than 78 participants. With such sample sizes, representing a population is quite difficult. This might be the reason why political scientist Robert Dahl, who first proposed the use of minipublics over three decades ago, suggested 1,000 participants as the sample size. And it is for this same reason that most surveys that attempt to represent a complex national population have a sample size of over 1,000 people. In the OECD database at the time of writing this article there are only 6 cases with a number of participants greater than 800 (OECD Database, 2023, accessed July 2024).

To illustrate the limitations of small samples, consider a sample size of 40 people, which is the median sample size found in the OECD database (2023). Suppose 60% of the general population is in favour of adopting a certain policy. If we survey a sample of 40 people about this policy, we will have a margin of error of about 15% at a 95% confidence level. This means that if we asked our 40 participants about their stance on the policy before the assembly, their preferences

could differ significantly from those of the majority of the population. If we drew many independent samples of 40 people from the population, their preferences would be found in the range of 45% to 75% with 95% probability.³

Small survey samples may suffice when there's a strong preference in the population and we are using the survey as a majority voting mechanism. For instance, if 70% of the general population favours option A, a small sample of 40 will reflect this preference with a 14% margin of error (95% confidence) that would not significantly alter a majority vote.

However, for more contentious issues, or if the survey is not being used as a majority voting mechanism but rather for a different purpose, such as generating a ranking of the top three issues to set a political agenda, the required level of precision might be significantly higher because small error margins might have large political consequences. Consequently, the requisite sample size would also increase. Ultimately, the acceptable margin of error is a political choice rather than a statistical one. In a complex citizens' assembly, a sample of more than 370 participants, yielding an error margin of approximately 5% for population parameters between 40% and 60%, might not be sufficient to claim representativeness.

3. Challenge II: Group Effects

A minipublic is not a survey, and to explain the difference we need to distinguish between input and output representativeness. Input representativeness describes how well the initial group preferences reflect the general population before the start of the minipublic, and output representativeness as how closely the minipublic's final decision aligns with what the general population would decide if it had gone through a similar process. So, the latter is an ideal benchmark, and it is useful to understand the key difference between the idea of representativeness applied to minipublics and surveys.

In a survey, participants do not interact with each other, allowing for the use of independent, identically distributed random samples to estimate general population attitudes. In a minipublic, however, this property applies only to input representativeness; as soon as participants interact, the data is no longer independent. Consequently, the output preferences cannot be used as easily to estimate the preferences of the original population.

Cristina Lafont has used a similar argument to criticize even an ideal minipublic that achieves perfect input representativeness (Lafont, 2015; 2020). According to her argument, the simple fact that participants change their opinions negates any possibility of claiming that the decision of a minipublic is representative of the preferences of the population. Proponents of minipublics respond to Lafont by clarifying that an ideal minipublic is not a mechanism to represent the

³ For those unfamiliar with statistics, various online calculators can help determine error margins for different sample sizes, population proportions, and confidence levels. And for a more rigorous treatment see the chapters on sampling in Bethlehem (2009, chapter 5), Groves *et al.* (2011, chapter 4.5), and Heeringa *et al.* (2017, chapter 2).

‘raw’ preferences of the general population on a topic (see, for example, Fishkin, 2020).

An ideal minipublic is designed to leverage a desirable emergent property of group deliberation and improve upon the capabilities of an aggregative mechanism such as a survey (or a prediction market). A vast body of literature on the desirable properties of carefully crafted group deliberation highlights its ability to leverage diversity to achieve better decisions or, at the very least, to identify problematic options, reduce the impact of biases, and promote perspective-taking and considered judgment (Landemore and Page, 2015; Mercier and Landemore, 2012; Druckman, 2004; Muradova, 2020).

Note however that another equally vast literature highlights the danger of badly crafted group decision-making architectures (Morrell, 2014; Sunstein, 2017), and a lot of research still needs to be done to fully understand the scope and conditions of some of the claims about diversity in the literature (Holou-Garcia, 2017; Sulik, Bahrami and Deroy, 2022) and the ideal ergonomics and design of these processes (Richards *et al.*, 2022).

But, setting aside the latter problems, and assuming we can implement an ideal minipublics process that achieves the desirable properties and is reproducible, then what sample size is needed to achieve or at least approximate this ideal output representativeness? Can we achieve this with the small sample sizes commonly used in the current generation of minipublics?

To explore this question empirically, we cannot rely on strategies developed for survey analysis; instead, we need to conceptualize minipublics as a partial randomized controlled trial. In this framework, a minipublic can be seen as one treatment group in a clustered randomized controlled trial (RCT), where the ‘treatment’ is the minipublics process itself, and our objective is to estimate the average outcome of the treatment group. Unlike in an RCT, we are not interested in the Average Treatment Effect — the difference between the average outcomes of the treatment and control groups — because there is no control group, and we are not measuring changes in outcomes caused by the minipublic. Instead, our sole focus is on using the minipublic as a treatment to estimate what an ideal deliberative population would have decided.

In this scenario, each cluster, i.e. each minipublic, would be affected by a treatment effect that may be slightly lower or higher than the ideal treatment effect due to group idiosyncrasies. Therefore, simply increasing the sample size of a minipublic to achieve input representativeness would not be sufficient to approximate the ideal decision. We would need to combine multiple identical minipublics and average their results.⁴

⁴ See the section of chapter 3 of Gerber and Green (2012) on estimating clustered randomized controlled trials for a detailed explanation. As noted in that chapter, increasing the number of clusters is a better strategy than increasing the sample size of each cluster for reducing the error in the estimate of the treatment effect, except in the unlikely scenario where each cluster has very similar potential outcomes. More research is needed to determine if a simple minipublics process, such as a deliberative poll, applied to a very large stratified representative sample of the population, would satisfy the conditions of such a scenario.

It is also important to note that while aggregating the results of many minipublics reduces idiosyncratic group effects, it does not guarantee output representativeness. For example, if all our minipublics systematically excluded a segment of the population that holds crucial information for the decision-making process, the average result would not reflect the ideal population decision. To address the issue of external validity, one potential solution would be to first recruit a large, representative sample of the population and then randomly divide it into smaller subgroups. This is the strategy used in large RCTs that aim to claim external validity.

In practice, another significant issue arises when considering the integration of multiple minipublics to reduce idiosyncratic group effects. This approach is relatively straightforward for simple minipublics without an ideation phase, such as deliberative polls, where participants rank a predetermined set of options. In these cases, results can be aggregated by combining each independent ranking. However, in a slightly more complex minipublic, where participants generate their own options to rank, aggregation becomes challenging, as each group will rank a different set of options. Developing a legitimate system to aggregate these diverse options might prove difficult. And in a complex citizens' assembly design that includes a sequence of ideation and ranking processes, the difficulty of maintaining independence across minipublics, while integrating their knowledge effectively, will increase even further.

In conclusion, scholars and practitioners who justify minipublics based on their representativeness of the population face significant challenges, particularly when making such claims for small minipublics. If we only consider input representativeness, larger sample sizes are needed to avoid significant estimation errors. If we aim for output representativeness, the current generation of single minipublics cannot adequately control for group effects, and we would need to adopt a new approach that integrates multiple independent minipublics. However, such an approach might only be feasible for relatively simple problems, and we would still need a combined large sample size to achieve output representativeness.

4. Challenge III: Biases Generated by Population Lists

Another consideration is the source population for creating a random sample. In some fields, the total population is well-defined with readily available data (e.g. students in a school). However, in other cases, this becomes more complex.

Collaboration with city or national governments can provide access to some of the most comprehensive lists available. Yet even official government records, such as electoral registers, may lack information about certain groups. These might include immigrants, homeless individuals, residents of informal settlements, or non-voters. In countries with limited state capacity, information about inhabitants in certain areas may be almost non-existent. To address these challenges, geographical sampling has been employed in some cases. This method involves randomly selecting areas that are then canvassed to invite participants for surveys or other processes. Recently, a geographical method

using satellite imagery has been utilized to randomly distribute invitations for a global citizens' assembly (<https://globalassembly.org/report>).

Many existing lists contain inherent biases. For instance, landline phone directories, once considered highly accurate for population sampling, have become increasingly unreliable due to the proliferation of alternative communication devices. When drawing from crowdsourced samples, such as those provided by typical service providers, or when conducting surveys through social media advertisements, these biases become particularly pronounced and well-documented. Such samples are significantly affected by digital divides, potentially excluding, or underrepresenting, certain demographic groups.

Many of these biases can be adjusted in a survey using a variety of statistical techniques, such as sample weights based on national demographic statistics that are more accurate. These weights enable the adjustment of estimates of a proportion in the population based on observable variables that can be matched. For example, if we know that the online population we are drawing from contains more educated conservative males, we can increase the weight of the answers of other strata of the population to generate a more accurate estimate of the variable of interest. These procedures allow for the elimination of bias generated by imbalances of observable variables but do not control for bias generated by unobservable variables, such as the distribution of voluntarism attitudes in the population (Heeringa, 2017, chapter 2.7).


Moreover, while these adjustments can be applied to the results of a survey, it is not immediately clear how to apply weights to a minipublic's decision in a legitimate manner. We will return to this at the end of this section, given that weighting is a general strategy that can be applied to all sample biases.

5. Challenge IV: Refusal to Participate (Non-Compliance)

A third factor, well-known among practitioners and community organizers, is that receiving an invitation to participate does not guarantee a person's involvement in the process. Any recruitment procedure faces issues of non-compliance, which might generate self-selection bias if the reasons for refusing to participate are systematically correlated with preferences relevant to the minipublic's discussion. If the reasons for not accepting participation are random, there would not be a problem.

In complex minipublics consisting of multi-event processes with high participation costs, such as citizens' assemblies, the initial acceptance rate is often quite low, sometimes less than 5%. By initial acceptance rate, we mean the percentage of contacted individuals who express willingness to participate and are included in the pool from which the final sample will be drawn. Simpler minipublics of shorter duration (e.g. one weekend) often achieve higher engagement. To our knowledge, a complete dataset on acceptance rates for minipublics does not exist, the OECD dataset has around 50% missing values and currently is the best

resource to explore acceptance rates. In this dataset at the time of this writing the median acceptance rate out of the 343⁵ reported cases is 4%, 90% of the cases have an acceptance rate below 11%, and only 4 cases report an acceptance rate above 40%.



*One person was unable to attend the second weekend due to illness.

Figure 4. Contact and recruitment numbers of UK's Citizens Assembly on Brexit (Renwick et al., 2017).

Note that it is not clear how comparable the data are. Different recruitment mechanisms collect and report slightly different information. For example, in the UK's Citizens' Assembly on Brexit in 2017, 2,742 individuals expressed interest in principle, and then 1,155 people agreed to enter the recruitment pool out of the 5,000 contacted, generating two possible acceptance rates that could be reported, the first one of 55% and the second one of 23%.

The OECD data show a high level of refusal to participate in minipublics, highlighting the probable presence of systematic, unobservable variables that explain why some people accept participation while others do not. Unfortunately, we know little about the reasons for this refusal. In one of the few empirical studies exploring why most people refuse to participate in minipublics, Jacquet (2017) identifies six explanatory logics of non-participation: concentration on the private sphere, internal political inefficacy, public meeting avoidance, scheduling conflicts, political alienation, and the perceived lack of impact of minipublics on the political system.

This result is intuitive — it is not difficult to imagine that introverts and those who find deliberation unpleasant or useless will tend to refuse participation. Additionally, a sense of efficacy, availability of free time, and various skills might play a role. If any of these variables is systematically correlated with refusal to participate and with opinions whose absence in deliberation might affect the final decision, then the minipublics output is biased. No matter what we do to improve the sample size, this bias will always remain. That is why this is one of the most difficult problems to solve.

⁵ We removed one case that reported 367% acceptance rate, and 26 cases that reported 0% to the dataset we accessed on 10th July 2024.

Some practitioners believe that increasing monetary incentives and institutionalizing minipublics will resolve the low acceptance rate. Other practitioners and scholars argue that our perspective is based on an outdated set of minipublics cases and does not consider recent data from institutionalized minipublics in Ireland and Belgium, which show increased acceptance rates above 10%.

The possibility of increasing engagement by providing better messaging and incentives is supported by a wealth of transferable strategies investigated in the expanding literature on voter mobilization, organizing, and engagement (Green and Gerber, 2019; Peixoto, Sjoberg and Mellon, 2020). We do not know, however, how much it is feasible and if there will be other biases generated by significantly increasing incentives. Dedicated research on minipublics' incentives strategies and acceptance rate is needed to explore such empirical questions.

We also agree that impact and meaningfulness are important drivers of engagement (Sjoberg, Mellon and Peixoto, 2017). However, extensive research shows that there is not a linear positive correlation between the institutionalization of democratic innovations and their impact, nor between the frequency of their adoption and impact (Caddy, Peixoto and McNeil, 2007; Bussu *et al.*, 2022). Therefore, the idea that institutionalization will solve the problem of minimal acceptance rates is, in our opinion, overly optimistic. For instance, emerging critiques of the Irish model and the political problems generated by the mishandling of the results of the latest Irish citizens' assembly suggest a more sobering view of the Irish system. The most recent OECD dataset shows that the assembly organized in Belgium had a 3% acceptance rate, the two organized in Ireland had 10% and 7.5%, while the one in Georgia had 44% acceptance rate, and the one in Bologna in Italy 17.5% — two countries with practically no previous experience in citizens' assemblies.

A possible solution for reducing non-compliance would be to impose mandatory participation. However, even mandatory mechanisms such as juries in the US often generate significant bias toward white, older, and more educated participants (Anwar, Bayer and Hjalmarsson, 2022).

It is evident from the four challenges that the claim that a minipublic, whether based on small samples or large samples with significant non-response, represents the broader population is questionable. The only way to start justifying such a claim would be to combine multiple parallel independent minipublics drawn from a large representative sample of the population, an approach that to our knowledge has not been implemented yet and might be applicable to only some problems.

The four sources of error discussed combine to create an overall error that cannot be easily addressed by statistical methods like reweighting. This is not only due to the difficulty of reweighting but also because a minipublic operates as a democratic decision-making process, where each participant's vote holds equal weight. For reweighting participants' votes to optimize overall representativeness, a fundamental shift in what we consider legitimate decision-making would be

required. In fact, we see the opposite trend: reweighting systems, like the often-cited *Matrix Orçamentária* used in some Brazilian cities' participatory budgeting — which favoured preferences of groups traditionally excluded from government decision-making — are being phased out.

In the following sections, we will review three possible responses to the critique of insufficient representativeness of minipublics: (1) using stratification, (2) introducing supermajority voting mechanisms, and (2) leveraging second-best arguments.

6. Could We Use Stratification to Boost the Representativeness of Small Minipublics?

In a survey, stratification can enhance the precision of estimates for population characteristics, particularly when the population is heterogeneous and can be divided into internally homogeneous subgroups. This sampling method ensures that samples are drawn from each subgroup in proportion to its representation in the entire population. However, it is critical to prevent overlap between strata; everyone in the population should belong to only one stratum, ensuring mutual exclusivity. Overlapping strata could lead to biased samples by increasing the likelihood of data duplication.

It is important to note that scholars do not typically use stratification to justify reducing sample sizes; rather, they use it to minimize extreme variability and improve estimates (Groves *et al.*, 2011, chapter 5). Therefore, in this subsection, we are not responding to a problem in the literature but addressing a growing confusion among some politicians, practitioners, and journalists who conflate stratification with enhanced representativeness and even the potential to achieve intersectional representation. Statisticians and academics might find the problem we want to address non-existent but given how frequently we have observed this confusion among practitioners, we believe it deserves clarification. As we will see, the requirements for using mutually exclusive strata prevent us from making any claims about enhanced representativeness in small samples.

For example, imagine we want to stratify for five key population strata — age, income, education, geographical location, and gender — because existing secondary data shows these groups hold different opinions on the topic the assembly will deliberate. Let's simplify by assuming each of the first four categories has five equal groups in society, and gender consists of two equal groups. In this example the minimal sample required to include the intersections of all the strata in the population is equal to $5^4 \times 2 = 1,250$. Note that we have maintained the somewhat unlikely assumption that all categories have equal size. If one stratum, for example ethnicity, includes a minority that is 1/10 of the population, then our multiplier would be 10 instead of 5, requiring a sample size of $5^3 \times 10 \times 2 = 2,500$. The latter is independent of the number of categories within the strata, so even if the strata have 2 categories, one comprising 90% (9/10) of the population, and one comprising 10% (1/10) of the population, the multiplier would be 10. When we want to represent a minority of 1% (1/100) of the population the multiplier becomes 100. Note that this would be a minimal sample size to include

the intersection of all the strata of such a population, but such a small sample will not be representative of each stratum. This simple math demonstrates how difficult it is to claim that even a large sample achieves intersectional representation of the population. It also helps explain why deliberative democracy scholars and advocates struggle to address intersectional critiques (Wojciechowska, 2019) and highlights the practical challenges of using random sampling in policy areas such as education and healthcare, where intersectionality may be an important variable.

In many practical applications, stratification is not used to justify a small sample size. Instead, it is employed to promote diversity and prevent extreme outcomes of a purely random sample, such as having all participants of the same gender. This approach is supported by the collective intelligence literature, which suggests that diverse groups can make better decisions (though we must be cautious about how we define diversity and the specific problems where diversity can enhance decision-making; see Houlou-Garcia, 2017).

In the presence of small samples, often organizations providing sampling support such as the Sortition Foundation or Yougov recommend choosing 2–3 ‘hard’ targets and more ‘soft’ targets. In the lingo of practitioners that offer sampling services, a hard target requires stratification with mutually exclusive subgroups, while a soft target involves a general rebalancing of the sample without stratifying the subgroups. As explained before, the soft targets might generate bias in the final estimate. This strategy could fit a sample size of 50 participants in our previous example if we chose as hard targets age, education, and gender ($5 \times 5 \times 2$) and we select the others as soft targets, but it has no clear benefits on representativity.

There are two other important issues to discuss regarding stratification. First: stratification on observable variables does not correct for unobservable ones. And as we discussed before in the section on acceptance rates this might generate systematic biases. Second: stratification opens the discussion on who should choose the strata and what should constitute the ideal set of strata. Should we include additional factors such as political attitudes (liberal vs. conservative), opinions on the topic of the assembly, personality types, education levels, income, or previous engagement in politics? Should we include decision-making styles to achieve the diversity suggested by the collective intelligence literature?

In conclusion, stratification cannot be used to justify very small sample sizes, while it can be used to avoid extreme outcomes and promote some diversity in the assembly. But stratification comes with some costs, it might introduce biases when soft targets are used, and it might reduce the legitimacy of the process if the choice of the strata is not properly justified. Additionally, stratification might increase the vulnerability of the assembly to manipulation if the criteria are known, given that certain combinations of strata in the population have a higher probability of being chosen to participate due to their rarity in the population. The latter is something that can be compensated with more advanced stratification algorithms, but it cannot be fully eliminated (Flanigan *et al.*, 2024).

7. Compensating the Lack of Input Representativeness with Supermajority Voting

One emerging approach to overcoming minipublics' lack of input representativeness is to impose a qualified majority voting rule on the final decisions, thereby strengthening the representativeness of the output. Many best practice guides, such as the OECD's 'Catching the Deliberative Wave' report (OECD, 2020, p. 102), the United Nation's 'Democracy Beyond Elections Handbook' (New-Democracy Foundation, 2018, p. 196), Marcin Gerwin's 'Citizens' Assemblies Guide' (Gerwin, 2018, p. 22), and Mosaic Lab's 'Facilitating Deliberation: A Practical Guide' (White, Hunter and Greaves, 2022, p. 116), recommend using decision thresholds of around 80%.

While it is not immediately clear if these recommendations aim to address emerging critiques on the representativeness of the assemblies — since the texts never mention error terms nor representativeness when discussing the recommendation⁶ — we observe that this argument has recently become more common as a defence against critiques of representativeness (e.g. Redman, Spada and Peixoto, 2023). The intuition behind this approach is that if more than 80% of the participants in a minipublic prefer a certain option, even the presence of an error of around 30% would not significantly alter the legitimacy of the decision because it would not alter the fact that more than 50% of people in the general population would have chosen the same.

While this perspective is intuitive at first glance, it is not free from implications. Before describing these issues, it is important not to confuse the imposition of an 80% supermajority rule with the task of estimating a parameter in the population that is 80% (for example, 80% of the population is in favour of introducing gender quotas in local elections). If we were in the latter situation and we asked a minipublic to decide whether to adopt or reject gender quotas in local elections, then the error terms generated by a small sample would be smaller because most people in the population would have the same preferences. But in most complex minipublics that include a co-design phase, we have very limited knowledge of the distribution of the population's preferences regarding the pro-

⁶ For example, this is the justification provided by the OECD report at page 102: 'A key difference between representative deliberative processes and other forms of citizen participation is that the outcome is not many individual views, but a collective and considered view. Citizens are tasked with finding consensus on the recommendations they provide to public decision makers. This does not mean that 100% of participants must agree with 100% of the proposals. This is highly unlikely and is arguably not desirable in a democracy that values pluralism. A common rule of thumb is that around 80% of the participants must agree that they would be fine with the recommendation. Sometimes the report with citizens' recommendations includes a minority report, where participants can include the proposals that garnered some support, but not enough to be accepted by the majority of the group.' While the justification provided by the Mosaic Lab guide is the following: 'In Deliberation a 50% level of support is usually not considered sufficient because it is important that most participants stand behind the group's recommendations. If a decision is passed by a 51% majority, 49% of participants will be left feeling dissatisfied and are likely to protest the decision.' Both definitions do not contain any hint about using supermajority rules to ensure representativeness of the assembly and instead derive the recommendation from deliberation or recommend it to avoid protest from the minority.

posals created by the assembly, and we cannot make any reasonable assumption over the population's preferences.

The first issue with requiring a supermajority is that prioritizing consensus can stifle critical dialogue, favouring shallow agreements and status-quo-friendly results. Status quo bias is particularly problematic and has been extensively studied (McGann, 2004). The literature on deliberative democracy has discussed the pros and cons of consensus for around thirty years. A recent special issue of the *Journal of Deliberative Democracy* offers an insightful overview of this issue (Friberg-Fernros, Schaffer and Holst, 2019).

The second issue is that supermajority rules might negatively impact participant engagement, particularly among minority groups or those with low political efficacy. Extensive literature dating back to the 1980s analyses the relationship between voting mechanisms and turnout, highlighting how mechanisms that restrict the likelihood of participants' preferences being represented reduce turnout, particularly among women (Skorge, 2023). Although this literature primarily contrasts plurality and proportional representation voting mechanisms, additional literature on deliberative settings (e.g. Lupien, 2018) warrants caution regarding the use of qualified majorities.

Third, imposing supermajority rules reduces the probability of change because almost everyone must agree for any proposal to be passed. This implies that more radical changes, which might be preferred by most of the general population, are not proposed by the assembly, which is forced to adopt only the least controversial issues. In other words, the capacity of the assembly to empower substantial changes is stifled. This effect is not that surprising: most constitutional democracies use qualified majority rules precisely to lock constitutions and prevent change.

While empirical evidence from assembly case studies demonstrating this reduction in radical demands due to the supermajority rule is not available — and we doubt it will ever be due to publication bias (Spada and Ryan, 2017) — some practitioners and some social movements that support citizens' assemblies are starting to realize that the citizens' assembly capacity to promote a radical agenda is unclear and possibly decreasing (XR Southwark, 2022). To be fair, in our opinion, the adoption of supermajority rules is only one factor that is contributing to this trend. This trend is also due to the emerging mainstream model of mini-publics, which are becoming more and more commodified products offered by various service providers to local authorities. Wilson and Mellier (2023) echo these emerging concerns, calling for a shift of perspective toward empowerment of citizens:

The current wave of deliberative processes focuses on helping governments increase their legitimacy and pays too little attention to helping citizens achieve the change they want to see. These often top-down invited spaces end up depoliticizing citizen engagement at best, and at worst they provide legitimacy for governments that are in effect opposed to the changes citizens want. (Wilson and Mellier, 2023)

Although anecdotal, this resonates with our experience with a European mini-public managed by a well-known service provider, a standardized, cost-effective

design including a low-budget sampling strategy, yielding an acceptance rate of less than 2%. The deliberation period was limited to two weekends and one online event, briefing packages were not included, and supermajority rules implemented. This cost-cutting approach yielded recommendations that aligned comfortably with existing government agendas and that reflected experts' interventions without adding any innovation. If the empirical evidence suggests that minipublics have been mostly inconsequential so far, with governments cherry-picking proposals (Giraudet *et al.*, 2022), a plausible hypothesis is that the recent popularity of minipublics also stems from this rather 'safe', commodified, low-cost adoption model. A similar trajectory is seen in Participatory Budgeting, which went mainstream in the mid-2000s and is now used in over 10,000 local governments. However, critics argue that its initial capacity to deepen democracy and drive reform has diminished (Baiocchi and Ganuza, 2014; Ganuza and Baiocchi, 2012; 2019).

Fourth, a minipublic is designed with the stated objective of allowing participants to change opinions when necessary. Therefore, the representativeness of the assembly cannot be derived from the qualified majority imposed on the final voting mechanism because the initial composition of the group might influence how they will vote. One of the core principles of minipublics is that participants deliberate, and their opinions may change because of the group process and composition. In other words, deliberation quality is a function of the representativeness of the participants. Thus, if we do not have sufficient input representativeness, the key process of transforming preferences, which is the basis of many positive claims in the deliberative and collective intelligence literature, is flawed. Such a flaw cannot be solved by imposing a supermajority rule at the end. For example, consider a minipublic tasked with choosing between options A and B. Suppose there is a sampling issue that excludes a segment of the population who would have contributed new arguments and perspectives (Zi), potentially persuading the entire group to favour option B. Without these perspectives, the group overwhelmingly supports option A, and the imposition of a supermajority decision rule does not address this underlying problem.

8. Second Best Arguments: Are Minipublics More Inclusive than Existing Alternatives?

Thus far, we have provided evidence that the claim of minipublics as representative of the population is problematic. But what about more cautious claims, such as minipublics being more inclusive than other participatory processes (e.g. participatory budgeting, e-petitions) that do not employ randomization? By inclusion, we refer to the practice or policy of providing equal access to opportunities and resources to participate effectively for people who might otherwise be excluded or marginalized.

Many would agree that traditional forms of consultation tend to attract 'usual suspects' — citizens who have a higher interest in politics, more spare time, higher education, enjoy talking in public, and sometimes enjoy any opportunity to criticize. In the US, for instance, these citizens are often older white males, or as

put by a practitioner once, ‘the male, pale and stale’. A typical minipublic instead manages to engage a more diverse set of participants than traditional consultations.

While this is an obvious reality, the engagement strategies of minipublics compared to traditional consultations based on self-selection have very different levels of sophistication and costs. Minipublics tend to invest more resources in engagement, sometimes tens of thousands of dollars, and thus it is not fair to compare traditional consultations to minipublics to exclude alternative methods of engagement (as it is not fair to compare minipublics that are not specifically designed to include marginalized populations to processes that are specifically designed for this purpose). This is because the amount dedicated to engagement is positively correlated with inclusion. For instance, processes specifically designed to include immigrants, youth, and native populations will be — if correctly executed — more successful at including such strata of the population than a general random selection strategy that does not have specific quotas and engagement strategies for these groups. We talk past one another when we try to rank processes with respect to their supposed inclusion performance without considering the impact of the resources dedicated to engagement or their intended effects (e.g. promoting diversity vs. representativeness).

Without a significant amount of research comparing different participatory methods with similar outreach and resources it is difficult to determine which approach is more inclusive. To our knowledge, the only study that compares two similar processes — one using random engagement and the other using an open-to-all invitation — found little difference in inclusiveness (Griffin *et al.*, 2015). The study also highlighted the importance of other factors such as the design of the process, potential political impact, and the topic of discussion. Many practitioners do not take these factors into account, and instead focus solely on recruitment strategies. While one study is not enough to make a conclusive judgment, it does suggest that the assumption that minipublics using randomly selected participants are automatically more inclusive than open-to-all processes is problematic. Even more so given that, in 90% of recorded cases in the OECD 2023 database, around 90% of the population refuses to enter the selection pool, it is very difficult to make any claim of greater inclusiveness. In these situations, the randomization occurs within a new type of usual suspects: the 10% of unusual people who accept the invitation. Even though one might try to adjust the results and use a lottery to randomly select a group of participants that is balanced across some demographic and attitudinal variables of interest, the probability that all the participants have a set of unique features that make them accept the invitation is very high. And if such unobservable common characteristics are correlated with preferences that matter for the deliberation, one is merely creating a facade of inclusiveness.

9. Concluding Remarks: Three Ways Forward

We hope that it is now clearer why we contend that it is very problematic to defend the application of minipublics based on their claimed ability to represent the population. So, should we abandon minipublics?

Our position is that we should not, and we argue that there seems to be three productive directions for minipublics: (1) scaling-up and integrating multiple independent minipublics, (2) solving previously detected failure of inclusion, and (3) leveraging non-domination claims.

First, we believe that there is a case for a novel minipublics design that integrates multiple smaller minipublics in a larger decision-making system because this approach would allow to better control for idiosyncratic group effects and approximate the ideal of output representativeness. But this approach is very costly and, so far, we have seen only a few attempts at large-scale deliberation (e.g. G1000), and to our knowledge no design has yet tried to integrate multiple independent minipublics. With sizes approaching 1,000 people, hundreds of moderators are normally required, and much of the exchange of information will occur not through synchronous exchanges in small groups, but through asynchronous transmission mechanisms across the groups. This is not necessarily a bad thing, but it will have the typical limitations of any type of aggregation mechanism that requires participant attention and effort. For example, in an ideation process with 100 groups of 10 people each, where each group proposes one idea and then discusses all other ideas, each group would have to discuss 100 ideas. This is a very intense task.

All else being equal, as the size of the assembly grows, and as the number of assemblies we need to integrate multiply, the logistical complexity and associated costs increases. At the same time, the ability to analyse and integrate all the information generated by participants diminishes. The question of whether artificial intelligence (AI) could help overcome the challenges associated with mass deliberation is an empirical one — and recent research shows promising avenues. For instance, recent research suggests AI-supported mediation can outperform human mediators in large-scale settings, potentially lowering the barriers to the implementation of larger minipublics (Tessler *et al.*, 2024). Still, AI solutions are — to date — unable to address the issue of non-compliance. Whether AI could help overcome non-compliance challenges to efficiently capture, predict, and aggregate individual preferences through mechanisms such as ‘agents for augmented democracy’ (e.g. Gudino-Rosero *et al.*, 2024) is a question that deserves both empirical and normative research.

From a less technological standpoint, recent designs of permanent minipublics such as the one adopted in Belgium (Ostbelgien, Brussels) that resample a small new group of participants every year could attempt to include over time a sufficiently large sample of the population to achieve a good level of representation, at least for some strata of the population, and if systematic sampling errors are corrected. In such a case, however, it might be better to justify the system of repeated minipublics with the idea of offering over time an equal chance to

participate to all members of the targeted community because each minipublic will not be representative in itself.

Second, another approach is to abandon the idea of achieving representativeness and instead target specific problems of inclusion. This represents a small shift from the current approach to minipublics but, in our opinion, could yield significant long-term legitimacy. Rather than justifying a minipublic through a broad claim of representation, this model would justify itself by addressing specific failures in inclusion. For example, consider neighbourhood-level urban planning meetings in a city that consistently fails to involve renters while disproportionately engaging developers, homeowners, and business owners. In such cases, a stratified random sample approach could be adopted, which reserves quotas for renters and includes specific incentives to attract them, excluding other participant types. This approach aims to prevent domination and ensure fair representation. However, it's important to note that this approach is only practical after clear inclusion failures have been identified. Mass LBP, a Canadian organization supporting various participatory and deliberative processes, implemented this design in the Grandview Woodland Citizens Assembly in 2014 following a previous participatory urban planning process organized by the City of Vancouver that had failed to recruit renters (Participedia case 4228).

Third, another approach is to completely abandon the claim of representativeness in favour of a weaker non-domination claim. Sortition-based recruitment systems offer a fair way to exclude certain groups from the minipublics. This is particularly so because, in certain cases, participatory mechanisms based on self-selection may be more easily captured by organized minorities to the detriment of disengaged majorities, making these processes less inclusive.

Consider, for instance, one of President Obama's first attempts to engage citizens at large scale, the White House's online town-hall. Through a platform named 'open for questions', citizens were able to submit questions to Obama and vote for which questions they would like to be answered by him. Over 92,000 people posted questions, and about 3.6 million votes were cast for and against those questions. Under the section 'budget' of the questions, seven of the ten most popular queries were about legalizing marijuana, many of which were about taxing it. The popularity of this issue was attributed to a campaign led by NORML, an organization advocating pot legalization. While the cause and ideas may be laudable, it is fair to assume that this was hardly the biggest budgetary concern of Americans in the aftermath of an economic downturn.

In a case like the White House's town-hall, the randomization of participants would be a fair and effective way to avoid the capture of the dialogue by organized groups. Randomization does not completely exclude the possibility of capture of a deliberative space, but it does increase the costs of doing so. The probability that members of an organized minority are randomly sampled to participate in a minipublic is minor, therefore the odds of their presence in the minipublic will be reduced. Thus, even if we had a technological solution capable of organizing large-scale deliberation in the millions, a randomization strategy could still be an effective means to protect deliberation from the capture by organized minorities. A legitimate method of exclusion will remain an asset — at

least until we have another legitimate way to mitigate the ability of small, organized minorities to bias deliberation.

When comparing these three ways forward, the first involves overcoming significant technical challenges and increased costs. Moreover, while we recognize the appeal of the idea of creating large-scale, integrated systems of minipublics, we approach it with caution. This perspective might lead to a practical implementation of deliberative democracy that, though intellectually engaging, risks assuming a somewhat solutionist and paternalistic stance, that can lead to a depoliticization of minipublics. As noted by Lafont and Urbinati (2024), this approach aligns with a technopopulist ideology, treating deliberation as a universal mechanism that neutralizes individual agency and political dynamics. Consequently, the political diversity and struggles of citizens become less relevant, as any representative sample, when subjected to this deliberative ‘treatment’, is expected to converge on the same decision (on average).

Thus, addressing specific cases of exclusion and preventing domination by organized minorities may be a more practical and desirable approach. The second and third approaches might not seem very appealing at first. But one should not be discouraged by our unglamorous example of fixing urban planning meetings or preventing interest groups from dominating a public agenda. In fact, these approaches are particularly attractive given that inclusion failures and domination attempts can be found across multiple spaces meant to be democratic — from neighbourhood meetings to parliaments around the globe. Moreover, they require a very limited change in the current set of procedures implemented in minipublics.

For minipublics practitioners and advocates like us, this should come as a comfort: there’s no shortage of work to be done. But we might be more successful if, in the meantime, we shift the focus away from the representative claim.

References

- Anwar, S., Bayer, P. & Hjalmarsson, R. (2022) Unequal jury representation and its consequences, *American Economic Review: Insights*, 4 (2), pp. 159–174.
- Baiocchi, G. & Ganuza, E. (2014) Participatory budgeting as if emancipation mattered, *Politics & Society*, 42 (1), pp. 29–50.
- Bethlehem, J. (2009) *Applied Survey Methods: A Statistical Perspective*, Hoboken, NJ: John Wiley & Sons.
- Bussu, S., Bua, A., Dean, R. & Smith, G. (2022) Introduction to symposium on embedding participatory governance, *Critical Policy Studies*, 16 (2).
- Caddy, J., Peixoto, T. & McNeil, M. (2007) *Beyond Public Scrutiny: Stocktaking of Social Accountability in OECD Countries*, Washington, DC: OECD.
- Carroll, R. (2024) Irish referendum fiasco puts future of lauded citizens’ assemblies in doubt, *Guardian*, [Online], <https://www.theguardian.com/world/2024/mar/20/irish-referendum-fiasco-puts-future-of-lauded-citizens-assemblies-in-doubt>.
- Curato, N., Farrell, D. & Geissel, B. (2021) *Deliberative Mini-Publics*, Bristol: Policy Press.
- Dahl, R.A. (1970) *After the Revolution? Authority in a Good Society*, Yale University Press.
- Dahl, R.A. (1989) *Democracy and Its Critics*, New Haven, CT: Yale University Press.
- Druckman, J.N. (2004) Political preference formation: Competition, deliberation, and the (ir)relevance of framing effects, *American Political Science Review*, 98 (4), pp. 671–686.
- Fishkin, J. (2020) Cristina Lafont’s challenge to deliberative minipublics, *Journal of Deliberative Democracy*, 16 (2), pp. 56–62. doi: doi.org/10.16997/jdd.394

- Flanigan, B., Liang, J., Procaccia, A.D. & Wang, S. (2024) Manipulation-robust selection of citizens' assemblies, in *Proceedings of the AAAI Conference on Artificial Intelligence*, **38** (9), pp. 9696–9703.
- Friberg-Fernros, H., Schaffer, J. & Holst, C. (2019) Deliberation after consensus: Introduction to the symposium, *Journal of Public Deliberation*, **15** (1). doi: [10.16997/jdd.325](https://doi.org/10.16997/jdd.325)
- Ganuza, E. & Baiocchi, G. (2012) The power of ambiguity: How participatory budgeting travels the globe, *Journal of Deliberative Democracy*, **8** (2).
- Ganuza, E. & Baiocchi, G. (2019) The long journey of participatory budgeting, in *Handbook of Democratic Innovation and Governance*, pp. 77–89, Cheltenham: Edward Elgar Publishing.
- Gerber, A.S. & Green, D.P. (2012) *Field Experiments — Design, Analysis, and Interpretation*, New York: W.W. Norton & Co.
- Gerwin, M. (2018) *Citizens' Assemblies Guide*, <https://citizensassemblies.org/download>
- Germann, M. (2024) Minipublics, (lack of) representativeness, and legitimacy beliefs, *British Journal of Political Science*, forthcoming.
- Giraudet, L.G., Apouey, B., Arab, H., Baeckelandt, S., Begout, P., Berghmans, N., Blanc, N., Boulou, J.Y., Buge, E., Courant, D. & Dahan, A. (2022) 'Co-construction' in deliberative democracy: Lessons from the French Citizens' Convention for Climate, *Humanities and Social Sciences Communications*, **9** (1), pp. 1–16.
- Goodin, R.E. & Dryzek, J.S. (2006) Deliberative impacts: The macro-political uptake of mini-publics, *Politics and Society*, **34** (2).
- Green, D.P. & Gerber, A.S. (2019) *Get Out the Vote: How to Increase Voter Turnout*, Washington, DC: Brookings Institution Press.
- Griffin, J., Abdel-Monem, T., Tomkins, A., Richardson, A. & Jorgensen S. (2015) Understanding participant representativeness in deliberative events: A case study comparing probability and non-probability recruitment strategies, *Journal of Public Deliberation*, **11** (1). doi: [10.16997/jdd.221](https://doi.org/10.16997/jdd.221)
- Groves, R.M., Fowler Jr, F.J., Couper, M.P., Lepkowski, J.M., Singer, E. & Tourangeau, R. (2011) *Survey Methodology*, Hoboken, NJ: John Wiley & Sons.
- Guardian (2023) Citizens' assemblies: Are they the future of democracy?, *Guardian*, [Online], <https://www.theguardian.com/us-news/2023/feb/01/citizens-assemblies-are-they-the-future-of-democracy>.
- Gudiño-Rosero, J., Grandi, U. & Hidalgo, C.A. (2024) Large Language Models (LLMs) as agents for augmented democracy, *arXiv*, preprint, arXiv:2405.03452.
- Heeringa, S.G., West, B.T. & Berglund, P.A. (2017) *Applied Survey Data Analysis*, London: Chapman and Hall/CRC.
- Houlou-Garcia, A. (2017) Collective wisdom, diversity and misuse of mathematics, *Revue française de science politique*, **67** (5), pp. 899–917.
- Jacquet, V. (2017) Explaining non-participation in deliberative minipublics, *European Journal of Political Research*, **56**, pp. 640–659. doi: [10.1111/1475-6765.12195](https://doi.org/10.1111/1475-6765.12195)
- Lafont, C. (2015) Deliberation, participation, and democratic legitimacy: Should deliberative mini-publics shape public policy?, *Journal of Political Philosophy*, **23** (1), pp. 40–63.
- Lafont, C. (2020) *Democracy Without Shortcuts*, Oxford: Oxford University Press.
- Lafont, C. & Urbinati, N. (2024) Defending democracy against lottocracy, in Grandjean, G. (ed.) *Against Sortition? The Problem with Citizens' Assemblies*, Exeter: Imprint Academic.
- Landemore, H. & Page, S.E. (2015) Deliberation and disagreement: Problem solving, prediction, and positive dissensus, *Politics, Philosophy & Economics*, **14** (3), pp. 229–254.
- Lupien, P. (2018) Participatory democracy and ethnic minorities: Opening inclusive new spaces or reproducing inequalities?, *Democratization*, **25** (7), pp. 1251–1269.
- McGann, A.J. (2004) The tyranny of the supermajority: How majority rule protects minorities, *Journal of Theoretical Politics*, **16** (1), pp. 53–77.
- Mercier, H. & Landemore, H. (2012) Reasoning is for arguing: Understanding the successes and failures of deliberation, *Political Psychology*, **33** (2), pp. 243–258.
- Morrell, M.E. (2014) Participant bias and success in deliberative mini-publics, *Deliberative Mini-publics: Involving Citizens in the Democratic Process*, pp. 157–176.

- Muradova, L. (2020) Seeing the other side? Perspective-taking and reflective political judgments in interpersonal deliberation, *Political Studies*, **69** (3), pp. 644–664. doi: 10.1177/0032321720916605
- New-Democracy Foundation (2018) *Enabling National Initiatives to Take Democracy beyond Elections*, [Online], <https://www.un.org/democracyfund/sites/www.un.org.democracyfund/files/newdemocracy-undef-handbook.pdf> [Accessed 10 July 2024].
- OECD (2020) *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*, Paris: OECD Publishing. doi: 10.1787/339306da-en
- OECD (2023) *OECD Deliberative Democracy Database*, [Online], <https://airtable.com/appP4czQLAU1My2M3/shrX048tmQLl8yzdc/tblm3C6n7vM6vPSCz/viw2EEqp8wOhaHACK/recB8dH9yxghQIsoX> [Accessed July 2024].
- Participedia case 4228, [Online], <https://participedia.net/case/4228>.
- Peixoto, T., Sjöberg, F.M. & Mellon, J. (2020) A get-out-the-vote experiment on the world's largest participatory budgeting vote in Brazil, *British Journal of Political Science*, **50** (1), pp. 381–389.
- Pogrebinschi, T. (2023) *Innovating Democracy? The Means and Ends of Citizen Participation in Latin America*, Cambridge: Cambridge University Press.
- Redman, K., Spada, P. & Peixoto, T. (2023) How representative is it really? A correspondence on sortition, *Deliberative Democracy Digest*, [Online], <https://www.publicdeliberation.net/how-representative-is-it-really-a-correspondence-on-sortition/> [Accessed on 10 July 2024].
- Renwick, A., Allan, S., Jennings, W., McKee, R., Russell, M. & Smith, G. (2017) *The Report of the Citizens' Assembly on Brexit*, London: UCL Constitution Unit.
- Richards, R., Morrell, M.E., Brinker, D. & Reedy, J. (2022) Psychological phenomena in democratic deliberation, *Journal of Deliberative Democracy*, **18** (2). doi: 10.16997/jdd.1277
- Saward, M. (2021) *Democratic Design*, Oxford: Oxford University Press.
- Sjöberg, F.M., Mellon, J. & Peixoto, T. (2017) The effect of bureaucratic responsiveness on citizen participation, *Public Administrative Review*, **77** (3), [Online], <https://www.oidp.net/docs/repo/doc381.pdf>.
- Skorge, Ø.S. (2023) Mobilizing the underrepresented: Electoral systems and gender inequality in political participation, *American Journal of Political Science*, **67** (3), pp. 538–552.
- Smith, G. (2009) *Democratic Innovations: Designing Institutions for Citizen Participation*, Cambridge: Cambridge University Press.
- Spada, P. & Ryan, M. (2017) The failure to examine failures in democratic innovations, *PS: Political Science and Politics*, **50** (3).
- Sulik, J., Bahrami, B. & Deroy, O. (2022) The diversity gap: When diversity matters for knowledge, *Perspectives on Psychological Science*, **17** (3), pp. 752–767.
- Sunstein, C.R. (2017) Deliberative trouble? Why groups go to extremes, in *Multi-Party Dispute Resolution, Democracy and Decision-Making*, pp. 65–95, London: Routledge.
- Tessler, M.H., Bakker, M.A., Jarrett, D., Sheahan, H., Chadwick, M.J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D.C. & Botvinick, M. (2024) AI can help humans find common ground in democratic deliberation, *Science*, **386** (6719), eadq2852.
- UN (n.d.) *UN Democracy Beyond Elections Handbook*, [Online], <https://www.un.org/democracyfund/sites/www.un.org.democracyfund/files/newdemocracy-undef-handbook.pdf>.
- White, K., Hunter, N. & Greaves, K. (2022) *Facilitating Deliberation — A Practical Guide*, Victoria, AU: Mosaic Lab.
- Wilson, R. & Mellier, C. (2023) *Getting Real About Citizens' Assemblies: A New Theory of Change for Citizens' Assemblies*, [Online], <https://europeandemocracyhub.epd.eu/getting-real-about-citizens-assemblies-a-new-theory-of-change-for-citizens-assemblies/>.
- Wojciechowska, M. (2019) Towards intersectional democratic innovations, *Political Studies*, **67** (4), pp. 895–911.
- XR Southwark (2022) *Citizen's Assemblies: Reflecting On Our Third Demand From First Hand Experience*, [Online], <https://xrsouthwark.earth/citizens-assemblies-how-participating-in-one-changed-our-perspective> [Accessed 10 July 2024].