POLICY BRIEFS SERIES

Mini-Public Selection: Ask What Randomness Can Do for You

Bailey Flanigan, Paul Gölz, Ariel Procaccia

NOVEMBER 2023

HARVARD Kennedy School
ASH CENTER
for Democratic Governance
and Innovation

POLICY BRIEFS SERIES

Mini-Public Selection: Ask What Randomness Can Do for You

Bailey Flanigan, Paul Gölz, Ariel Procaccia

NOVEMBER 2023

About the Authors

Bailey Flanigan is a PhD student in the Carnegie Mellon University Department of Computer Science. Her research is interdisciplinary, combining tools from theoretical computer science and political science to design more equitable and effective democratic decision processes. One major focus of her research is deliberative processes: she has worked on fairer and more transparent algorithms for selecting the participants of citizens' assemblies, and she is now studying how the conditions promoted by deliberation can improve democratic outcomes. Outside of research, Bailey is passionate about equitable education: she led the development and teaching of CS-JEDI, an introductory DEI course for first-year CS PhD students at CMU.

Paul Gölz is a postdoctoral researcher at the School of Engineering and Applied Sciences at Harvard. He received his PhD in computer science from Carnegie Mellon University under the supervision of Ariel Procaccia. Paul studies democratic decision-making and the fair allocation of resources, using tools from algorithms, optimization, and artificial intelligence. Algorithms developed in his work are now deployed to select citizens' assemblies around the world and to allocate refugees for a major US resettlement agency. In 2024, Paul will join Cornell as assistant professor in Operations Research and Information Engineering.

Ariel Procaccia is Gordon McKay Professor of Computer Science at Harvard University. He works on a broad and dynamic set of problems related to AI, algorithms, economics, and society. He has helped create systems and platforms that are widely used to solve everyday fair division problems, resettle refugees, mitigate bias in peer review and select citizens' assemblies. To make his research accessible to the public, he regularly writes opinion and exposition pieces for publications such as the *Washington Post*, *Bloomberg*, *Wired* and *Scientific American*.

About the Ash Center

The Mission of the Roy and Lila Ash Center for Democratic Governance and Innovation at is to develop ideas and foster practices for equal and inclusive, multi-racial and multi-ethnic democracy and self-government.

This paper is one in a series published by the Ash Center for Democratic Governance and Innovation at Harvard University's John F. Kennedy School of Government. The views expressed in the Ash Center Policy Briefs Series are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. The papers in this series are intended to elicit feedback and to encourage debate on important public policy challenges.

A PUBLICATION OF THE

Ash Center for Democratic Governance and Innovation
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138

617-495-0557
ash.harvard.edu

Contents

Abstract

Deliberative mini-publics convene a panel of randomly selected constituents to deeply engage with complex policy issues. This essay considers the process of selecting the members of this panel, with a particular focus on the central role of *randomness* in this process. Among other benefits, this randomness confers legitimacy by affording all members of society some chance of participation. The main constraint placed on this selection process is *representativeness*—the requirement that the resulting panel accurately reflects the demographic and ideological makeup of the population. Within this constraint, there is a vast space of possible ways to randomly choose a panel, and each method spreads the random chance of selection differently across willing participants. In this essay, we explore approaches for intentionally designing this randomness to promote specific goals, such as increasing fairness, transparency, non-manipulability, and richness of representation in mini-publics.

1. Introduction

Suppose you want to elicit public opinion on a controversial issue in order to inform an upcoming political decision. For example, you may want feedback on regulating interaction on online platforms, or whether people have the "right to die."[1] You could run a simple poll, but you don't want to survey the kind of uninformed attitudes that polls are known to elicit.[2] Given the topic's sensitivity, complexity, and polarization, you want respondents to rigorously consider the issue. To facilitate such deeper consideration, you decide to run a *deliberative mini-public*, where participants will convene for multiple days. During this time, they will hear expert testimony and deliberate with each other before weighing in on the issue, e.g., through a poll or by collectively producing a policy recommendation. Ample evidence indicates that this process enables "high-quality deliberation," helps "overcome polarization," and produces proposals rooted in "considered judgment."[3] Because this process is involved and participants are to be compensated, you decide to limit the panel to, say, 100 participants.

The question then becomes: how should you choose the participants of the mini-public? This is an important decision for many reasons, perhaps most notably that it must convince decision-makers and the public that the mini-public speaks for the community as a whole. This property, known as *legitimacy*, is particularly challenging to achieve because opinions expressed by mini-public participants after the deliberation process may differ from those of the general population.[4] This divergence is natural, as participants' opinions have been filtered through an intensive process of learning, questioning, and hearing others' perspectives. It then depends in large part on the legitimacy of the *selection process* to determine whether the mini-public's output is perceived as "what *the people* would think under good conditions,"[5] instead of just another voice in the crowd of opinion groups.

A central tenet of a *legitimate* selection process is to always produce a panel that is *representative* of key subdivisions of society. These subdivisions are usually defined by demographic or ideological categories like race/ethnicity, gender, geographic region, education level, and political leaning (as we will discuss in more detail later). Among other reasons, enforcing representation is important for the inclusion of underserved stakeholders. As a negative example, suppose you selected participants for a panel on artificial intelligence regulation via an application process that filtered for individuals with knowledge or expertise in technological tools. The resulting panel might skew toward young, tech-savvy, and educated participants, potentially underrepresenting the perspectives of older people, those with less access to technology, and those with less formal education—all of whom have an important stake in the issue.

Even when representation is achieved, there remains a significant threat to legitimacy: public suspicion about *why* the particular mini-public participants were chosen. Are they truly everyday people, or were they hand-picked behind the scenes? This threat is of real concern: between 2017 and 2018, a rogue employee recruited seven participants for the Irish Citizens' Assembly among his personal connections rather than following procedures.[6] This prompted activists opposed to the mini-public's recommendations to raise questions about whether the assembly was biased toward or against certain outcomes.[7]

To avoid such concerns about organizers "stacking the panel," mini-public selection methods rely heavily on *randomness*, leaving who participates up to chance rather than filters or personal connections. For instance, suppose you selected mini-public participants via a simple lottery, giving each constituent the same chance of being selected. If this lottery-based process were feasible, it would achieve the conditions for legitimacy we have laid out, producing a representative panel[8] while clearly demonstrating that no individual is preferred over another.[9] Unlike in the selection of a U.S. jury, however, you cannot compel those selected by lottery to participate in a mini-public. As a result, organizers typically contact thousands of constituents at random and invite them to *opt into* the mini-public selection process.[10] Then, panel members are chosen from among those who opted in.[11]

It is tempting to execute this final selection step via a simple lottery among those who opted in. Unfortunately, doing so almost always proves unsatisfactory because of *self-selection bias*—i.e., different groups in the population opt in at different rates. When there is self-selection bias, a simple lottery would yield a panel whose makeup reproduces the demographic skew among those who opted in. As an example, consider an assembly on climate change: if those least concerned about climate change are less likely to volunteer their time to discuss it, a simple lottery would result in a panel that grossly underrepresents this perspective.

With a simple lottery off the table, how should we randomly choose our representative panel? We call this task the *panel selection problem*: given a pool of volunteers whose demographic composition may differ arbitrarily from the underlying population, we must choose a pre-specified number of panelists (e.g., 100) whose demographic composition closely resembles that of the underlying population. We call any fixed procedure that solves this task, whether executed by a computer or person, a *selection algorithm*. Even within the restriction that the panel is representative, there is a massive design space of selection algorithms, each of which uses randomness differently.

In this essay, we focus on defining and exploring the design space of this randomness, aiming to start a conversation about what properties of randomness can most effectively promote legitimacy. This conversation will require input from political scientists, practitioners, and mathematicians. For example, we will explore questions around how to quantify conceptual normative properties of this randomness, how to trade off these desiderata in practice (and which trade-offs are mathematically feasible), and which properties can be achieved by efficient algorithms. Ultimately, we hope that the outcomes of this conversation can encourage the development of new selection algorithms and guide practitioners in choosing the selection algorithm best suited to their mini-public context.

To launch this conversation, Section 2 first defines the design space of randomness in a selection algorithm; that is, it addresses the question of what randomness remains *available* after representation constraints are imposed. In Section 3, we brainstorm a preliminary list of several potential benefits that can result from carefully designed randomness. Then, in Section 4, we describe selection algorithms that we, along with other co-authors, have developed, which have been deployed in practice. These algorithms' defining feature is that they satisfy representation constraints while enabling *optimal* design of the randomness with respect to certain goals.[12] Since many goals still remain out of reach algorithmically—and there are undoubtedly more goals to be proposed—we conclude the essay by laying out challenges and proposing opportunities for future work.

2. The Design Space of Randomness

2.1 Representativeness Limits the Ways We Can Design Randomness

To delineate the design space of randomness, we now describe the constraints that restrict how we can use it. These constraints are specified by the mini-public's organizer and enforce representation requirements through *quotas* on *stratified attributes*.[13] It's worth noting that while there are other methods to guarantee representation, quotas are widely used in practice, and representation requirements of other forms would restrict the randomness in similar ways.

The panel selection problem begins with our given group of willing participants, called a *pool*, whose composition may differ arbitrarily from the underlying population. For each member of the pool, we have information about their *stratified attributes:* a pre-specified set of characteristics that we want to ensure are properly represented. These attributes can be based on demographics (e.g., age, gender, socioeconomic status, location of residence, etc.) or attitudes (e.g., political leaning or opinions on a policy question). We can also define stratified groups by *intersections* of attributes, ensuring representation for different subsets, like urban Democrats, rural Democrats, urban Republicans, rural Republicans, and so on.

To make sure our panel reflects the population on these margins, we will require the panel to satisfy hard quotas. For example, if 18 percent of constituents are "not very concerned" about climate change, we might require our panel of 100 to contain between 16 and 20 panelists who enter the deliberation with this attitude. Typically, mini-public organizers set quotas for multiple overlapping categories of stratified attributes (e.g., age, gender, and location), so each volunteer possesses one stratified attribute per category.

In the introduction, we already informally discussed how these quotas restrict our usage of randomness by preventing us from using a simple lottery over the pool members. To illustrate this point within our 100-person panel example, suppose that only 5 percent of our pool is "not very concerned" about climate change. If we run a simple lottery, the resulting panel will contain approximately five people with such opinions, grossly violating our lower quota of 16. In other words, with a simple lottery, we would likely choose a panel that does not satisfy our quotas. This example underscores how quotas fundamentally limit the ways we can randomize: the random process *must not, with any probability, produce a panel that does not satisfy the quotas.* As we will see next, this idea allows us to succinctly summarize *all possible* ways to design the randomness.

2.2 Capturing Randomness with Panel Distributions

Continuing our example, suppose our pool contains 1,000 participants. Then, the number of possible *panels*—i.e., groups of 100 pool members—is astonishingly large: "1,000 choose 100," a number with 140 digits. Unintimidated by this abundance, we can describe the behavior of *any* given selection algorithm on this pool by specifying the probability (between 0 and 1) that the algorithm outputs each possible panel. If the algorithm is *random*, it will distribute the total amount of probability, 1, across multiple different panels; if it is not, it will choose a single panel with probability 1 and all other panels with probability 0. In any case, these probabilities amount to the algorithm's *panel distribution*: the probability the algorithm places on each possible panel.

Using the idea of a panel distribution, we can exactly characterize what randomness is available after imposing the quotas. We illustrate this in Figure 1, below, which shows two different panel distributions, corresponding to two different selection algorithms. Each panel distribution specifies the probabilities of the corresponding algorithm outputting each of the seven possible panels (an unrealistically small number of panels, for the sake of simplicity). Among these panels, only some satisfy all the quotas; we call these the *representative panels.*

Of the two selection algorithms pictured in Figure 1, the one represented by the panel distribution on the left places some probability on panels that are not representative (red). This means there is some chance that this algorithm will output a panel that fails to satisfy the quotas. By contrast, the algorithm on the right places probability only on representative panels, meaning it is *guaranteed* to output a panel that satisfies the quotas.

Figure 1. The panel distributions of two different selection algorithms. The horizontal axis lists all possible panels we can select from the pool, where black labels indicate representative panels and red labels indicate nonrepresentative ones. The vertical axis represents the probability each algorithm chooses each possible panel. (left) The panel distribution of a selection algorithm that does not necessarily select a representative panel from the pool at hand. (right) The panel distribution of a selection algorithm that is guaranteed to choose a representative panel from the pool at hand.

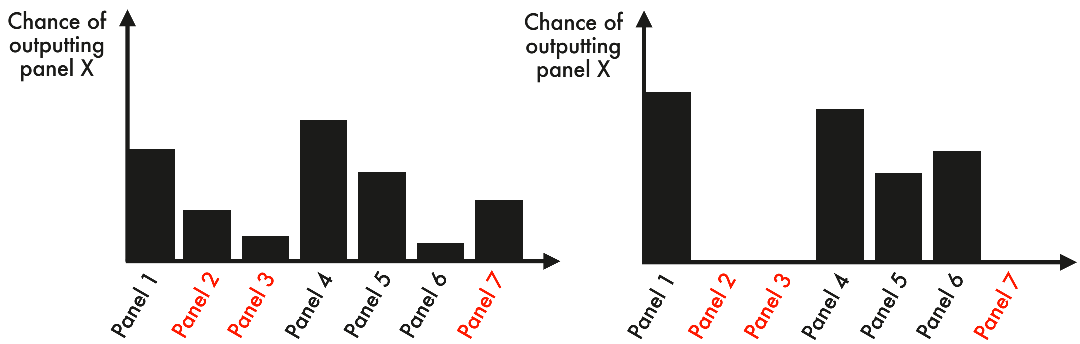

From this example, we can see that requiring quotas to be met means restricting our focus to panel distributions that place 0 probability on all panels that do not satisfy the quotas. Using this correspondence, we can now more specifically define the task of designing the randomness subject to quotas. Our goal is to design the *panel distribution*, considering only those distributions that give positive probability solely to representative panels.

3. Possible Goals for the Design of Randomness

An important observation from the previous section is that *any selection algorithm* is committing to some distribution over panels, whether intentionally or not. This motivates the question: *among the panel distributions that only assign probability to representative panels, which one should we choose?* Below, we outline several potentially desirable properties of a panel distribution, all motivated by the legitimacy-conferring properties of a simple lottery.[14] Of course, due to quotas, there might not be a panel distribution that perfectly achieves any of these ideals; however, we can try to pick the panel distribution that achieves a given ideal to the *highest possible degree*. In practice, we may want multiple of these properties at once. Since each panel distribution entails different trade-offs among these desiderata, one must determine the relative importance of these (and potentially other) considerations to their project, and then choose a selection algorithm whose panel distribution aligns most closely with their preferences. Below, we outline some considerations and properties that might guide the choice of panel distribution.

3.1 Individual Fairness

One important property of a panel distribution is how it treats *individual pool members*—i.e., the probability it gives each pool member of being chosen for the final panel. For example, suppose that a panel distribution chooses between two panels, each with probability ½ (thus placing probability 0 on all other panels). Then, pool members contained in both of these panels would be chosen with probability 1, pool members contained in only one of the two panels with probability ½, and all remaining pool members with probability 0. This example illustrates how a panel distribution translates to each pool member's chance of being on the chosen panel: a pool member's *selection probability* can be calculated by summing up the probabilities of all panels that include that pool member.

Given this relationship between the panel distribution and individuals' selection probabilities, the space of available panel distributions gives the selection algorithm ample freedom in how it allocates selection probabilities across the pool members. However, there are some restrictions. First, selection probability is not an unlimited resource. Since any panel contains a fixed number of panelists (100, in our example), it holds for any panel distribution that the selection probabilities across all pool members sum up to the number of available panel seats (i.e., 100). Second, it is typically impossible to give all pool members *equal* selection probabilities, even though that would be normatively attractive.[15] The climate-change example in Section 2.1 illustrates this limitation: when there are groups that are substantially over- or under-represented in the pool relative to the population (as is virtually certain in practice), a panel distribution that gives all pool members equal probabilities would, just like the simple lottery, place probability on unrepresentative panels.

Without the possibility of giving all pool members exactly the same probabilities, the next-best goal is to give them probabilities that are *as close to equal as possible*. In fact, this stated goal, when implemented precisely, turns into a myriad of related goals, as there are many ways to measure deviation from equal selection probabilities. A mini-public organizer might want to trade off different kinds of such deviations. For example, an organizer determined to avoid very low selection probabilities might get closest to their ideal of equality by maximizing the minimum selection probability in exchange for allowing some pool members to receive selection probabilities close to 1, whereas an organizer who wants to avoid large selection probabilities might achieve their goal by choosing some pool members with probability 0. Ultimately, one's notion of closeness to equal selection probabilities reflects their priority over different ways in which a similarity of selection probabilities can support legitimacy. For instance, here are some conceptual goals with which one might pursue maximally equal selection probabilities:

Ensuring that everyone has an opportunity to participate. Suppose that a selection algorithm gives some pool members selection probabilities that are 0 (or very close to 0). Further, suppose this is not just a fluke of the current pool; rather, the selection algorithm *systematically* gives people with certain characteristics very low selection probabilities. In this case, constituents with these characteristics would be systematically excluded from the mini-public and have a strong argument against the mini-public's legitimacy. By contrast, if a selection algorithm consistently gives no pool member a very small selection probability, it avoids this danger altogether.

Reducing envy between pool members. If the selection algorithm is transparent enough for pool members to observe their selection probabilities—an outcome already known to be possible[16]—those with low selection probabilities might feel envious of those with higher selection probabilities. Admittedly, since equal selection probabilities are typically not possible, this potential for envy cannot be eliminated. Still, the selection probabilities that result from our choice of panel distribution determine how many pool members experience envy, how many others they envy, and the magnitude of their envy. Therefore, the selection algorithm can be designed to minimize the disruption caused by envy within the pool. For example, when choosing 100 panelists out of 1,000 pool members, the average pool member has a selection probability of 100/1000 or 10 percent. If certain pool members are chosen with a probability close to 1, they might be natural targets of envy; in this case, a selection algorithm that avoids high selection probabilities would be desirable. For a different conception of the risks of envy, the selection algorithm might aim for a panel distribution where the Gini inequality coefficient of the selection probabilities is small.[17]

Reducing incentives for misrepresentation. To achieve representation, any selection algorithm must know pool members' attributes in order to ensure it produces a representative panel. However, in practice, pool members' attributes are typically *self-reported*, and some key types of attributes—e.g., someone's beliefs about a topic—cannot be externally verified. As a result, pool members' selection probabilities must, to some degree, depend on the attributes they report. Of course, if a selection algorithm could always give all pool members equal selection probabilities, this would not be an issue; however, since larger differences in selection probabilities are sometimes necessary to ensure representation, pool members with lower selection probabilities may be tempted to lie about their attributes to increase their own selection probability.[18] Note that this concern not only suggests pursuing panel distributions that give participants maximally equal selection probabilities but also underscores that the selection algorithm should not drastically change the panel distribution based on a single pool member's attributes.

3.2 Complementary Forms of Representativeness

As a way of enforcing representativeness, quotas are particularly firm, since they impose constraints on *all* panels the algorithm might produce, no matter the algorithm's randomness. To complement these deterministic constraints, one can use the available randomness to promote the representation of additional groups in random ways, i.e., by shaping the *likelihood* of these groups being well-represented. We discuss two such ways below.

Representing unknown groups. A unique advantage of the ideal of a simple lottery is that the panels it generates are not only representative for the stratified attributes but for any other group in the constituency as well.[19] For instance, if the political identity of 8 percent of constituents was shaped by Bruce Springsteen's lyrics, or if 17 percent of the population have had their identity stolen in the past, then panels produced by a simple lottery of 100 constituents will contain about eight Springsteen fans and about 17 victims of identity theft.

Unfortunately, the opt-in step of practical assembly recruitment implies that some groups will inevitably be inaccurately represented. If invitations are sent by physical mail, for example, those who do not read their mail will invariably be underrepresented on the panel. But if membership in a given group is statistically *independent* from opting into the pool when controlling for the stratified attributes, then a selection algorithm with the right kind of randomness might be able to recover representativeness for that group. It is not obvious which properties of a panel distribution make it best suited for this purpose, and the answer depends on which kinds of groups an organizer is most concerned about. Still, as for most of the other objectives, we would expect selection algorithms that spread selection probability "more equally" to perform better toward this goal.

Ex-ante representation. One caveat about quotas is that, if for instance, we set quotas requiring the inclusion of between 16 and 20 panelists who are "not very concerned" about climate change, it does not necessarily mean that the panel distribution will include panels across the entire range between 16 and 20. For instance, the panel distribution might place non-zero probability only on panels with exactly 16 members of this group, resulting in consistently lopsided representation of this group. To rule out such lopsidedness, we can control the randomness through an *ex-ante requirement*. In this example, we can require that the total probability given to pool members in this group is 18, thereby ensuring that the *average* panel in the distribution must contain 18 panelists who are "not very concerned." This would mean that the number of such panelists would vary across different draws from the panel distribution, sometimes being above 18 and sometimes below it.

But if we believe that 16 panelists from this group are too few, why don't we simply impose tighter quotas instead of an ex-ante requirement? We could set both the lower and upper quota to 18, which would mean that every panel must have exactly 18 "not very concerned" panelists. The main reason is that imposing quotas can come at a cost. As we impose more quotas, the number of representative panels decreases, limiting the available randomness. At the extreme, we can impose so many quotas that there will no longer be *any* panels within our pool that satisfy all our quotas. Adding some slack between lower and upper quotas and imposing 18 only as an ex-ante requirement is less restrictive. While it is true that too many ex-ante representation requirements can also become impossible for any panel distribution to satisfy, we can design the panel distribution with the goal of violating these requirements by as little as possible. On a technical level, this goal is closely related to the goal of making selection probabilities as equal as possible, as discussed in Section 3.1.

4. Selection Algorithms: Possibilities and Future Work

Suppose you, as the mini-public organizer, have decided which benefits of randomness you value most. Now, you want a selection algorithm that uses randomness to implement these priorities. While some goals can be achieved via existing selection algorithms, other goals would require new algorithms. In Section 4.1, we describe the state-of-the-art selection algorithms for achieving maximal individual fairness, as discussed in Section 3.1. Although these algorithms can be adapted toward some other goals, there currently exist no selection algorithms that are specifically tailored toward other goals. We attribute this largely to challenges in three areas: goal quantification, algorithm design, and data availability. Each of these areas presents a fruitful opportunity for future work, and we discuss how work from multiple disciplines can contribute to these three areas in Sections 4.2, 4.3, and 4.4, respectively.

4.1 Existing Possibilities for Individual Fairness

If your goal for your mini-public is individual fairness, you can use an algorithmic approach developed by the authors in collaboration with Anupam Gupta at Carnegie Mellon University and Brett Hennig at the Sortition Foundation.[20] These algorithms permit you to select *any measure* of closeness to equal selection probabilities (so long as it satisfies a relatively weak mathematical property). Then, given this measure, our algorithm will find a panel distribution over only representative panels that makes pool members' selection probabilities *as fair as possible* according to the measure you chose. For example, if the specific measure of fairness aims to avoid small selection probabilities to the largest extent possible (as discussed in Section 3.1), and our algorithm gives any pool member a selection probability of 5 percent, then, for this pool and set of quotas, you can be sure that no distribution over representative panels—and therefore *no selection algorithm*—can give every pool member a selection probability higher than 5 percent. We call this objective *maximin* because it *maximizes* the *minimum* selection probability.

If you want to achieve the particular objective *maximin*, a software implementation of our algorithm is available for free on our online platform Panelot.[21] This implementation, as well as an implementation that measures closeness to equality by maximizing the geometric mean (an objective commonly known as *Nash Welfare*), is also implemented in the open-source program StratifySelect.[22] To our knowledge, software implementations of this class of algorithms for other fairness measures do not currently exist. However, extending the algorithmic approach to many other notions of individual fairness, and further to the goal of ex-ante representation, as discussed in Section 3.2, should be algorithmically feasible.

4.2 Future Work in Quantifying Goals

A key step in our approach to maximizing individual fairness was *formulating these fairness goals in precise, quantitative ways* (e.g., maximizing the lowest selection probability). Similarly, to design a selection algorithm that aims to, say, represent unknown groups (as discussed in Section 3.2), it would first be necessary to specify a measure that quantifies a panel distribution's performance with respect to this desideratum. We enthusiastically welcome proposals for quantitative measures that capture desirable normative properties of selection algorithms beyond fairness, such as those described in Section 3. The work to be done here is only in part quantitative: one needs to not only specify a measure but also argue that the proposed measure indeed captures a desirable normative property of randomness.

4.3 Future Work in Algorithm Design

To the best of our knowledge, no selection algorithm currently used in practice—except for those mentioned above—maximizes any particular goal.[23] Instead, most existing algorithms simply focus on finding any representative panel, adding in some randomness along the way. These procedures implicitly give rise to a panel distribution with no particular guarantees other than ensuring representation. This is reasonable, as controlling selection algorithms' randomness while ensuring representation is algorithmically challenging for two reasons. First, among the enormous number of possible panels, identifying the (still substantial number of) *representative* panels is a formidable task. Second, any panel distribution that achieves a given goal may need to place probability over more panels than can be found in a practical timeframe.

The algorithms in Section 4 sidestep these challenges by targeting fairness to *individual pool members*, a goal that has a special structure: the fairest panel distribution only randomizes over a fairly small number of panels, which means that almost all panels have 0 probability in this panel distribution. As a result, the algorithmic problem boils down to finding the right set of panels; from there on, finding the optimal panel distribution is comparatively easy.

This presents a promising direction for future work: the development of selection algorithms for goals beyond individual fairness. Unfortunately, the trick of constructing only a few panels does not generalize to all goals beyond individual fairness. In cases where this approach does not generalize, achieving any "good" panel distribution with respect to a given objective might fundamentally require allocating non-zero probability to a larger set of panels. In such cases, selection algorithms may be unable to explicitly compute the optimal panel distribution. This motivates the exploration of new techniques to *implicitly* determine the right panel distributions, new algorithms with approximation guarantees, and/or heuristic algorithms distinguished by good empirical performance rather than theoretical guarantees.

4.4 Future Work in Data Availability

A final challenge in mini-public selection is the lack of publicly available data and shared infrastructure, which poses a challenge to accessible and standardized algorithmic evaluation. The lack of publicly available data sets (a "data set" consisting of the pool and quotas of a real mini-public) is understandable, given the relative infrequency with which most organizations convene mini-publics and privacy concerns about sharing (even anonymized) data sets. A second challenge is that few of the selection algorithms used in practice are publicly implemented. For many of the selection algorithms developed in-house by practitioners, only fragmentary descriptions are publicly available. This makes it difficult to evaluate their advantages and disadvantages.

As a step toward addressing these challenges, we propose the creation of an online platform that collects quantitative performance measures, implementations of selection algorithms, and mini-public data sets for the purpose of evaluating these algorithms. If required to accommodate privacy concerns, the platform could keep mini-public data sets inaccessible to the public, and just allow researchers and practitioners to request benchmarking results for new measures and algorithms.

In conclusion, randomness is arguably *the* central feature of deliberative mini-public selection, permitting their defining claim of including "everyday citizens." Principled design of this randomness is crucial: after all, *any* selection process inherently commits to some design of this randomness, whether intentional or not. Our initial understanding of how to design randomness for the goal of individual fairness can serve as a foundation for the pursuit of new goals, though new technical ideas will be needed. Challenges remain in precisely defining desirable goals, developing algorithms to achieve them, and evaluating these algorithms in a standardized way. We see an opportunity to address these challenges through a richer and more collaborative conversation across disciplines, organizations, and countries. To start this conversation, we pose the following question to the global deliberative mini-public community: *what can randomness do for you–and what can be done to make it happen?*

Endnotes

1. Both of these issues have been considered in high-profile deliberative mini-publics in the past year. The former was considered at a global deliberative forum sponsored by Meta (Harris, "Improving People's Experiences"), and the latter was considered in a citizens' assembly at the French national level (Jérôme, "Assisted Dying").
2. Converse, "The Nature of Belief Systems," 206–261.
3. Dryzek, "The Crisis of Democracy," 1144–1146.
4. Farrar et al., "Disaggregating Deliberation's Effects," 333–347.
   Fishkin, *When the People Speak.*
   Ingham and Levin, "Effects of Deliberative Minipublics," 51–78.
5. Fishkin, *Democracy When the People Are Thinking.* Emphasis ours.
6. Gleeson, "Seven Citizens' Assembly Members."
7. Leogue, "Chair Rejects Criticism."
8. If we were to draw a large number of panels using a simple lottery, a group making up *x* percent of constituents would, in the average panel, compose *x* many of the 100 panelists. Moreover, the law of large numbers says that, in the vast majority of panels, the number of panelists from this group will not deviate much from *x*.
9. Carson and Martin, *Random Selection.*
10. Throughout this paper, we will assume that the mini-public organizer is directly involved in panel recruitment. Other mini-publics outsource recruitment to a polling company, which seeks to solve a similar problem but tends to be less transparent about its process.
11. How this recruitment process works in practice varies, to some degree, across types of mini-publics and organizing groups. For example, sometimes pool members are recruited through processes that more heavily target groups that are less likely to respond. Some processes adhere more closely to standard polling recruitment methods and expend more effort to ensure that those originally invited participate rather than assembling a large pool of willing participants and selecting from among them. In this essay, we commit to the model described above for the sake of specificity and because its core attributes are common across a wide range of deliberative mini-publics.
12. Flanigan et al, "Fair Algorithms," 548–552.
13. We emphasize the distinction between *stratified attributes* and *strata* (as in stratified sampling). Here, "stratified attributes" are single attributes, like "female" or "lives in a rural area." In contrast, "strata" are mutually exclusive segments of the population that can be defined by arbitrary combinations of attributes.
14. Carson and Martin, *Random Selection.*
15. If the selection algorithm could choose all pool members with equal probability, and all constituents were given an equal opportunity to join the pool, all constituents would have equal opportunity of participation (at least in theory).
16. Flanigan et al., "Fair Sortition Made Transparent".
17. Ceriani and Verme, "The Origins."

18. Flanigan et al., "Manipulation-Robust."
19. Benadè et al., "No Representation."
20. Flanigan et al., "Fair Algorithms," 548–552.
21. https://panelot.org.
22. https://github.com/sortitionfoundation/stratification-app.
23. Flanigan et al., "Fair Algorithms," supplementary information 12.

Bibliography

Benadè, Gerdus, Paul Gölz, and Ariel D. Procaccia. "No Stratification without Representation," *Proceedings of the ACM Conference on Economics and Computation* (2019): 281–314.

Converse, Philip E., "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, edited by David Apter: 206–261. Free Press, 1964.

Dryzek, John S. *et al*. "The Crisis of Democracy and the Science of Deliberation." *Science* 363, no. 6432 (March 2019): 1144–1146.

Carson, Lyn and Brian Martin. *Random Selection in Politics*. Westport, CT: Praeger, 1999.

Ceriani, Lidia and Paolo Verme. "The origins of the Gini index: extracts from Variabilità e Mutabilità (1912) by Corrado Gini." *Journal of Economic Inequality* 10 (2012): 421–443.

Flanigan, Bailey, Paul Gölz, Anupam Gupta, Brett Hennig, and Ariel D. Procaccia. "Fair Algorithms for Selecting Citizens' Assemblies." *Nature* 596, (2021): 548–552.

Flanigan, Bailey, Gregory Kehne, and Ariel D. Procaccia*l*. "Fair Sortition Made Transparent." *Advances in Neural Information Processing Systems* 34 (2021): 25720–25731.

Flanigan, Bailey, Jennifer Liang, Ariel D. Procaccia, and Sven Wang. "Manipulation-Robust Citizens' Assembly Selection." *Preprint*.

Farrar, Cynthia, James S. Fishkin, Donald P. Green, Christian List, Robert C. Luskin, and Elizabeth Levy Paluck. "Disaggregating Deliberation's Effects: An Experiment within a Deliberative Poll." *British Journal of Political Science 40, no.* 2 (2010).), 333–347.

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

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

Harris, Brent. "Improving People's Experiences Through Community Forums." Meta, November 16, 2022, https://about.fb.com/news/2022/11/improving-peoples-experiences-through-community-forums/.

Ingham, Sean and Ines Levin. "Effects of Deliberative Minipublics on Public Opinion: Experimental Evidence from a Survey on Social Security Reform." *International Journal of Public Opinion Research* 30, no.1 (2018): 51–78.

Leogue, Joe. "Citizens' Assembly Chair Rejects Criticism of 8th Report." *Irish Examiner*, February 23, 2018, https://www.irishexaminer.com/news/arid-20467523.html.

Gleeson, Colin. "Seven Citizens' Assembly Members Not Randomly Recruited." *Irish Times*, February 21, 2018, https://www.irishtimes.com/news/social-affairs/seven-citizens-assembly-members-not-randomly-recruited-1.3400734.

Jérôme, Béatrice. "Assisted Dying: French Citizen's Convention Starts Brainstorming with Possible New Legislation in Sight." *Le Monde*, December 12, 2022, https://www.lemonde.fr/en/france/article/2022/12/12/assisted-dying-french-citizen-s-convention-starts-brainstorming-with-possible-new-legislation-in-sight_6007444_7.html.

"Stratify Select," Accessed August 20, 2023, https://github.com/sortitionfoundation/stratification-app.

A PUBLICATION OF THE

Ash Center for Democratic Governance and Innovation
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138

617-495-0557
ash.harvard.edu

HARVARD Kennedy School
ASH CENTER
for Democratic Governance
and Innovation