Local Evidence and Diversity in Minipublics*

Arjada Bardhi† Nina Bobkova‡

January 16, 2021

Abstract

We study optimal minipublic design with endogenous evidence. A policymaker selects a group of citizens—a minipublic—for advice on the desirability of a policy. Citizens can discover local evidence but might be deterred by uncertainty about the policymaker's adoption standard. We show that such uncertainty can be detrimental to evidence discovery even with costless evidence, civic-minded citizens, and ex ante aligned players. Evidence discovery is hardest to sustain under moderate uncertainty. The optimal minipublic has low diversity: it overrepresents citizens around the median citizen and underrepresents those at the margins. Our findings bear implications for the French Citizens' Convention on Climate.

JEL: D71, D72, D83

Keywords: minipublic, evidence discovery, political uncertainty, informational diversity, demographic diversity

1 Introduction

Evaluating the impact of a novel public policy is a complex task, especially in diverse societies. Citizens of different socio-economic backgrounds are affected by the policy in varied, far-reaching, and uncertain ways. What is more, the evidence critical for evaluation is often quintessentially local: lay citizens can claim privileged insight into how the policy is likely to impact them and fellow

*We thank Attila Ambrus, Alina Arefeva, Steven Callander, Gent Carrabregu, Alessandra Casella, Miaomiao Dong, Laura Doval, Hülya Eraslan, Marina Halac, Daniel Hauser, Dimitri Landa, Stephan Lauermann, Jacopo Perego, Anja Prummer, João Ramos, Kareen Rozen, Bruno Strulovici, Leeat Yariv, Huseyin Yildirim, and seminar audiences at Princeton, Notre Dame, Columbia, Manchester, Queen Mary, UTexas Austin, UToronto, and conference audiences at the Bonn CRC Junior Workshop on Economic Theory, SEA Annual Meeting, and SITE Political Economic Theory Workshop for valuable feedback. Bardhi gratefully acknowledges financial support from the National Science Foundation (SES-1919329).

†Department of Economics, Duke University; email: arjada.bardhi@duke.edu.

‡Department of Economics, Rice University; email: nina.bobkova@rice.edu.

citizens of similar background.[1] Accordingly, policymakers have to rely on citizens' willingness to discover and bring forth local evidence, which in turn feeds into the policy decision. Indeed, over the last few decades, policymakers have actively sought ways to induce higher citizen participation in evidence-based policymaking (Michels and De Graaf (2010), Jacobs and Kaufmann (2019)).[2]

This paper studies a policy environment in which (i) a small number of citizens is targeted by the policymaker to contribute their local evidence, and (ii) citizens face prior uncertainty about the standard that the policymaker follows when making the policy decision. First, we examine how the presence of uncertainty about how demanding the policymaker will be toward the policy—which we refer to as *political uncertainty*—distorts citizens' incentives to contribute evidence. Are citizens less willing to contribute evidence in the presence of greater political uncertainty? Second, we analyze the policymaker's problem of assembling a group of citizens—which, for reasons that become clear below, we refer to as a *minipublic*—to learn as much as possible about the policy in question. To what extent is the optimal minipublic demographically representative of the citizenry at large? How much evidence does it endogenously produce? Does the policymaker target citizens from more diverse backgrounds in the presence of greater political uncertainty?

Our framework is motivated by the increased popularity of a particular form of citizen participation: minipublics.[3] Conceptualized first by Dahl (1989), a minipublic is a small group of ordinary citizens selected from the entire citizenry, tasked with gathering information on an issue of public interest in order to inform policymaking. A recent illustrious minipublic is the Citizens' Convention on Climate (CCC), implemented in France between October 2019 and June 2020. The CCC was an initiative of French President Emmanuel Macron to engage citizens in formulating France's climate policy. It consisted of 150 citizens, drawn by lot and representative of the French population along six criteria: gender, age, education, occupation, residence, and geographical area. Strikingly, political uncertainty accompanied the CCC throughout its proceedings. At the onset, Macron vowed that the citizens' policy recommendations would be enacted "without filter." In January 2020, he invoked provisions of the French Constitution to narrow down the meaning of "without filter" by reserving the option to not act upon some of the recommendations. In June 2020, after the CCC put forth its

[1] This is especially the case for policy decisions that involve normative and societal values, rather than mere technical knowledge, and that entail complex tradeoffs, e.g., tradeoffs across generations of citizens, as in the case of climate policy.

[2] Such efforts have taken many forms, from collaborative governance, citizens' advisory committees, and participatory budgeting, to deliberative bodies of citizens such as consensus conferences, deliberative polls, citizen juries, and planning cells. In 2017, the OECD launched the Innovative Citizen Participation Project, aiming to record and analyze all cases worldwide of deliberative, collaborative, and participatory decision making.

[3] Since the early 1990s, minipublics, as an innovative form of citizen participation, have been drawing increased attention from policymakers and academics alike. Breckon, Hopkins and Rickey (2019) reviews case studies of minipublics around the world.

149 recommendations, Macron said he would not pass along three of the recommendations. There is ongoing concern among the participants of the CCC that even more recommendations might not be passed along, and the uncertainty about Macron's response is heightened by the ongoing pandemic.[4] A key takeaway of this example is that political uncertainty is a first-order consideration when the minipublic's role is strictly advisory.

Our main insight is that political uncertainty can be detrimental to the production of evidence within small minipublics and leads to the policymaker choosing minipublics that are of inefficiently low diversity. We formalize this insight in a setting that is arguably conducive of more, rather than less, evidence production: (i) the policymaker and the citizens agree ex ante on the decision threshold to be followed; (ii) minipublic citizens are not self-interested, i.e., just like the policymaker, they evaluate the policy based on its average outcome across all citizens; and (iii) evidence is costless to discover. The policymaker controls the citizens' incentives to discover evidence through designing the size and the composition of the minipublic. We first propose a benchmark for the ideal minipublic of a given size—the one that maximizes the welfare of both the policymaker and the citizens in the absence of political uncertainty—which is also the most informative minipublic of that size. Then we show that when this ideal minipublic is unfeasible in the presence of political uncertainty, the policymaker's only instrument for inducing citizens to discover evidence is to make them less marginally informative within the minipublic. This diminishes the diversity of the minipublic below the level of the ideal minipublic: the policymaker ends up overly informed about citizens around the median citizens and rather uninformed about more peripheral citizens.

Even though our results speak directly to the optimal design of minipublics, the framework applies beyond this specific application. The assignment of members to legislative committees, the appointment of faculty to university committees, and the drawing by lot of the Council of 500 (the Boule) in ancient Athens are other natural applications of our framework. We essentially study the design of an advisory committee in which the designer and the committee members share a common value for a multi-attribute object that is available for adoption. Decision-making power rests with the designer rather than the committee. Each member is an expert on a single attribute and chooses whether to discover evidence about it. Therefore, by choosing members, the designer chooses which attributes to learn about.

Framework. The model features a unit mass of citizens and a single policymaker. A citizen's position in this unit mass indexes his demographic background. The policymaker chooses a finite

[4] For early coverage of the CCC, see https://www.theguardian.com/world/2020/jan/10/citizens-panels-ready-help-macron-french-climate-policies. On the criticism that Macron is facing about his lack of commitment, see https://www.france24.com/en/europe/20201207-facing-criticism-macron-defends-his-government-s-climate-policies.

set of citizens, a minipublic, from this unit mass. She is constrained by a maximal capacity for the minipublic: she can costlessly target up to, but no more than, n citizens. The game proceeds in three stages. First, the policymaker chooses a minipublic. Second, each minipublic citizen observes who else has been chosen in the minipublic and then decides whether to discover his local evidence about the policy. Third, the policymaker observes the citizens' evidence and decides whether to adopt the policy.

The local evidence of citizen i consists of his realized policy outcome $\beta(i)$. Each citizen can become an "expert" on the impact of the policy for citizens from his demographic background only. Moreover, if the citizens discover evidence, it becomes publicly observable—that is, evidence is fully transparent and the players are symmetrically informed at all times. This structure captures in a stylized way the intricate processes of information gathering and deliberation that take place in real-world minipublics. We model policy outcomes as given by the sample path of an Ornstein-Uhlenbeck process, which has several appealing properties. First, all policy outcomes are positively correlated. Second, citizens are ex ante equally uncertain about their policy outcomes. Third, citizens who are closer in the unit interval have more strongly correlated outcomes. In this sense, distance captures the similarity of the citizens' backgrounds.

Put simply, our game is one of local evidence but global value for the policy. That is, both the policymaker and the minipublic citizens care about the average outcome of the policy across all citizens. The only friction between the citizens and the policymaker is due to a random threshold of adoption, which captures implementation costs that only the policymaker incurs. This threshold, whose expected value is zero, is realized after the policymaker chooses the minipublic but before she makes a decision. The more variable this threshold is, the greater is the political uncertainty that the citizens face.

Main results. Our first contribution is to establish that the conflict between the policymaker and the citizens boils down to whether a more informative minipublic is preferred. The analysis identifies a sufficient *informativeness* statistic for any minipublic. The players' expected payoffs from a given minipublic depend only on this statistic, which equals the variance in the posterior value of the policy given the minipublic's evidence. Notably, the policymaker's payoff strictly increases in informativeness: the more informative the evidence provided by the minipublic, the more precise the estimate of the policy's value. Hence, the problem of the policymaker reduces to identifying the most informative minipublic of at most n citizens in which each citizen discovers his evidence.

By contrast, the citizen's payoff is quasiconvex in informativeness. This is due to two opposing effects. On the one hand, a more informative minipublic leads to a better-informed adoption decision, which benefits the citizen. On the other hand, more evidence might make ex post disagreement with the policymaker more likely, in which case the policymaker misuses evidence. When does the possibility of such misuse dominate the citizen's calculus, discouraging him from discovering evidence? We show that two conditions must be met. First, when starting from no information, the citizen's marginal value from an infinitesimally small amount of information must be negative. Second, the total informativeness of the minipublic must be sufficiently low. We refer to these two conditions in tandem as *the curse of too little information*, which is a key force in this framework.

Disregarding at first the incentives of minipublic citizens to discover evidence, we characterize the unique first-best minipublic. This minipublic has the highest informativeness among all minipublics of size at most n and provides a benchmark for the informativeness-maximizing level of diversity. This characterization is a special case of the single-player attribute sampling characterization in Bardhi (2020).[5] The first-best minipublic consists of exactly n citizens, distributed symmetrically about the median citizen.

Our second contribution is to characterize how and when the optimal minipublic is distorted relative to the first-best one. Whenever some citizens in the first-best minipublic prefer not to discover evidence, the policymaker has two instruments available to incentivize evidence discovery: distortion in minipublic size and distortion in its composition. We establish that the first instrument—distortion in size—is counterproductive for the policymaker. The optimal minipublic, if nonempty, consists of exactly n citizens. If the policymaker is able to incentivize evidence discovery among $n' < n$ citizens, then these citizens must have already escaped the curse of too little information. Hence, adding more citizens to this minipublic strictly improves informativeness, while all citizens—both those in the original minipublic and the added ones—prefer to discover evidence. In light of this observation, a distorted optimal minipublic is either empty or it consists of n citizens, at least some of which are not in the first-best minipublic.

How does the degree of political uncertainty affect the optimal choice of a minipublic? On a cautionary note, our analysis reveals that what discourages evidence discovery and forces the policymaker toward a minipublic other than the first-best one is the presence of moderate, rather than high, political uncertainty. Against a first intuition that greater political uncertainty makes evidence discovery more challenging, we show that the first-best minipublic is optimal if political uncertainty is either sufficiently high or sufficiently low. Two conditions must be met for a citizen to be unwilling to discover evidence: the likelihood of ex post misalignment should be sufficiently

[5]Moreover, our model differs from the strategic attribute sampling framework in Bardhi (2020) in two important aspects. First, we have a continuum of citizens, each of whom can discover only their own outcome, whereas in Bardhi (2020) there is one agent who can discover any outcome. Second, our citizens face political uncertainty, whereas the decision maker in Bardhi (2020) has a known threshold.

high and the citizen's evidence should significantly influence the probability of such misalignment. Both conditions are satisfied *only* when political uncertainty is moderate. Importantly, this implies that higher political uncertainty is not always detrimental to evidence production in minipublics.

Our third contribution consists in formalizing the diversity of a minipublic and establishing that the optimal minipublic is less diverse than the first-best minipublic. To make possible such comparison, we formalize two notions of diversity within a minipublic: *informational diversity* and *demographic diversity.* Informational diversity quantifies the novelty of the evidence of individual citizens relative to the evidence provided by the rest of the minipublic. It is thus based on the marginal informativeness of each citizen's evidence: the more novel the evidence that each individual citizen contributes to the minipublic, the more informationally diverse the minipublic is. Demographic diversity, in contrast, is defined in terms of the pairwise distances between any two citizens in the minipublic: the farther apart the citizens' backgrounds are, the more demographically diverse the minipublic is.

The optimal minipublic is less informationally diverse than the first-best one for any minipublic size. When the first-best minipublic is not feasible, all minipublics in which citizens discover evidence—and hence, the optimal minipublic as well— are strictly less informationally diverse than the first-best one. Citizens contribute less novel evidence in the minipublic. Moreover, we establish that the optimal minipublic is also less demographically diverse for small minipublics ($n = 2$). A small minipublic is a natural setting in which to study demographic diversity, because, as we show, if a distorted minipublic ever arises, it is for a sufficiently small capacity. The optimal minipublic consists of two citizens who, while still symmetric about the median citizen, are closer to each other than the citizens in the first-best minipublic. The optimal minipublic, hence, overrepresents citizens around the median citizen and underrepresents those at the margins, turning into an echo chamber in which the local evidence of minipublic citizens is too correlated.

Practical implications for minipublic design. A recent report by OECD (2020) identifies *policy impact* and *representativeness* as two desiderata for minipublic design, where *impact* means that "the commissioning public authority should publicly commit to responding to or acting on participants' recommendations," and *representativeness* means that "the participants should be a microcosm of the general public." Our analysis, which proxies policy impact by political uncertainty and representativeness by diversity, sheds light on the intricate interaction between these two desiderata.

On the one hand, either very high impact or almost no impact on the policy decision can both sustain evidence discovery in a representative minipublic. In such instances there is no tension between the two desiderata. Arguably, the CCC faced low political uncertainty when it was formed, as it was preceded by the Grand Débat National, a national effort that confirmed Macron's commitment to a greener economy, and as it was backed by his promise to enact the CCC's proposals "without filter." Our findings imply that such low uncertainty might have helped engage citizens in a sufficiently representative CCC.

On the other hand, it is when the minipublic's impact is moderately uncertain that at least part of a representative minipublic—starting with citizens at the margins of the society—self-selects out by not bringing local evidence to the table, ultimately leading to a poorly informed policy decision. Citizens prefer to engage in evidence discovery only if the background of other minipublic citizens is similar to theirs. This implies that if the policymaker can guarantee some, but not sufficient, impact, she has to sacrifice representativeness. The two desiderata must go hand in hand. Our findings further imply that this tension is easier to overcome the more resources the society has to afford more citizens in a minipublic, the more homogeneous the society is, and the weaker is the public sentiment about the policy.

The rest of this section reviews the related literature. Section 2 presents and discusses the model. Section 3 simplifies the policymaker's problem and solves for the first-best minipublic. The characterization of the optimal minipublic is to be found in section 4, whereas section 5 discusses the robustness of our findings to alternative assumptions. Section 6 concludes.

Related literature

Our work connects to several strands of the existing literature. First, it builds on a large literature on the optimal choice of statistical experiments. The evidence discovery game among minipublic citizens relates to models of *Bayesian persuasion with one and multiple senders* (Kamenica and Gentzkow, 2011; Gentzkow and Kamenica, 2016, 2017; Li and Norman, 2018), because citizens do not have private information about the policy outcomes but commit simultaneously and costlessly to experiments. Our focus, however, is on a specific class of experiments that are not Blackwell-connected: each citizen's experiment is either perfectly informative about only his outcome or entirely uninformative.[6] In this respect, our model shares with Koessler, Laclau and Tomala (2018), Boleslavsky and Cotton (2018), and Au and Kawai (2019, 2020) the premise that each sender can design information on only one dimension of a multidimensional state (in our case, infinite-dimensional).[7] As in Au and Kawai (2019), the dimensions in our model are positively correlated.

[6]The experiments are not Blackwell-connected because each citizen can produce information only about his outcome, not those of others. Therefore, our setting is not a special case of Gentzkow and Kamenica (2016).

[7]It also shares this feature with the model of advocacy in Dewatripont and Tirole (1999), where each agent explores only one cause. Our model differs in that evidence is costless and transfers are assumed away.

Despite key modeling differences, one recurring focus that we share with this literature is whether the equilibrium becomes more informative as the number of senders increases.

The policymaker too designs a statistical experiment through her choice of minipublic. In the absence of evidence discovery considerations, her problem is one of *selective learning about multiple correlated attributes*, where each citizen is an attribute and his outcome the attribute realization. The literature on this problem is thin. Our first-best benchmark is a special case of the single-player benchmark in Bardhi (2020). Liang, Mu and Syrgkanis (2020) focuses on gradual learning of finitely many attributes. In these papers, as in ours, attributes are jointly Gaussian. However, our main difference with this line of work is that each attribute is a strategic player in our framework: each attribute needs to be incentivized in order to reveal its realization.

Because each citizen's outcome is predictive of the outcomes of others, minipublic evidence is an example of social data as defined in Bergemann, Bonatti and Gan (2020), which studies data intermediation in a product market with correlated consumer preferences. Despite our vastly different settings, their design of a data policy is similar to our design of a minipublic insofar as consumers must be induced to volunteer their data while being aware of such data externality. As in our analysis, they also leverage the tractability of Gaussian data.

Second, in modeling correlation in policy outcomes through an Ornstein-Uhlenbeck process, our work is methodologically connected to a literature starting with Jovanovic and Rob (1990) and Callander (2011) that adopts *Brownian motion to model uncertainty over a continuum of correlated alternatives* in a search framework. Agents choose which alternatives to explore sequentially so as to identify the best one. Callander and Clark (2017) is closer to us in that it studies the optimal selection of legal cases by a higher court under resource constraints, so as to guide decisions about all possible cases by a lower court. In contrast to the friction between citizens and the policymaker in our model, frictions between the two courts are assumed away. Callander, Lambert and Matouschek (2018) models expert advice over a large space of uncertain decisions through disclosure of hard evidence. In contrast, we model binary policy decisions and commitment in evidence discovery. Notably, both the Ornstein-Uhlenbeck process, which is a better fit for our application, and Brownian motion are special cases of a large class of Gaussian processes, recently explored in Bardhi (2020).

Third, our questions closely relate to those in the literature on the *optimal composition of a team of experts*. What differentiates our paper from this literature is (i) our focus on the citizens' incentives to discover evidence, and (ii) our rich modeling of the correlation between citizens' outcomes as distance. Lamberson and Page (2012) considers the optimal composition of a team of forecasters from one of two statistical groups. In contrast to our model, the group forecast is assumed to weigh individual forecasts equally. Hong and Page (2001) studies the optimal diversity of a problem-solving team, in which agents differ in both perspectives and heuristics. Although our citizens also vary in perspectives, their task is one of evaluation rather than problem-solving. Prat (2002) also studies the optimal diversity of a team in the context of workforce recruitment and complementarities across workers. Chade and Eeckhout (2018) studies sorting of experts into teams in a model in which the precision of an expert's signal is his expertise and the correlation between any two experts is a constant. This is in sharp contrast to our correlation structure, in which each citizen's evidence has the same precision, but the correlation between any two citizens' evidence weakens with distance.

Fourth, our paper shares themes with a large literature on *deliberation in voting committees* and *collective evaluation of multi-attribute proposals*, and in particular on committee models with interdependent values and heterogeneous information (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1997; Visser and Swank, 1999; Moldovanu and Shi, 2013; Gradwohl and Feddersen, 2018; Name Correa and Yildirim, 2020, among others). The presence of private information and voting are two key differences with this line of work. Related work on endogenous information acquisition in committees includes Gerardi and Yariv (2008), Cai (2009), Chan et al. (2018), and Tan and Wen (2020).

Our paper also adds to a vast *social science literature on minipublics* (Dahl, 1989; Fishkin, 2011; Warren and Gastil, 2015), in particular to work that studies non-participation in minipublics (Jacquet, 2017). For a recent model of information acquisition in citizens' assemblies that consist of a single rationally inattentive representative citizen, see Kwiek (2020). To the best of our knowledge, we offer the first formal model of optimal minipublic composition. More broadly, our work relates to a growing economics literature on direct democracy and citizen participation (see, among others, Matsusaka and McCarty (2001), Matsusaka (2005), and Prato and Strulovici (2017)).

2 Environment

2.1 Model

Players. A policymaker ("she") and a unit mass of citizens, each ("he") indexed by $i \in [0,1]$, evaluate the desirability of an uncertain policy. This evaluation takes place within a minipublic. A minipublic is a finite set of distinct citizens, denoted by $\mathbf{m} = \{i_1, i_2, ..., i_n\}$, where $0 \leq i_1 < ... < i_n \leq 1$. The set of all minipublics of size at most n is denoted by $\mathcal{M}_n$.

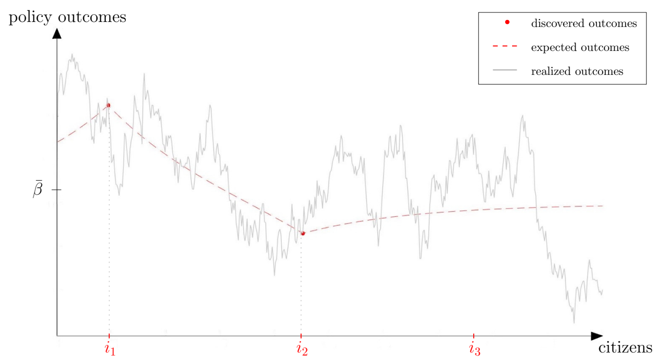

Figure 1: Minipublic $\{i_1, i_2, i_3\}$, in which i_1 and i_2 are active whereas i_3 is passive. The dashed red line depicts how the policymaker extrapolates from minipublic outcomes if $\bar{\beta}(i) = \bar{\beta}$ for all $i \in [0, 1]$.

Policy. The outcome of the policy is given by a randomly drawn mapping $\beta : [0, 1] \to \mathbb{R}$, where $\beta(i)$ denotes the realized policy outcome (hereafter, outcome) for citizen i. This mapping is initially unobserved by all players. The outcomes $\beta(i)$ and $\beta(j)$ of any two citizens i and j are correlated, and the correlation structure is common knowledge. Assumption 1 introduces the distribution over possible mappings from which β is drawn.[8]

Assumption 1 (Distribution of the outcome mapping)**.** *The outcome mapping β is drawn from the space of sample paths of an Ornstein-Uhlenbeck process on* $[0, 1]$*, where*

(i) for each i, $\beta(i) \sim \mathcal{N}(\bar{\beta}(i), 1)$ and $\bar{\beta}(i)$ is integrable on $[0, 1]$*; and*

(ii) for any two citizens $i, j \in [0, 1]$, the correlation between $\beta(i)$ and $\beta(j)$ is given by $e^{-|i-j|/\ell}$, where $\ell \in (0, +\infty)$.

Figure 1 depicts in grey one such realization of the outcome mapping. All citizens' outcomes are equally uncertain ex ante. Moreover, the correlation between any two citizens depends only on the distance between them. For a given fixed distance, the greater ℓ is, the stronger the correlation between the citizens' outcomes. Hence, ℓ measures the degree of *heterogeneity* among citizens. As $\ell \to 0$ (or $\ell \to \infty$), outcomes become almost independent (or almost perfectly correlated), so the citizenry becomes perfectly heterogeneous (or perfectly homogeneous).

Actions and strategies. The game proceeds in three stages: (i) policymaker's choice of a minipublic, (ii) citizens' evidence discovery, and (iii) policymaker's policy adoption. In the first stage, the policymaker chooses a minipublic $\mathbf{m}$ subject to the constraint that $\mathbf{m} \in \mathcal{M}_n$, where

[8] Appendix A presents axioms on the correlation structure which are satisfied if and only if Assumption 1 holds. Our results generalize to richer Gaussian processes, as shown in section 5.

$n \in \mathbb{N}$ is an exogenous capacity constraint on the size of the minipublic. That is, the marginal cost of adding an additional citizen to the minipublic is zero below capacity n, and infinite otherwise. The policymaker's minipublic choice strategy consists of a lottery over feasible minipublics $\Delta(\mathcal{M}_n)$. Each citizen in a realized minipublic $\mathbf{m}$ observes the entire $\mathbf{m}$.

Let $\mathbf{m}$ be the minipublic realized in the first stage. In the second stage, each citizen $i \in \mathbf{m}$ decides whether to costlessly and publicly discover his outcome. If i is *active* in evidence discovery, then $\beta(i)$ is observed by all players. Otherwise, i is *passive*, and $\beta(i)$ remains unobserved by all. Hence, each minipublic citizen has an evidence discovery strategy $\delta_i : \mathcal{M}_n \to \Delta\{0, 1\}$. All discovery decisions are taken simultaneously within $\mathbf{m}$. An *active minipublic* $\hat{\mathbf{m}}$ is a minipublic in which all citizens are active. Figure 1 depicts an example in which $\hat{\mathbf{m}} = \{i_1, i_2\}$ is an active minipublic.

In the third stage, the policymaker decides whether to adopt the policy or keep the status quo. She makes this decision based on $(\hat{\mathbf{m}}, \beta(\hat{\mathbf{m}}))$, where $\beta(\hat{\mathbf{m}}) := \{\beta(i) : i \in \hat{\mathbf{m}}\}$ is the set of discovered outcomes in the active minipublic. We solve for the set of policymaker-preferred Perfect Bayesian equilibria of this game.[9]

Payoffs. For all players, the *value* of the policy equals the average outcome of the policy across the unit mass of citizens:

$$B := \int_0^1 \beta(i)\mathrm{d}i.$$

The payoff of each citizen i is given by B if the policy is adopted and 0 otherwise. The payoff of the policymaker is given by $(B - c)$ if the policy is adopted and 0 otherwise, where $c \sim \mathcal{N}(0, \tau^2)$ is the policymaker's realized threshold of adoption. This threshold is realized at the beginning of the third stage. We refer to the variance τ^2 of this adoption threshold as the degree of *political uncertainty*: the more variable the threshold, the more uncertain the adoption decision of the policymaker.

We let $\bar{B} := \mathbb{E}[B]$ denote the prior value of the policy. This determines the players' ex ante sentiment about the policy: the more positive (negative) $\bar{B}$ is, the more strongly they prefer adopting the policy (keeping the status quo) and the less likely any new evidence is to overturn this ex ante preferred decision. We refer to $|\bar{B}|$ as the degree of *policy sentiment*. High (low) $|\bar{B}|$ indicates strong (weak) policy sentiment.

2.2 Discussion of assumptions

Our stylized model is guided by key features of real-world minipublics, as documented by Elstub (2014) among others. First, a defining feature of a minipublic is that citizens cannot self-select into it, but rather are targeted by the minipublic organizer. The organizer usually aims for the

[9] The focus on policymaker-preferred equilibria is common in the literature on information and mechanism design.

minipublic to be representative—namely, a microcosm—of the larger population. The selection procedure typically consists of stratified sampling from specific demographic groups. For instance, citizen i can be interpreted as belonging to the i^{th} income percentile of the population and the policymaker as selecting the income composition of the minipublic. Second, the primary role of minipublics consists in making nonbinding recommendations to decision makers. Elstub (2014) documents that "[n]ot only do minipublics not make decisions but they have little influence over what they make decisions about because citizens usually have little control over the agendas for minipublics." Third, even though a minipublic is minuscule in size relative to the overall population, as a microcosm of the larger population it can be informative about the population-wide effects of policies under consideration. Fourth, minipublics are intended as a mechanism for producing public evidence about novel policy issues rather than for aggregating existing private information. Simultaneous public learning is a reduced-form way of modeling the deliberation process within the minipublic. This is why we focus on symmetrically informed players.

In our model, the value of the policy is global and common across players, whereas evidence is local and positively correlated. The correlation structure of Assumption 1 has properties that make it a natural choice for modeling correlation across a large population (see appendix A). Each citizen's outcome has the same unit variance ex ante. Therefore, what drives the policymaker's choice of the minipublic is how much she can extrapolate from its local evidence rather than the variability of the outcomes of its citizens. Moreover, the distance between citizens determines how correlated the citizens' outcomes are.[10] In addition, extrapolation from observed outcomes is local: in order to guess the expected outcome of a citizen, the policymaker looks only at the observed outcomes of the most similar citizens.

Our setting assumes civic-minded minipublic citizens who care about the common good rather than their own outcome. More than a perfect description of minipublic citizens, this assumption captures the ideal of a perfectly non-partisan minipublic.[11] By abstracting away from citizens' self-interest, we zoom into the future conflict between the policymaker and the citizenry rather than conflict among citizens.

This possibility of future conflict between the citizens and the policymaker is captured by political uncertainty τ^2. Such uncertainty, which is resolved after the minipublic issues its final report

[10]If β were instead the realization of a Brownian motion, these properties would not hold: the variance of $\beta(i)$ would increase with i, and the correlation between $\beta(i)$ and $\beta(j)$ would depend on both the distance $|i-j|$ and i.

[11]Recent survey evidence presented in Fabre et al. (2020) and at the PSE workshop "The French Citizens' Convention on Climate: A forerunner to future democracy?" suggests that participants in CCC were much more civic-minded than the overall French population. For instance, CCC participants report that they less often "have the feeling that they belong to their neighborhood, village or city" and more often "have the feeling that they belong to the world."

but before the policymaker decides about adoption, captures the time lag that is commonly required for the recommendations of a minipublic to translate into a binding decision. During this time lag, the political priorities of the policymaker might shift or new budgetary considerations might arise, as was arguably the case in the CCC example.

3 Policymaker's problem

This section derives the informativeness of a minipublic, simplifies the policymaker's problem of choosing a minipublic, and characterizes the unique first-best minipublic.

3.1 Minipublic informativeness

Fix a minipublic $\mathbf{m}$ and an evidence discovery strategy profile $\delta := (\delta_i)_{i\in\mathbf{m}}$. This strategy profile induces a lottery over active minipublics $\hat{\mathbf{m}} \subseteq \mathbf{m}$. The policymaker extrapolates from the observed outcomes $\beta(\hat{\mathbf{m}})$ to the rest of the citizenry, as in figure 1, and updates the expected value of the policy from $\bar{B}$ to

$$B_{\hat{\mathbf{m}}} \coloneqq \mathbb{E}\left[\int_0^1 \beta(i)\,\mathrm{d}i | \hat{\mathbf{m}}, \beta(\hat{\mathbf{m}})\right].$$

We refer to $B_{\hat{\mathbf{m}}}$ as the *post-minipublic value.* Passive citizens within the minipublic and citizens outside the minipublic do not enter into the policymaker's updating.

For any active minipublic $\hat{\mathbf{m}}$, the post-minipublic value $B_{\hat{\mathbf{m}}}$ is Gaussian and centered at $\bar{B}$. What varies with $\hat{\mathbf{m}}$—and thus with minipublic choice $\mathbf{m}$—is the variance of the post-minipublic value. The greater this variance, the lower the residual uncertainty about the policy value B. Therefore, minipublics that result in a more variable post-minipublic value are more informative. We refer to the variance of $B_{\hat{\mathbf{m}}}$, denoted by $\sigma^2_{\hat{\mathbf{m}}}$, as the *informativeness of minipublic* $\hat{\mathbf{m}}$. Lemma 3.1 characterizes this minipublic informativeness.

Lemma 3.1 (Informativeness of a minipublic)**.** *Consider an active minipublic* $\hat{\mathbf{m}} = \{i_1, \ldots, i_k\}$, *where* $2 \leqslant k \leqslant n$ *and* $0 \leqslant i_1 < \ldots < i_k \leqslant 1$. *The post-minipublic value is distributed according to* $B_{\hat{\mathbf{m}}} \sim \mathcal{N}\left(\bar{B}, \sigma^2_{\hat{\mathbf{m}}}\right)$, *and the informativeness of minipublic* $\hat{\mathbf{m}}$ *is given by*

$$\sigma^2_{\hat{\mathbf{m}}} = \sum_{j=1}^{k} \gamma_j(\hat{\mathbf{m}})\sigma_{i_j}. \tag{1}$$

Here, $\sigma^2_{i_j} := \ell^2\left(2 - e^{-i_j/\ell} - e^{-(1-i_j)/\ell}\right)^2$ *is the informativeness of the singleton minipublic* $\{i_j\}$, *and* *the weights $\gamma_j(\hat{\mathbf{m}})$ are given by*

$$\gamma_j(\hat{\mathbf{m}}) = \begin{cases} \ell\left(1 - e^{-i_1/\ell} + \tanh\left(\frac{i_2 - i_1}{2\ell}\right)\right) & \text{if } j = 1; \\ \ell\left(\tanh\left(\frac{i_j - i_{j-1}}{2\ell}\right) + \tanh\left(\frac{i_{j+1} - i_j}{2\ell}\right)\right) & \text{if } j \in \{2, \ldots, k-1\}; \\ \ell\left(1 - e^{-(1-i_k)/\ell} + \tanh\left(\frac{i_k - i_{k-1}}{2\ell}\right)\right) & \text{if } j = k. \end{cases} \tag{2}$$

The informativeness of the active minipublic $\hat{\mathbf{m}}$ is a linear combination of the informativeness of k active singleton minipublics, one for each of its citizens. The singleton minipublic is most informative when it consists of the median citizen $i = 1/2$. The further away a citizen is from the median citizen, the less informative his singleton minipublic is.

Moreover, the weight assigned to the informativeness of singleton minipublic $\{i_j\}$ depends only on i_j's distance from his closest active minipublic neighbors i_{j-1} and i_{j+1}. This weight depends neither on i_j's absolute position in $[0, 1]$, nor on the rest of the active citizens. This is because the outcome $\beta(i_j)$ is useful in inferring only the outcomes of citizens in (i_{j-1}, i_{j+1}).[12] Second, these weights reflect the fact that the farther i_j is from either active neighbor, the greater is the weight assigned to his singleton minipublic.

3.2 Conflict between the policymaker and the citizens

Players' expected payoffs. Lemma 3.2 characterizes the expected payoffs of the policymaker and of the citizens in an active minipublic $\hat{\mathbf{m}}$. The expectations are taken with respect to both the minipublic outcomes and the threshold of adoption: the policymaker adopts the policy if $B_{\hat{\mathbf{m}}} > c$ and keeps the status quo if $B_{\hat{\mathbf{m}}} < c$.[13] This payoff characterization establishes that minipublic informativeness $\sigma^2_{\hat{\mathbf{m}}}$ is a sufficient statistic for players' payoffs. This is a direct implication of the Gaussian distribution of the post-minipublic value, as derived in lemma 3.1.

Lemma 3.2 (Players' payoffs)**.** *Fix an active minipublic $\hat{\mathbf{m}}$ with informativeness $\sigma^2_{\hat{\mathbf{m}}}$. The expected payoff of the policymaker is*

$$V_P(\sigma^2_{\hat{\mathbf{m}}}) := \bar{B}\Phi\left(\frac{\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right) + \sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}\,\phi\left(\frac{\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right), \tag{3}$$

and the expected payoff of every citizen $i \in \hat{\mathbf{m}}$ is

[12] This follows from the Markov property of the Ornstein-Uhlenbeck process. See also Axiom 5 in Appendix A.
[13] We allow for any tie-breaking rule in the zero-probability event $B_{\hat{\mathbf{m}}} = c$.

$$V_C(\sigma^2_{\hat{\mathbf{m}}}) := \bar{B}\Phi\left(\frac{\bar{B}}{\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}}\right) + \frac{\sigma^2_{\hat{\mathbf{m}}}}{\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}}\phi\left(\frac{\bar{B}}{\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}}\right). \tag{4}$$

If the policy is not adopted, all players obtain a payoff of zero. However, the citizens' expected payoff from adoption is different from that of the policymaker, because the citizens do not internalize the realized cost of adoption c. That is, the expected payoff of the policymaker is $\Pr[B_{\hat{\mathbf{m}}} > c]\mathbb{E}[B_{\hat{\mathbf{m}}} - c | B_{\hat{\mathbf{m}}} > c]$ whereas that of the citizens is $\Pr[B_{\hat{\mathbf{m}}} > c]\mathbb{E}[B_{\hat{\mathbf{m}}} | B_{\hat{\mathbf{m}}} > c]$. The expected payoff of the policymaker (3) and that of the citizen (4) share the same first term, which is determined by the prior value $\bar{B}$ and the probability with which the policymaker adopts the policy, $\Pr[B_{\hat{\mathbf{m}}} > c] = \Phi\left(\bar{B}/\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}\right)$. The difference in their expected payoffs is reflected in the second terms of (3) and (4), respectively. The citizen's payoff in (4) can be rewritten so as to highlight the misalignment between him and the policymaker:

$$\begin{aligned} V_C(\sigma^2_{\hat{\mathbf{m}}}) &= V_P(\sigma^2_{\hat{\mathbf{m}}}) + \Pr\left[B_{\hat{\mathbf{m}}} > c\right]\mathbb{E}\left[c | B_{\hat{\mathbf{m}}} > c\right] \\ &= V_P(\sigma^2_{\hat{\mathbf{m}}}) - \frac{\tau^2}{\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}}\phi\left(\frac{\bar{B}}{\sqrt{\tau^2+\sigma^2_{\hat{\mathbf{m}}}}}\right) < V_P(\sigma^2_{\hat{\mathbf{m}}}). \end{aligned} \tag{5}$$

For any level of political uncertainty $\tau^2 > 0$, the policymaker's expected threshold of adoption *conditional on adoption* is strictly less than $\mathbb{E}[c] = 0$. The citizens perceive the policymaker as being too positively disposed toward the policy. Hence, the payoff of the citizens is less than that of the policymaker for any informativeness $\sigma^2_{\hat{\mathbf{m}}}$. As $\tau^2 \to 0$, the misalignment term vanishes because $\mathbb{E}[c \mid B_{\hat{\mathbf{m}}} > c] \to 0$, so $V_C(\sigma^2_{\hat{\mathbf{m}}}) \to V_P(\sigma^2_{\hat{\mathbf{m}}})$. As $\tau^2 \to +\infty$, by contrast, V_C collapses only to the first term in (4). As political uncertainty becomes arbitrarily large, adoption becomes perfectly uninformative; hence, the expected value of an adopted policy approaches the prior value $\bar{B}$.

We next examine the conflict between the policymaker and the citizens over the level of informativeness. Lemma 3.3 shows that the policymaker always prefers higher informativeness, whereas the minipublic citizens might not.

Lemma 3.3 (Dependence of payoffs on informativeness)**.**

(i) The expected payoff of the policymaker, $V_P(\sigma^2)$, is strictly increasing in $\sigma^2 \in [0, \infty)$.

(ii) The expected payoff of citizen i, $V_C(\sigma^2)$, is strictly quasiconvex in σ^2, with a minimum at

$$\underline{\sigma}^2 = \max\left\{0, \frac{1}{2}\left(\sqrt{\tau^4 + 4\bar{B}^2\tau^2} - 3\tau^2\right)\right\}. \tag{6}$$

The policymaker aims to get as precise an estimate of the value of the policy as she can. For any realization of the threshold c, higher informativeness leads to a more accurate decision, which strictly benefits her. Hence, the policymaker prefers all citizens in the minipublic to be active, as each citizen strictly adds to informativeness.

In contrast, the citizen's payoff need not be increasing in informativeness. Higher informativeness has two opposing effects on the citizen's payoff. On the one hand, a more informative minipublic leads to a more precise estimate of the policy's value. This in turn leads to a better informed adoption decision, which benefits the citizen. With no political uncertainty this effect is the only one present. On the other hand, in the presence of political uncertainty, higher informativeness might increase the probability of ex post misalignment between the policymaker and the citizens.

To see what drives the quasiconvexity of the citizen's payoff, suppose the status quo is preferred ex ante: $\bar{B} < 0$.[14] More information unequivocally increases the probability that the policy is adopted. For low informativeness, the citizens might be better off with an entirely uninformed policymaker. Increasing informativeness only marginally above zero strictly increases the probability that a bad policy is adopted, because the expected value of an adopted policy is $\mathbb{E}[B_{\hat{\mathbf{m}}} \mid B_{\hat{\mathbf{m}}} > c] \approx \bar{B} < 0$ for $\sigma^2 \approx 0$. The citizens would rather let the policymaker keep the status quo based on $\bar{B}$ than risk a policy adoption. On the other hand, when informativeness is high, $\mathbb{E}[B_{\hat{\mathbf{m}}} \mid B_{\hat{\mathbf{m}}} > c]$ is positive and increasing in informativeness. If the policymaker is already well informed, the chance of ex post misalignment is negligible. Providing more evidence to the policymaker primarily contributes to the adoption of a rare right-tail policy of high value. Hence, the citizen's payoff strictly increases for sufficiently high informativeness.

The curse of too little information. At the evidence discovery stage, citizen $i \in \hat{\mathbf{m}}$ induces *active informativeness* $\sigma^2_{\hat{\mathbf{m}}}$ by being active, or *passive informativeness* $\sigma^2_{\hat{\mathbf{m}}\setminus i}$ by being passive. Because citizens' outcomes are imperfectly correlated, active informativeness is strictly higher than the passive one for any $\hat{\mathbf{m}}$ and $i \in \hat{\mathbf{m}}$.[15] The difference $(\sigma^2_{\hat{\mathbf{m}}} - \sigma^2_{\hat{\mathbf{m}}\setminus i})$ corresponds to the *marginal informativeness* of citizen i. A citizen's decision whether or not to be active, therefore, boils down to whether or not he prefers his active informativeness over his passive informativeness. We refer to this as the *evidence discovery* constraint (ED): citizen $i \in \hat{\mathbf{m}}$ prefers to be active if and only if $V_C(\sigma^2_{\hat{\mathbf{m}}}) \geqslant V_C(\sigma^2_{\hat{\mathbf{m}}\setminus i})$.

The citizen fears that, by becoming active, he might sway the policymaker toward the wrong adoption decision. The probability of landing in such a disagreement region, as well as the expected cost that citizens would bear from the wrong decision in this region, varies with minipublic

[14] A similar argument holds for $\bar{B} > 0$, as the misalignment term in (5) is identical for $\bar{B}$ and $-\bar{B}$.

[15] While active informativeness is the same for all citizens in $\hat{\mathbf{m}}$, passive informativeness might vary.

informativeness. The first effect dominates when the informativeness of the rest of the minipublic $\hat{\mathbf{m}} \setminus i$ is high, since citizen i's evidence primarily improves the quality of the adoption decision. By contrast, the second effect dominates when the informativeness of the rest of the minipublic is low: by providing evidence, citizen i risks rocking the boat too much.

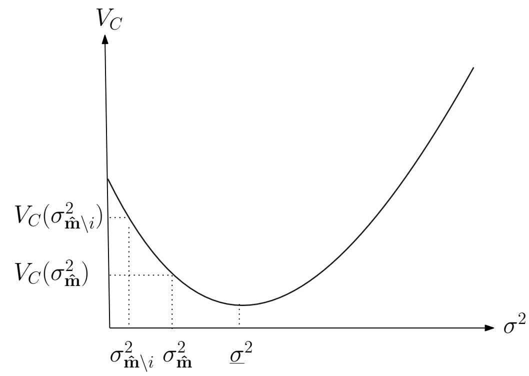

Figure 2: The curse of too little information.

Formally, when does the possibility of such misuse dominate the citizen's calculus, thus leading him to be passive in the minipublic? It is immediate from lemma 3.3(ii) that two conditions must be met for this to arise, as illustrated in figure 2. First, the minimum of the citizen's payoff must be to the right of zero. That is, the citizen's payoff must be strictly decreasing in informativeness at $\sigma^2 = 0$: contributing a small amount of information to an otherwise entirely uninformed policymaker is bound to harm the citizen.[16] Second, both the citizen's passive informativeness and his marginal informativeness in the minipublic must be sufficiently low. That is, $\sigma^2_{\hat{\mathbf{m}}\setminus i}$ must be lower than $\underline{\sigma}^2$, and $\sigma^2_{\hat{\mathbf{m}}}$ must not be too large relative to $\underline{\sigma}^2$. Figure 2 illustrates how (ED) is violated if $\sigma^2_{\hat{\mathbf{m}}} < \underline{\sigma}^2$.

We refer to these two conditions—namely, the negative marginal value of information at $\sigma^2 = 0$ and the sufficiently low informativeness of the minipublic—taken in tandem as the *curse of too little information*. This will be a driving force in the rest of our analysis.

3.3 The simplified problem of the policymaker

We next show that it is without loss for the policymaker's problem to restrict attention to degenerate lotteries over active minipublics of size at most n.

[16]Even though evidence discovery is costless, the marginal value of a small amount of informativeness can be strictly negative due to the implicit cost of the misuse of information. For a classic result on the negative marginal value of a small amount of information in single-agent settings with costly learning, see Radner and Stiglitz (1984) and Chade and Schlee (2002). In contrast to Martinelli (2006) and Strulovici (2010), the negative value of experimentation in our framework is not due to costly information or pivotality considerations in voting.

Proposition 3.1. *In the class of policymaker-preferred equilibria, it is without loss to restrict attention to deterministic minipublics* $\mathbf{m} \in \mathcal{M}_n$ *in which all citizens are active with probability one.*

Proposition 3.1 marks a significant step in simplifying the policymaker's problem, and it follows from three observations. First, the policymaker cannot encourage more evidence discovery by using lotteries over different minipublics; this is because each citizen observes the realized minipublic before deciding whether to be active. Second, for any minipublic and any equilibrium in which a subset of citizens mix between being active and being passive, we can construct another equilibrium in pure evidence discovery strategies that guarantees a higher expected payoff to the policymaker. Third, the policymaker cannot benefit from including passive citizens in the minipublic. Not only is her expected payoff V_P independent of the presence of passive citizens, but they also do not affect the incentives of other minipublic citizens.

From lemma 3.3, maximizing the expected payoff $V_P(\sigma^2_{\mathbf{m}})$ is equivalent to choosing the minipublic with the highest informativeness $\sigma^2_{\mathbf{m}}$ among all active minipublics of size at most n. (ED) guarantees that each citizen in the minipublic is active. Thus, we can write the policymaker's minipublic choice problem as the following simplified problem:

$$\max_{\mathbf{m} \in \mathcal{M}_n} \quad \sigma^2_{\mathbf{m}} \tag{P}$$

$$\text{s.t.} \quad V_C(\sigma^2_{\mathbf{m}}) \geq V_C(\sigma^2_{\mathbf{m}\setminus i}) \quad \forall i \in \mathbf{m}. \tag{ED}$$

We call a minipublic $\mathbf{m}$ *feasible* if (i) it consists of at most n citizens, and (ii) every $i \in \mathbf{m}$ is active.

3.4 Benchmark: The first-best minipublic

Ignoring (ED) for now, the policymaker's unconstrained problem (P) consists of choosing a feasible minipublic that maximizes informativeness $\sigma^2_{\mathbf{m}}$. The unconstrained problem of the policymaker is equivalent to one in which the policymaker can herself discover any n outcomes in the outcome mapping β. The next result follows from Bardhi (2020), which provides a full characterization of the solution to this problem.

Corollary 3.1. *For any* n*, the first-best minipublic* $\mathbf{m}^f_n := \{i^f_1, ..., i^f_n\}$ *is unique and satisfies the following three properties:*

(i) It is symmetric about the median citizen: $i^f_k = 1 - i^f_{n-k+1}$ *for every* $k \in \{1, \dots, n\}$.

(ii) It has equal distance between adjacent minipublic citizens: $i^f_k - i^f_{k-1}$ *is constant for* $k \in \{2, ..., n\}$.

(iii) It has the peripheral citizens $0 < i_1^f = 1 - i_n^f < 1$ *pinned down by*

$$1 - e^{-i_1^f/\ell} = \tanh\left(\frac{1 - 2i_1^f}{2\ell(n-1)}\right). \tag{7}$$

This first-best minipublic, which in the absence of political uncertainty also maximizes the payoff of each player, provides a benchmark for the informativeness-maximizing level of diversity. In particular, it pins down how far apart adjacent citizens should be, whether the minipublic should be skewed relative to the median citizen, and how large the minipublic range $(i_n - i_1)$ should be.

The informativeness of $\mathbf{m}_n^f$ gives an upper bound on the informativeness that can be attained with capacity n. We let $\bar{\sigma}_n^2$ denote the informativeness attained by the first-best minipublic of size n. Because the outcomes of any two citizens are only imperfectly correlated, $\bar{\sigma}_n^2$ strictly increases in n. Moreover, the first-best informativeness $\bar{\sigma}_n^2$ approaches the ex ante variance of the policy value B as $n \to \infty$. With sufficiently large capacity, almost the entire uncertainty about the policy can be resolved through the minipublic. This limit informativeness is

$$\bar{\sigma}_\infty^2 = 2\ell\left(1 - \ell\left(1 - e^{-1/\ell}\right)\right) < 1.$$ [17]

Because outcomes are positively correlated, the ex ante variance of the policy $\bar{\sigma}_\infty^2$ is strictly lower than the unit variance of any citizen's outcome.

Let us now return to the constrained problem. Whenever $\mathbf{m}_n^f$ is not feasible, (ED) of at least one citizen in it must be violated. Which citizens are more likely to be passive in the first-best minipublic? We show that the structure of $\mathbf{m}_n^f$ is such that either (i) only the most peripheral citizens i_1^f and i_n^f are passive, or (ii) all citizens are passive. To establish this, lemma 3.4 ranks the passive informativeness of all citizens in the first-best minipublic.

Lemma 3.4 (Ranking passive informativeness in $\mathbf{m}_n^f$)**.** *In the first-best minipublic* $\mathbf{m}_n^f$*, for any* $j, k \in \{2, \ldots, n-1\}$,

$$\sigma^2_{\mathbf{m}_n^f \setminus i_1^f} = \sigma^2_{\mathbf{m}_n^f \setminus i_n^f} < \sigma^2_{\mathbf{m}_n^f \setminus i_j^f} = \sigma^2_{\mathbf{m}_n^f \setminus i_k^f}.$$

Inner citizens $i_2^f, \ldots, i_{n-1}^f$ all have the same marginal informativeness. This is because for each such citizen i_k^f, the mass of citizens to his immediate left $(i_k^f - i_{k-1}^f)$ is equal to the mass of citizens to his immediate right $(i_{k+1}^f - i_k^f)$. If one inner citizen is passive, they are all passive. Second, the peripheral citizens i_1^f and i_n^f are the most marginally informative in the first-best minipublic because, of all citizens in $\mathbf{m}_n^f$, only they are informative about the periphery of the citizenry $[0, i_1^f]$

[17] Lemma B.5 provides the details.

and $[i_n^f, 1]$, respectively. By the symmetry of $\mathbf{m}_n^f$, their marginal informativeness is the same. So if a subset of citizens is passive in the first-best minipublic, then i_1^f and i_n^f are in that subset.

4 The optimal minipublic

4.1 No distortion in minipublic size

In designing a minipublic, the policymaker chooses both how large the minipublic is and which citizens it consists of. If the first-best minipublic $\mathbf{m}_n^f$ is unfeasible, the policymaker has two instruments through which to satisfy (ED): distorting the minipublic size and distorting its composition. Our first observation concerns the size of the optimal minipublic. Does the policymaker ever sample fewer citizens than what her capacity allows? Proposition 4.1 shows that the optimal minipublic either uses the entire capacity n or is empty. The latter is the case if the policymaker cannot incentivize any feasible minipublic to be active. But whenever it is possible to incentivize *some* evidence discovery, the policymaker is strictly worse off from reducing the minipublic size below capacity.

Proposition 4.1 (Optimal minipublic size). *Given capacity n, the optimal minipublic either is empty or consists of exactly n citizens.*

The proof of proposition 4.1 establishes that for any minipublic $\mathbf{m}$ of size $n' < n$, adding a new citizen $j \notin \mathbf{m}$ satisfies the (ED) of each citizen while strictly increasing minipublic informativeness from $\sigma^2_{\mathbf{m}}$ to $\sigma^2_{\mathbf{m}\cup j}$. One crucial observation is that if a citizen is active for a given pair of active and passive informativeness (i.e., if he already escaped the curse of too little information), then he remains active if both his passive and active informativeness increase further.[18] This observation, illustrated in figure 3, serves us in two ways. On the one hand, due to the addition of citizen j, both the passive and the active informativeness of each original citizen $i \in \mathbf{m}$ have increased. Hence, citizen i continues to be active in $\mathbf{m} \cup j$. On the other hand, because each original citizen $i \in \mathbf{m}$ prefers $\sigma^2_{\mathbf{m}}$ to his passive informativeness $\sigma^2_{\mathbf{m}\setminus i}$, the new citizen j prefers $\sigma^2_{\mathbf{m}\cup j}$ to his own passive informativeness $\sigma^2_{\mathbf{m}}$ in the new minipublic. Hence, j is also active in $\mathbf{m} \cup j$.

Proposition 4.1 thus narrows down the types of the optimal minipublic $\mathbf{m}^*$ that can arise to these three: (i) the first-best minipublic $\mathbf{m}^* = \mathbf{m}_n^f$, (ii) a minipublic of n citizens $\mathbf{m}^* \neq \mathbf{m}_n^f$, and (iii) the empty minipublic $\mathbf{m}^* = \emptyset$. The optimal minipublic is distorted when it is of either type (ii) or type (iii). The rest of this section identifies conditions on the primitives of our environment (political uncertainty, policy sentiment, citizenry's heterogeneity, and minipublic capacity) for each type to

[18] See lemma B.2 in Appendix B.

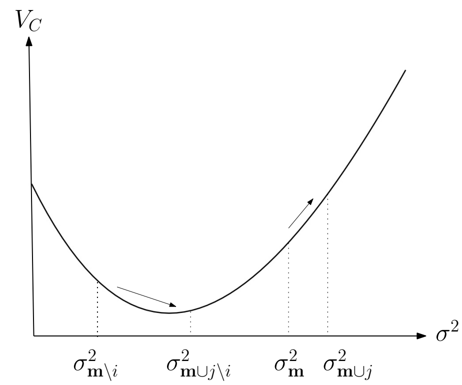

Figure 3: Citizen $i \in \mathbf{m}$ prefers to be active if citizen $j \notin \mathbf{m}$ is added to $\mathbf{m}$.

arise as optimal. In particular, section 4.2 clarifies when the policymaker attains the first-best benchmark and when no active minipublic is feasible. Section 4.3, in turn, analyzes the diversity within an optimal minipublic of type (ii).

4.2 Types of the optimal minipublic

Distortions arise only for moderate political uncertainty. We first argue that political uncertainty has a non-monotonic effect on the optimal minipublic. Intuitively, greater political uncertainty implies greater agency loss, and hence a lower payoff, for the citizens. For any level of informativeness σ^2, the citizens' payoff (4) strictly decreases in τ^2. A first intuition might suggest that the greater the political uncertainty that the citizens face, the more distorted the optimal minipublic is relative to the first-best one. With greater political uncertainty, the citizens are more likely to fear the misuse of the evidence that they provide, which makes it more challenging for the policymaker to motivate evidence discovery. However, contrary to this intuition, proposition 4.2 establishes that the optimal minipublic is the first-best one for either sufficiently low or sufficiently high political uncertainty. The curse of too little information disappears at either extreme.

Proposition 4.2 (No distortions under high or low political uncertainty)**.** *Fix all parameters other than τ^2. There exist cutoffs $0 < \underline{\tau}^2 \leq \bar{\tau}^2 < \infty$ such that $\mathbf{m}^* = \mathbf{m}_n^f$ if political uncertainty is either (i) sufficiently low (i.e., $\tau^2 \leqslant \underline{\tau}^2$) or (ii) sufficiently high (i.e., $\tau^2 \geqslant \bar{\tau}^2$).*

Political uncertainty determines the level of informativeness $\underline{\sigma}^2$ at which the citizens' payoff is at its lowest, as given in (6). This minimum $\underline{\sigma}^2$ is single-peaked in τ^2: when τ^2 is either sufficiently close to zero or sufficiently high, $\underline{\sigma}^2$ is exactly zero. But if $\underline{\sigma}^2$ equals zero, then the marginal value of information is positive for any level of informativeness; therefore, all citizens prefer to be active. As a consequence, the first-best minipublic is active for either a sufficiently low or a sufficiently high level of political uncertainty. However, the reason that the citizens' fear of the misuse of evidence does not dominate their calculus is fundamentally different across the two cases.

Under low political uncertainty (i.e., as $\tau^2 \to 0$), the policymaker is increasingly likely to prefer the same adoption decision as the citizens for any post-minipublic value $B_{\mathbf{m}}$. In the limit, the misalignment term in (5) vanishes, and $V_C(\sigma^2)$ gets arbitrarily close to $V_P(\sigma^2)$. The value of a more accurate adoption decision encourages citizens to discover evidence, while their fear of the policymaker misusing the evidence disappears.

By contrast, under high political uncertainty (i.e., as $\tau^2 \to +\infty$), citizens cannot affect the probability of an ex post misalignment by being passive. The policymaker's decision becomes fully unpredictable from the citizens' perspective: for any minipublic informativeness, the policy is adopted with probability approaching $1/2$. The expected value of the adopted policy approaches the prior value $\bar{B}$. Yet evidence still contributes to a slightly more accurate adoption decision. Therefore, citizens in $\mathbf{m}_n^f$ are willing to be active.

A key implication of proposition 4.2 is that what hampers evidence discovery and thus forces the policymaker to choose a minipublic other than the first-best one is the presence of *moderate*, rather than high, political uncertainty. In contrast to when $\tau^2 \to \infty$, moderate political uncertainty guarantees that each citizen's evidence can significantly influence the probability of ex post misalignment. Yet, in contrast to when $\tau^2 \to 0$, the likelihood of ex post misalignment is sufficiently high as to discourage evidence discovery. Hence, under moderate political uncertainty, the optimal minipublic might be distorted away from the first-best minipublic, or even be empty.

Distortions arise for strong policy sentiment. Our next insight is that stronger policy sentiment makes it more challenging for the policymaker to incentivize evidence discovery within a minipublic. In fact, when policy sentiment is extremely strong, no minipublic, no matter its size and composition, is ever active. In such a case, citizens find it too costly to risk rocking the boat by discovering evidence. Strikingly, no minipublic is feasible precisely for policies for which minipublic evidence is unlikely to overturn the decision that is ex ante preferred by all players.

Proposition 4.3 (Distortions arise under strong policy sentiment)**.** *Fix capacity n. There exist unique cutoffs $0 < \underline{b} \leqslant \bar{b} < \infty$ such that:*

(i) the optimal minipublic is the first-best minipublic $\mathbf{m}_n^f$ if and only if $|\bar{B}| \leqslant \underline{b}$;

(ii) the optimal minipublic consists of n citizens but is distinct from the first-best minipublic if and only if $|\bar{B}| \in (\underline{b}, \bar{b}]$;

(iii) the optimal minipublic is empty if and only if $|\bar{B}| > \bar{b}$.

If policy sentiment is sufficiently weak, any minipublic is active. To see this, consider $\bar{B} = 0$, for which the citizens are ex ante indifferent between the policy and the status quo. For such $\bar{B}$, the citizen's payoff globally increases in informativeness: the probability of adoption is exactly $1/2$ for any level of informativeness, but the more evidence is available to the policymaker, the higher is the expected value conditional on adoption. This is why the misalignment term in (5) strictly decreases in informativeness. By providing evidence to the policymaker, not only do citizens contribute toward a more precise decision, but they also hedge against political uncertainty. Intuitively, because both the distribution of c and that of the post-minipublic value are centered at zero, the likelihood of ex post misalignment is high. A more informative minipublic increases the likelihood that more extreme post-minipublic values are realized, for which the policymaker and the citizens prefer the same adoption decision ex post. Because any citizen is willing to be active in any minipublic, the policymaker attains the first-best level of informativeness.

This reasoning extends to when $|\bar{B}|$ is small but nonzero. The true value of the policy B is then almost as likely to be below zero as above zero. Learning more about the policy is therefore highly valuable. Moreover, minipublic informativeness continues to serve a hedging role as long as policy sentiment is weak relative to the level of political uncertainty.

In contrast, when $|\bar{B}|$ is sufficiently high, evidence is unlikely to overturn the ex ante preferred decision. Although both the value of learning new evidence and the probability of ex post misalignment are small, the latter dominates the citizen's payoff. Evidence discovery contributes more to an increase in the likelihood of ex post misalignment than to an increase in the accuracy of adoption. Therefore, no citizen is active in any minipublic when policy sentiment is extremely strong.

Distortions arise only for a sufficiently heterogeneous citizenry. We next seek to understand whether greater heterogeneity of the citizenry helps or hampers evidence discovery. To do so, we characterize the optimal minipublic in two polar cases: (i) when the citizenry is extremely heterogeneous, i.e., $\ell \to 0$; and (ii) when the citizenry is extremely homogeneous, i.e., $\ell \to +\infty$. The policymaker attains the first-best minipublic when facing a sufficiently homogeneous citizenry, but is unable to incentivize any minipublic to be active when the citizenry is sufficiently heterogeneous.

As starkly different as these two polar cases are, one important commonality between them is that the marginal informativeness of any citizen approaches zero in either limit. When citizens' outcomes are almost independent, it becomes increasing difficult to extrapolate from any one citizen to the rest of the minipublic. Thus, the marginal informativeness of a single citizen approaches zero as $\ell \to 0$. By contrast, when outcomes are almost perfectly correlated, the outcome of one citizen suffices to resolve almost all uncertainty about the policy. Therefore, the marginal informativeness of any additional citizen beyond the first citizen also vanishes as $\ell \to \infty$.

Despite this commonality, what makes these two cases fundamentally different is the active informativeness of any given minipublic, which is infinitesimally small when citizens' outcomes are almost independent ($\ell \to 0$) but close to $\bar{\sigma}^2_\infty \to 1$ when they are almost perfectly correlated ($\ell \to \infty$). This distinction determines whether the curse of too little information is likely to arise and leads to drastically different optimal minipublics, as the following result establishes.[19]

Proposition 4.4. *Let $0 < \underline{\sigma}^2 < 1$. There exist $\underline{\ell}$ and $\overline{\ell}$ with $0 < \underline{\ell} < \overline{\ell} < \infty$ such that*

(i) in a sufficiently heterogeneous citizenry with $\ell \leqslant \underline{\ell}$, the optimal minipublic is empty;

(ii) in a sufficiently homogeneous citizenry with $\ell \geqslant \overline{\ell}$, the optimal minipublic is the first-best minipublic for any $n > 1$.

Because $\underline{\sigma}^2 > 0$ by assumption, providing the policymaker with little information can harm the citizen. In a heterogeneous citizenry in which outcomes are almost independent, any minipublic carries a negligible amount of information. Moreover, each citizen contributes little to an infinitesimally low passive informativeness. Hence, as the citizenry becomes increasingly heterogeneous, the curse of too little information discourages any citizen from being active in any minipublic.

In contrast, for a sufficiently homogeneous citizenry, the active informativeness in any minipublic approaches the unit variance of any citizen's outcome, because outcomes become perfectly correlated. Any citizen reveals almost all there is to know about the policy; hence, each citizen in a minipublic consisting of more than one citizen adds little to a passive informativeness very close to one. At this level of informativeness, the citizen has already escaped the curse of too little information because $\underline{\sigma}^2 < 1$. Therefore, (ED) is satisfied for every citizen in the first-best minipublic.

As the citizenry becomes highly homogeneous, the optimal minipublic, which is the first-best minipublic, approaches an intuitive form: it minimizes the sum of the squared distance from each citizen in $[0, 1]$ to his closest minipublic citizen. The minipublic $\mathbf{m}^d_n$ with $i^d_k = (2k - 1)/2n$ for every $k = 1, ..., n$ minimizes such average distance to the closest minipublic representative. Each minipublic citizen represents a mass $1/n$ of citizens, equally split between either side of the minipublic citizen. For instance, $\mathbf{m}^d_2 = \{1/4, 3/4\}$ for $n = 2$. Citizens $i^d_1 = 1/4$ and $i^d_2 = 3/4$ are, respectively, the representatives for citizens to the left and to the right of the median citizen.

Corollary 4.1. *As $\ell \to \infty$, the optimal minipublic approaches the distance^2-minimizing minipublic $\mathbf{m}^d_n$, which consists of $i^d_k = (2k - 1)/(2n)$ for $k \in \{1, \ldots, n\}$.*

[19] The parameter ℓ does not enter V_C and $\underline{\sigma}^2$, but only $\sigma^2_{\mathbf{m}}$ and $\sigma^2_{\mathbf{m}\setminus i}$. To restrict attention to nontrivial cases, we impose that $\underline{\sigma}^2 \in (0, 1)$. Otherwise, the citizens' payoff is either globally increasing or globally decreasing in informativeness on $[0, \bar{\sigma}^2_\infty]$ and citizens in any minipublic are either always active or always passive for any ℓ.

Distortions arise only in small minipublics. Finally, we show that if the first-best minipublic is ever optimal, it is when minipublic capacity is large. That is, distortions are more likely to arise when the policymaker can sample only a few citizens. This observation motivates our focus on the composition of small minipublics in section 4.3 below.

As already shown, the curse of too little information arises in minipublics that have low informativeness. Because informativeness increases with each additional citizen, low informativeness is characteristic of small, rather than large, minipublics. In a large minipublic each citizen has a strong incentive to be active because both his passive and his active informativeness are high. Proposition 4.5 asserts that if it is at all possible to have an active minipublic of some size, then the first-best minipublic prevails as optimal for sufficiently large capacity.

Proposition 4.5 (No distortion for large n)**.** *If $\bar{\sigma}^2_\infty > \underline{\sigma}^2$, then there exists $\bar{n} \in \mathbb{N}$ such that for all $n > \bar{n}$, the optimal minipublic is $\mathbf{m}^f_n$. If $\bar{\sigma}^2_\infty \leqslant \underline{\sigma}^2$, then the optimal minipublic is empty for any n.*

If the informativeness corresponding to minimum payoff $\underline{\sigma}^2$ is to the left of $\bar{\sigma}^2_\infty$, the citizens' payoff is not globally decreasing over the relevant domain $[0, \bar{\sigma}^2_\infty]$. That is, when the rest of the environment is fixed, it is possible to evaluate the policy through *some* active minipublic, however large that might be. In a sufficiently large first-best minipublic, both the informativeness of $\mathbf{m}^f_n \setminus i^f_k$ for any i^f_k and that of $\mathbf{m}^f_n$ are eventually larger than the payoff-minimizing informativeness $\underline{\sigma}^2 < \bar{\sigma}^2_\infty$. In fact, as n gets large, both of them approach $\bar{\sigma}^2_\infty$ and the marginal informativeness of each citizen vanishes: minipublic $\mathbf{m}^f_n$, with or without citizen i^f_k, discovers almost all there is to know about the policy. This guarantees that the first-best minipublic is active for large capacities.

However, if the minimum $\underline{\sigma}^2$ is to the right of $\bar{\sigma}^2_\infty$—so that V_C strictly decreases in informativeness over the relevant informativeness interval—then every citizen is passive, no matter how large the minipublic. This leads to a total collapse of minipublic choice. Even if the minipublic capacity is so large that it affords very extensive learning about the policy, the prospect of future misalignment deters all evidence discovery and the policymaker remains uninformed.

Implications for the CCC. We draw two main takeaways from the results above. First, political uncertainty is detrimental to minipublic design only when such uncertainty is neither too low nor too high. Arguably, in the case of the CCC, political uncertainty was low. The CCC was preceded by the Grand Débat National, a national effort that showcased Macron's commitment to a greener economy, and was backed by his promise to enact the CCC's proposals "without filter". Had political uncertainty been moderately higher, citizens would have likely been more concerned about providing evidence. The resulting CCC would have to be less representative of the French society in order to incentivize citizens to bring in their local evidence. A second takeaway is that the less clear it is ex ante what the correct adoption decision is, the stronger are the citizens' incentives to provide local evidence. Indeed, climate policy is one such instance in which evidence is highly likely to overturn any prior attitudes, as it relies on frontier scientific and technological advances. Therefore, our results imply that the citizens' incentives within the CCC were likely sustained by both the low political uncertainty and the weak policy sentiment given the nature of climate policy.

4.3 Distortion in minipublic composition

We now turn to characterizing the optimal minipublic when the set of feasible minipublics is nonempty but does not include $\mathbf{m}_n^f$. Per proposition 4.1, any such optimal minipublic consists of exactly n citizens. To guarantee that (ED) is satisfied for each minipublic citizen, the policymaker distorts the composition rather than the size of the minipublic.

To address how diverse the optimal minipublic is relative to the first-best one, we introduce two notions of diversity within a minipublic: *informational diversity* and *demographic diversity*. Whereas informational diversity captures the diversity of a minipublic in terms of the marginal informativeness of each citizen's evidence, demographic diversity captures it as the linear distance between citizens' positions in $[0, 1]$. Notably, the optimal minipublic is strictly less diverse than the first-best one with respect to both notions.

Informational diversity. Our first notion of diversity quantifies the novelty of the evidence of individual citizens relative to the evidence provided by the rest of the minipublic. All else fixed, the more novel the evidence that an individual citizen contributes to the minipublic, the more diverse the minipublic is. Formally, we define the informational diversity of a minipublic $\mathbf{m}$ to be

$$\max_{i \in \mathbf{m}} \qquad \frac{\sigma_{\mathbf{m}}^2 - \sigma_{\mathbf{m}\setminus i}^2}{\sigma_{\mathbf{m}\setminus i}^2}.$$

This definition has two building blocks. First, we compute each citizen's marginal informativeness normalized by the informativeness of the rest of the minipublic. This measures the relative increase in minipublic informativeness that each citizen can bring about by being active. The more novel the citizen's evidence is relative to that of other citizens in the minipublic, the higher is his marginal informativeness. Second, the informational diversity of a minipublic is determined by the citizens with the highest such index, i.e., the citizens in the minipublic with the most novel evidence.

The key insight upon which proposition 4.6 builds is that if the first-best minipublic is not feasible, then all feasible minipublics are strictly less informationally diverse than the first-best one. Hence, the optimal minipublic is strictly less informationally diverse as well.

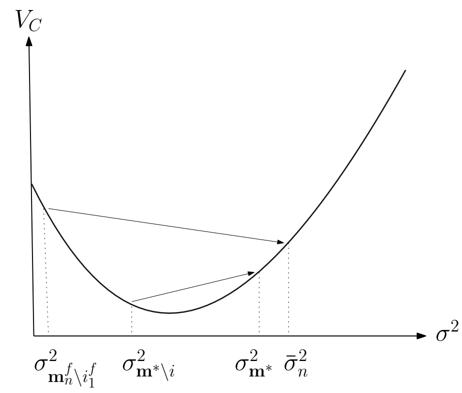

Figure 4: The optimal minipublic $\mathbf{m}^*$ has strictly lower informational diversity than $\mathbf{m}^f_n$. (ED) is violated for i^f_1 and i^f_n in $\mathbf{m}^f_n$, but it holds for each $i \in \mathbf{m}^*$.

Proposition 4.6 (Lower informational diversity)**.** *Fix capacity n. If the optimal minipublic is neither empty nor $\mathbf{m}^f_n$, then it is strictly less informationally diverse than the first-best minipublic.*

Lemma 3.4 implies that the informational diversity of the first-best minipublic is determined by citizens i^f_1 and i^f_n. They have the lowest passive informativeness, and thus the highest marginal informativeness, in $\mathbf{m}^f_n$. Hence, whenever $\mathbf{m}^f_n$ is not feasible, it must be the case that i^f_1 and i^f_n are passive, as shown in figure 4. This means that if the optimal minipublic is nonempty, the passive informativeness of *each citizen* in it must be higher than the passive informativeness of i^f_1 in $\mathbf{m}^f_n$. Coupled with the observation that the informativeness of any minipublic of size n is lower than that of $\mathbf{m}^f_n$, this implies that each citizen in the optimal minipublic has lower marginal informativeness than i^f_1 in the first-best minipublic. Therefore, the citizens with the most novel evidence in $\mathbf{m}^*$ add less, in relative terms, to minipublic informativeness than do the citizens with the most novel evidence in $\mathbf{m}^f_n$. That is, $\mathbf{m}^*$ is less informationally diverse than $\mathbf{m}^f_n$.

Demographic diversity. Next, we analyze the implications of lower informational diversity for demographic diversity. For conceptual clarity, we do so in the context of a small minipublic with $n = 2$, in which demographic diversity corresponds to the distance between the two minipublic citizens. This is a natural context in which to study demographic diversity because, per proposition 4.5, if a distorted minipublic ever arises, it is for sufficiently small n. For $n = 1$, the optimal minipublic is either the first-best $\mathbf{m}^f_1 = \{1/2\}$ or empty.[20] So the smallest capacity for which a distorted minipublic arises is $n = 2$.

[20]The active informativeness of each citizen is lower than that of the median citizen, whereas passive informativeness is zero for all. By the quasiconvexity of V_C, it follows that if the median citizen prefers to be passive, then so does any other citizen.

Proposition 4.7 establishes that if the optimal minipublic is a distorted minipublic, the distance between the two citizens in $\mathbf{m}^*$ is smaller than that between i_1^f and i_2^f in the first-best minipublic. This means that the citizens in the optimal minipublic are more similar to each other—i.e., their respective outcomes are more strongly correlated than the outcomes of the first-best minipublic—so the optimal minipublic is less demographically diverse. The optimal minipublic overrepresents citizens around the median citizen and underrepresents those at the peripheries, risking to be an echo chamber in which the local evidence of minipublic citizens is too correlated. Lower informational diversity, therefore, goes hand in hand with lower demographic diversity in the optimal minipublic.

Proposition 4.7 (Lower demographic diversity)**.** *Fix $n = 2$ and let $|\bar{B}| \in (\underline{b}, \bar{b})$, where $\underline{b}$ and $\bar{b}$ are as defined in proposition 4.3. The unique optimal minipublic is the symmetric minipublic $\mathbf{m}^* = \{i^*, 1 - i^*\}$ such that $i_1^f < i^* < 1/2$ and (ED) binds for both citizens.*

For moderate policy sentiment, the policymaker resorts to a distorted minipublic of smaller distance between citizens because (ED) is violated for each citizen in the first-best minipublic. The passive informativeness of citizens i_1^f and i_2^f in $\mathbf{m}_2^f$, which corresponds exactly to the informativeness of singleton minipublics $\{i_2^f\}$ and $\{i_1^f\}$, respectively, is too low to encourage them to be active. The only way to increase such passive informativeness is to choose citizens who are closer to the median citizen than i_1^f and i_2^f are. This is because the informativeness of a singleton minipublic increases in proximity to the median citizen. Therefore, all minipublics with at least one citizen further away from the median citizen than i_1^f will have at least one passive citizen.

As is true for the first-best minipublic, this distorted minipublic continues to be symmetric about the median citizen. Given an active asymmetric minipublic with a fixed distance between citizens, the policymaker attains strictly higher minipublic informativeness by shifting the citizens toward the median citizen while keeping the distance between them fixed. Such a shift continues to satisfy (ED). As a result, the symmetry of the optimal minipublic reduces the policymaker's problem to one of choosing a minipublic of the form $\{i^*, 1 - i^*\}$. Intuitively, the policymaker prefers i^* to be as close to i_1^f as possible, because the closer i^* gets to the median citizen, the more redundant is the evidence of the minipublic citizens and, hence, the lower the informativeness of the minipublic. That is, the policymaker prefers a more demographically diverse minipublic. But the first-best minipublic is too diverse to incentivize both citizens to discover evidence. Therefore, the policymaker chooses the citizen i^* closest to i_1^f who is exactly indifferent between being passive and being active. By the symmetry of the correlation structure, if (ED) binds for i^*, it also binds for $1 - i^*$.

One implication of proposition 4.7 is that as policy sentiment becomes stronger, the optimal minipublic gets increasingly less demographically diverse.[21] A higher $|\bar{B}|$ makes it more challenging to induce evidence discovery. Thus, only symmetric minipublics of sufficiently low demographic diversity are feasible, and the policymaker chooses the most diverse among them. When policy sentiment grows too strong, suppressing demographic diversity is no longer effective in encouraging evidence discovery, and so only the empty minipublic remains feasible.

5 Discussion and extensions

Noisy evidence discovery. Our baseline model assumes that minipublic citizens discover either $\beta(i)$ perfectly or nothing at all. However, assuming away partial evidence might stack the deck against evidence discovery: citizens who are otherwise willing to discover some partial evidence might remain passive when only perfect evidence is available for discovery. To evaluate this conjecture, we consider an alternative setting in which each minipublic citizen $i \in \mathbf{m}$ publicly discovers a noisy signal $s(i) = \beta(i) + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \xi_i^2)$ and the noise $\xi_i^2 \in \mathbb{R}_0^+ \cup \{+\infty\}$ is chosen by citizen i. In the baseline model, $\xi_i^2 = 0$ (or $\xi_i^2 = +\infty$) corresponds to the citizen being active (or passive). Crucially, we continue to assume that any choice of ξ_i^2 is entirely costless to the citizen.

This richer discovery model reveals that the all-or-nothing nature of evidence in the baseline model is in fact without loss. The informativeness of any minipublic continues to be the same as in the baseline model, and therefore the optimal minipublic is the same as well. This is due to the quasiconvexity of the citizens' payoff in informativeness. If a citizen has already escaped the curse of too little information for some finite $\xi_i^2 > 0$, then he prefers to increase informativeness even further by choosing the least noisy signal available $\xi_i^2 = 0$. Citizens never prefer to discover their outcomes partially if they can discover them perfectly.

General Gaussian outcomes. The payoff characterization in lemmata 3.2 and 3.3 relies on the Gaussian distribution of the citizens' outcomes and the policymaker's threshold. Hence, the policymaker's problem, as derived in section 3.3, as well as the curse of too little information arise more generally for a rich class of Gaussian processes over policy outcomes.[22] The arguments behind propositions 4.1 (no distortion in minipublic size), 4.2 (distortions for moderate uncertainty), 4.3 (distortions for strong policy sentiment), and 4.6 (lower informational diversity) can be generalized as well. We adopt the Ornstein-Uhlenbeck process for two reasons. First, among all Gaussian

[21] By lemma C.1, if a minipublic is feasible under $|\bar{B}|$, then it is feasible under $|\bar{B}'| < |\bar{B}|$ as well. Hence, the optimal minipublic $\mathbf{m}^*(|\bar{B}|)$ is more informative than $\mathbf{m}^*(|\bar{B}'|)$; for $n = 2$ this implies that the citizens in $\mathbf{m}^*(|\bar{B}|)$ are farther away from the median citizen than are those in $\mathbf{m}^*(|\bar{B}'|)$.

[22] With a continuum of citizens, for the policy value B to be well-defined, the Gaussian process must satisfy almost sure Lebesgue integrability and $\bar{\sigma}_\infty^2 < \infty$.

processes, it is a natural candidate to model correlation in local evidence, as argued in the axiomatization in appendix A. Second, and more importantly, the Ornstein-Uhlenbeck process provides a tractable framework for studying demographic diversity (proposition 4.7) and the dependence of the optimal minipublic on the heterogeneity of the citizenry (proposition 4.4). This is because the minipublic informativeness and the first-best minipublic take a strikingly simple form.

No commitment in evidence disclosure. The game of section 2.1, which the discussion below refers to as the *commitment game*, assumes commitment in evidence disclosure: the outcome of each active citizen is disclosed publicly regardless of its realization. The citizen cannot withhold unfavorable outcome realizations. We examine here the robustness of our analysis to this commitment assumption. To do so, we consider the following *no-commitment game* which differs from the commitment game only at the evidence discovery stage: (i) each minipublic citizen simultaneously decides whether to discover evidence,[23] (ii) each citizen who discovers evidence observes his outcome privately, and (iii) citizens decide simultaneously whether to disclose or conceal their privately observed outcomes. That is, citizens' evidence is verifiable (e.g., as in Milgrom and Roberts (1986)).

Proposition 5.1 establishes that for any feasible minipublic in the commitment game, there exists an equilibrium in the no-commitment game in which the policymaker perfectly infers all minipublic outcomes. In particular, this equilibrium guarantees that in the no-commitment game the policymaker can attain at least the same level of informativeness as that of the optimal minipublic in the commitment game.

Proposition 5.1. *Let* $\mathbf{m}$ *be any feasible minipublic in the commitment game. Then, in the no-commitment game, there exists an equilibrium for minipublic* $\mathbf{m}$ *in which (i) all citizens in* $\mathbf{m}$ *discover evidence, and (ii) the policymaker infers all outcomes* $\beta(\mathbf{m})$ *perfectly.*

In this equilibrium, each minipublic citizen $i \in \mathbf{m}$ discloses all but a single outcome realization $\beta(i) = x_i$, which is pinned down by $\mathbb{E}[B \mid \beta(i) = x_i] = 0$. This is the unique realization that leaves him indifferent between the policy and the status quo. To see that this is indeed an equilibrium, consider a minipublic of two citizens $\{i, j\}$. If citizen i conceals evidence in favor of the policy $\beta(i) > x_i$, this encourages the policymaker to be more demanding on $\beta(j)$ for adoption because she incorrectly believes that $\beta(i) = x_i$. This has two opposing effects on i's payoff: the expected value of the policy conditional on its being adopted increases, but the probability of such an adoption decreases. The latter effect dominates. Some policies that are preferable to the status quo, given i's

[23] We assume that discovery decisions are observable to the policymaker: she can distinguish between a citizen with no evidence and one who conceals evidence. Yet this assumption is without loss for proposition 5.1. If the decision were instead unobservable, it would be weakly dominant for each $i \in \mathbf{m}$ to discover evidence.

favorable evidence, are forgone. Because citizen i's preference is aligned with the policymaker's ex ante, he does not benefit from inducing false pessimism by concealing favorable evidence about the policy. The reasoning is analogous if citizen i holds unfavorable evidence $\beta(i) < x_i$ instead. False optimism from concealing $\beta(i)$ would lead to policy adoption with too high of a probability.

Thus, for any minipublic that is feasible in the commitment game, the policymaker is not worse off if citizens lack commitment in disclosure. However, there might exist minipublics which are not feasible in the commitment game but are informative in the no-commitment game. Therefore, the policymaker attains weakly higher informativeness in the no-commitment game.

Delegation of decisional authority. The curse of too little information arises because minipublic citizens fear the misuse of evidence by the policymaker. If the minipublic were in charge of making the adoption decision instead, there would be no downside to discovering evidence, so any minipublic would be active. Does the policymaker benefit from delegating the adoption decision to a minipublic? We argue that delegation is a useful but limited instrument in resolving the conflict between the policymaker and the citizens.

To address this, we consider the extreme case in which no minipublic is active if the policymaker holds decisional authority. In such a case, the policymaker stands to benefit greatly from delegation, given that without delegation she is bound to act without further evidence. If she were to delegate the decision, her minipublic of choice would be the first-best one $\mathbf{m}_n^f$ and the attained informativeness would be $\bar{\sigma}_n^2$. The policymaker's payoff if the decision is not delegated is

$$V_P(0) = \bar{B}\Phi\left(\frac{\bar{B}}{\tau}\right) + \tau\phi\left(\frac{\bar{B}}{\tau}\right).$$

If the policymaker delegates the decision to the first-best minipublic, the policy gets adopted if and only if $B_{\mathbf{m}} \geqslant 0$, so her payoff is

$$\begin{aligned} V_P^{delegation}(\bar{\sigma}_n^2) &= \mathbb{E}\left[B_{\mathbf{m}} - c | B_{\mathbf{m}} \geq 0\right] \Pr\left[B_{\mathbf{m}} \geq 0\right] \\ &= \bar{B}\Phi\left(\frac{\bar{B}}{\bar{\sigma}_n}\right) + \bar{\sigma}_n\phi\left(\frac{\bar{B}}{\bar{\sigma}_n}\right). \end{aligned}$$

Note the similarity in how these expected payoffs depend, respectively, on τ^2, the political uncertainty, and $\bar{\sigma}_n^2$, the highest feasible informativeness with the given capacity. The relative ranking of τ^2 and $\bar{\sigma}_n^2$ determines whether the policymaker benefits from delegating decisional authority. By delegating, the policymaker trades her ability to fine-tune the adoption decision to her realized threshold for a more informed decision. The ability to fine-tune is more valuable when political uncertainty is high relative to how informative the first-best minipublic is. Therefore, the policymaker prefers not to delegate when $\tau^2 > \bar{\sigma}_n^2$. By contrast, when political uncertainty pales in comparison to the informativeness of the first-best minipublic, i.e., when $\tau^2 < \bar{\sigma}_n^2$, delegation is beneficial.[24] For instance, when the policymaker's threshold is as uncertain as each citizen's outcome, i.e., when $\tau^2 = 1$, she never prefers to delegate because $\bar{\sigma}_n^2 < 1$ for any n.

Biased policymaker. So far, we focused on the presence of political uncertainty rather than the presence of an ex ante bias in the policymaker's threshold. The ex ante alignment assumed in the baseline model is natural in many applications in which the realized threshold of the policymaker is due to idiosyncratic random shocks observed after the minipublic discovers its evidence. In other applications, an ex ante biased policymaker might be more natural. That is, $c \sim \mathcal{N}(\bar{c}, \tau^2)$, where $\bar{c} \neq 0$. Our technical arguments extend to this case.

For a stark illustration, let $\tau^2 = 0$ and $\bar{c} < 0$: the policymaker is with certainty less demanding toward the policy than the citizens. Informativeness continues to be a sufficient statistic for players' payoffs, and the policymaker's problem once again reduces to maximizing minipublic informativeness. Importantly, the citizen's payoff continues to be quasiconvex in informativeness.[25] In this case, however, there is one important difference from our analysis. The citizen's payoff reaches a minimum at $\max\{0, \bar{c}(\bar{B} - \bar{c})\}$, which depends on the direction of policy sentiment $\bar{B}$ rather than just its magnitude $|\bar{B}|$.

The citizens are always active for policies with $\bar{B} > \bar{c}$, so the first-best minipublic is feasible. The policymaker favors adoption ex ante, and evidence discovery is the citizens' way of sowing skepticism and swaying her toward the status quo. For policies with a sufficiently negative sentiment $\bar{B} < \bar{c}$, citizens might have an incentive to forgo evidence discovery because the policymaker is less demanding than the citizens. That is, citizens do not want to risk rocking the boat since they fear inducing the policymaker to adopt a policy with a post-minipublic value in $(\bar{c}, 0)$, to which the status quo is preferable in their eyes.

Policymaker turnover. The baseline model assumed that, when choosing a minipublic, the policymaker is certain that she will be in charge of the adoption decision after evidence is discovered and her threshold of adoption is realized. But this might well not be the case if there is turnover in policymakers, e.g., if the policymaker faces term limits and there is a substantial delay between the minipublic choice stage and the adoption stage. In such a case the policymaker herself faces

[24] Per proposition 4.2, for τ^2 sufficiently low or sufficiently high, the first-best minipublic is feasible even without delegation.

[25] For a deterministic threshold $\bar{c} \neq 0$, the citizens' expected payoff is $V_C(\sigma^2; \bar{c}) = \bar{B}\Phi\left(\frac{\bar{B}-\bar{c}}{\sigma}\right) + \sigma\phi\left(\frac{\bar{B}-\bar{c}}{\sigma}\right)$.

political uncertainty at the minipublic choice stage. If she evaluates the policy with the same zero threshold as the citizens, how is the optimal minipublic affected by this uncertainty?

In the presence of policymaker turnover, the optimal minipublic is never distorted in its composition. The policymaker chooses either the first-best minipublic or the empty one. This is because the policymaker's interests are perfectly aligned with those of the citizens —it is, in fact, as if one of the citizens is choosing the minipublic. Hence, the policymaker's payoff at the minipublic choice stage is exactly $V_C(\cdot)$. The quasiconvexity of the payoff implies that the policymaker prefers either the most or the least informative minipublic. Incentivizing citizens to be active in the first-best minipublic is also no longer a concern: if the policymaker prefers the first-best minipublic to the empty one, then so does each citizen in the first-best minipublic.[26]

6 Concluding remarks

This paper studies how political uncertainty dampens the incentives of minipublic citizens to discover local evidence. It does so in a framework with costless and transparent evidence, civic-minded citizens, and no initial conflict of interest between the citizens and the policymaker. We have found that incentivizing evidence discovery is more challenging (i) for moderate political uncertainty, (ii) for strong policy sentiment, (iii) in smaller minipublics, and (iv) for heterogeneous citizenries. To restore incentives for evidence discovery, the policymaker resorts to minipublics that are not sufficiently diverse, either informationally or demographically. The policymaker relies too heavily on citizens close to the median citizen, at the expense of more peripheral citizens. At its extreme, no evidence discovery is possible in any minipublic.

Our framework can prove useful for answering several policy-relevant questions that are beyond the scope of this paper. First, in many minipublics, citizens vote on the policy recommendation. How does this voting stage impact the incentives of citizens to discover their local evidence, if minipublic citizens face uncertainty regarding each other's thresholds for adoption? Second, a policymaker's consultation with a minipublic often precedes a citizenry-wide referendum in which all citizens vote on the policy decision. Examples include the Irish referendums on same-sex marriage in 2015 and on abortion in 2018, as well as a referendum tentatively planned for 2021 in France on a recommendation of the CCC to add environmental protection to the French constitution.[27] It would

[26] Every citizen $i^f \in \mathbf{m}_n^f$ prefers to be active because $V_C(\sigma^2_{\mathbf{m}_n^f \setminus i^f}) \leqslant \max\{V_C(0), V_C(\sigma^2_{\mathbf{m}_n^f})\} = V_C(\sigma^2_{\mathbf{m}_n^f})$.

[27] The Irish referendums were preceded, respectively, by the Constitutional Convention of 2012-14 and the Citizens' Assembly of 2016-18, set up by the Irish government to examine various constitutional reforms. For the planned French referendum, see https://www.politico.eu/article/macron-agrees-to-add-environmental-protection-as-a-constitutional-duty/.

be fruitful to study how the format of the referendum, as well as potential conflict among citizens of different backgrounds regarding the threshold for policy adoption, might curtail a minipublic's ability to discover useful evidence. Third, in our framework, the policymaker does not benefit from randomizing over minipublics. Yet, most sampling protocols for minipublics involve some form of randomization. Are there features of real-world minipublics that would rationalize randomization by the policymaker, or is randomization observed in practice because the policymaker cannot target citizens with arbitrary precision? We leave these questions for future work.

References

Au, Pak Hung, and Keiichi Kawai. 2019. "Competitive Disclosure of Correlated Information." *Economic Theory*, 1–33.

Au, Pak Hung, and Keiichi Kawai. 2020. "Competitive Information Disclosure by Multiple Senders." *Games and Economic Behavior*, 119: 56–78.

Austen-Smith, David, and Jeffrey Banks. 1996. "Information Aggregation, Rationality and the Condorcet Jury Theorem." *American Political Science Review*, 90: 34–45.

Bardhi, Arjada. 2020. "Attributes: Selective Learning and Influence." Working Paper.

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan. 2020. "The Economics of Social Data." Working Paper.

Boleslavsky, Raphael, and Christopher Cotton. 2018. "Limited Capacity in Project Selection: Competition Through Evidence Production." *Economic Theory*, 65(2): 385–421.

Breckon, Jonathan, Anna Hopkins, and Ben Rickey. 2019. "Evidence vs Democracy: How 'Mini-publics' Can Traverse the Gap Between Citizens, Experts, and Evidence." Alliance for Useful Evidence, NESTA.

Cai, Hongbin. 2009. "Costly Participation and Heterogeneous Preferences in Informational Committees." *The RAND Journal of Economics*, 40: 173–189.

Callander, Steven. 2011. "Searching and Learning by Trial and Error." *American Economic Review*, 101(6): 2277–2308.

Callander, Steven, and Tom S. Clark. 2017. "Precedent and Doctrine in a Complicated World." *American Political Science Review*, 111(1): 184–203.

Callander, Steven, Nicolas S. Lambert, and Niko Matouschek. 2018. "The Power of Referential Advice." Working Paper.

Chade, Hector, and Edward E. Schlee. 2002. "Another Look at the Radner–Stiglitz Nonconcavity in the Value of Information." *Journal of Economic Theory*, 107: 421–452.

Chade, Hector, and Jan Eeckhout. 2018. "Matching Information." *Theoretical Economics*, 13: 377–414.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv. 2018. "Deliberating Collective Decisions." *The Review of Economic Studies*, 85(2): 929–963.

Dahl, Robert A. 1989. *Democracy and its Critics.* Yale University Press.

Dewatripont, Mathias, and Jean Tirole. 1999. "Advocates." *Journal of Political Economy*, 107(1): 1–39.

Doob, J. L. 1942. "The Brownian Movement and Stochastic Equations." *Annals of Mathematics*, 43(2): 351–369.

Elstub, Stephen. 2014. "Mini-publics: Issues and Cases." In *Deliberative Democracy: Issues and Cases.* , ed. Stephen Elstub and Peter McLaverty, 166–188. Edinburgh University Press.

Fabre, Adrien, Bénédicte H. Apouey, Thomas Douenne, Louis-Gaëtan Giraudet, Jean-François Laslier, and Antonin Macé. 2020. "Convention Citoyenne pour le Climat: Les Citoyens de la Convention Comparés à des Échantillons Représentatifs de la Population Française." Working Paper.

Feddersen, Timothy, and Wolfgang Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections with Private Information." *Econometrica*, 65(5): 1029–1058.

Fishkin, James S. 2011. *When the People Speak: Deliberative Democracy and Public Consultation.* Oxford University Press.

Gentzkow, Matthew, and Emir Kamenica. 2016. "Competition in Persuasion." *The Review of Economic Studies*, 84(1): 300–322.

Gentzkow, Matthew, and Emir Kamenica. 2017. "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces." *Games and Economic Behavior*, 104: 411–429.

Gerardi, Dino, and Leeat Yariv. 2008. "Costly expertise." *American Economic Review*, 98(2): 187–93.

Gradwohl, Ronen, and Timothy Feddersen. 2018. "Persuasion and Transparency." *The Journal of Politics*, 80(3): 903–915.

Hong, Lu, and Scott E. Page. 2001. "Problem Solving by Heterogeneous Agents." *Journal of Economic Theory*, 97: 123–163.

Jacobs, Daan, and Wesley Kaufmann. 2019. "The Right Kind of Participation? The Effect of a Deliberative Mini-public on the Perceived Legitimacy of Public Decision-making." *Public Management Review*, 1–21.

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

Jovanovic, Boyan, and Rafael Rob. 1990. "Long Waves and Short Waves: Growth Through Intensive and Extensive Search." *Econometrica*, 1391–1409.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.

Koessler, Frédéric, Marie Laclau, and Tristan Tomala. 2018. "Interactive Information Design." *HEC Paris Research Paper No. ECO/SCD-2018-1260.*

Kwiek, Maksymilian. 2020. "Information Acquisition in Citizen and Representative Assemblies." Working Paper.

Lamberson, P.J., and Scott E. Page. 2012. "Optimal Forecasting Groups." *Management Science*, 58(4): 805–810.

Liang, Annie, Xiaosheng Mu, and Vasilis Syrgkanis. 2020. "Dynamically Aggregating Diverse Information." Working Paper.

Li, Fei, and Peter Norman. 2018. "On Bayesian Persuasion with Multiple Senders." *Economics Letters*, 170: 66–70.

Martinelli, César. 2006. "Would rational voters acquire costly information?" *Journal of Economic Theory*, 129(1): 225–251.

Matsusaka, John G. 2005. "Direct Democracy Works." *Journal of Economic Perspectives*, 19(2): 185–206.

Matsusaka, John G., and Nolan M. McCarty. 2001. "Political Resource Allocation: Benefits and Costs of Voter Initiatives." *Journal of Law, Economics, and Organization*, 17(2): 413–448.

Michels, Ank, and Laurens De Graaf. 2010. "Examining Citizen Participation: Local Participatory Policy Making and Democracy." *Local Government Studies*, 36(4): 477–491.

Milgrom, Paul, and John Roberts. 1986. "Relying on the Information of Interested Parties." *The RAND Journal of Economics*, 17: 18–32.

Moldovanu, Benny, and Xianwen Shi. 2013. "Specialization and Partisanship in Committee Search." *Theoretical Economics*, 8(3): 751–774.

Name Correa, Àlvaro, and Huseyin Yildirim. 2020. "Biased Experts, Majority Rule, and the Optimal Composition of Committee." Working Paper.

OECD. 2020. "Innovative Citizen Participation and New Democratic Institutions." OECD Publishing. https://www.oecd-ilibrary.org/sites/339306da-en/index.html?itemId=/content/publication/339306da-en.

Prat, Andrea. 2002. "Should a Team be Homogeneous?" *European Economic Review*, 46: 1187–1207.

Prato, Carlo, and Bruno Strulovici. 2017. "The Hidden Cost of Direct Democracy: How Ballot Initiatives Affect Politicians' Selection and Incentives." *Journal of Theoretical Politics*, 29(3): 440–466.

Radner, Roy, and Joseph E. Stiglitz. 1984. "A Non-Concavity in the Value of Information." In *Bayesian Models in Economic Theory.* Vol. 5 of *Studies in Bayesian Econometrics.* North-Holland.

Strulovici, Bruno. 2010. "Learning while voting: Determinants of collective experimentation." *Econometrica*, 78(3): 933–971.

Tan, Xu, and Quan Wen. 2020. "Information Acquisition and Voting with Heterogeneous Experts." *The RAND Journal of Economics*, 51(4): 1063–1092.

Visser, Bauke, and Otto H. Swank. 1999. "On Committees of Experts." *The Quarterly Journal of Economics*, 122(1): 337–372.

Warren, Mark E, and John Gastil. 2015. "Can Deliberative Minipublics Address the Cognitive Challenges of Democratic Citizenship?" *The Journal of Politics*, 77(2): 562–574.

A Axiomatization of the correlation structure

The Ornstein-Uhlenbeck process introduced in Assumption 1 uniquely satisfies the following set of natural axioms on the outcome mapping β.

A.1 *(Principle of maximal ignorance)* For any group of citizens $\{i_1, \ldots, i_n\}$, outcomes $\{\beta(i_1), \ldots, \beta(i_n)\}$ follow a multivariate Gaussian distribution.

A.2 *(Similar citizens, similar outcomes)* $\beta(\cdot)$ is almost surely continuous.

A.3 *(Identical outcome uncertainty)* For each $i \in [0, 1]$, $\beta(i) - \bar{\beta}(i) \sim \mathcal{N}(0, 1)$.

A.4 *(Distance-based correlation)* For any two $i_1, i_2 \in [0, 1]$, the correlation between $\beta(i_1)$ and $\beta(i_2)$ depends only on the distance $|i_1 - i_2|$.

A.5 *(Look to your left, look to your right)* For any $i_1 < \ldots < i_k < \ldots i_n$, the distribution of $\beta(i_k)$ depends on the outcomes of other citizens in the set only through $\beta(i_{k-1})$ and $\beta(i_{k+1})$.

Axiom A.1 imposes a general Gaussian structure, whereas Axioms A.2-A.5 specify additional properties. A.1 can also be interpreted as a maximal-ignorance desideratum. The Gaussian distribution maximizes entropy among all unbounded distributions of a fixed mean and variance, therefore the Gaussian structure allows one to draw the weakest conclusions possible from a set of outcomes. A.2 requires that for any two citizens that are arbitrarily close to each other, their realized outcomes are also close. A.3 requires that all citizens face the same uncertainty about their outcomes. In understanding how informative a citizen's outcome is for the rest of the citizenry, this axiom allows us to isolate the role of the citizen's position in $[0, 1]$ from the role of the outcome uncertainty that he faces. Axioms A.4 and A.5 specify how a citizen's position determines his correlation to other citizens. Correlation between any two citizens' outcomes depends only on how far the two citizens are from each other (A.4). Moreover, given a set of citizens the outcomes of which are observed, the best conjecture for the outcome of any citizen outside this set depends only on the outcomes of his closest neighbors in this set (A.5).

Corollary A.1, which follows from Theorem 1.1 of Doob (1942), establishes that not only does the Ornstein-Uhlenbeck process satisfy this set of natural axioms A.1-A.5, but it is the *only* nontrivial stochastic process that does so.

Corollary A.1 (Doob (1942))**.** *Suppose $\bar{\beta}(i)$ is continuous. Then, the Ornstein-Uhlenbeck process on domain $[0, 1]$ uniquely satisfies assumptions A.1-A.5.*

B Proofs for section 3

Proof for Lemma 3.1. From Lemma 3.3 in Bardhi (2020), the post-minipublic value $B(\hat{\mathbf{m}})$ is

$$B_{\hat{\mathbf{m}}} = \bar{B} + \sum_{j=1}^{k} \gamma_j(\hat{\mathbf{m}}) \left(\beta(i_j) - \bar{\beta}(i_j)\right), \tag{8}$$

where $\gamma_j(\hat{\mathbf{m}})$ is as stated in lemma 3.1. By Theorem 3.1 in Bardhi (2020), the distribution of $B_{\hat{\mathbf{m}}}$ is

$$\sigma^2_{\hat{\mathbf{m}}} = \sum_{j=1}^{k} \sum_{h=1}^{k} \gamma_j(\hat{\mathbf{m}}) \gamma_h(\hat{\mathbf{m}}) e^{-|i_j - i_h|/\ell}.$$

The informativeness induced by the active minipublic $\hat{\mathbf{m}}$ can be rewritten as

$$\sigma^2_{\hat{\mathbf{m}}} = \sum_{j=1}^{k} \gamma_j(\hat{\mathbf{m}}) \left(\gamma_j(\hat{\mathbf{m}}) + \sum_{h \neq j} \gamma_h(\hat{\mathbf{m}}) e^{-|i_h - i_j|/\ell} \right).$$

We distinguish two cases, based on whether $j \in \{1, k\}$ or $j \in \{2, \ldots, k-1\}$. First, let $j = 1$.

$$\begin{aligned}
\left(\gamma_1(\hat{\mathbf{m}}) + \sum_{h \geqslant 2} \gamma_h(\hat{\mathbf{m}}) e^{-|i_h - i_1|/\ell} \right) &= \ell \left(1 - e^{-i_1/\ell} + \tanh\left(\frac{i_2 - i_1}{2\ell} \right) \right) \\
&\quad + \sum_{h=2}^{k-1} e^{-(i_h - i_1)/\ell} \ell \left(\tanh\left(\frac{i_h - i_{h-1}}{2\ell} \right) + \tanh\left(\frac{i_{h+1} - i_h}{2\ell} \right) \right) \\
&\quad + e^{-(i_k - i_1)/\ell} \ell \left(1 - e^{-(1 - i_k)/\ell} + \tanh\left(\frac{i_k - i_{k-1}}{2\ell} \right) \right) \\
&= \ell \left(1 - e^{-i_1/\ell} - e^{-(1 - i_1)/\ell} \right) + e^{-(i_k - i_1)/\ell} \ell \\
&\quad + \sum_{h=2}^{k} \ell \left(e^{-(i_{h-1} - i_1)/\ell} + e^{-(i_h - i_1)/\ell} \right) \tanh\left(\frac{i_h - i_{h-1}}{2\ell} \right) \\
&= \ell \left(1 - e^{-i_1/\ell} - e^{-(1 - i_1)/\ell} \right) + e^{-(i_k - i_1)/\ell} \ell \\
&\quad + \sum_{h=2}^{k} \ell \left(e^{-(i_{h-1} - i_1)/\ell} - e^{-(i_h - i_1)/\ell} \right) \\
&= \ell \left(2 - e^{-i_1/\ell} - e^{-(1 - i_1)/\ell} \right),
\end{aligned}$$

where the second equality rearranges the last additive term $\gamma_k(\hat{\mathbf{m}}) e^{-|i_k - i_1|/\ell}$ and the third equality uses the observation that for any $h \geqslant 2$,

$$\left(e^{-(i_{h-1} - i_1)/\ell} + e^{-(i_h - i_1)/\ell} \right) \tanh\left(\frac{i_h - i_{h-1}}{2\ell} \right) = \left(e^{-(i_{h-1} - i_1)/\ell} - e^{-(i_h - i_1)/\ell} \right).$$

The last equality follows from cancelling opposite-sign terms. For any singleton minipublic $\{i_1\}$, the informativeness is

$$\sigma^2_{i_1} = \ell^2 \left(2 - e^{-i_1/\ell} - e^{-(1 - i_1)/\ell} \right)^2.$$

This gives us the result. The case of $j = k$ is similar and hence omitted.

Next, suppose $j \in \{2, \dots, k-1\}$. Rearranging terms in a similar way to the case of $j = 1$, we obtain

$$\begin{aligned}
&\left(\gamma_j(\hat{\mathbf{m}}) + \sum_{h \neq j} \gamma_h(\hat{\mathbf{m}}) e^{-|i_h - i_j|/\ell}\right) \\
&= \ell(1 - e^{-i_1/\ell}) e^{-(i_j - i_1)/\ell} + \ell \left(1 - e^{-(1-i_k)/\ell}\right) e^{-(i_k - i_j)/\ell} + \sum_{h=2}^{k} \ell \left(e^{-|i_h - i_j|/\ell} + e^{-|i_{h-1} - i_j|/\ell}\right) \tanh\left(\frac{i_h - i_{h-1}}{2\ell}\right) \\
&= \ell \left(e^{-(i_j - i_1)/\ell} + e^{-(i_k - i_j)/\ell} - e^{-i_j/\ell} - e^{-(1-i_j)/\ell}\right) + \\
&\qquad + \sum_{h=2}^{j} \ell \left(e^{-(i_j - i_h)/\ell} - e^{-(i_j - i_{h-1})/\ell}\right) + \sum_{h=j}^{k-1} \ell \left(e^{-(i_h - i_j)/\ell} - e^{-(i_{h+1} - i_j)/\ell}\right) \\
&= \ell \left(2 - e^{-i_j/\ell} - e^{-(1-i_j)/\ell}\right).
\end{aligned}$$

■

Proof for Lemma 3.2. Because c and $B_{\hat{\mathbf{m}}}$ are independent Gaussian variables, the policymaker's ex post payoff $B_{\hat{\mathbf{m}}} - c$ is distributed according to $B_{\hat{\mathbf{m}}} - c \sim \mathcal{N}(\bar{B}, \tau^2 + \sigma^2_{\hat{\mathbf{m}}})$. The policymaker observes $B_{\hat{\mathbf{m}}} - c$ and adopts the policy if $B_{\hat{\mathbf{m}}} - c > 0$. Hence, the probability of adoption is

$$\Pr\left[B_{\hat{\mathbf{m}}} - c > 0\right] = 1 - \Phi\left(\frac{-\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right) = \Phi\left(\frac{\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right).$$

Let $\lambda(x) \coloneqq \phi(x)/(1 - \Phi(x))$ denote the inverse Mills ratio. The following is a standard result about the conditional expectation of a joint Gaussian distribution.

Lemma B.1. *Let X, Y be two jointly Gaussian random variables with respective means μ_x, μ_y, respective variances σ_x^2, σ_y^2, and covariance $\mathrm{Cov}[X,Y]$. Then, $\mathbb{E}[X|Y > y] = \mu_x + \frac{\mathrm{Cov}[X,Y]}{\sigma_y}\lambda\left(\frac{y-\mu_y}{\sigma_y}\right)$.*

Applying lemma B.1 for $X = Y = B_{\hat{\mathbf{m}}} - c$ and $y = 0$, the expected payoff of the policymaker conditional on adoption is

$$\mathbb{E}[B_{\hat{\mathbf{m}}} - c | B_{\hat{\mathbf{m}}} - c > 0] = \bar{B} + \sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}\lambda\left(\frac{-\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right).$$

Applying lemma B.1 for $X = B_{\hat{\mathbf{m}}}$, $Y = B_{\hat{\mathbf{m}}} - c$ and $y = 0$ (and thus, $\mathrm{Cov}[B_{\hat{\mathbf{m}}}, B_{\hat{\mathbf{m}}} - c] = \sigma^2_{\hat{\mathbf{m}}}$), the expected payoff of the citizens conditional on adoption is

$$\mathbb{E}[B_{\hat{\mathbf{m}}} | B_{\hat{\mathbf{m}}} - c > 0] = \bar{B} + \frac{\sigma^2_{\hat{\mathbf{m}}}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\lambda\left(\frac{-\bar{B}}{\sqrt{\tau^2 + \sigma^2_{\hat{\mathbf{m}}}}}\right).$$

The unconditional expected payoff is $\Pr\left[B_{\hat{\mathbf{m}}} - c > 0\right] \mathbb{E}[B_{\hat{\mathbf{m}}} - c | B_{\hat{\mathbf{m}}} - c > 0]$ for the policymaker, and $\Pr(B_{\hat{\mathbf{m}}} - c > 0)\mathbb{E}[B_{\hat{\mathbf{m}}} | B_{\hat{\mathbf{m}}} - c > 0]$ for the citizens. Plugging in the above expressions yields the result. ■

Proof of Lemma 3.3. (i) From lemma 3.2, payoffs V_P and V_C are continuous and differentiable. Differentiating the expected payoff of the policymaker with respect to σ^2, we obtain

$$\frac{\partial V_P(\sigma^2)}{\partial \sigma^2} = \frac{\phi\left(\frac{\bar{B}}{\sqrt{\tau^2+\sigma^2}}\right)}{2\sqrt{\sigma^2+\tau^2}} > 0.$$

(ii) Differentiating the expected payoff of the citizen with respect to σ^2, we obtain

$$\frac{\partial V_C(\sigma^2)}{\partial \sigma^2} = \frac{\phi\left(\frac{\bar{B}}{\sqrt{\tau^2+\sigma^2}}\right)}{2\sqrt{\sigma^2+\tau^2}} \left(\sigma^4 - (\bar{B}^2 - 3\sigma^2)\tau^2 + 2\tau^4\right).$$

Because $\sigma^2 > 0$, the only admissible root of the quadratic $\left(\sigma^4 - (\bar{B}^2 - 3\sigma^2)\tau^2 + 2\tau^4\right)$ is

$$\sigma^2_{root} = \frac{1}{2}\left(\sqrt{\tau^4 + 4\bar{B}^2\tau^2} - 3\tau^2\right),$$

which is positive if and only if $\bar{B}^2 \geqslant 2\tau^2$. Therefore, the payoff minimum is reached at 0 if $\sigma^2_{root} < 0$ and σ^2_{root} otherwise. V_C is strictly decreasing over $[0, \underline{\sigma}^2)$ and strictly increasing over $(\underline{\sigma}^2, \infty)$. ■

Lemma B.2. *Fix σ^2 and $\tilde{\sigma}^2$ such that $\sigma^2 < \tilde{\sigma}^2$.*

(i) Suppose that $V_C(\sigma^2) > V_C(\tilde{\sigma}^2)$. Then, for any σ_1^2 and σ_2^2 such that $\sigma_1^2 \leqslant \sigma^2$, $\sigma_2^2 \leqslant \tilde{\sigma}^2$ and $\sigma_1^2 < \sigma_2^2$, it holds that $V_C(\sigma_1^2) > V_C(\sigma_2^2)$.

(ii) Suppose that $V_C(\sigma^2) \leqslant V_C(\tilde{\sigma}^2)$. Then, for any σ_1^2 and σ_2^2 such that $\sigma_1^2 \geqslant \sigma^2$, $\sigma_2^2 \geqslant \tilde{\sigma}^2$ and $\sigma_1^2 < \sigma_2^2$, it holds that $V_C(\sigma_1^2) \leqslant V_C(\sigma_2^2)$.

Proof. Let $V_C(\sigma^2) > V_C(\tilde{\sigma}^2)$. As established in lemma 3.3, $V_C(\cdot)$ is strictly quasiconvex with a minimum at $\underline{\sigma}^2$. Because $\sigma^2 < \tilde{\sigma}^2$ and $V_C(\sigma^2) > V_C(\tilde{\sigma}^2)$, it must be that $\underline{\sigma}^2 \geqslant \sigma^2$. If $\sigma_2^2 \leq \sigma^2$, the statement follows by the fact that V_C is strictly decreasing to the left of $\underline{\sigma}^2$. If $\sigma_2^2 \in (\sigma^2, \tilde{\sigma}^2]$ instead, then by the strict quasiconvexity of V_C, it follows that

$$V_C(\sigma_1^2) \geq V_C(\sigma^2) = \max\left\{V_C(\sigma^2), V_C(\tilde{\sigma}^2)\right\} > V_C(\sigma_2^2).$$

This establishes part (i). Part (ii) follows by a similar argument. ■

Proof of Proposition 3.1. First, we prove that the restriction to deterministic minipublics $\mathbf{m} \in \mathcal{M}_n$ is without loss. Suppose the policymaker chooses a random lottery with at least two distinct minipublics $\mathbf{m} \neq \mathbf{m}'$ in its support. Let $\hat{\mathbf{m}}$ and $\hat{\mathbf{m}}'$ denote the random subset of active citizens in $\mathbf{m}$ and $\mathbf{m}'$ respectively.[28] Then, the policymaker's expected payoff from $\mathbf{m}$ and $\mathbf{m}'$ respectively is $\mathbb{E}_{\hat{\mathbf{m}}}[V_p(\hat{\mathbf{m}})]$ and $\mathbb{E}_{\hat{\mathbf{m}}'}[V_p(\hat{\mathbf{m}}')]$. Citizens observe the realized minipublic from the policymaker's lottery. Hence, the policymaker randomizes over minipublics only if she is indifferent between them, i.e., $\mathbb{E}_{\hat{\mathbf{m}}}[V_p(\hat{\mathbf{m}})] = \mathbb{E}_{\hat{\mathbf{m}}'}[V_p(\hat{\mathbf{m}}')]$. But then, the policymaker obtains the same payoff by choosing either $\mathbf{m}$ or $\mathbf{m}'$ deterministically.

Next, we prove that the restriction to pure strategies at the evidence discovery stage, $\delta_i(\mathbf{m}) \in \{0, 1\}$ for all $\mathbf{m} \in \mathcal{M}_n$ and for all $i \in \mathbf{m}$, is without loss. Consider a minipublic $\mathbf{m} \in \mathcal{M}_n$ and an equilibrium δ at the evidence discovery stage in which $\delta_i \in (0, 1)$ for at least some $i \in \mathbf{m}$. Let $\mathbf{m}^A := \{i \in \mathbf{m} : \delta_i = 1\}$ be the set of active citizens, $\mathbf{m}^P := \{i \in \mathbf{m} : \delta_i = 0\}$ the set of passive citizens, and $\mathbf{m}^M := \{i \in \mathbf{m} : \delta_i \in (0, 1)\}$ the remaining set of citizens who strictly mix between being active and passive, such that $\mathbf{m} = \mathbf{m}^A \cup \mathbf{m}^P \cup \mathbf{m}^M$.

[28]Given a minipublic $\mathbf{m}$, the citizen's equilibrium strategy δ induces a lottery $\Delta(\mathcal{P}(\mathbf{m}))$ over the power set of $\mathbf{m}$.

For the same minipublic $\mathbf{m}$, we construct an alternative equilibrium $\tilde{\delta}$ at the evidence discovery stage in which $\tilde{\delta}_i = 1$ for all $i \in \mathbf{m}^A \cup \mathbf{m}^M$. The set of active citizens in equilibrium $\tilde{\delta}$ is a strict superset of the active citizens in δ with positive probability, therefore $\tilde{\delta}$ guarantees a strictly higher minipublic informativeness with positive probability and thus a strictly higher expected payoff for the policymaker.

Lemma B.3 below establishes that in $\tilde{\delta}$, every citizen i in the active set of citizens $\mathbf{m}^A \cup \mathbf{m}^M$ prefers to be active if he expects the set $\mathbf{m}^A \cup \mathbf{m}^M \backslash i$ to be active and the set $\mathbf{m}^P$ to be passive.

Lemma B.3. *For a minipublic $\mathbf{m}$, let δ be an equilibrium in which the nonempty set $\mathbf{m}^M$ mixes between being active and passive. Then, $V_C(\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M}) \geq V_C(\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M \backslash i})$ for every $i \in \mathbf{m}^A \cup \mathbf{m}^M$.*

Proof. By contradiction, suppose there exists $i \in \mathbf{m}^A \cup \mathbf{m}^M$ such that $V_C(\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M}) < V_C(\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M \backslash i})$. For any realized subset of citizens $S \subset \mathbf{m}^M$ in the original equilibrium δ, it holds that (i) $\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M} \geq \sigma^2_{\mathbf{m}^A \cup S}$, (ii) $\sigma^2_{\mathbf{m}^A \cup \mathbf{m}^M \backslash i} \geq \sigma^2_{\mathbf{m}^A \cup S \backslash i}$, and (iii) $\sigma^2_{\mathbf{m}^A \cup S \backslash i} < \sigma^2_{\mathbf{m}^A \cup S}$. Then, lemma B.2(i) implies that for any $S \subset \mathbf{m}^M$, $V_C(\sigma^2_{\mathbf{m}^A \cup S \backslash i}) > V_C(\sigma^2_{\mathbf{m}^A \cup S})$. This yields the contradiction that citizen i has a strictly profitable deviation in δ, because being passive yields a strictly higher payoff than being active for every realization of S. ■

Thus, in the alternative candidate equilibrium $\tilde{\delta}$ where $\tilde{\delta}_j = 1$ if $\delta_j \in (0, 1]$ and $\tilde{\delta}_i = 0$ if $\delta_i = 0$, none of the active citizen can profit from being passive. However, it might be that the passive citizens can now strictly gain from being active and distort the incentives of the active citizens. We prove that if this is the case, no active citizen can gain from being passive.

Lemma B.4. *Suppose citizen i prefers to be active when the set of active citizens is $\hat{\mathbf{m}}$. Then, citizen i prefers to be active for any larger set of active citizens $\hat{\mathbf{m}}'$ such that $\hat{\mathbf{m}}' \supset \hat{\mathbf{m}}$.*

Proof. By assumption, $V_C(\sigma^2_{\hat{\mathbf{m}}}) \geq V_C(\sigma^2_{\hat{\mathbf{m}} \backslash i})$. For any $\hat{\mathbf{m}}'$, it holds that (i) $\sigma^2_{\hat{\mathbf{m}}'} \geq \sigma^2_{\hat{\mathbf{m}}}$, (ii) $\sigma^2_{\hat{\mathbf{m}}' \backslash i} \geq \sigma^2_{\hat{\mathbf{m}} \backslash i}$, and (iii) $\sigma^2_{\hat{\mathbf{m}}' \backslash i} < \sigma^2_{\hat{\mathbf{m}}'}$. Then, lemma B.2(ii) implies that for any $\hat{\mathbf{m}}' \supset \hat{\mathbf{m}}$, $V_C(\hat{\mathbf{m}}' \backslash i) \leq V_C(\hat{\mathbf{m}}')$. ■

Lemma B.4 implies that even if a subset of passive citizens in δ prefer to be active in $\tilde{\delta}$, this will not distort the incentives of the active citizens in $\tilde{\delta}$. Therefore, we have identified an equilibrium $\tilde{\delta}$ that is strictly preferred by the policymaker. This contradicts the optimality of the original equilibrium δ.

Finally, we prove that it is without loss to include only active citizen in the minipublic. Let $\mathbf{m}$ be an optimal minipublic in which a deterministic subset $\hat{\mathbf{m}} \subsetneq \mathbf{m}$ is active. Consider an alternative minipublic $\mathbf{m}' := \hat{\mathbf{m}}$ that consists only of the active citizens in $\mathbf{m}$. When every citizen in $\hat{\mathbf{m}}$ is active, then the passive and active informativeness of each $i \in \hat{\mathbf{m}}$ coincide, $\sigma^2_{\mathbf{m}'} = \sigma^2_{\mathbf{m}}$ and $\sigma^2_{\mathbf{m}' \backslash i} = \sigma^2_{\mathbf{m} \backslash i}$. Hence, the incentives to be active are the same for every $i \in \hat{\mathbf{m}}$ in the policymaker-preferred equilibrium in both minipublics, and the policymaker obtains the same expected payoff from $\mathbf{m}$ and $\mathbf{m}'$. ■

Proof for Corollary 3.1. This follows immediately from Proposition 3.4 in Bardhi (2020). The proof is therefore omitted. ■

Lemma B.5 (First-best informativeness $\bar{\sigma}^2_n$)**.**

(i) $\bar{\sigma}^2_1 > 0$;

(ii) $\bar{\sigma}^2_n$ *strictly increases in* n;

(iii) as $n \to \infty$, $\bar{\sigma}^2_n \to 2\ell\left(1 - \ell\left(1 - e^{-1/\ell}\right)\right)$.

Proof. (i) The informativeness of any single citizen $i \in [0,1]$ is $\sigma_i^2 = \ell^2 \left(2 - e^{-i/\ell} - e^{-(1-i)/\ell}\right)^2 > 0$ (lemma 3.1). Hence, $\bar{\sigma}_1^2 \geqslant \sigma_i^2 > 0$.

(ii) Suppose there exist $n < n'$ such that $\bar{\sigma}_n^2 = \bar{\sigma}_{n'}^2$. Consider a minipublic $\tilde{\mathbf{m}} := \mathbf{m}_n^f \cup M$, where M is a set of $(n'-n)$ distinct citizens such that $\mathbf{m}_n^f \cap M = \emptyset$. Because any two citizens are imperfectly correlated, $\sigma_{\tilde{\mathbf{m}}}^2 > \bar{\sigma}_n^2 = \bar{\sigma}_{n'}^2 = \sigma^2_{\mathbf{m}^f_{n'}}$, which contradicts the fact that $\mathbf{m}_{n'}^f$ is the first-best minipublic.

(iii) By the characterization of the first-best minipublic in corollary 3.1, as $n \to \infty$ the first-best minipublic converges to $i_1^f \to 0$, $i_n^f \to 1$, and the distance between any two adjacent citizens goes to zero. Therefore, $\mathbf{m}_n^f$ defines a partition of $[0,1]$, the mesh of which converges to zero as $n \to \infty$. Moreover, $\gamma_k(\mathbf{m}_n^f) \to 0$ for all $k = 1, \ldots, n$ as $n \to \infty$. By the mean-square continuity of the Ornstein-Uhlenbeck process, this limit-partition property of the first-best minipublic implies that

$$B = \lim_{n\to\infty} B_{\mathbf{m}_n^f}.$$

Hence, the variance of $B_{\mathbf{m}_n^f}$ approaches that of B as $n \to \infty$.

Next we solve for the variance of B.

$$\begin{aligned}
\mathbb{E}[B^2] &= \mathbb{E}\left[\left(\int_0^1 \beta(i)\mathrm{d}i\right)\left(\int_0^1 \beta(j)\mathrm{d}j\right)\right] \\
&= \int_0^1\int_0^1 \Big(\mathrm{Cov}[\beta(i),\beta(j)] + \mathbb{E}[\beta(i)]\mathbb{E}[\beta(j)]\Big)\mathrm{d}i\mathrm{d}j \\
&= \int_0^1\int_0^1 \mathrm{Cov}[\beta(i),\beta(j)]\mathrm{d}i\mathrm{d}j + \left(\mathbb{E}\left[\int_0^1 \beta(i)\mathrm{d}i\right]\right)^2.
\end{aligned}$$

Using the fact that $\mathrm{Var}[B] = \mathbb{E}[B^2] - (\mathbb{E}[B])^2$ and that $\mathrm{Cov}[\beta(i),\beta(j)] = e^{-|i-j|/\ell}$, we obtain

$$\mathrm{Var}[B] = \int_0^1\int_0^1 e^{-|i-j|/\ell}\mathrm{d}i\mathrm{d}j = 2\ell\left(1 - \ell\left(1 - e^{-1/\ell}\right)\right).$$

■

Proof for Lemma 3.4. We first show that for any $j, k \in \{2, \ldots, n-1\}$, the passive informativeness of i_j^f is equal to that of i_k^f. From corollary 3.1, distance $\Delta := i_{j+1}^f - i_j^f$ is the same for all $j \in \{1, \ldots, n-1\}$. The claim holds vacuously for $n \leqslant 3$. Let $n > 3$. Pick any four consecutive citizens $\{i_j^f, i_j^f + \Delta, i_j^f + 2\Delta, i_j^f + 3\Delta\}$, where $j \geqslant 1$. We compare the informativeness of the minipublic $\mathbf{m}_n^f \setminus (i_j^f + \Delta)$ with that of $\mathbf{m}_n^f \setminus (i_j^f + 2\Delta)$. Note that removing $i_j^f + \Delta$ or $i_j^f + 2\Delta$ from $\mathbf{m}_n^f$ does not affect γ_k for citizens outside this set of four consecutive citizens $k \notin \{j, j+1, j+2, j+3\}$. Therefore, using the characterization in lemma 3.1,

$$\begin{aligned}
\sigma^2_{\mathbf{m}_n^f\setminus(i_j^f+\Delta)} - \sigma^2_{\mathbf{m}_n^f\setminus(i_j^f+2\Delta)} &= \ell\left(\tanh\left(\frac{2\Delta}{2\ell}\right) - \tanh\left(\frac{\Delta}{2\ell}\right)\right)\left(\sigma_{i_j^f} - \sigma_{i_j^f+3\Delta}\right) \\
&\qquad + \ell\left(\tanh\left(\frac{2\Delta}{2\ell}\right) + \tanh\left(\frac{\Delta}{2\ell}\right)\right)\left(\sigma_{i_j^f+2\Delta} - \sigma_{i_j^f+\Delta}\right) \\
&= \ell^2 e^{-(1+3\Delta+i_j^f)/\ell}\left(e^{3\Delta/\ell} - 1\right)\left(e^{(3\Delta+2i_j^f)/\ell} - e^{1/\ell}\right)\mathrm{sech}\left(\frac{\Delta}{\ell}\right)\tanh\left(\frac{\Delta}{2\ell}\right) \\
&\quad - \ell^2 e^{-(1+2\Delta+i_j^f)/\ell}\left(e^{\Delta/\ell} - 1\right)\left(e^{(3\Delta+2i_j^f)/\ell} - e^{1/\ell}\right)\left(\tanh\left(\frac{\Delta}{2\ell}\right) + \tanh\left(\frac{2\Delta}{2\ell}\right)\right) \\
&= 0,
\end{aligned}$$

where the second equality uses $\sigma_i = \ell\left(2 - e^{-i/\ell} - e^{-(1-i)/\ell}\right)$ and the last equality follows from the trigonometric identities $\operatorname{sech}(x) = 2/(e^x + e^{-x})$ and $\tanh(x) = (e^x - e^{-x})/(e^x + e^{-x})$. We proved that the passive informativeness of any two adjacent citizens k and $k+1$ such that $1 < k < n-1$ is the same. Therefore, any two non-adjacent citizens $j, k \in \{2, \dots, n-1\}$ have the same passive informativeness as well.

The fact that $\sigma^2_{\mathbf{m}_n^f \setminus i_1^f} = \sigma^2_{\mathbf{m}_n^f \setminus i_n^f}$ follows from the symmetry of $\mathbf{m}_n^f$ about $1/2$. To rank the passive informativeness of a peripheral citizen and that of an inner citizen, without loss we consider the difference $\sigma^2_{\mathbf{m}^f \setminus (i_1^f + \Delta)} - \sigma^2_{\mathbf{m}^f \setminus i_1^f}$. For this difference, we need to only consider the subset of citizens $\{i_1^f, i_1^f + \Delta, i_1^f + 2\Delta\}$. That is,

$$
\begin{aligned}
\sigma^2_{\mathbf{m}_n^f \setminus (i_1^f + \Delta)} - \sigma^2_{\mathbf{m}_n^f \setminus i_1^f} &= \ell\left(1 - e^{-i_1^f/\ell} + \tanh\left(\frac{2\Delta}{2\ell}\right)\right)\sigma_{i_1^f} + \ell\left(\tanh\left(\frac{\Delta}{2\ell}\right) + \tanh\left(\frac{2\Delta}{2\ell}\right)\right)\sigma_{i_1^f + 2\Delta} \\
&\quad - \ell\left(1 - e^{-(i_1^f + \Delta)/\ell} + \tanh\left(\frac{\Delta}{2\ell}\right)\right)\sigma_{i_1^f + \Delta} - \ell\left(\tanh\left(\frac{\Delta}{2\ell}\right) + \tanh\left(\frac{\Delta}{2\ell}\right)\right)\sigma_{i_1^f + 2\Delta}
\end{aligned}
$$

From corollary 3.1, the first-best i_1^f and Δ are given by

$$
1 - e^{-i_1^f/\ell} = \tanh\left(\frac{1 - 2i_1^f}{2\ell(n-1)}\right), \qquad \Delta = \frac{1 - 2i_1^f}{n-1}.
$$

Therefore,

$$
e^{-i_1^f/\ell} = 1 - \tanh\left(\frac{\Delta}{2\ell}\right).
$$

This allows us to simplify the expression for $\sigma^2_{\mathbf{m}_n^f \setminus (i_1^f + \Delta)} - \sigma^2_{\mathbf{m}_n^f \setminus i_1^f}$. First, the coefficient in front of $\sigma_{i_1^f + \Delta}$ becomes

$$
\ell\left(1 - e^{-(i_1^f + \Delta)/\ell} + \tanh\left(\frac{\Delta}{2\ell}\right)\right) = 2\ell(1 - e^{-\Delta/\ell}).
$$

Second, the coefficient in front of $\sigma_{i_1^f}$ becomes

$$
\ell\left(1 - e^{-i_1^f/\ell} + \tanh\left(\frac{2\Delta}{2\ell}\right)\right) = \ell\left(\tanh\left(\frac{\Delta}{2\ell}\right) + \tanh\left(\frac{2\Delta}{2\ell}\right)\right).
$$

Therefore,

$$
\begin{aligned}
\sigma^2_{\mathbf{m}_n^f \setminus (i_1^f + \Delta)} - \sigma^2_{\mathbf{m}_n^f \setminus i_1^f} &= \ell\left(\tanh\left(\frac{\Delta}{2\ell}\right) + \tanh\left(\frac{2\Delta}{2\ell}\right)\right)\left(\sigma_{\{i_1^f\}} + \sigma_{\{i_1^f + 2\Delta\}}\right) \\
&\quad -2\ell(1 - e^{-\Delta/\ell})\sigma_{\{i_1^f + \Delta\}} - 2\ell\tanh\left(\frac{\Delta}{2\ell}\right)\sigma_{\{i_1^f + 2\Delta\}} \\
&= \frac{2\ell^2 e^{-(2\Delta + i_1^f)/\ell}(e^{\Delta/\ell} - 1)^2(1 + e^{2\Delta/\ell} - 2e^{(\Delta + i_1^f)/\ell})}{1 + e^{2\Delta/\ell}}.
\end{aligned}
$$

This difference is strictly positive because $1 + e^{2\Delta/\ell} - 2e^{(\Delta + i_1^f)/\ell} = 1 + e^{2\Delta/\ell} - \frac{2e^{\Delta/\ell}}{1 - \tanh\left(\frac{\Delta}{2\ell}\right)} > 0$. This implies that $\sigma^2_{\mathbf{m}_n^f \setminus i_j^f} > \sigma^2_{\mathbf{m}_n^f \setminus i_1^f}$ for any $j = 2, \dots, n-1$. ■

C Proofs for section 4

Proof of Proposition 4.1. First, observe that for $n = 1$, the optimal minipublic is either empty or it consists of one citizen, so the statement holds trivially. Next, let $n \geqslant 2$. By contradiction, suppose that the optimal minipublic consists of $n' < n$ distinct citizens, $\mathbf{m}^* = \{i_1, \dots, i_{n'}\}$, where $i_1 < \dots < i_{n'}$ and $n' \geqslant 1$. Consider a modified minipublic $\tilde{\mathbf{m}} := \mathbf{m}^* \cup \tilde{i}$, where $\tilde{i} \notin \mathbf{m}^*$. First, this modified minipublic is strictly more informative because $\beta(\tilde{i})$ is imperfectly correlated with $\beta(\mathbf{m}^*)$, i.e., $\sigma^2_{\tilde{\mathbf{m}}} > \sigma^2_{\mathbf{m}^*}$. Moreover, for any $i_k \in \mathbf{m}^*$, $\sigma^2_{\tilde{\mathbf{m}}\setminus i_k} > \sigma^2_{\mathbf{m}^*\setminus i_k}$ because $\tilde{\mathbf{m}} \setminus i_k = (\mathbf{m}^* \setminus i_k) \cup \{\tilde{i}\}$. Therefore, by lemma B.2(ii), all citizens in $\mathbf{m}^*$ continue to be active in $\tilde{\mathbf{m}}$. Consider (ED) for the newly added citizen $\tilde{i}$:

$$\sigma^2_{\mathbf{m}^*\setminus i_k} < \sigma^2_{\tilde{\mathbf{m}}\setminus \tilde{i}} = \sigma^2_{\mathbf{m}^*} < \sigma^2_{\tilde{\mathbf{m}}}.$$

Again invoking lemma B.2(ii), because $i_k \in \mathbf{m}^*$ is active in $\mathbf{m}^*$, $\tilde{i}$ is also active in $\tilde{\mathbf{m}}$. Hence, all citizens are active in $\tilde{\mathbf{m}}$. Hence, the policymaker strictly prefers $\mathbf{m}$ to $\mathbf{m}^*$, which contradicts the optimality of $\mathbf{m}^*$. ■

Proof of Proposition 4.3. The following lemma will be invoked in the proof.

Lemma C.1 (Single crossing in $|\bar{B}|$)**.** *For any given $0 \leqslant \underset{\sim}{\sigma}^2 < \tilde{\sigma}^2$, there exists a unique cutoff $b > 0$ such that $V_C(\tilde{\sigma}^2) \geqslant V_C(\underset{\sim}{\sigma}^2)$ if and only if $|\bar{B}| \leqslant b$.*

Proof. Without loss, let $\bar{B} \geqslant 0$. Consider the function $D(\bar{B}) := V_C(\tilde{\sigma}^2) - V_C(\underset{\sim}{\sigma}^2)$, which is continuously differentiable in $\bar{B}$. We claim that D crosses zero exactly once from above.

The payoff-minimizing informativeness $\underline{\sigma}^2$ in (6) is increasing in $\bar{B}$. Fix $0 \leqslant \underset{\sim}{\sigma}^2 < \tilde{\sigma}^2$, and define $\underline{b} := \{\bar{B} : \underline{\sigma}^2 = \underset{\sim}{\sigma}^2\}$, and $\bar{b} := \{\bar{B} : \underline{\sigma}^2 = \tilde{\sigma}^2\}$. From equation (6), it immediately follows that $\underline{b} > 0$.

By Lemma 3.3, for any $0 \leq \bar{B} \leq \underline{b}$, the payoff-minimizing $\underline{\sigma}^2$ is weakly to the left of both $\underset{\sim}{\sigma}^2$ and $\tilde{\sigma}^2$. Hence, a higher informativeness yields a higher expected payoff, so $D(\bar{B}) > 0$. Similarly, for any $\bar{B} \geq \bar{b}$, $\underline{\sigma}^2$ is weakly to the right of both $\underset{\sim}{\sigma}^2$ and $\tilde{\sigma}^2$, and hence, $D(\bar{B}) < 0$. This implies that the function D crosses zero k times in the interval $(\underline{b}, \bar{b})$, where $k \geq 1$ and k odd.

We claim that $k = 1$. By contradiction, suppose that $k > 1$. By the continuous differentiability of D, the derivative D' must be switching sign at least three times, so there exist at least two different values of $\bar{B} \in (\underline{b}, \bar{b})$ at which $D'(\bar{B}) = 0$. We next show that at most one such value can exist, which generates the desired contradiction. The first derivative of $D(\bar{B})$ is

$$D'(\bar{B}) = h(\bar{B}; \tilde{\sigma}^2) - h(\bar{B}; \underset{\sim}{\sigma}^2)$$

where $h(\bar{B}; \sigma^2) := \Phi(\frac{\bar{B}}{\sqrt{\sigma^2+\tau^2}}) + \phi(\frac{\bar{B}}{\sqrt{\sigma^2+\tau^2}})\frac{\bar{B}\tau^2}{\sqrt{(\sigma^2+\tau^2)^3}}$. The partial derivative of h with respect to $\bar{B}$ is

$$\begin{aligned}
\frac{\partial h(\bar{B}; \sigma^2)}{\partial \bar{B}} &= \phi(\frac{\bar{B}}{\sqrt{\sigma^2+\tau^2}}) \left[\frac{1}{\sqrt{\sigma^2+\tau^2}} + \frac{\tau^2}{\sqrt{(\sigma^2+\tau^2)^3}} - \frac{\bar{B}^2\tau^2}{\sqrt{(\sigma^2+\tau^2)^5}} \right] \\
&= \phi(\frac{\bar{B}}{\sqrt{\sigma^2+\tau^2}}) \left[\frac{\sigma^4 - \tau^2(\bar{B}^2 - 3\sigma^2) + 2\tau^4}{\sqrt{(\sigma^2+\tau^2)^5}} \right].
\end{aligned}$$

From (6), V_C increases in σ^2 if and only $\sigma^2 \geq \underline{\sigma}^2$, or equivalently if $\sigma^4 - (\bar{B}^2 - 3\sigma^2)\tau^2 + 2\tau^4 \geqslant 0$. As $\underset{\sim}{\sigma}^2 < \underline{\sigma}^2 < \tilde{\sigma}^2$ for every $\bar{B} \in (\underline{b}, \bar{b})$, it holds that $\underset{\sim}{\sigma}^4 - (\bar{B}^2 - 3\underset{\sim}{\sigma}^2)\tau^2 + 2\tau^4 < 0$ and $\tilde{\sigma}^4 - (\bar{B}^2 - 3\tilde{\sigma}^2)\tau^2 + 2\tau^4 \geq 0$. Hence, $\frac{\partial h(\bar{B};\tilde{\sigma}^2)}{\partial \bar{B}} \geq 0$ and $\frac{\partial h(\bar{B};\underset{\sim}{\sigma}^2)}{\partial \bar{B}} < 0$.

To sum up, on the entire domain $\bar{B} \in (\underline{b}, \bar{b})$, $h(\bar{B}; \tilde{\sigma}^2)$ is increasing in $\bar{B}$ and $h(\bar{B}; \underset{\sim}{\sigma}^2)$ is strictly decreasing in it. Hence, there is at most one $\bar{B}'$ that satisfies $D'(\bar{B}') = 0$. We have reached a contradiction, so $k = 1$. ■

(i) The first-best minipublic is optimal if and only if each $i \in \mathbf{m}_n^f$ is active, i.e., $V_C(\sigma^2_{\mathbf{m}_n^f}) \geqslant V_C(\sigma^2_{\mathbf{m}_n^f \setminus i})$. By lemma C.1 and the fact that $\sigma^2_{\mathbf{m}_n^f \setminus i_k^f} < \sigma^2_{\mathbf{m}_n^f}$ for any $k \in \{1, \ldots, n\}$, there exists a threshold $b_k > 0$ such that i_k^f is active if and only if $|\bar{B}| \leqslant b_k$. Define $\underline{b}$ to be the lowest among all such b_k for $k \in \{1, \ldots, n\}$. By construction, each $i_k^f \in \mathbf{m}_n^f$ is active if $|\bar{B}| \leqslant \underline{b}$. On the other hand, if $|\bar{B}| > \underline{b}$, there exists at least one citizen in $\mathbf{m}_n^f$ who prefers to be passive.

(ii) The optimal minipublic is empty if and only if no minipublic $\mathbf{m} \in \mathcal{M}_n$ is active. First, we establish the sufficiency of $|\bar{B}| \geqslant \bar{b}$ for the optimal minipublic to be empty, then its necessity.

As $|\bar{B}| \to +\infty$, the limit payoff V_C is strictly decreasing over the relevant domain $[0, \bar{\sigma}_n^2]$ because $\underline{\sigma}^2 \to +\infty$. That is, for $|\bar{B}|$ sufficiently large, for all $\mathbf{m} \in \mathcal{M}_n$ each $i \in \mathbf{m}$ is passive. Therefore, there exists a cutoff $\bar{b} < +\infty$ such that the optimal minipublic is empty if $|\bar{B}| \geqslant \bar{b}$.

Suppose there exist $b < \tilde{b} < \bar{b}$ such that (i) if $|\bar{B}| = b$ the optimal minipublic is empty, and (ii) if $|\bar{B}| = \tilde{b}$ the optimal minipublic, denoted by $\tilde{\mathbf{m}}$, is nonempty. Then, each $i \in \tilde{\mathbf{m}}$ is active, which means that $D(\tilde{b}) > 0$, and hence D is strictly decreasing at $|\bar{B}| = \tilde{b}$. Because $b < \tilde{b}$, it must be that $D(b) > D(\tilde{b}) > 0$, hence $\tilde{\mathbf{m}}$ is active when $|\bar{B}| = b$ as well. This contradicts the empty optimal minipublic for $|\bar{B}| = b$. Hence, the optimal minipublic is empty only if $|\bar{B}| \geqslant \bar{b}$. ■

Proof of Proposition 4.2. First, we examine how the V_C-minimizing informativeness $\underline{\sigma}^2$ varies with τ^2.

Lemma C.2. *Let $\tau_x^2 := (\frac{3}{\sqrt{2}} - 2)\bar{B}^2$ and $\tau_y^2 := \frac{1}{2}\bar{B}^2$. The informativeness at which V_c reaches its minimum, $\underline{\sigma}^2$, is*

(i) strictly increasing in $\tau^2 \in (0, \tau_x^2)$, and $\underline{\sigma}^2 \to 0$ as $\tau^2 \to 0$;

(ii) strictly decreasing in $\tau^2 \in (\tau_x^2, \tau_y^2)$, and for $\tau^2 \in [\tau_y^2, \infty)$, $\underline{\sigma}^2 = 0$.

Proof. Let $g(\tau^2) := \frac{1}{2}\sqrt{\tau^4 + 4\bar{B}^2\tau^2} - 3\tau^2$ and $\underline{\sigma}^2 = \max\{0, g(\tau^2)\}$ as in (6). The first derivative of $g(\tau^2)$ with respect to τ^2 is

$$\frac{1}{2}\left(\frac{4\bar{B}^2 + 2\tau^2}{2\sqrt{4\bar{B}^2\tau^2 + \tau^4}} - 3\right),$$

which is negative (positive) if and only if $\tau^2 > (<)(\frac{3}{\sqrt{2}} - 2)\bar{B}^2 =: \tau_x^2$. First, as $\tau^2 \to 0$, $g(\tau^2) \to 0$ from above, hence $\underline{\sigma}^2 \to 0$. Moreover, the positive derivative of g implies that $\underline{\sigma}^2$ strictly increases in $\tau^2 \in [0, \tau_x^2)$. Second, g is strictly positive over $[\tau_x^2, \tau_y^2)$ and strictly decreasing for $\tau^2 > \tau_x^2$. Hence, $\underline{\sigma}^2$ strictly decreases for $\tau^2 \in (\tau_x^2, \tau_y^2)$. For $\tau^2 = \tau_y^2$, $\underline{\sigma}^2 = 0$ because $g(\tau_y^2) = 0$. Therefore, because $g(\tau_y^2) \leqslant 0$ for $\tau^2 \in [\tau_y^2, \infty)$, $\underline{\sigma}^2 = 0$. ■

If $\bar{B} = 0$, then $\underline{\sigma}^2 = 0$ for any τ^2. This implies any citizen in any minipublic is active. Hence, $\mathbf{m}^* = \mathbf{m}_n^f$ for any τ^2, so $\underline{\tau}^2 = \bar{\tau}^2$. If $\bar{B} \neq 0$, let $\bar{\tau}^2 = \tau_y^2$ as in lemma C.2. By lemma C.2(ii), for any $\tau^2 \geqslant \bar{\tau}^2$, $\underline{\sigma}^2 = 0$ so any citizen is active in any minipublic. Therefore, $\mathbf{m}_n^f$ is feasible, so $\mathbf{m}^* = \mathbf{m}_n^f$.

By lemma 3.4, if i_1^f is active then $\mathbf{m}_n^f$ is feasible. By lemma C.2(i), for any $\sigma^2_{\mathbf{m}_n^f \setminus i_1^f} > 0$, there exists $\underline{\tau}^2$ sufficiently small so that $0 < \underline{\sigma}^2 < \sigma^2_{\mathbf{m}_n^f \setminus i_1^f}$ for all $\tau^2 \leq \underline{\tau}^2$. Thus, for τ^2 sufficiently small, $\mathbf{m}^* = \mathbf{m}_n^f$. ■

Proof of Proposition 4.4. (i) We first show that the active informativeness $\sigma^2_{\mathbf{m}}$ of any minipublic $\mathbf{m} \in \mathcal{M}_n$ converges to zero as $\ell \to 0^+$. For any $i_k \in \mathbf{m}$, the informativeness of the singleton minipublic $\{i_k\}$ vanishes: $\sigma^2_{i_k} \to 0$ as $\ell \to 0^+$. Furthermore, for any distance $d > 0$, $\tanh(d/\ell) \to 1$ as $\ell \to 0^+$, and hence, $\gamma_k(\mathbf{m}) \to 0$ for any $k \leqslant n$. Thus, by (1) in lemma 3.1, $\sigma^2_{\mathbf{m}} \to 0$ as $\ell \to 0^+$.

Next, because $0 < \underline{\sigma}^2$ by assumption and $\sigma^2_{\mathbf{m}} \to 0$ as $\ell \to 0^+$ for any $\mathbf{m} \in \mathcal{M}_n$, there exists ℓ sufficiently small such that for any $\mathbf{m} \in \mathcal{M}_n$, $\sigma^2_{\mathbf{m}} < \underline{\sigma}^2$. By the quasiconvexity of V_C from lemma 3.3, each minipublic citizen in any $\mathbf{m}$ then strictly prefers to be passive. Hence, the set of feasible minipublics is empty.

(ii) Fix an arbitrary $\mathbf{m} \in \mathcal{M}_n$. As $\ell \to +\infty$, using L'Hôpital's rule, $\sigma^2_{i_k} \to 1$ for any $i_k \in \mathbf{m}$. Furthermore, $\ell \tanh(\frac{d}{2\ell}) \to \frac{d}{2}$ as $\ell \to \infty$ for any distance $d > 0$. Using this observation and L'Hôpital's rule, the limit of γ_k as given by (2) is

$$\lim_{\ell \to +\infty} \gamma_k(\mathbf{m}) = \begin{cases} \dfrac{i_1}{2} + \dfrac{i_2}{2} & \text{if } k = 1, \\ \dfrac{i_{k+1}}{2} - \dfrac{i_{k-1}}{2} & \text{if } k = 2, ..., n-1, \\ 1 - \dfrac{i_{n-1}}{2} - \dfrac{i_n}{2} & \text{if } k = n. \end{cases} \tag{9}$$

Thus, the limit of informativeness of any minipublic with $n \geq 1$ is

$$\lim_{\ell \to +\infty} \sigma^2_{\mathbf{m}} = \lim_{\ell \to +\infty} \sum_k \gamma_k(\mathbf{m}) \sigma_{i_k} = 1.$$

Let $n > 1$. By the argument above, in the first-best minipublic, both the passive informativeness $\sigma^2_{\mathbf{m}^f_n \setminus i^f_k}$ for any $i^f_k \in \mathbf{m}^f_n$ and the active informativeness $\sigma^2_{\mathbf{m}^f_n}$ converge to 1 as $\ell \to +\infty$. Hence, for ℓ sufficiently high, $\underline{\sigma}^2 < \sigma^2_{\mathbf{m}^f_n \setminus i^f_k} < \sigma^2_{\mathbf{m}^f_n}$ because $\underline{\sigma}^2 < 1$ by assumption. As V_C is strictly increasing in informativeness for $\sigma^2 > \underline{\sigma}^2$, (ED) of each citizen in the first-best minipublic is satisfied. ∎

Proof of Corollary 4.1. The distance^2-minimizing minipublic $\mathbf{m}^d_n = \{i^d_1, \ldots, i^d_n\}$ solves

$$\min_{\{i_1, \ldots, i_n\}} \quad \int_0^{(i_1+i_2)/2} (i - i_1)^2 \mathrm{d}i + \int_{(i_1+i_2)/2}^{(i_2+i_3)/2} (i - i_2)^2 \mathrm{d}i + \ldots + \int_{(i_{n-1}+i_n)/2}^{1} (i - i_n)^2 \mathrm{d}i.$$

Setting the partial derivatives of the objective with respect to each minipublic citizen i_k equal to zero and solving this system of n equations for each citizen, we obtain $i^d_k = (2k-1)/(2n)$ for every $k \in \{1, \ldots, n\}$.

From proposition 4.4, for ℓ sufficiently large, the optimal minipublic is $\mathbf{m}^f_n$. From Proposition G.4(iii) in Bardhi (2020), each first-best minipublic citizen i^f_k approaches $i^d_k = (2k-1)/(2n)$. ∎

Proof of Proposition 4.5. First, let $\underline{\sigma}^2 < \bar{\sigma}^2_\infty$. Therefore, V_C is strictly increasing in informativeness for $\sigma^2 \in (\underline{\sigma}^2, \bar{\sigma}^2_\infty]$. In any minipublic of the form $\mathbf{m}^f_n \setminus i^f_k$, where $k \in \{1, \ldots, n\}$, the mesh of the partition of $[0, 1]$ corresponding to this minipublic goes to zero as $n \to \infty$. By an argument similar to the proof of lemma B.5, $\sigma^2_{\mathbf{m}^f_n \setminus i^f_k} \to \bar{\sigma}^2_\infty$ as $n \to \infty$. Hence, for n sufficiently large, $V_C(\sigma^2)$ is strictly increasing at $\sigma^2 = \sigma^2_{\mathbf{m}^f_n \setminus i^f_k}$. Therefore, (ED) is satisfied for every $i^f_k \in \mathbf{m}^f_n$.

Second, let $\underline{\sigma}^2 \geqslant \bar{\sigma}^2_\infty$. This implies that V_C is strictly decreasing for every $\sigma^2 \in (0, \bar{\sigma}^2_\infty)$. Therefore, (ED) for any citizen in any minipublic is violated. Only the empty minipublic is feasible. ∎

Proof of Proposition 4.6. By the premise that $\mathbf{m}^* \notin \{\emptyset, \mathbf{m}^f_n\}$ and lemma 3.4, it must be that the minimum of the citizens' payoff function is reached at $\underline{\sigma}^2 \in (\sigma^2_{\mathbf{m}^f_n \setminus i^f_1}, \sigma^2_{\mathbf{m}^f_n})$. The informational diversity of the first-best minipublic is

$$\frac{\sigma^2_{\mathbf{m}^f_n} - \sigma^2_{\mathbf{m}^f_n \setminus i^f_1}}{\sigma^2_{\mathbf{m}^f_n \setminus i^f_1}}.$$

Because each citizen is active in the optimal minipublic, it must be that for any $i \in \mathbf{m}^*$, $\sigma^2_{\mathbf{m}^*\setminus i} > \sigma^2_{\mathbf{m}^f_n\setminus i^f_1}$. Moreover, by the definition of the first-best minipublic and the fact that $\mathbf{m}^f_n$ is infeasible, $\sigma^2_{\mathbf{m}^*} < \sigma^2_{\mathbf{m}^f_n}$. Hence, for any $i \in \mathbf{m}^*$, $\sigma^2_{\mathbf{m}^*} - \sigma^2_{\mathbf{m}^*\setminus i} < \sigma^2_{\mathbf{m}^f_n} - \sigma^2_{\mathbf{m}^f_n\setminus i^f_1}$, and therefore,

$$\frac{\sigma^2_{\mathbf{m}^*} - \sigma^2_{\mathbf{m}^*\setminus i}}{\sigma^2_{\mathbf{m}^*\setminus i}} \leqslant \max_{i\in\mathbf{m}^*} \quad \frac{\sigma^2_{\mathbf{m}^*} - \sigma^2_{\mathbf{m}^*\setminus i}}{\sigma^2_{\mathbf{m}^*\setminus i}} < \frac{\sigma^2_{\mathbf{m}^f_n} - \sigma^2_{\mathbf{m}^f_n\setminus i^f_1}}{\sigma^2_{\mathbf{m}^f_n\setminus i^f_1}}.$$

■

Proof of Proposition 4.7. Because $|\bar{B}| > \underline{b}$, $\mathbf{m}^f_2$ is infeasible. By lemma 3.4, this implies that both i^f_1 and i^f_2 are passive because $\sigma^2_{\mathbf{m}^f_2\setminus i^f_1} = \sigma^2_{i^f_2} = \sigma^2_{\mathbf{m}^f_2\setminus i^f_2} = \sigma^2_{i^f_1}$. Moreover, because $|\bar{B}| < \bar{b}$, the optimal minipublic is nonempty. Therefore, $\underline{\sigma}^2 \in (\sigma^2_{i^f_1}, \bar{\sigma}^2_2)$. By proposition 4.1, the optimal minipublic consists of two citizens. Let $\mathbf{m}^* = \{i_1, i_2\}$, where $i_2 \neq i_1$.

Claim 1. *In any optimal minipublic, $i^f_1 < i_1 < 1/2 < i_2 < i^f_2$.*

Proof. In order for each citizen i_1 and i_2 to be active in $\mathbf{m}^*$, it must be that $\sigma^2_{i^f_1} < \sigma^2_{i_2} < \sigma^2_{1/2}$ and $\sigma^2_{i^f_1} < \sigma^2_{i_1} < \sigma^2_{1/2}$. Because the informativeness of a singleton minipublic σ^2_i is single-peaked in i with a maximum at $i = 1/2$ (see lemma 3.1 for informativeness σ^2_i), it must be that $|1/2 - i_1| < 1/2 - i^f_1 = i^f_2 - 1/2$ and $|1/2 - i_2| < 1/2 - i^f_1 = i^f_2 - 1/2$. Hence, $i^f_1 < i_1 < i_2 < i^f_2$.

Suppose, by contradiction, that $i_1 < i_2 < 1/2$. Then, there exists a modified minipublic $\tilde{\mathbf{m}} = \{i_1, 1 - i_2\}$ which is feasible and attains a strictly higher informativeness. First, $\sigma^2_{\tilde{\mathbf{m}}} > \sigma^2_{\mathbf{m}^*}$ because $\sigma^2_{\{i_1,i\}}$ is strictly increasing in i over the region $i \in (i_1, i^f_2)$. This is because (i) $\sigma^2_{\{i^f_1,i\}}$ is maximized at $i = i^f_2$, (ii) $\sigma^2_{\{i_1,i\}}$ is increasing in i_1 for each i, and (iii) $i_1 > i^f_1$ as established above. Second, both i_1 and $1 - i_2$ are active in $\tilde{\mathbf{m}}$ if i_1 and i_2 are active in $\mathbf{m}^*$. Citizen $1 - i_2$ is active because the passive informativeness of $1 - i_2$ in $\tilde{\mathbf{m}}$, i.e., $\sigma^2_{i_1}$, is the same as that of i_2 in $\mathbf{m}^*$, whereas the active informativeness of $1 - i_2$ in $\tilde{\mathbf{m}}$ is strictly higher than that of i_2 in $\mathbf{m}^*$. Citizen i_1 is also active in $\tilde{\mathbf{m}}$ because his passive informativeness remains unchanged, i.e., by the symmetry of the informativeness of a singleton minipublic $\sigma^2_{i_2} = \sigma^2_{1-i_2}$, whereas his active informativeness strictly increases from $\mathbf{m}^*$ to $\tilde{\mathbf{m}}^*$. Therefore, we have reached a contradiction to the optimality of $\mathbf{m}^*$. The case of $1/2 < i_1 < i_2$ leads to a similar contradiction and is omitted. ■

Claim 2. *In any optimal minipublic, $i_1 = 1 - i_2 = i^*$ and $V_C(\sigma^2_{i^*}) = V_C(\sigma^2_{\{i^*,1-i^*\}})$.*

Proof. We first show that the optimal minipublic is symmetric. By way of contradiction, suppose that $i_1 < 1 - i_2$. Consider a modified minipublic $\tilde{\mathbf{m}}_\epsilon = \{i_1 + \epsilon, i_2 + \epsilon\}$ where $\epsilon > 0$ small. This minipublic attains strictly higher informativeness for ϵ sufficiently small because for $i_1 + i_2 < 1$,

$$\left.\frac{\partial\sigma^2_{\tilde{\mathbf{m}}_\epsilon}}{\partial\epsilon}\right|_{\epsilon=0} = 2\ell e^{-\frac{2(1+i_2)}{\ell}} \left(e^{\frac{i_1+i_2}{\ell}} - e^{\frac{1}{\ell}}\right)\left(e^{\frac{1}{\ell}} + e^{\frac{i_1+i_2}{\ell}} - 2e^{\frac{1+i_1}{\ell}}\right) > 0.$$

Moreover, for ϵ sufficiently small the passive informativeness of the rightmost citizen $\sigma^2_{i_1+\epsilon}$ strictly increases. Hence, $i_2 + \epsilon$ is active in $\tilde{\mathbf{m}}$ if i_2 is active in $\mathbf{m}^*$. The only remaining (ED) is that of $i_1 + \epsilon$: note that for ϵ sufficiently small, the passive informativeness of the leftmost citizen $\sigma^2_{i_2+\epsilon}$ decreases in ϵ. But, $\sigma^2_{i_1} < \sigma^2_{i_2}$ because $i_1 < 1 - i_2$, so for ϵ sufficiently small $\sigma^2_{i_1+\epsilon} < \sigma^2_{i_2+\epsilon}$. Hence, if $i_2 + \epsilon$ is active, then $i_1 + \epsilon$ is active as well in $\tilde{\mathbf{m}}$. We have thus reached a contradiction to the optimality of $\mathbf{m}^*$. This implies that $\mathbf{m}^* = \{i_1, 1 - i_1\}$.

Next we show that (ED) binds for both citizens in $\mathbf{m}^*$. Suppose by contradiction that (ED) for i_1 is slack. Because the passive informativeness and the active informativeness of $1-i_1$ are identical to those of i_1, then (ED) for $1-i_1$ must be slack as well. Consider a modification $\tilde{\mathbf{m}}_\delta = \{i_1 - \delta, 1 - i_1 + \delta\}$ for $\delta > 0$ small, i.e., both citizens are shifted further away from the median citizen. Note that the informativeness of a symmetric minipublic $\sigma^2_{\{i_1, 1-i_1\}}$ strictly decreases in i_1 for $i_1 \in (i_1^f, 1/2]$. Hence, because $i_1 > i_1^f$,

$$\left.\frac{\partial \sigma^2_{\tilde{\mathbf{m}}_\delta}}{\partial \delta}\right|_{\delta=0} > 0.$$

For δ sufficiently small, the minipublic informativeness strictly increases. Moreover, given that each (ED) is slack, it continues to be satisfied for δ sufficiently small. Hence, $\tilde{\mathbf{m}}_\delta$ is a strict improvement over $\mathbf{m}^*$, which contradicts the optimality of $\mathbf{m}^*$. Therefore, (ED) binds for both citizens in the optimal minipublic. ■

Claim 3. *The optimal minipublic is unique.*

Proof. Suppose that there exist more than one optimal minipublic. By Claim 1 and Claim 2, they are of the form $\{i^*, 1-i^*\}$ and $\{i^{**}, 1-i^{**}\}$, where $i_1^f < i^* < i^{**}$. But, by an argument similar to that of Claim 2, $\sigma^2_{\{i^*,1-i^*\}} > \sigma^2_{\{i^{**},1-i^{**}\}}$. Hence, if both these minipublics are feasible, then $\{i^{**}, 1-i^{**}\}$ is suboptimal. We have reached a contradiction. ■

■

D Proofs for section 5

Proof of Proposition 5.1. We prove that the following constitutes an equilibrium: (i) each citizen $i \in \mathbf{m}$ discovers evidence, (ii) each citizen $i \in \mathbf{m}$ discloses $\beta(i) \neq x_i$ and conceals $\beta(i) = x_i$, where x_i uniquely solves $\mathbb{E}[B \mid \beta(i) = x_i] = 0$, and (iii) the policymaker adopts the policy if and only if her post-minipublic value is higher than her realized threshold of adoption c.

First, (iii) is a best response for the policymaker. Because there is a single outcome realization $\beta(i) = x_i$ which citizen $i \in \mathbf{m}$ conceals and both disclosure and no disclosure are on the equilibrium path, the policymaker perfectly infers $\beta(\mathbf{m})$ and the post-minipublic value $B_{\mathbf{m}} = \mathbb{E}[B \mid \beta(i), \beta(\mathbf{m} \setminus i)]$. Therefore, she best responds as in the commitment game.

Second, we show that it is a best response for citizen i to disclose $\beta(i)$ if $\beta(i) \neq x_i$ conditional on i having discovered $\beta(i)$ and all other citizens following strategy (ii). Consider first $\beta(i) \neq x_i$. For simplicity of notation, let $\mu := \mathbb{E}[B|\beta(i)]$. The distribution of the post-minipublic value from the perspective of citizen i with evidence $\beta(i)$ is denoted by $B_{\mathbf{m}}|\beta(i)$. Using the law of iterated expectations, we have

$$\mathbb{E}[B_{\mathbf{m}}|\beta(i)] = \mathbb{E}[\mathbb{E}[B|\beta(\mathbf{m}\backslash i), \beta(i)]|\beta(i)] = \mathbb{E}[B|\beta(i)] = \mu.$$

Let σ^2 be the variance of $B_{\mathbf{m}}|\beta(i)$, which does not depend on the realization of $\beta(i)$.[29] The random variable $B_{\mathbf{m}}|\beta(i)$ is distributed according to

$$B_{\mathbf{m}}|\beta(i) \sim \mathcal{N}(\mu, \sigma^2).$$

If citizen i discloses $\beta(i)$, the policymaker's post-minipublic value is $B_{\mathbf{m}}$. If he conceals $\beta(i)$, due to linearity of the post-minipublic value in (8), the policymaker's post-minipublic value is $\hat{B}_{\mathbf{m}} = B_{\mathbf{m}} - \lambda$, where

[29] The variance σ^2 is independent of $\beta(i)$ because the joint distribution of B and $\beta(i)$ is Gaussian. The functional form of σ^2 is inconsequential for this proof and therefore omitted.

$\lambda \coloneqq \gamma_i(\mathbf{m})(\beta(i) - x_i)$. That is, concealing $\beta(i)$ shifts the policymaker's post-minipublic value either up or down by λ.

By similar calculations as in the proof of lemma 3.2, we calculate the expected payoff of citizen i if he discloses or conceals $\beta(i)$. If he discloses $\beta(i)$, he obtains

$$\begin{aligned} V_C(\beta(i)) &\coloneqq \mathbb{E}[B_{\mathbf{m}} | B_{\mathbf{m}} \geq c, \beta(i)] \Pr[B_{\mathbf{m}} \geq c | \beta(i)] \\ &= \mu\Phi\left(\frac{\mu}{\sqrt{\tau^2+\sigma^2}}\right) + \frac{\sigma^2}{\sqrt{\tau^2+\sigma^2}}\phi\left(\frac{\mu}{\sqrt{\tau^2+\sigma^2}}\right) \end{aligned}$$

Similarly, if he conceals $\beta(i) \neq x_i$, he obtains

$$\begin{aligned} \tilde{V}_C(\beta(i)) &\coloneqq \mathbb{E}[B_{\mathbf{m}} | \hat{B}_{\mathbf{m}} \geq c, \beta(i)] \Pr\left[\hat{B}_{\mathbf{m}} \geq c | \beta(i)\right] \\ &= \mu\Phi\left(\frac{\mu-\lambda}{\sqrt{\tau^2+\sigma^2}}\right) + \frac{\sigma^2}{\sqrt{\tau^2+\sigma^2}}\phi\left(\frac{\mu-\lambda}{\sqrt{\tau^2+\sigma^2}}\right) \end{aligned}$$

Note that the function $f(a) \coloneqq \mu\Phi(\frac{a}{\sqrt{\tau^2+\sigma^2}}) + \frac{\sigma^2}{\sqrt{\tau^2+\sigma^2}}\phi(\frac{a}{\sqrt{\tau^2+\sigma^2}})$ varies in a depending on the sign of μ and the relation between μ and a:

$$\frac{\partial f(a)}{\partial a} = \phi(\frac{a}{\sqrt{\tau^2+\sigma^2}})\frac{\mu(\tau^2+\sigma^2) - a\sigma^2}{(\tau^2+\sigma^2)^{\frac{3}{2}}} \begin{cases} > 0 \text{ if } \mu > 0 \text{ and } a < \mu, \\ < 0 \text{ if } \mu < 0 \text{ and } a > \mu. \end{cases}$$

Next, we show that $V_C(\beta(i)) > \tilde{V}_C(\beta(i))$ for every $\beta(i) \neq x_i$. Let $\beta(i) > x_i$. Then, by the definition of x_i and the monotonicity of $B_{\mathbf{m}}$ in $\beta(i)$ in (8), $\mu > 0$ and $\lambda > 0$. This means that disclosing evidence in favor of the policy yields a higher payoff than concealing it, as $f(a)$ is increasing in a for these parameters and $\mu - \lambda < \mu$. Similarly, let $\beta(i) < x_i$. In this case, $\mu < 0$ and $\lambda < 0$. Disclosing evidence in favor of the status quo yields a higher payoff than concealing it. If $\beta(i) = x_i$, then $\lambda = 0$. Then, concealing is a weak best response because $V_C(\beta(i)) = \tilde{V}_C(\beta(i))$.

Finally, we show that (i) holds: it is a best response for citizen i to (privately) discover $\beta(i)$ if all other citizens in $\mathbf{m}$ discover their respective outcomes. Because in the continuation equilibrium (ii) and (iii) all minipublic outcomes $\beta(\mathbf{m})$ are perfectly inferred by the policymaker, citizen i discovers $\beta(i)$ if and only if his (ED) in the commitment game holds. By the premise, every $i \in \mathbf{m}$ in the commitment game is active, hence citizen i prefers to discover $\beta(i)$ in the no-commitment game as well. ∎