What Do Mini-Publics Prove?

Abstract: Lottocracy is a novel political arrangement that replaces elected politicians with ordinary citizens selected at random. It is unlike any political system that currently exists, so it is natural to have anxiety about how it would work in practice. To mollify these concerns, many lottocrats point to the successful performance of mini-publics as evidence that lottocracy would work well, if tried. This inference resembles a claim about external validity, I argue, and develop a framework for evaluating it by relying on work from the social sciences and philosophy of science about when such claims are warranted. With this framework in hand, I show that mini-publics teach us little about how lottocracy would work in practice, although the strength of this conclusion depends on which specific version of lottocracy is under consideration. Mini-publics, however, do provide evidence about the abilities of ordinary citizens more generally, undermining some of the strongest claims about their lack of political competence. The framework developed here applies to a broader class of arguments in political philosophy: arguments that use the performance of existing institutions to predict how similar but untried institutions would perform.

Keywords: sortition; lottocracy; mini-publics; deliberative democracy; external validity.

Word count (excluding title page and works cited): 11,881

1. Introduction

A growing body of work in contemporary political philosophy argues that elected legislators should be replaced by ordinary citizens selected at random. The most prominent examples of this view are Alexander Guerrero's (2024) *Lottocracy: Democracy Without Elections* and Hélène Landemore's (2026) *Politics Without Politicians: The Case for Citizen Rule*, but they are not alone (Zakaras 2010; Van Reybrouck 2018; Vandamme and Verret-Hamelin 2017; Allen 2018; Gastil and Wright 2019; Abizadeh 2021; Shoemaker 2024). For the time being, define *lottocracy* as a system of government identical to ordinary representative democracies, except all legislators are chosen by lot rather than election. This definition, of course, obscures a great deal of variation among lottocratic proposals. These nuances will be introduced later in the analysis (§5.7).

Lottocracy is unlike any existing political system. It is therefore natural to have some anxiety about how it might work in practice. Are ordinary citizens selected at random really up to the task of governing? Lottocrats try to ease this anxiety by pointing to the performance of real-world mini-publics as evidence that lottocracy, if tried, would work well. Prominent examples include the British Columbia Citizens' Assembly, where a randomly selected group of citizens was tasked with reforming the province's electoral system, and the 2016–2018 Irish Citizens' Assembly, where randomly selected citizens addressed major policy questions ranging from abortion to climate change. Pointing to mini-publics as evidence that lottocracy would work well is an extraordinarily common argumentative move. Indeed, I can identify at least twelve instances of it (§2).

Mini-publics, by all accounts, have performed admirably. That is not what is at issue here. The question is how much their performance can really tell us about how lottocracy would

work in practice. It cannot tell us very much, I argue. Inferring the likely performance of lottocracy from the actual performance of mini-publics resembles, I contend, a claim about external validity. But external validity, as social scientists know all too well, can never simply be assumed. Policies and interventions that succeed in one setting often fail when transported to another. We are warranted in believing that an intervention will travel only when certain conditions are met.

I outline these conditions in detail below (§3), but the basic idea is simple. Before concluding that a policy that worked in one setting will work in another, we must first have a sense of *why* it worked in the original setting—that is, we must know something about the underlying causal factors at play. We must then have good reason to believe those causal factors will also be present in the setting to which we hope to transport it. Neither condition is satisfied in the case of mini-publics. Mini-publics seem to work well, but we do not really understand why (§4). Moreover, many of the most plausible causal factors behind their successful performance would not be present in lottocracy (§5). Some versions of lottocracy, however, perform better across this second desideratum, making the inference from mini-publics slightly more warranted in those cases (§5.7).

Does this mean mini-publics teach us nothing? No. They undermine the strongest skeptical claims about ordinary citizens' competence to engage in politics by showing that, under certain circumstances, they can deliberate quite well about policy (§6). But demonstrating this is *not* the same as demonstrating that ordinary citizens can competently perform the tasks that a fully institutionalized lottocracy calls on them to perform. Mini-publics provide little evidence for this latter claim, though they do support a weaker, more qualified claim about ordinary citizens' political competence.

My focus in this paper is on a very common argument in the lottocracy literature, but the framework I develop has far broader implications. Political philosophers frequently defend their proposed but untried institutions by pointing to existing institutions that resemble them in certain respects. Once you begin looking for this argumentative move, you see it everywhere. Some examples:

- Advocates of workplace democracy who appeal to the Mondragón cooperatives in Spain's Basque Country as evidence that democratically governed firms can operate efficiently (Wright 2010: 240-246).
- Anarcho-capitalists who appeal to medieval Iceland and other historical cases in which law was privately enforced as evidence that a society without centralized law enforcement need not collapse into a Hobbesian state of war (Benson 1990: ch. 2; Skoble 2026: ch. 4).
- Confucian meritocrats who point to Singapore's success as evidence for political meritocracy more broadly (Bell 2015: 31–36).
- Advocates of universal basic income who appeal to the Alaska Permanent Fund for evidence about how a society-wide basic income might work (Widerquist 2024: ch. 5).
- A recent paper proposes that political parties select candidates by sortition, who then compete in ordinary elections. To gain insight into how this arrangement might work in practice, the authors examine the Mexican political party Morena, which has used a similar procedure (Bosworth et al. 2026).

These inferences resemble the one lottocrats draw from mini-publics. Are they warranted? I cannot answer that in the current paper. But the framework developed here provides a way to

answer it, because these all resemble, at bottom, inferences about external validity. According to the framework I develop, they are warranted when we understand what caused the initial institution to succeed in the first place and also have reason to believe those causal factors will also be present in the proposed institution.

2. The Inference from Mini-Publics

Lottocrats argue that professional politicians should be replaced by ordinary citizens selected at random. A natural concern is that ordinary citizens are not competent enough to govern (Landa and Pevnick 2025a: 174). Competence can mean many things, but what most people who raise this objection probably have in mind is something like intelligence (Kogelmann 2026). It need not mean this, however. One could worry that politics involves certain skills and that randomly selected individuals lack these skills, not because they are dumb, but because they lack practice. Nothing much hangs on how we understand competence, for my focus in this paper is on how lottocrats respond to these sorts of objections, not the objections themselves. So, let's say that a *competence objection* to lottocracy worries that ordinary citizens are less likely to have some trait *t* when compared to elected legislators, and that *t* likely has a positive impact on governance.

How do lottocrats respond to competence objections? Many point to the successful performance of mini-publics as evidence that ordinary citizens are, in fact, sufficiently competent to govern. Consider a few examples. Hélène Landemore and Théophile Pénigaud (2026: 1) note that arguments for lottocracy "rest on the assumption that ordinary citizens can competently engage in policy making and lawmaking. Yet a long tradition in political science … has

entrenched skepticism about this assumption." Mini-publics like the French Citizens' Convention for Climate, however, act as "a proof of concept for involving ordinary citizens qua legislators in more inclusive forms of democracy" (Landemore and Pénigaud 2026: 2). As another example, responding to the question of "whether a randomly selected assembly will be *competent*," Arash Abizadeh (2021: 801) points to several real-world mini-publics—such as the Ontario Citizens' Assembly—and concludes "that we have little reason to expect them to be any worse than the many incompetent representatives who regularly populate existing chambers, whether elected or appointed."

The same move appears throughout the lottocratic literature. After proposing his own version of lottocracy, David Van Reybrouck (2018: 153) imagines an interlocutor objecting that "citizens can't do this!" He responds by observing that "all experiments with citizens' forums show that, in general, the behaviour of participants chosen by lot is dedicated and constructive and their recommendations surprisingly sophisticated" (Van Reybrouck 2018: 155-156). Erik Olin Wright (2010: 172) likewise considers the worry that "members of the Citizens' Assembly will generally not be very well informed about the issues under discussion at the time they are chosen." In response, he closely examines deliberative polling and the British Columbia Citizens' Assembly, concluding that "there is good evidence that with suitable conditions, ordinary citizens are capable of assimilating large amounts of information, evaluating it in a reasonable manner, and using that information to make well reasoned collective decisions" (Wright 2010: 173).

This argumentative move is incredibly common throughout the lottocracy literature (e.g., Zakaras 2010: 465–466; Hennig 2017: 148; Allen 2018: 77; Gastil and Wright 2019: 18n28, 352;

Landemore 2020a: 193; Benson 2024: 212; Bouricius 2026: 178; Landemore 2026: 14). Here is the general structure:

> *Inference from Mini-Publics (IMP)*: The successful performance of mini-publics provides substantial evidence that ordinary citizens can competently perform the tasks that lottocracy calls on them to perform.

The IMP is the most common rejoinder in the literature to competence objections. The aim of this paper is to determine whether it succeeds. But before turning to that, we must unpack it a bit more carefully.

First, what are mini-publics? They are "small groups of randomly selected citizens, operating according to deliberative principles (including facilitated small-group discussions) and tasked with considering one or a number of important policy, institutional, or constitutional reform issues" (Farrell and Stone 2020: 229). This broad definition encompasses a wide variety of deliberative bodies, including citizens' juries, deliberative polls, citizens' assemblies, and more. These different kinds of mini-publics vary along several dimensions, including the number of participants (12–26 in the case of citizens' juries; 100–500 in the case of deliberative polls), the number of meetings (4–5 days in the case of planning cells; 20–26 in the case of citizens' assemblies), and the nature of their outputs (for example, a policy recommendation to parliament, a referendum to be voted on by citizens, a position report, or an opinion survey) (Farrell and Stone 2020: Table 11.1).

Which kinds of mini-publics do IMP proponents appeal to? A wide range of examples are invoked, some of which do not even qualify as mini-publics under the definition offered above. Mini-publics, as just defined, are constituted by random selection, yet some who advance the

IMP invoke the Icelandic Constitutional Council, which was constituted by election (e.g., Landemore 2026: 36-37). Others invoke the 2012–2014 Irish Constitutional Convention, which was only partly constituted by random selection (e.g., Van Reybrouck 2018: 116). Here are the examples IMP proponents explicitly appeal to:

- **The British Columbia Citizens' Assembly**: 160 randomly selected citizens were tasked with deciding whether British Columbia should retain or change its electoral system. The assembly recommended a change, and its proposal was put to a province-wide referendum, where it failed to meet the supermajority threshold required for adoption (Fournier et al. 2011: 6–8).
- **The Dutch *Burgerforum***: 143 randomly selected individuals were asked to consider the Dutch electoral system. The assembly recommended minor changes, which were rejected by parliament (Fournier et al. 2011: 8-9).
- **The Ontario Citizens' Assembly**: 103 randomly selected citizens were asked to consider the province's electoral system. The assembly offered a recommendation which was put to a province-wide referendum. It did not pass; only 37 percent of voters favored it (Fournier et al. 2011: 9).
- **The 2012-2014 Irish Constitutional Convention**: 66 randomly selected citizens and 33 professional politicians were asked to consider eight potential changes to the Constitution of Ireland. Two of the convention's recommendations (on marriage equality and blasphemy laws) went up as referenda, the former passing in 2015 and the latter in 2018 (Farrell et al. 2021: 628).
- **The 2016-2018 Irish Citizens' Assembly**: 99 randomly selected citizens were given five policy questions to consider (Farrell et al. 2019: 114). Notably, one of the assembly's

recommendations led to a referendum on abortion, which passed in 2018 (Farrell et al. 2019: 115).

- **The Icelandic Constitutional Council**: Initiated by an act of parliament, 25 individuals were elected to the body, but no professional politicians were allowed to run. Members drafted a new constitution, which received two-thirds approval in a low-turnout, non-binding referendum. Parliament never held a vote on whether to implement it (Landemore 2020b: 184).
- **The French Citizens' Convention for Climate**: Initiated in 2019, 159 randomly selected individuals were tasked with "defining structuring measures to achieve, in a spirit of social justice, a reduction of greenhouse gas emissions of at least 40% by 2030" (Giraudet et al. 2022: 2). 149 measures were proposed to the government. 20 percent were implemented as proposed and 51 percent in modified form (Averchenkova et al. 2025).

For the remainder of this paper, I will refer to this set as our *sample*. I will also refer to all the members of our sample as "mini-publics" even though, again, some do not meet the technical definition of the term.

The IMP says the successful performance of mini-publics provides evidence that ordinary citizens are sufficiently competent to perform the tasks that lottocracy calls on them to perform. What is meant by "successful performance"? The idea seems to be that the outputs of these deliberative bodies—in particular, their policy recommendations—meet or exceed some relevant normative threshold. I grant this for the sake of argument.

Some who advance the IMP do so only with caution. Landemore (2020a: 181), for example, writes that "the empirical lesson of Iceland as this book describes it—that more truly

democratic forms of representation (or more participation on some other people's readings) is a good idea—needs to be made with caution." The IMP must be advanced with caution because there are many differences between the mini-publics in our sample and lottocracy, making generalization hazardous. Landemore (2020a: 180-191) addresses this by arguing that unique features of the Icelandic case—like the country's small population and homogeneity—do not prevent its lessons from generalizing. This, I argue below, is not the right way to establish that lessons drawn from mini-publics generalize.

Finally, not all lottocrats respond to competence objections with the IMP. Although Guerrero (2024: 275–284) responds to competence objections at length, he makes no appeal to mini-publics. That said, Guerrero does think mini-publics can teach us something about how lottocracy might work in practice. He writes:

> There are important differences between these institutions and full lottocratic institutions, but this kind of experimentation with institutional forms and powers is useful for gathering evidence and building a broader foundation of understanding to reduce some uncertainty about how lottocratic institutions might work in practice (Guerrero 2024: 412).

Guerrero is careful to note the differences between lottocracy and mini-publics, but he ultimately believes that the latter can teach us something about the former. Thus, although he does not advance the IMP, he appears open, at least in principle, to drawing inferences from mini-publics about how lottocracy might work in practice. This matters because, although the paper focuses primarily on the IMP, the framework I develop applies to any attempt to draw inferences from real-world cases to institutions that have not yet been tried.

3. The Problem of External Validity

Is the IMP a warranted inference? We need evaluative standards to answer that. My contention is that the literature on external validity is the most promising place to find them. In particular, we can understand the IMP as resembling an inference about the external validity of the mini-publics in our sample. We can thus evaluate it according to the relevant standards. I now lay the groundwork for this approach. I introduce the concept of external validity and, drawing on the work of both philosophers of science and social scientists, develop standards for determining when external validity inferences are warranted. I then argue that the IMP has the structure of an external validity inference. The remainder of the paper then evaluates the IMP according to the standards developed here.

Policy interventions sometimes achieve great success in one setting, only to fail when tried elsewhere. This is the problem of *external validity*. External validity is "the extent to which inferences drawn from a given study's sample apply to a broader population or other target populations" (Findley et al. 2021: 368). Call the initial case in which an intervention succeeds the *source case*, and the case to which we hope it travels the *target case*. When an intervention is externally valid, its effects travel from the source to the target case. When it is not externally valid, its effects do not travel from the source to the target case. Some examples of failures of external validity:

- In Tamil Nadu, an intervention combining nutrition education, child growth monitoring, and supplemental feeding produced large declines in malnutrition (White and Masset

2007: 628). When a similar intervention was implemented in Bangladesh, it had no effect (Hossain et al. 2005: 38).

- In Kenya, adding teachers on fixed-term contracts to schools staffed by permanent government teachers improved student performance when an NGO administered the program. When the Ministry of Education implemented the same intervention at scale, the effect vanished (Bold et al. 2018: 2).
- In Tennessee, reducing class sizes in early grades substantially increased student achievement. The same intervention was tried in California, but researchers could find no conclusive link between class-size reduction and student achievement (Bohrnstedt and Stecher 2002: 34).
- After producing positive effects in the United States, a nurse home-visiting program for teenage mothers was implemented in England. Adding it "to the usually provided health and social care provided no additional short-term benefit to our primary outcomes" (Robling et al. 2016: 146).
- Teacher incentive programs improved student achievement in Kenya and India. When a similar program was implemented in New York City, however, researchers found "no evidence that teacher incentives increase student performance, attendance, or graduation" (Fryer 2013: 373).

External validity, then, can never be assumed. Even effects demonstrated in randomized controlled trials—the gold standard of experimental research—"cannot automatically be extrapolated outside the context in which they were obtained" (Deaton 2010: 449). Indeed, doing so can be "dangerous" (Banerjee et al. 2017: 96). Under what circumstances, then, are scholars warranted in inferring that a policy that "worked here" will also "work there"? Here I primarily

follow the work of Nancy Cartwright and Jeremy Hardie (2012), although, as I demonstrate below, their general approach to external validity is widely embraced. To begin, why do interventions sometimes fail to travel?

Interventions fail to travel because causal factors present in the source case are absent in the target case. One way to understand causal factors is in terms of what philosophers of science call *INUS conditions*. An INUS condition is an *I*nsufficient but *N*ecessary part of an *U*nnecessary but *S*ufficient condition (Cartwright and Hardie 2012: 63). Less formally, we can think of causal factors as ingredients in a "causal cake" (Cartwright and Hardie 2012: 63). Flour alone does not make a cake; eggs, milk, sugar, and other ingredients are also required. If one of the ingredients is missing, the cake will not turn out. The causal factors responsible for the positive result in the source case include the policy itself—reduced class sizes in Tennessee, nutrition education programs in Tamil Nadu, and so on. But the policy is only one ingredient (i.e., causal factor) in the causal cake. Other ingredients must also be present. If they are not, the policy won't produce the desired effect.

For example, the best explanation for why smaller class sizes succeeded in Tennessee but not California is that important causal factors present in the former were absent in the latter. Schools need suitable spare rooms to hold the additional classes in. Tennessee had many available; California did not. Consequently, many of California's new classes were held in spaces poorly suited to the needs of students and teachers. Smaller classes also require more teachers. Tennessee had an ample supply, but California had to hire an additional 12,000, many of whom were unqualified (Cartwright and Hardie 2012: 64-67). Thus, the intervention failed to travel successfully from Tennessee to California because two causal factors crucial for its success in

Tennessee—adequate classroom space and a sufficient supply of qualified teachers—were absent in California.

As another example, consider what went wrong in Kenya. The intervention supplemented teachers who enjoyed civil service protections with teachers hired on fixed-term contracts. One reason the policy initially worked was because teachers hired on fixed-term contracts believed they needed to perform well to get rehired (Bold et al. 2018: 3). The credible threat of nonrenewal was a causal factor. But when the government—unlike the NGO—began hiring teachers on fixed-term contracts, this credible threat vanished, because unions pressed the government to grant new teachers civil service protections (Bold et al. 2018: 16). Hence, an important causal factor—the credible threat of nonrenewal—was absent.

Now that we understand why external validity sometimes fails, we can turn to the circumstances under which we are warranted in inferring it holds. Philosophers of science have developed accounts of when inferences about external validity are warranted. These accounts differ in important ways, but many share a common theme. In particular, we are warranted in inferring that an effect observed in a source case will travel to a target case when we understand the causal factors—sometimes "causal principles" or "causal roles" or "causal mechanisms" are used instead—responsible for the effect in the source case, and also have reason to believe they will be present in the target case.

According to Cartwright and Hardie (2012: 54), to infer that a policy that worked "there" will also work "here," we must know that "the policy can play the *same causal role* here as there." In a similar vein, Daniel Steel (2008: 85) argues that "knowledge of mechanisms and factors capable of interfering with them can provide a basis for extrapolation." As another example, Caterina Marchionni and Samuli Reijula (2019: 62) argue that "extrapolating

behavioral policies requires evidence about the invariance of the causal relationship between the policy lever and the outcome."

Social scientists agree that the key to warranted external validity inferences is causal knowledge. For example, a large team of experimental economists led by Nobel laureates Abhijit Banerjee and Esther Duflo argue that moving from a successful study to a large-scale policy "requires elaborating a model of the mechanism at play" (Banerjee et al. 2017: 96). Angus Deaton (2010: 448), another Nobel laureate, writes that "for an RCT to produce 'useful knowledge' beyond its local context, it must illustrate some general tendency, some effect that is the result of a mechanism that is likely to apply more broadly." Finally, a group of political scientists argues that external validity inferences are credible when they "are tied to mechanisms," where "the mechanism is clearly specified, which entails an articulation of the causal principles" (Findley et al. 2021: 377).

Let me try to pull these related ideas together in the form of two conditions that, when jointly satisfied, are sufficient for a minimally warranted external validity inference (what I mean by "minimally warranted" I say below). To begin, before we can know whether causal factors transfer from the source to the target case, we must first know what the causal factors in the source case are. Hence:

> *Articulation condition*: We have a well-supported account of the causal factors that produced the desired effect in the source case.

How do we develop a well-supported account of causal factors in the source case? There is no single answer. The gold standard is further experimental work (Banerjee et al. 2017: 96), but that is not always possible. Researchers might instead employ process tracing (Deaton 2010: 450) or

similar strategies (Cartwright and Hardie 2012: ch. III.B). Causal factors can be inferred through careful study of small-*N* case studies (Collier 1993: 108). A recent book develops a framework for drawing causal inferences from qualitative data (Waldner 2026). There are many credible ways to identify causal factors. I take no stance on how we arrive at a well-supported account of the causal factors at work.

Understanding the relevant causal factors in the source case is, of course, insufficient to warrant an inference about external validity. We must also show that these factors will be present in the target case. More specifically:

> *Transfer condition*: The causal factors that produced the desired effect in the source case are present in the target case.

For example, if we know that small classes work only when there are enough good teachers, then before we are warranted in believing that small classes will work in California, we must know something about the supply of qualified teachers in the state. If we know that small classes work only when coupled with classrooms of a certain quality, then before we are warranted in believing that they will work in California, we must know something about the condition of the state's school facilities.

When both the articulation and transfer conditions are satisfied, we are *minimally warranted* in believing the effect will travel. That is, we should have limited confidence in the belief. Why? Because this is a domain of inquiry in which mistakes come cheap. Despite our best efforts to identify the causal factors present in the source case, we may overlook a crucial causal factor that also happens to be absent in the target case. Likewise, despite our best efforts to determine whether the relevant causal factors will be present in the target case, we may misjudge

the matter. Indeed, causal factors can vanish unexpectedly. Return to the Kenya case. The credible threat of contract nonrenewal was a relevant causal factor. When unions pressed the government to grant contract teachers permanent protections, that threat disappeared. But who could have predicted this *ex ante*? Those preparing to scale the program may well have recognized the importance of nonrenewal, and they may well have expected that threat to remain. Politics had other plans.

The two conditions therefore support only a minimally warranted belief that an effect will travel. Perhaps further conditions could strengthen the inference, but the two developed here are good enough for my purposes. Even when the two conditions are satisfied, the effect may still fail to travel. Conversely, when at least one of them is not satisfied and we therefore lack minimal warrant, the effect may travel nonetheless. Such is the nature of inductive inference. Still, we can say this much: absent some alternative inferential strategy, if the articulation and transfer conditions are not satisfied, then we have very little reason to believe that the intervention's effects will travel.

Back to the IMP. It says that the successful performance of mini-publics provides substantial evidence that ordinary citizens can competently perform the tasks that lottocracy calls on them to perform. This resembles an external validity inference. Rather than a single source case, we have many: the mini-publics in our sample. The intervention is randomly selected individuals deliberating about policy. Across our source cases, the intervention produces an effect: participants make good policy recommendations. The question is whether there is even minimal warrant to believe this effect will travel to the target case, lottocracy. Understanding the IMP as an external validity inference allows us to apply the criteria developed in this section. If

the articulation and transfer conditions are satisfied, then the IMP is minimally warranted. If not, then not.

To close, let me briefly return to a point from the prior section. There, I noted that Landemore is the theorist who spends the most time reflecting on whether the IMP is a warranted inference. Her strategy is to identify two potential differences between the Icelandic Constitutional Council and lottocracy—size and homogeneity—and argue that neither prevents the result from generalizing. This is not a promising approach. As philosophers of science note, any two source and target cases will both resemble and differ from one another along an indefinite number of dimensions (Guala 2003: 1201). If we try to explain away every difference between the Icelandic Constitutional Council and lottocracy, we will run out of ink. A better approach is through the articulation and transfer conditions. Landemore should first ask why, exactly, the Icelandic Constitutional Council—as well as all the other mini-publics in our sample—worked as well as it did. Once she has identified the relevant causal factors, she can then ask whether those factors would also be present in lottocracy. This is what I now turn my attention to.

4. The Articulation Condition

For the IMP to satisfy the articulation condition, we need a well-supported account of the causal factors that lead mini-publics to produce good policy recommendations. Do we have such an account? Landemore (2020a: 170-171) argues that the Icelandic Constitutional Council performed well because it included a diversity of perspectives. A severely disabled participant, for example, purportedly exerted significant influence on the proposed constitution's progressive

human rights provisions. More generally, there is a large literature on epistemic democracy that highlights how ordinary citizens can make good decisions through a variety of mechanisms such as the Condorcet Jury Theorem, the wisdom of crowds, the Hong-Page theorem, and more (Schwartzberg 2015). These are, no doubt, plausible causal factors that could explain the success of mini-publics. But there are other causal factors that are equally plausible. I now consider five of them.

Mini-publics may perform well due to *selection bias*. Social scientists must constantly worry about this: those who voluntarily participate in an experiment may differ from the broader population in ways that contribute to the observed effect (Banerjee et al. 2017: 80-82). Indeed, this has led to failures of external validity in some cases. For example, the Opower energy conservation program ran trials in which consumers received tips for reducing energy use, along with comparisons between their energy consumption and that of their neighbors. The initial trials produced much larger reductions in energy consumption than later trials, partly because the utility companies that first volunteered for the program served populations that were already environmentally conscious (Allcott 2015). The initial trials thus overestimated the effects of the program when implemented more broadly. Something similar may be happening with mini-publics. Those who accept an invitation to participate may be atypical in ways that impact the overall performance.

We know that mini-public participants are atypical in at least one potentially relevant respect: they are substantially more educated than the populations from which they are drawn. In the Dutch *Burgerforum* case, for example, 9 percent of the population did not have a high school degree, but only 1 percent of the assembly did not. Meanwhile, 26 percent of the population held a bachelor's degree, compared to 57 percent of the assembly (Fournier et al. 2011: 55). Most of

the mini-publics in our sample have similar disparities (e.g., Fournier et al. 2011: 55; Landemore 2015: 178; Farrell et al. 2021: 634). The closest any mini-public in our sample has come to educational parity was the French Citizens' Convention for Climate, where 40.9 percent of participants lacked the baccalauréat, compared with 49.6 percent of the population (Giraudet et al. 2022: Table 1). But even this difference is not insubstantial. With 159 total participants, the 8.7-percentage-point gap amounts to roughly fourteen participants, which can make a difference in small group settings.[1]

Education is simply the trait we happen to measure. That mini-publics are demonstrably unrepresentative on it should raise the worry that they are unrepresentative on traits that are harder to observe (Landa and Pevnick 2025b: 443). For example, it would not be terribly surprising if those who agree to participate possess greater political knowledge than the broader population. This is a plausible causal factor in mini-public performance: those who pay close attention to politics and thus know more about it may enter mini-publics with a reservoir of information that helps them understand the issues and evaluate the arguments. It would also not be too surprising if participants were more conscientious than the average citizen or had a stronger sense of civic duty. These are also candidate causal factors, since the quality of mini-public decisions will plausibly be higher if participants carefully read and reflect on the materials they are given.

[1] It's also worth noting that the French Citizens' Convention for Climate contained a "plethora of supervisory bodies" which was "unparalleled in other citizens' and climate assemblies" (Giraudet et al. 2022: 9). This unusually extensive structure and oversight should be kept in mind when interpreting its performance.

Selection bias is not the only potential causal factor. Mini-publics may perform well because they deliberate over a *narrow policy agenda*. The British Columbia, Dutch, and Ontario assemblies focused only on electoral reform, while the Icelandic Constitutional Council on a new constitution. The two Irish bodies considered a somewhat wider array of questions, including abortion, climate change, electoral rules, marriage equality, and constitutional reform (Farrell et al. 2019: 114; Farrell et al. 2021: 630). But even these are narrow mandates compared to ordinary legislatures, which must address everything from foreign affairs and monetary policy to healthcare, immigration, and trade. This narrow mandate may partially explain their success. Participants can focus their attention on a bounded set of questions and develop a level of expertise that would be much harder to achieve across the full range of issues that confront ordinary legislatures.

Another possible causal factor is the *absence of lobbying*. The stakes are low in mini-publics. They typically produce recommendations for parliament or for the electorate via referendum. They never exercise decisive authority, so there is little incentive to lobby. Actors with intense preferences over the issues addressed by mini-publics have better options for influence. If a recommendation proceeds directly to referendum, they can launch a public campaign to convince voters; if it goes to parliament—as happened in the modal case in our sample—they can lobby politicians. Participants in mini-publics thus face far less lobbying than ordinary legislators do. This may partially explain their success.

Indeed, mini-public participants were quite insulated from outside influence. In the British Columbia, Dutch, and Ontario cases, expert testimony "did not rely on advocates for specific systems. Rather, the staff outlined the strengths and weaknesses of each [voting system] in a neutral fashion" (Fournier et al. 2011: 35). As another example, in the 2012-2014 Irish

Constitutional Convention, "the advisory and legal group team played a key role in identifying academics who were the leaders in their fields and were noted for their independence and not being identified with a partisan position" (Arnold et al. 2019: 113). That said, both Irish assemblies permitted testimony from advocacy groups (Farrell et al. 2019: 116; Farrell et al. 2021: 631). But even here, facilitators prescreened which ones were allowed to testify (Arnold et al. 2019: 113-114).

Yet another possible causal factor is the *highly structured procedures*. Participants are told which policy questions to consider (Farrell et al. 2021: 630). Their deliberation is managed by trained moderators (Farrell et al. 2021: 631). The experts from whom they hear are selected and vetted by those organizing the mini-public (Fournier et al. 2011: 30–31). The entire sequence of events—how much time participants spend learning, hearing testimony, deliberating, and so on—is planned (Fournier et al. 2011: 29). In the French Citizens' Convention for Climate, "between sessions, the citizens' proposals would be compiled by the facilitators, assessed by the technical advisors and reformulated by the legal advisors in legal terms. At the next session, the citizens would start with their reworked proposals" (Giraudet et al. 2022: 6). Such structure plausibly improves the quality of decisions by directing participants' attention toward relevant evidence and arguments while preventing their deliberation from becoming disorganized and unproductive.

Finally, *piloting bias* may be another causal factor. Piloting bias arises when those who run an intervention are motivated to see it succeed. This may lead them to implement it with unusual care, closely monitor whether participants adhere to the intervention, and correct problems if they arise. When the intervention is implemented on a larger scale, by contrast, these tasks typically fall to bureaucrats who may lack the same motivation (Banerjee et al. 2017: 82–

83). Indeed, there is empirical evidence that interventions produce larger effects when implemented by academics or non-profits than when implemented by bureaucracies (Vivalt 2020). To be clear, the concern here is not that facilitators behave corruptly, but that small differences in implementation—such as greater care in interacting with participants—can impact outcomes.

The mini-publics in our sample were plausibly subject to piloting bias. The British Columbia Citizens' Assembly was run by an academic from Simon Fraser University who founded a research center dedicated to civil discourse (Lang 2007: 39). The 2012–2014 Irish Constitutional Convention grew out of *We the Citizens*, a pilot study organized by academics who studied deliberative democracy with the explicit aim of persuading the government to adopt the model (Farrell et al. 2013: 111). When the government did, the same academics helped run the convention (Convention on the Constitution 2014: 6; Courant 2021: 6-7). They remained involved in the 2016–2018 Irish Citizens' Assembly, surveying participants after every meeting. Their research was expressly intended to provide feedback that could contribute to the assembly's "ongoing development and improvements" (Citizens' Assembly 2017: 57–58). The French Citizens' Convention for Climate was co-chaired by Thierry Pech, the president of a progressive think tank, and Laurence Tubiana, an architect of the Paris Agreement (Giraudet et al. 2022: 4).

Here, then, are some candidate causal factors: the epistemic benefits of deliberation, various forms of selection bias, narrow policy agendas, the absence of lobbying, highly structured procedures, and piloting bias. I am not claiming that all—or even any—of these causal factors *actually* explain the strong performance of the mini-publics in our sample. My claim is that they are *plausible*. Some, all, or none of them may be actual causal factors. We simply do

not know. This uncertainty is a problem for the articulation condition. The condition requires a well-supported account of the causal factors that produced the desired effect in the source case. *But we have no such account for the mini-publics in our sample*. We instead have several candidate causal factors, but little work that attempts to adjudicate which of them are actual causal factors. Proponents of the IMP, we can conclude, have not demonstrated that the articulation condition is satisfied.

The articulation condition is not satisfied, but this failure need not be permanent. As I noted in the prior section, there are many ways to acquire knowledge of causal factors. In the case of mini-publics, further experimental work with greater variation strikes me as the most promising path forward.

For example, researchers can test whether positive selection matters by convening mini-publics disproportionately composed of participants with less education, low levels of political knowledge, or a low propensity to vote (the latter acting as a proxy for conscientiousness). They can test the importance of narrow policy agendas by varying the number and range of topics participants must consider. Why not ask a single mini-public to address monetary policy, climate change, trade deficits, welfare reform, healthcare reform, the national debt, AI regulation, and the Russia–Ukraine War? They can test the importance of absent lobbying by exposing participants to unfiltered, biased, or even deliberately misleading testimony. They can test the importance of highly structured procedures by allowing participants greater control over which issues to discuss, how to allocate their time, and which experts to invite. Indeed, why not drop the moderators altogether? And they can test for piloting bias by comparing mini-publics administered by experienced advocates of deliberative democracy with those administered by neutral facilitators.

None of these studies would, by itself, provide a complete account of why mini-publics succeed. Together, however, they could help us distinguish the causal factors that actually contribute to their success from features that are merely incidental. In principle, then, we could develop a well-supported account of the causal factors that produce good outcomes in mini-publics. If we did, then the articulation condition would be satisfied. But at present, we simply lack such an account.

5. The Transfer Condition

The articulation and transfer conditions are jointly sufficient for a minimally warranted belief that an intervention's effects will travel from a source to a target case. Because proponents of the IMP have not yet established the articulation condition, the IMP remains unwarranted under the inferential framework adopted in this paper. Strictly speaking, then, there is no need to ask whether the transfer condition is satisfied. Yet IMP proponents may eventually provide a well-supported account of the causal factors responsible for the successful performance of mini-publics. Indeed, I sketched how they might do so at the end of the prior section. If they succeed, the transfer condition becomes relevant. It is therefore worth asking whether the IMP could satisfy that condition.

The transfer condition requires that the causal factors responsible for the effect in the source case are also present in the target case. Does the IMP satisfy this condition? We cannot answer that question because we do not yet know which of the possible causal factors are actual. We can, however, ask whether the candidate causal factors identified in the prior section would likely also be present in lottocracy. If all of them would be, that would be encouraging: even if

every candidate causal factor contributed to the successful performance of mini-publics, none would create a transfer problem for the IMP. If, however, some candidate causal factors would most likely be absent in lottocracy, then the IMP faces a potentially serious problem. Should any of those candidates turn out to be actual causal factors, the transfer condition would not be satisfied.

Answering this question requires we speculate about how lottocracy would work in practice. This is far more difficult than the cases discussed in §3. To determine whether those kinds of interventions satisfy the transfer condition, we must ask whether California has enough classroom space or qualified teachers. Here, by contrast, we must ask whether lobbying would arise in an untried political system, what kinds of selection bias might emerge, and so on. Some regard questions of this kind as essentially unanswerable (Gaus 2016: ch. 2). If correct, then the IMP—and analogous inferences from real-world to untried institutions—is a nonstarter. I am less skeptical. We can gain imperfect insight into how untried institutions might work by carefully thinking through the incentives they create. That is what I do below. The conclusions are of course tentative and open to debate.

One final point before I begin: I grant, without argument, that the various epistemic mechanisms that may explain mini-public performance would also be present in lottocracy. I make this assumption partly to give proponents of the IMP the benefit of the doubt (these mechanisms are how many of them explain the success of mini-publics), and partly because I am short on space.

5.1 *Selection bias*.

Will lottocracy exhibit the kinds of selection bias that may be present in mini-publics? Start with positive selection on education. This occurs in mini-publics because less educated citizens decline invitations to participate at higher rates. Why they do so is hard to determine; people decline invitations to participate in mini-publics for many reasons (Jacquet 2017: 648; Kübler and Rochat 2025). Two explanations seem especially plausible. First, most mini-publics meet on weekends, and white-collar workers are more likely than blue-collar workers to have weekends free. This is the *availability gap*. Second, more educated citizens may simply be more interested in politics than less educated citizens. This is the *interest gap*. Would these gaps persist in lottocracy?

The availability gap would vanish because service in a randomly selected chamber would not be performed on top of one's regular job; it would replace that job. The interest gap, however, may persist: less educated citizens may remain less interested in politics than their more educated counterparts. But in lottocracy, a lack of interest may not translate into a declined invitation because compensation for service would be generous (Zakaras 2010: 458; Van Reybrouck 2018: 138; Gastil and Wright 2019: 19; Abizadeh 2021: 800). Guerrero (2024: 159), for example, recommends paying participants 110 percent of their current salary, with a floor of $100,000. Given that the median U.S. worker with a high school diploma but no college degree earns about $49,000 a year (Bureau of Labor Statistics 2025), serving in a randomly selected chamber would nearly double the income of many blue-collar workers. This financial incentive would likely lead many citizens who lack an intrinsic interest in politics to nonetheless accept an invitation.

There are good reasons to think that less educated citizens will not select out of lottocracy at the same rate they do for mini-publics. There are also good reasons to think that more

educated citizens will not select into lottocracy at the same rate they do into mini-publics. Brian Kogelmann (2026: §4) makes this argument in a recent paper. He argues that highly educated citizens may decline to serve in a randomly selected legislature because doing so imposes especially high opportunity costs on them. Since education is associated with higher earnings, temporarily leaving one's career may involve substantial financial sacrifice. Paying members 110 percent of their current salaries may not solve this problem, because it only compensates them for present income, but not for the future earnings they may forgo by interrupting their careers. Mini-publics do not face this problem, because participation does not require interrupting one's career. It is a weekend activity.

Turn now to positive selection on political knowledge. If mini-publics do in fact contain individuals with unusually high levels of political information, what explains this? The interest gap is again relevant. Those with an intrinsic interest in politics are more likely to possess high levels of political knowledge (Luskin 1990: 348). Presumably, these individuals are also more likely to accept invitations to participate in mini-publics. But generous compensation, I just argued, can induce people with a weak interest in politics to participate in lottocracy. Receiving 110 percent of one's current salary, with a floor of $100,000, will likely persuade many people who neither follow nor care much about public affairs to serve. Hence, just as generous compensation likely mitigates positive selection on education, it probably also mitigates positive selection on political knowledge.

What about positive selection on conscientiousness? Once again, generous compensation may encourage people who lack a strong sense of civic duty to participate in lottocracy. Not only this, but it may also encourage some highly conscientious citizens to *not* participate. A large body of empirical work shows that paying people to perform activities they would otherwise

voluntarily undertake can reduce their willingness to perform them. The most famous example comes from Richard Titmuss's (1997) work on blood donation, which showed that paying people to donate blood reduced donations relative to uncompensated giving. Another example comes from a daycare in Israel. To encourage parents to pick up their children on time, the daycare began fining those who arrived late. Contrary to the intended effect, tardiness increased after the fine was introduced (Gneezy and Rustichini 2000). More examples are available (e.g., Frey et al. 1996; Frey and Jegen 2001; Besley and Ghatak 2018).

Why does this happen? One explanation says that introducing monetary incentives "crowds out" civic duty. If an activity is framed as something done for money, people who were willing to do it out of civic duty become less willing to do so (Bowles 2016). If correct, then some conscientious citizens who would voluntarily participate in an unpaid or low-paid mini-public may decline a highly paid position in lottocracy, just as some individuals who would donate blood for free won't do it for money. Lottocracy may therefore reduce positive selection on conscientiousness from both sides: by attracting those who are less conscientious, and repelling some who are more.

5.2 *Narrow policy agendas*.

Will lottocratic legislatures have more extensive policy agendas than mini-publics? Almost certainly. A legislature must decide questions across the full range of public policy, so its agenda will be far broader than any mini-public. Most mini-publics focus on a single issue—such as voting rules, a new constitution, or climate change. The widest agenda among the mini-publics in our sample belonged to the 2012–2014 Irish Constitutional Convention, which

considered eight separate questions. A legislature for an entire nation must address more than eight questions.

5.3 *Absence of lobbying*.

Lottocratic chambers, unlike mini-publics, can make legislation. The stakes will be much higher. There will, as a result, be tremendous incentive to lobby them, just as there is for elected legislatures. But though the incentive will be present, will those who wish to lobby be able to do so effectively?

To think through this question more carefully, it is helpful to distinguish between different types of lobbying. First, *exchange lobbying* happens when private interests offer politicians campaign contributions and, in return, expect legislators to provide favorable legislation (Morton and Cameron 1992: 88). Lottocracy, one might think, eliminates this form of lobbying. Because there are no elections, private interests cannot offer campaign contributions as part of a *quid pro quo*. Some argue, however, that exchange lobbying will persist. Corporations might, for example, establish a reputation for offering lucrative jobs to members of a lottocratic chamber who vote in ways they find salutary (Landa and Pevnick 2021: 55–56). Lottocrats have pushed back, arguing that such practices would be relatively easy to regulate (Guerrero 2024: 290). I tend to agree with the lottocrats that exchange lobbying is not a significant threat in lottocracy.

Far more likely to arise in lottocracy is *informational lobbying*. Private interests attempt to influence legislators by providing biased information that they hope will shape policy (Hall and Deardorff 2006). For example, a fossil fuel company might supply legislators with

information suggesting that climate change is less serious than the scientific consensus says. Informational lobbying is especially dangerous in lottocracy. Because ordinary citizens with no expertise in the relevant domains are asked to legislate, they will require extensive testimony from experts and advocacy groups (Abizadeh 2021: 802; Guerrero 2024: 160–161). We should expect private interests to do everything in their power to bias this pool of testimony to secure favorable policies.

Some lottocrats recognize this danger and propose ways to mitigate it. Guerrero (2024: 204–209), for example, suggests establishing a database of experts who meet some credential threshold and then randomly selecting from it when expert testimony is needed. This may reduce some forms of informational lobbying, but it could open others. Private interests would be incentivized to influence the institutions that produce the experts. We might, for instance, see a flood of money enter academia to bias those who are in the database. Fossil fuel companies might begin lavishly funding earth science departments in hopes of influencing experts' beliefs. Indeed, funding earth science departments might be cheaper than current spending on elections we see in the United States.

We should also not be surprised to find *grassroots lobbying* in lottocracy. In elected legislatures, grassroots lobbying occurs when interest groups inform their members about a legislator's voting record or anticipated vote in hopes of generating pressure (Lohmann 1998). An interest group representing corn farmers might alert its members in Iowa that their senator is about to vote on a bill slashing corn subsidies. The farmers flood the senator's office with angry calls to change her vote. It's not hard to imagine analogues to this in lottocracy. Suppose a randomly selected legislature is considering environmental regulations that threaten a major employer in a legislator's hometown. The employer might launch a campaign portraying the

legislator as personally responsible for the potential loss of local jobs. It could run advertisements featuring her name and face, organize demonstrations in the town, and encourage those who know her personally to contact her.

We should expect significant informational and grassroots lobbying in lottocracy. We should also worry about novel forms of lobbying. One lesson from the long effort to regulate campaign finance in the United States is that when one outlet for money in politics is closed, another is often found. Some legal scholars liken this dynamic to water always finding its level: "political money, like water, has to go somewhere. It never really disappears into thin air" (Issacharoff and Karlan 1999: 1708). Trillions of dollars are at stake in politics. Interested parties will not sit idly by. Lottocracy may eliminate some avenues for influence, but this may simply lead to the development of new ones.

5.4 *Highly structured procedures*.

Will lottocratic chambers be as highly structured as the mini-publics in our sample? Lottocrats want them to be (e.g., Zakaras 2010: 458; Abizadeh 2021: 802; Guerrero 2024: 160–161), but I am not convinced this is entirely possible. To see why, let's begin with a question: if members of a randomly selected legislature did not adhere to the prescribed procedures—if, for example, they chose to skip the learning phase, invited their own expert witnesses rather than those selected for them, decided to deliberate without moderators present, and so on—who would stop them? In mini-publics, the body can be dissolved, or disruptive members removed, by those who administer it. But it is less clear who would have the authority to do this in a lottocracy. Perhaps the structure and procedures governing randomly selected legislatures would

be written into a constitution and enforced by courts. The constitution might specify that before drafting a bill, members of the legislature must hear from experts, those experts must be selected in a specific way, etc.

A constitution could impose some structure on randomly selected legislatures, but limits remain. First, the constitution cannot determine in advance which issues the legislature must consider, since the pressing political questions of the day cannot be known *ex ante*. Nor can it specify how much time should be devoted to each phase of the process. Some highly technical issues might require extensive expert testimony, while others might call for greater input from affected stakeholders. Decisions of this kind would have to be made as issues arise and are thus inappropriate for constitutional entrenchment.

Second, although courts can ensure that legislators follow certain procedures in a technical sense, they cannot force them to follow the spirit of the procedures. Suppose the ideal procedure requires legislators to hear expert testimony and deliberate before they begin drafting legislation. What if they don't want to? Perhaps a court could sanction the legislature if it called no witnesses. But it probably could not sanction the legislature for calling witnesses and then not paying attention to them. Moreover, how would a court determine whether the legislature deliberated sincerely before making policy? What if legislators wanted to hear additional testimony from experts who were not selected through whatever filtering mechanism was in place? Would a judge invalidate the resulting legislation because they brought in additional experts? What if they heard from those experts off-site, perhaps at a bar? Formal compliance with a procedure is not the same as adhering to its spirit. Courts may be able to enforce the former, but not the latter.

All of this suggests that we should be somewhat skeptical that lottocracy would exhibit the same rigid structure found in mini-publics. Some features of that structure simply could not be reproduced at all, such as prespecifying the questions participants must consider. Other features could in principle be reproduced—such as a designated learning phase or specific way of selecting experts—but enforcing these procedures would be difficult if legislators did not wish to comply. We should thus understand many of the structural proposals offered by defenders of lottocracy as closer to recommendations than binding constraints. Some may be present; others may not.

5.5 *Piloting bias*.

Mini-publics may benefit from especially attentive facilitators. Will they be present in lottocracy? I'm skeptical, for three reasons.

First, the presumption should be that whatever benefits mini-publics derive from piloting bias will weaken in the transition to lottocracy, for that is what often happens in the real world. A large literature finds that interventions with substantial effects in pilot studies often produce smaller effects when scaled up (Al-Ubaydli et al. 2017; Banerjee et al. 2017: 82–83; Vivalt 2020; DellaVigna and Linos 2022). The pattern is familiar: academics run a small study and find a large effect; the government adopts the intervention on a large scale, only to find a smaller one. Unless we are given a compelling reason to believe that scaling mini-publics into lottocracy will escape this pattern, we should expect it to hold.

Second, it may not be politically feasible for the unusually motivated facilitators found in our sample mini-publics to occupy the same positions in lottocracy. Facilitators exercise

immense influence over how deliberative bodies proceed. They set the rules, select experts, structure deliberation, moderate discussion, and make countless procedural decisions. How they exercise this authority can impact outcomes. And so, we should expect political backlash if those exercising this power appear partisan. Such backlash was absent in mini-publics because the stakes were so low that few had reason to care. But things will be different in lottocracy, where the stakes are immense.

Many of the academics and non-profit professionals who helped run mini-publics would probably appear partisan to the average citizen. Do we really expect people to accept a professor of philosophy or politics, who has published controversial views, enforcing the rules of the lottocratic chamber? I imagine there will be strong pressure for lottocratic chambers to be administered by neutral career civil servants, lest the process appear tainted. Instead of academics who write about deliberative democracy, heads of progressive think tanks, or architects of international climate agreements, facilitators in lottocracy will probably more closely resemble IRS employees.

Finally, even if lottocracy employs highly motivated facilitators, their motivation will not last. Helping facilitate a mini-public is a major event in an academic's career, especially for someone who studies deliberative democracy. Such opportunities are rare. If you are lucky, you may get it once in a lifetime. We should not be surprised if those involved approach the task with unusual care and enthusiasm. But some of that care and enthusiasm will surely fade once facilitating a lottocratic chamber becomes your day job. What is exciting when done once in a lifetime becomes drudgery when done five days a week for thirty years. Thus, even if the same kinds of people who facilitate mini-publics also facilitate lottocracy—and I seriously doubt that

will prove acceptable—the piloting bias generated by their unusual motivation may nonetheless dissipate.

Candidate causal factor	Likelihood of transfer to lottocracy
Epistemic benefits of deliberation	High
Selection bias	Low
Narrow policy agenda	Low
Absence of lobbying	Low
Highly structured procedures	For some procedures, low; for others, moderate
Piloting bias	Low

Table 1

5.6 *Summary and analysis*.

Table 1 summarizes the analysis. If selection bias, narrow policy agendas, the absence of lobbying, or piloting bias turn out to be actual causal factors, then the transfer condition will not be satisfied. Highly structured procedures are more complicated. Some procedures found in mini-publics, such as prespecified questions, cannot be present in lottocracy. Others in principle could, but it is hard to say whether they will. A more complete analysis might individuate these procedures and assess separately the likelihood they will be present in lottocracy. For now, we can say that, if highly structured procedures turn out to be an actual causal factor, then we should have doubts about the transfer condition. The only candidate causal factor that would not create a

transfer problem is the epistemic benefits of deliberation, which I granted would be present in lottocracy *arguendo*.

What, then, can we conclude? We cannot determine whether the transfer condition is satisfied until we know which candidate causal factors are actual. We can nevertheless say that the prospects are not good. Several candidate causal factors are unlikely to be present in lottocracy. If any of them turn out to be actual causal factors, then the transfer condition will not be satisfied, and the IMP will not achieve minimal warrant. The prior section established that the IMP is currently not minimally warranted because proponents have not established the articulation condition. The analysis in this section gives us reason to worry that satisfying that condition may not do much good.

5.7 *Extensions*.

Thus far, the analysis has understood lottocracy as a political system identical to ordinary representative democracies, except all legislators are selected by lot rather than election. Yet this definition obscures considerable variation among lottocratic proposals. Moreover, scholars have proposed novel uses of random selection that do not involve substituting elected legislators at all. Some examples:

- Many think we should replace only one chamber of a bicameral legislature with randomly selected citizens, leaving the other chamber elected (Gastil and Wright 2019; Abizadeh 2021). Among those who adopt this view, there is disagreement about the role of the randomly selected chamber. Some think it should have only review power (Zakaras

2010: 457). Others think it should have the power to initiate legislation (Vandamme and Verret-Hamelin 2017: 2).

- Guerrero (2024: 151–153) proposes replacing a single generalist legislature with multiple specialized ones. These single-issue legislatures (SILLs) would focus on distinct policy areas—one legislature for agriculture and food safety, another for housing, another for foreign affairs, and so on. He recommends establishing twenty to thirty permanent bodies overall (Guerrero 2024: 238).
- Some advocate using random selection to constitute temporary bodies that review, and potentially veto, specific policies or regulatory rulings (Bagg 2024: 100; Benson 2024: 220).
- Some recommend "constituency juries": bodies of randomly selected citizens drawn from an elected politician's district charged with overseeing the politician's performance in office (Leipold 2025).
- Some advocate using randomly selected bodies in a merely advisory role: to propose policies to legislatures or the public (Setälä 2017: 853). This, however, would be done more systematically (i.e., they would be constituted at regular intervals), rather than through the *ad hoc* initiatives we see today.

Many other uses of sortition have been proposed. Indeed, something of a cottage industry has emerged in political philosophy devoted to finding new uses for random selection. I cannot list or discuss them all. The point I want to make is that, compared to lottocracy as I initially defined it, some of these alternative arrangements are more likely to include the possible causal factors we've been focusing on. And so, some of these proposed arrangements are more likely to satisfy the transfer condition.

Consider Guerrero's SILL proposal. If lottocracy takes this form, the narrow agendas found in mini-publics do seem to transfer over to the target case. Indeed, that's the very point of SILLs: legislators can focus all their attention on a single policy area. The French Citizens' Convention for Climate, which focused on a single policy area but issued 149 recommendations, offers a reasonable approximation of what we might expect from a SILL: many proposals, but within a narrow domain.

Or consider the use of many short-term, one-off randomly selected chambers to review proposed bills. These, too, seem to transfer over the narrow agendas found in mini-publics to the target case. Jonathan Benson (2024: 223) also argues that such arrangements reduce the threat of lobbying. Because there are so many of these chambers and because they meet for such a short period of time, the expected payoff of lobbying any one of them is small, making special interests less likely to expend resources to do so. We might further add that if these assemblies have only veto power—that is, if they cannot add novel provisions to bills—the potential benefits of lobbying are further reduced. Only some groups—those opposed to the bill—would have incentive to lobby. Those seeking to alter the bill would have no incentive. Thus, a case can be made that this use of sortition transfers over both narrow agendas and the absence of lobbying to the target case.

Indeed, we can imagine a political system identical to modern democracies but with advisory mini-publics that are given a more permanent role. Rather than being used only on an *ad hoc* basis, mini-publics could be formed regularly, with their proposals submitted to the public in referendums. The bodies themselves have no power to make binding decisions. Depending on how the details are worked out, all the potential causal factors we've been focusing on may transfer to the target case.

So, although the general conclusion of §5 is pessimistic, there are grounds for qualified optimism: alternative versions of lottocracy may be more likely to satisfy the transfer condition. Two qualifications are important, however. First, we cannot conclude that *any* particular form of lottocracy satisfies the transfer condition until we know which causal factors explain the success of the mini-publics in our sample. The articulation condition always comes first. Second, although some variants of lottocracy may include more of the causal factors we've been considering, I doubt that any of them—with the exception of the variant discussed in the preceding paragraph—will include *all* of them. Relevant causal factors, moreover, may have escaped our analysis. The pessimistic conclusion of this section should therefore be tempered, but not abandoned.

6. Weakening the IMP

The IMP says that the successful performance of mini-publics provides substantial evidence that ordinary citizens can competently perform the tasks that lottocracy calls on them to perform. Perhaps those lottocrats who appeal to mini-publics when confronted with competence objections are not actually making this strong claim, though. Perhaps they are actually making a weaker claim: not that ordinary citizens are competent enough for lottocracy to work well, but only that they possess some kind of general competence. This leads to a weaker version of the IMP:

> *Weak Inference from Mini-Publics (WIMP)*: The successful performance of mini-publics provides substantial evidence that ordinary citizens, in some settings, can competently deliberate about public policy.

Is this inference warranted? Note first that it does not resemble an external validity inference: it does not claim that a result will travel from a source to target case. Rather, the WIMP resembles something more like a possibility proof: mini-publics show that ordinary citizens can deliberate competently, at least in some cases, *contra* those skeptics who doubt their capacity to do so. Because the WIMP does not claim that a result will travel, the articulation and transfer conditions are not relevant to its evaluation.

Whether the WIMP is warranted depends on who "ordinary citizens" refers to. If the term refers to something like a random sample of the public, then the WIMP is unwarranted. As I have shown, there are compelling reasons to think mini-publics are subject to selection bias. We know that participants are far more educated than the public. There may also be positive selection on unobservable traits. Thus, if "ordinary citizens" refers to something like a random sample, then we cannot infer from mini-publics that ordinary citizens can deliberate competently, for we have reason to doubt that mini-public participants are ordinary in the relevant sense. But if "ordinary citizens" refers to something like "people who do not work professionally in politics," then the WIMP is warranted, for mini-public participants fit that description. From mini-publics we *can* infer that people with little professional experience in politics can deliberate well about policy, at least in some settings.

Beyond stating that ordinary citizens have certain capacities in certain settings, what is the significance of the WIMP? It does not respond to competence objections to lottocracy. Competence objections worry that ordinary citizens lack traits needed to legislate well, at least when compared to professional politicians. But showing that ordinary citizens are competent in some settings does not show that they will be competent in a legislature. And so, the WIMP does not answer competence objections.

The WIMP, however, does undermine more general concerns about citizen competence. Consider the following claim:

> *Unconditional Elitism*: In general, most citizens lack the ability to competently deliberate about public policy.

The WIMP undermines this claim, but it's not clear who actually defends it. Landemore and Pénigaud (2026: 1), when reviewing those who worry about citizen competence, cite Joseph Schumpeter, Jason Brennan, and Bryan Caplan, among others. These thinkers do not defend Unconditional Elitism. Indeed, Brennan (2016: 35) and Caplan (2007: 123) are clear that the incompetence we find in mass electorates is a function of the incentives imposed by the environment. Because one's vote does not make a difference, it is rational to be uninformed and/or irrational. Schumpeter (2008: 262, emphasis mine) does not state the logic of rational ignorance (the idea would be developed by Anthony Downs a bit later), but he does say that "the typical citizen drops down to a lower level of mental performance *as soon as he enters the political field*." I'm not saying that Unconditional Elitism has no defenders, only that many of the usual suspects do not defend it.

Many who worry about the competence of ordinary citizens in democracies advance something like:

> *Conditional Elitism:* Under the typical conditions of democratic electoral politics, most citizens are unable to deliberate competently about public policy.

The WIMP does not undermine this claim, for the specific settings in which it finds ordinary citizens competent are distinct from the typical conditions of electoral politics. In particular, mini-public participants deliberate in settings characterized by a high level of decisiveness,

among other features. And, we know, a high level of decisiveness in a deliberative body such as a mini-public or a legislature can induce participants to deliberate carefully when they otherwise would not (Landa and Pevnick 2025a: ch. 7).

So, lottocrats can advance the WIMP instead of the IMP. Depending on how we interpret "ordinary citizens," the WIMP may be warranted. It does not, however, undermine competence objections to lottocracy. It can undermine Unconditional Elitism, but few actually embrace that claim. Many who worry about citizen competence in democracies instead advance Conditional Elitism. The WIMP does not undermine it.

7. Conclusion

Mini-publics have performed admirably. But from this it does not follow that lottocracy would perform admirably. That inference is warranted only if we have a well-supported account of why mini-publics succeed and good reason to believe that the relevant causal factors would also be present in lottocracy. At present, neither condition is satisfied. We do not know whether mini-publics succeed because of the epistemic benefits of deliberation, selection bias, their narrow agendas, the absence of lobbying, their highly structured procedures, piloting bias, some other causal factor not considered here, or some combination. Moreover, several of the candidate causal factors considered in this paper are unlikely to survive the transition to lottocracy, although different versions of lottocracy perform to varying degrees across this desideratum. The IMP is therefore unwarranted, but I did note that mini-publics do support more circumscribed claims about citizen competence.

The central claim of this paper does not show that lottocracy would fail. Nor does it show that ordinary citizens are incompetent to govern. It shows only that the successful performance of mini-publics does not provide evidence to the contrary. Further research may eventually identify the causal factors responsible for successful mini-public performance, and these factors may end up transferring to lottocracy. But until then, mini-publics should not be treated as proof of concept for lottocracy, as many claim.

The broader lesson of this paper extends beyond mini-publics and lottocracy. Political philosophers often defend untried institutions by pointing to existing institutions that resemble them in certain respects. But resemblance is not enough. Before drawing lessons from an existing institution, we must understand why it produced the desired result. We must also know whether those causal factors would also be present in the novel institution. The articulation and transfer conditions provide a framework for answering these questions. They also provide a reason for greater caution when moving from institutions that have worked somewhere to institutions that have never existed anywhere.

Works Cited

Abizadeh, Arash. 2021. "Representation, Bicameralism, Political Equality, and Sortition: Reconstituting the Second Chamber as a Randomly Selected Assembly." *Perspectives on Politics* 19: 791-806.

Allcott, Hunt. 2015. "Site Selection Bias." *Quarterly Journal of Economics* 130: 1117-1165.

Allen, Peter. 2018. *The Political Class*. Oxford University Press.

Al-Ubaydli, Omar, John A. List, Danielle LoRe, and Dana Suskind. 2017. "Scaling for Non-Economists: Lessons from the Non-Adherence Problem in the Medical Literature." *Journal of Economic Perspectives* 31: 125-144.

Arnold, Tom, David M. Farrell, and Jane Suiter. 2019. "Lessons from a Hybrid Sortition Chamber: The 2012-2014 Irish Constitutional Convention." In *Legislature by Lot*: 101-121. Verso Books.

Averchenkova, Alina, Arnaud Koehl, and Graham Smith. 2025. "Policy Impact of the French Citizens' Convention for the Climate: Untangling the Fate of the Citizens' Recommendations." *KNOCA Briefing No. 13*. URL: https://www.knoca.eu/briefings/policy-impact-of-the-french-citizens-convention-for-the-climate-untangling-the-fate-of-the-citizens-recommendations.

Bagg, Samuel. 2024. "Sortition as Anti-Corruption: Popular Oversight against Elite Capture." *American Journal of Political Science* 68: 93-105.

Banerjee, Abhijit, Rukmini Banerji, James Berry, Esther Duflo, Harini Kannan, Shobhini Mukerji, Marc Shotland, and Michael Walton. 2017. "From Proof of Concept to Scalable Policies: Challenges and Solutions, with an Application." *Journal of Economic Perspectives* 31: 73-102.

Bell, Daniel A. 2015. *The China Model*. Princeton University Press.

Benson, Bruce L. 1990. *The Enterprise of Law*. Pacific Research Institute for Public Policy.

Benson, Jonathan. 2024. *Intelligent Democracy*. Oxford University Press.

Besley, Timothy and Maitreesh Ghatak. 2018. “Prosocial Motivation and Incentives.” *Annual Review of Economics* 10: 411-438.

Bohrnstedt, George W. and Brian M. Stecher. 2002. *What We Have Learned about Class Size Reduction in California*. California Department of Education.

Bold, Tessa, Mwangi Kimenyi, Germano Mwabu, Alice Ng’ang’a, and Justin Sandefur. 2018. “Experimental Evidence on Scaling Up Education Reforms in Kenya.” *Journal of Public Economics* 168: 1-20.

Bosworth, William, Keith Dowding, and Adriano Giuliani. 2026. “Sortition Within Political Parties.” *Political Studies* 74: 1293-1309.

Bouricius, Terry. 2026. *Democracy Without Politicians*. Routledge.

Bowles, Samuel. 2016. *The Moral Economy*. Yale University Press.

Brennan, Jason. 2016. *Against Democracy*. Princeton University Press.

Bureau of Labor Statistics. 2025. “Usual Weekly Earnings of Wage and Salary Workers: First Quarter 2025.” URL: https://www.bls.gov/news.release/wkyeng.htm.

Caplan, Bryan. 2007. *The Myth of the Rational Voter*. Princeton University Press.

Cartwright, Nancy and Jeremy Hardie. 2012. *Evidence-Based Policy*. Oxford University Press.

Citizens’ Assembly. 2017. *First Report and Recommendations of the Citizens’ Assembly: The Eighth Amendment of the Constitution*. 29 June 2017. URL: https://citizensassembly.ie/wp-content/uploads/2023/02/FirstReport_EIGHTAMENDMENT.pdf.

Collier, David. 1993. “The Comparative Method.” In *Political Science: The State of the Discipline II*, edited by Ada F. Finifter: 105-119. Washington, DC: American Political Science Association.

Convention on the Constitution. 2014. *Ninth Report of the Convention on the Constitution: Conclusions and Final Recommendations*. March 2014. URL: https://citizensassembly.ie/wp-content/uploads/Conclusions-and-final-recommendations.pdf.

Courant, Dimitri. 2021. “Citizens’ Assemblies for Referendums and Constitutional Reforms: Is There an ‘Irish Model’ for Deliberative Democracy?” *Frontiers in Political Science* 2: 591983.

Deaton, Angus. 2010. “Instruments, Randomization, and Learning about Development.” *Journal of Economic Literature* 48: 424-455.

DellaVigna, Stefano and Elizabeth Linos. 2022. “RCTs to Scale: Comprehensive Evidence from Two Nudge Units.” *Econometrica* 90: 81-116.

Farrell, David M., Eoin O’Malley, and Jane Suiter. 2013. “Deliberative Democracy in Action Irish-Style: The 2011 *We the Citizens* Pilot Citizens’ Assembly.” *Irish Political Studies* 28: 99-113.

Farrell, David M. and Peter Stone. 2020. “Sortition and Mini-Publics: A Different Kind of Representation.” In *The Oxford Handbook of Political Representation in Liberal Democracies*, edited by Robert Rohrschneider and Jacques Thomassen: 228-246. Oxford University Press.

Farrell, David M., Jane Suiter, and Clodagh Harris. 2019. "'Systematizing' Constitutional Deliberation: The 2016-2018 Citizens' Assembly in Ireland." *Irish Political Studies* 34: 113-123.

Farrell, David M., Jane Suiter, Clodagh Harris, and Kevin Cunningham. 2021. "Ireland's Deliberative Mini-Publics." In *The Oxford Handbook of Irish Politics*, edited by David M. Farrell, Niamh Hardiman, Eoin Carolan, John Coakley, Aidan Regan, Colin Scott, Ben Tonra, and Alexa Zellentin: 627-643. Oxford University Press.

Findley, Michael G., Kyosuke Kikuta, and Michael Denly. 2021. "External Validity." *Annual Review of Political Science* 24: 365-393.

Fournier, Patrick, Henk van der Kolk, R. Kenneth Carty, André Blais, and Jonathan Rose. 2011. *When Citizens Decide: Lessons from Citizens' Assemblies on Electoral Reform*. Oxford: Oxford University Press.

Frey, Bruno S. and Reto Jegen. 2001. "Motivation Crowding Theory." *Journal of Economic Surveys* 15: 589-611.

Frey, Bruno S., Felix Oberholzer-Gee, and Reiner Eichenberger. 1996. "The Old Lady Visits Your Backyard: A Tale of Morals and Markets." *Journal of Political Economy* 104: 1297-1313.

Fryer, Roland G. 2013. "Teacher Incentives and Student Achievement: Evidence from New York City Public Schools." *Journal of Labor Economics* 31: 373-407.

Gastil, John and Erik Olin Wright. 2019. *Legislature by Lot*. Verso Books.

Gaus, Gerald. 2016. *The Tyranny of the Ideal*. Princeton University Press.

Giraudet, Louis-Gaëtan, Bénédicte Apouey, Hazem Arab, Simon Baeckelandt, Philippe Bégout, Nicolas Berghmans, Nathalie Blanc et al. 2022. "'Co-construction' in Deliberative Democracy: Lessons from the French Citizens' Convention for Climate." *Humanities and Social Sciences Communications* 9: 207.

Gneezy, Uri and Aldo Rustichini. 2000. "A Fine is a Price." *Journal of Legal Studies* 29: 1-17.

Guala, Francesco. 2003. "Experimental Localism and External Validity." *Philosophy of Science* 70: 1195-1205.

Guerrero, Alexander. 2024. *Lottocracy: Democracy Without Elections*. Oxford University Press.

Hall, Richard L. and Alan V. Deardorff. 2006. "Lobbying as Legislative Subsidy." *American Political Science Review* 100: 69-84.

Hennig, Brett. 2017. *The End of Politicians.* Unbound.

Hossain, S.M. Moazzem, Arabella Duffield, and Anna Taylor. 2005. "An Evaluation of the Impact of a US$60 Million Nutrition Programme in Bangladesh." *Health Policy and Planning* 20: 35-40.

Issacharoff, Samuel and Pamela Karlan. 1999. "The Hydraulics of Campaign Finance Reform." *Texas Law Review* 77: 1705-1738.

Jacquet, Vincent. 2017. "Explaining Non-Participation in Deliberative Mini-Publics." *European Journal of Political Research* 56: 640-659.

Kogelmann, Brian. 2026. "Sortition and Cognitive Ability." *Politics, Philosophy & Economics* (Online First).

Kübler, Daniel and Philippe E. Rochat. 2025. "Nonparticipation in Democratic Assemblies: Factors, Reasons and Suggestions." *European Political Science Review* 17: 569-588.

Landa, Dimitri and Ryan Pevnick. 2021. "Is Random Selection a Cure for the Ills of Electoral Representation?" *Journal of Political Philosophy* 29: 46-72.

Landa, Dimitri and Ryan Pevnick. 2025a. *Representative Democracy*. Oxford University Press.

Landa, Dimitri and Ryan Pevnick. 2025b. "Should We Be Lottocrats?" *Free & Equal* 1: 424-444.

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

Landemore, Hélène. 2020a. *Open Democracy*. Princeton University Press.

Landemore, Hélène. 2020b. "When Public Participation Matters: The 2010-2013 Icelandic Constitutional Process." *International Journal of Constitutional Law* 18: 179-205.

Landemore, Hélène. 2026. *Politics Without Politicians*. Thesis.

Landemore, Hélène and Théophile Pénigaud. 2026. "Citizen Legislators: Citizens on Top and Experts on Tap in the French Convention for Climate (2019-20)." *Perspectives on Politics* (First View).

Lang, Amy. 2007. "But Is It for Real? The British Columbia Citizens' Assembly as a Model of State-Sponsored Citizen Empowerment." *Politics & Society* 35: 35-69.

Leipold, Bruno. 2025. "Constituency Juries: Holding Elected Representatives Accountable through Sortition." *Perspectives on Politics* 24: 1077-1095.

Lohmann, Susanne. 1998. "An Information Rationale for the Power of Special Interests." *American Political Science Review* 92: 809-827.

Luskin, Robert C. 1990. “Explaining Political Sophistication.” *Political Behavior* 12: 331-361.

Marchionni, Caterina and Samuli Reijula. 2019. “What Is Mechanistic Evidence, and Why Do We Need It for Evidence-Based Policy?” *Studies in the History and Philosophy of Science Part A* 73: 54-63.

Morton, Rebecca and Charles Cameron. 1992. “Elections and the Theory of Campaign Contributions.” *Economics and Politics* 4: 79-108.

Robling, Michael, Marie-Jet Bekkers, Kerry Bell, Christopher C. Butler, Rebecca Cannings-John, Sue Channon, Belen Corbacho Martin et al. 2016. “Effectiveness of a Nurse-Led Intensive Home-Visitation Programme for First-Time Teenage Mothers (Building Blocks): A Pragmatic Randomised Controlled Trial.” *The Lancet* 387: 146-155.

Schumpeter, Joseph. 2008. *Capitalism, Socialism and Democracy*. Harper Perennial.

Schwartzberg, Melissa. 2015. “Epistemic Democracy and its Challenges.” *Annual Review of Political Science* 18: 187-203.

Setälä, Maija. 2017. “Connecting Deliberative Mini-Publics to Representative Decision Making.” *European Journal of Political Research* 56: 846–863.

Shoemaker, Eric. 2024. “Democratic Equality Requires Randomly Selecting Legislators.” *Public Affairs Quarterly* 38: 132-152.

Skoble, Aeon. 2026. *Deleting the State*. Independent Institute.

Steel, Daniel P. 2008. *Across the Boundaries*. Oxford University Press.

Titmuss, Richard M. 1997. *The Gift Relationship*. The New Press.

Vandamme, Pierre-Étienne and Antoine Verret-Hamelin. 2017. “A Randomly Selected Chamber: Promises and Challenges.” *Journal of Public Deliberation* 13: 1-24.

Van Reybrouck, David. 2018. *Against Elections*. Seven Stories Press.

Vivalt, Eva. 2020. “How Much Can We Generalize from Impact Evaluations?” *Journal of the European Economic Association* 18: 3045-3089.

Waldner, David. 2026. *Qualitative Causal Inference and Explanation*. Cambridge: Cambridge University Press.

White, Howard and Edoardo Masset. 2007. “Assessing Interventions to Improve Child Nutrition: A Theory-Based Impact Evaluation of the Bangladesh Integrated Nutrition Project.” *Journal of International Development* 19: 627-652.

Widerquist, Karl. 2024. *Universal Basic Income*. MIT Press.

Wright, Erik Olin. 2010. *Envisioning Real Utopias*. Verso Books.

Zakaras, Alex. 2010. “Lot and Democratic Representation: A Modest Proposal.” *Constellations* 17: 455-471.