Experimentation in Democratic Mechanisms*

Volker Britz
CER-ETH – Center of Economic
Research at ETH Zurich
Zuerichbergstrasse 18
8092 Zurich, Switzerland
vbritz@ethz.ch

Hans Gersbach
CER-ETH – Center of Economic
Research at ETH Zurich and CEPR
Zuerichbergstrasse 18
8092 Zurich, Switzerland
hgersbach@ethz.ch

This version, February 12, 2015

Abstract

We examine whether and how democratic procedures can achieve socially desirable public good provision in the presence of profound uncertainty about the benefits of the public good, i.e., when citizens are able to identify the distribution of benefits only if they aggregate their private information. Some members of the society, however, are harmed by socially desirable policies and aim at manipulating information aggregation by misrepresenting their private information. We show that information can be aggregated and a socially desirable policy implemented under a new class of democratic mechanisms involving an experimentation group. Those mechanisms reflect the principles of liberal democracy, are prior–free, and involve a differential tax treatment of experimentation group members which motivates them to reveal their private information truthfully. Conditional experimentation ensures that all citizens are treated equally on the equilibrium path. Conversely, we show that standard democratic mechanisms with an arbitrary number of voting rounds without experimentation do not generally lead to socially desirable policies.

*We are grateful to Carlos Alos–Ferrer, Johannes Becker, Patrick Bolton, Ulrich Erlenmaier, Volker Hahn, Hans Haller, Martin Hellweg, Mark Machina, Herve Moulin, Klaus Schmidt, Urs Schweizer, seminar participants in Davis, Heidelberg, Irvine, Munich, San Diego, and Zurich for valuable comments and suggestions on democratic mechanisms. A first round of research in this area was conducted jointly by Volker Britz, Hans Gersbach, and Vitalijs Butenko and was included in the doctoral thesis by Butenko (2013).

1 Introduction

The ability of democratic decision–making procedures to achieve socially optimal outcomes is the topic of a long–standing and complex debate with many unresolved issues. In particular, one open question is whether democratic procedures can resolve profound uncertainty. We mean by profound uncertainty a situation in which members of the society perceive the individual benefits (or costs) of a policy as realizations of some probability distribution, while this probability distribution itself is also unknown. Such a situation features not only uncertainty but ambiguity. In this paper, we consider mechanisms which mimic decision–making in a liberal democracy. We show that such democratic mechanisms can be used to implement the socially optimal level of a public good under profound uncertainty about individual benefits. However, the implementation of the socially optimal choice requires the use of experimentation.

In the presence of profound uncertainty about benefits of the public good, the implementation of a socially desirable policy requires aggregation of individual citizens' private information. We mean by a socially desirable policy a public good level which would be preferred by a majority of citizens to any feasible alternative if the underlying distribution of benefits was known. Put another way, one of the possible alternatives in our setup is a Condorcet winner, but citizens can identify this Condorcet winner only through information aggregation. Although such information aggregation is socially desirable, it is harmful for some members of the society. The core issue of this paper is whether democratic procedures can lead to information revelation even when a subset of the society is interested in concealing the relevant information. In order to approach this question, we adopt the notion of a democratic mechanism. Democratic mechanisms are collective choice procedures which differ from standard mechanisms in three respects. First, they fulfill requirements of liberal democracy such as equal voting rights of all citizens and anonymous voting. In fact, only Yes–or–No votes (or abstentions) are allowed in democratic mechanisms. Second, all citizens are equally likely to be in the position of the agenda–setter who has the opportunity to make proposals which are then voted upon. There is no benevolent social planner. Third, all citizens are coerced to participate in the implementation of the collective decision. The idea is that there is a governmental authority which can enforce the law and, in particular, tax citizens. The collective decision need not be "individually rational" for each citizen, and thus need not fulfill a "participation constraint." While the first two features above make it harder to achieve socially desirable outcomes, the third feature obviously facilitates the implementation of socially desirable collective decisions. In the present paper, we investigate whether socially desirable outcomes can be achieved with such mechanisms in the presence of ambiguity about the underlying distribution of benefits from public goods. The problem is that some citizens (acting as agenda–setters or voters) may not be interested the resolution of this ambiguity through information aggregation. The question of interest is how information aggregation can nevertheless be accomplished by a suitably designed democratic mechanism.

In this paper, we study the following model. A society chooses which quantity of a public good to provide. The set of possible public good levels is discrete, and it is also possible to provide zero public good. The provision of the public good is financed by the uniform taxation of all citizens. Each citizen is privately informed about his own valuation of the public good. The cross–section distribution depends on the unobservable realization of a state of nature. The individual valuation serves as a signal from which the individual citizen infers the valuations of other citizens. A citizen who values the public good highly tends to believe that it is highly beneficial to the society as well. The set of possible public good levels and their valuations are such that exactly one of the public good levels is a Condorcet winner in each state of nature. However, due to the profound uncertainty, it is unknown which possible public good level is the Condorcet winner. This uncertainty can only be resolved by the aggregation of private information.

First, we consider a class of mechanisms called two–stage voting mechanisms. In such a mechanism, two rounds of majority voting are held, and the intention is that the voting behavior in the first round reveals information about the socially desirable public good level, while the second round of voting serves to make the decision. We demonstrate, however, that two–stage voting mechanisms are prone to manipulation and the Condorcet may not be discovered. This occurs when the agenda–setters is not interested in information revelation and chooses proposals in such a way that a subgroup of citizens can tacitly coordinate their votes on these proposals such that the state of nature is concealed. This impossibility result persists if voting is repeated any finite number of times.

We then show that this impossibility result can be overcome using a new class of democratic mechanisms, which we call "democratic mechanisms with experimentation." The essence of a democratic mechanism with experimentation is that a small subset of the society acts as an experimentation group. Members of the experimentation group are exempted from taxation and their votes in the first round are counted separately from those of the rest of the society. We show that this group can reveal the Condorcet winner on behalf of the society, independently of the interests of the agenda–setter. One important property of the new class of democratic mechanisms with experimentation is that they do not depend on citizens' prior beliefs about the state of nature and the associated Condorcet winner. In that sense, the newly introduced mechanisms are robust.

The formation of an experimentation group and the concomitant tax exemption can be seen as a challenge to the idea of equal treatment in the liberal–democracy constraint. We address this problem either by requiring that ex ante every citizen can become an experimentation group member with equal probability, or through conditional experimentation where tax exemptions occur only off the equilibrium path. In the latter case, all citizens receive equal tax treatment ex post as well as ex ante on the equilibrium path.

Overall, the results of our present paper describe environments in which voting – supplemented by experimentation – achieves socially desirable outcomes. Our results add to a body of literature pioneered by the seminal papers of Ledyard and Palfrey (1994, 2002). They discuss how efficient public good provision can be approximated by voting schemes ("referenda") with a simple yes–or–no message space. The main insights of this literature have been summarized and further developed by Ledyard (2005). We propose the use of experimentation groups in order to enhance the capability of voting procedures to achieve socially desirable outcomes. In particular, experimentation groups can be helpful in environments with ambiguity about the distribution of benefits from the public good. Moreover, the use of experimentation groups can be used to neutralize the bad incentives of agenda–setters who are interested in blocking information aggregation.

2 Related literature

Our paper also contributes to the literature on incomplete social contracts and democratic mechanisms. Since the classic work of Buchanan and Tullock (1962), a vast literature on optimal constitutions has been developed. Aghion and Bolton (2003) have introduced incomplete social contracts and have explored how simple or qualified majority rules balance the need to overcome vested interests and respect the preferences of a majority of citizens. Gersbach (2009) introduces the notion of democratic mechanisms and shows how increasingly sophisticated combinations of agenda, treatment, and decision rules can yield first–best allocations when each citizen only faces two possible realizations, namely being either a winner or a loser of a public project. In this paper, we extend the democratic mechanism approach to the case of profound uncertainty about the distribution of valuations. More specifically, neither individual valuations nor the underlying distributions are common knowledge. We continue this line of research and explore the scope of "simple" mechanisms when there is profound uncertainty about the distribution of benefits. Following Ledyard and Palfrey (1994, 2002), we consider mechanisms which are "simple" in the sense that their message space is minimal: Only yes–or–no messages can be sent. We explore how incentives of agenda–setters to conceal information can be overcome in such simple democratic mechanisms.

Moreover, our paper relates to several strands of literature on experimentation in single–agent decision problems and games. A sizable literature dating back at least to Rothschild (1974) deals with experimentation in the context of the famous bandit problems; a survey can be found in Bergemann and Välimäki (2006). Uur paper is concerned with experimentation in collective decisions, and with ways to ensure that experimentation does take place. At least since the seminal paper by Rose–Ackermann (1980), it is well known that the rules which govern collective decisions also govern the incentives of office–holders as to whether to experiment. She showed that free–riding in federal systems reduces the incentives of candidates for office to undertake policy experiments. This line of research has been extended and deepened by Strumpf (2002), Cai and Treisman (2009), Volden, Ting and Carpenter (2008), and Kollman, Miller, and Page (2000), as well as Bednar (2011). Callander and Harstad (2013) study quantity and quality of policy experimentation and characterize advantages and drawbacks of federal systems in this respect.

The importance of experimentation in actual choice procedures is well established. For instance, Volden (2006) and Shipan and Volden (2006) study policy experimentation and diffusion across jurisdictions in the United States. Buera, Monge–Naranjo, and Primiceri (2011) provide a theory of policy diffusion at a global level and find empirically that learning from experience across countries is an important factor behind changes in economic policy. Our approach is complementary to this literature. We adopt a constitutional approach and develop a set of rules that together induce effective experimentation in the polity. Our approach is related to the logic outlined in Bendor and Mookherjee (1987) according to which repetition of collective decisions improves efficiency. In our setup, a combination of repetitive voting and experimentation can efficiently counteract the attempts of agenda–setters to make proposals which deter information revelation. Repetitive voting alone, however, cannot accomplish this result.[1]

3 The public good problem

We consider a society which consists of a continuum of risk–neutral citizens of unit mass. Citizens collectively decide how much of a public good to provide. The public good is not arbitrarily divisible, that is, the set of possible public good levels is discrete. We are interested in the decision–making process by which the society chooses from such a set of public good levels.

We will assume that the per capita cost of providing a quantity $q \in \mathbb{R}_+$ of the public good is given by a twice continuously differentiable, strictly increasing, and strictly convex function[2] $c : \mathbb{R}_+ \to \mathbb{R}_+$ with $c(0) = 0$. This implies in particular that average cost $c(q)/q$

[1]Experimentation differs from opinion polling since it is not ensured in the latter that polled citizens report honestly. See Bernhardt, Duggan and Squintani (2008) for a recent survey of rational choices of polling.

[2]Although we will eventually consider a discrete set of possible quantities, it is convenient to start by defining continuous cost and utility functions on $\mathbb{R}_+$.

is strictly increasing. Public good provision is financed by uniform taxation so that every citizen pays $c(q)$ when q is provided. Typically, one assumes that each citizen is initially endowed with w ($w > 0$) units of a private consumption good which can either be consumed or transformed into the public good. The per capita costs $c(q)$ are the utility losses due to foregone private consumption.

While the cost of public good provision is common knowledge among all citizens, there is uncertainty with regard to the benefit of the public good. Every citizen obtains a benefit from public good provision which is proportional to q, with the factor of proportionality given by the citizen's *type.* The type, in turn, is a realization of a random variable z taking values in a non–empty, non–degenerate, and connected interval Z, which is closed in $\mathbb{R}_{++}$. The probability distribution from which the types are drawn will be specified later. Each citizen is privately informed about his type. We will henceforth refer to the citizen of type z as *citizen* z. If the public good level q is provided, citizen z obtains a utility of

$$u(z, q) = zq - c(q). \tag{1}$$

Moreover, we assume that one out of finitely many *states of nature* has been realized, and that no citizen is informed about this state of nature. We index the states of nature by $k = 1, \ldots, n$ (or by i and j when necessary) and write $N = \{1, \ldots, n\}$. The aggregate and individual uncertainties are related as follows. When the state of nature is k, then the types z are drawn from Z by a probability distribution associated with the cumulative distribution function $F_k : Z \to [0, 1]$. The cumulative distribution functions $(F_k)_{k=1,\ldots,n}$ are each twice continuously differentiable. We denote the concomitant probability density functions by $(f_k)_{k=1,\ldots,n}$. We assume that $f_k(z) > 0$ for all $k \in N$ and all $z \in Z$ and, moreover, that

$$\frac{f'_{k+1}(z)}{f_{k+1}(z)} > \frac{f'_k(z)}{f_k(z)}, \ \forall z \in Z, \ \forall k \in N \setminus \{n\}. \tag{2}$$

The above inequality reflects a property of the family of probability distributions which is known as the *monotonicity of likelihood ratios.* This property has three key implications. First, the probability distribution associated with F_{k+1} first–order stochastically dominates the one associated with F_k for every $k \in N \setminus \{n\}$. In that sense, the benefits from the public good are higher in state $k+1$ than in state k. Second, the monotonicity of likelihood ratios implies a single–crossing property of the probability density functions, which will be crucial for our analysis. Finally, the monotonicity of likelihood ratios imposes a monotonicity property on citizens' posterior beliefs. To be more precise, denote by Δ^n the unit simplex in $\mathbb{R}^n$ and by Δ^n_{++} the intersection of Δ^n with $\mathbb{R}^n_{++}$. We assume that all citizens share a common prior $p \in \Delta^n_{++}$ about the states of nature. Upon observing his type, citizen z updates the prior belief p with his type z, thus obtaining the posterior belief $\beta(z)$ about the state of nature. According to Bayes' rule, we can write the k^{th} component of $\beta(z)$ as follows,

$$\beta_k(z) = \frac{f_k(z)p_k}{\sum_{j=1}^{n} f_j(z)p_j}, \ k = 1, \ldots, n. \tag{3}$$

Since $p \in \Delta^n_{++}$, we also have $\beta(z) \in \Delta^n_{++}$ for all $z \in Z$. Now the monotonicity of likelihood ratios implies, loosely speaking, that a higher type tends to believe with higher probability in higher states of nature. We can interpret $F_k(z)$ as the cross–sectional distribution of z in the population when the state is k. It is well known that this interpretation requires the application of a suitable version of the law of large numbers for a continuum of random variables. One implication is that the state of nature can be inferred by aggregating information about the realized types. Therefore, voting may serve as an information aggregation device.

Recall that we have defined the cost and utility functions on $q \in \mathbb{R}_+$. For what follows, we need to impose the technical assumption ensuring that the type space Z and the cost function c with first derivative c' are such that Z is included in the image of $\mathbb{R}_+$ under c'.[3] Together with the strict convexity of the cost function, this implies that the preferences of each type $z \in Z$ are single–peaked. As a result, we can define for each state $k \in N$ a quantity $q_k \in \mathbb{R}_+$ which satisfies the equality $F_k\left(c'(q_k)\right) = \frac{1}{2}$. The quantity q_k is a *Condorcet winner* in the sense that in state k, a simple majority of citizens prefers q_k over any other quantity $q \in \mathbb{R}_+ \setminus \{q_k\}$. For the remainder of the paper, we assume a discrete set Q of possible public good levels which consists of the Condorcet winners $\{q_1, \ldots, q_n\}$ and of the default quantity $q_0 = 0$. A non–zero public good quantity is only possible if it is the Condorcet winner in some state. This assumption rules out the possibility that our decision–making procedure might replace the default quantity q_0 by some public good level $q > 0$ which is itself not stable to majority voting and would thus provide incentives for "renegotiation."

Now let us define for every $k \in N$ a quantity $\tilde{q}_k$ which satisfies the equality $F_k\left(\frac{c(\tilde{q}_k)}{\tilde{q}_k}\right) = \frac{1}{2}$. In state k, a majority of citizens prefers any quantity $q < \tilde{q}_k$ to zero public good, but prefers zero public good to any quantity $q > \tilde{q}_k$. Given that average cost $c(q)/q$ is increasing, it is straight–forward that $\tilde{q}_k \geq q_k$ for every $k \in N$. Assumption 3.1 below is the key assumption of our model; loosely speaking, its role is to ensure that the states are "sufficiently different" from each other.

Assumption 3.1. *The type space Z, the cost function c, and the cumulative distribution functions $(F_k)_{k \in N}$ are such that $q_{k+1} > \tilde{q}_k$ for every $k \in N \setminus \{n\}$.*

[3] This assumption is weaker than imposing the Inada condition that $c'(q)$ tends to zero as $q \to 0$.

If this assumption were not satisfied and the number of states were large, then our model would become similar to a model with a continuous state space. While such a setup would be of interest in its own right, it is fundamentally different from the problem of interest in the present paper.

The aforementioned assumptions on cumulative distribution functions, the type space, and the set Q imply the following three statements, which we will work with extensively in the analysis:

$$\left\{ \frac{c(q_1)}{q_1}, \frac{c(q_n) - c(q_{n-1})}{q_n - q_{n-1}} \right\} \subset int(Z), \tag{4}$$

$$F_k \left(\frac{c(q_k) - c(q_j)}{q_k - q_j} \right) < \frac{1}{2}, \; k \in N, \; j \in \{0, 1, \ldots, k-1\}, \tag{5}$$

$$F_k \left(\frac{c(q_{k+1})}{q_{k+1}} \right) > \frac{1}{2}, \; k \in N \setminus \{n\}. \tag{6}$$

The set inclusion (4) says that non–degenerate subsets of the society prefer q_1 over q_0 and q_n over q_{n-1}. This is a direct consequence of the fact that both q_1 and q_n are Condorcet winners in states 1 and n, respectively. The inequality (5) says that in state k, the quantity q_k is preferred by the simple majority to any smaller element of Q; this is an obvious implication of the construction of q_k as the Condorcet winner in state k. Finally, inequality (6) says that in state k, the simple majority prefers q_0 over q_{k+1}. This property is due to Assumption 3.1, which says that any two consecutive states are sufficiently different from each other.

4 The decisive voting round

In this paper, we study a class of democratic mechanisms which allow for repeated voting on public good provision. In particular, each of these mechanisms consists of several rounds, and the decision for the actual public good level is taken by a vote in the final round of each mechanism. The rules governing this final round are identical in all the mechanisms we are going to consider. We refer to the final round of each mechanism as the *decisive voting round.* In this section, we study the decisive voting round in isolation as a strategic game among the citizens. It can be described in extensive form as follows.

Decisive Voting Round. *At the beginning of this round, all citizens decide simultaneously whether or not to apply for the role of the agenda–setter. Then, the agenda–setter is chosen by fair randomization from all citizens who have applied. If no citizen has applied, then no public good is provided and all citizens obtain zero utility. The citizen who is* *chosen to be the agenda–setter makes a proposal $q \in Q$. Finally, all voters simultaneously cast votes in favor of or against this proposal. If the simple majority votes in favor of the proposal q, then q is implemented. Otherwise, the status quo is implemented. The status quo is determined in previous rounds of the mechanism and denoted by $\bar{q} \in Q$.*[4] *All citizens are taxed uniformly, except for the agenda–setter, who is tax–exempt.*

The decisive voting round as described above has the following properties.

1. Every citizen has the right to abstain from proposal–making.
2. Every citizen who does not abstain from proposal–making has the same probability of being chosen as the agenda–setter.
3. Every citizen has the right to vote, and all votes count equally.
4. Voting is binary, only yes–or–no–approval is allowed.

The decisive voting round complies with the definition of a *democratic mechanism,* as given by Gersbach (2009). We note that in the decisive voting round, the agenda–setter is exempted from taxation while all other citizens are subject to uniform taxation. Equal treatment of citizens with regard to taxation can be viewed as a further desirable feature of democratic mechanisms. We stress that without the tax–exemption of the agenda–setter in the decisive voting round, there is no chance that democratic mechanisms in consideration here can yield socially desirable public good levels. For instance, if the status quo was q_0 and the agenda–setter had a very low valuation, then he could simply propose q_0 in order to minimize his tax burden. Notice that in our setup with a continuum of citizens, a tax–exemption for one (or finitely many) citizens does not change the tax burden for the rest of the society. As a result of the tax–exemption, the agenda–setter of type z obtains the utility

$$u_e(z, q) = zq. \tag{7}$$

For tractability and ease of presentation, we work with a continuum of voters. One problem we face, however, is that an individual vote has no influence and thus any outcome in the decisive voting round can be rationalized as an equilibrium. To exclude implausible outcomes, we mimic voting behavior in a large but finite society where individuals eliminate weakly dominated strategies.[5] Thus, no citizen abstains. Moreover, since there is only a

[4] At this stage of the analysis, $\bar{q}$ can be any element of Q.

[5] This is a standard procedure for the examination of voting in a continuum economy, see for instance Gersbach (2009).

binary decision in the decisive voting round, there are no gains from voting strategically. Therefore, the voting behavior of each citizen in the decisive voting round is *sincere,* which means that they vote in favor of the proposed alternative if and only if they strictly prefer it to the status quo prevailing in the decisive voting round. More formally, we obtain the following lemma.

Lemma 4.1. *In a decisive voting round with status quo $\bar{q}$, citizen z who is not tax–exempt votes in favor of a proposal $q \in Q$ if and only if $u(z,q) > u(z,\bar{q})$. In particular, in a decisive voting round with status quo q_0, citizen z who is not tax–exempt votes in favor of $q \in Q$ if and only if $z > c(q)/q$.*

For the rest of this section, we denote the true state of nature by k^*, and we write q^* for the quantity q_{k^*}.

If the status quo in the decisive voting round is q_0, then the inequalities (5) and (6) imply that the proposal $q_k \in Q \setminus \{q_0\}$ will be accepted if and only if $k \leq k^*$, where $k^* \in N$ is the true state of nature.

If the proposal $q \in Q$ is accepted in a decisive voting round, we say that q is the *outcome* of the decisive voting round. If a proposal is rejected in a decisive voting round with status quo $\bar{q}$, we say that $\bar{q}$ is the outcome of the decisive voting round. Sincere voting implies the following statement.

Lemma 4.2. *If the quantity q^* is either the proposal or the status quo in the decisive voting round, then it will be the outcome of the decisive voting round.*

Because citizens vote sincerely, the agenda–setter of the decisive voting round faces a simple decision problem when choosing the proposal. Since the agenda–setter is tax–exempt, whatever his type is, he has the preferences $q_n \succ \ldots \succ q_1 \succ q_0$ over the set Q, that is, he strictly prefers more to less of the public good. As noted above, the democratic mechanisms under consideration here will never stand a chance to achieve the socially desirable solution in the absence of a tax–exemption for the agenda–setter.

We are now going to describe in more detail the decision problem of the agenda–setter in a decisive voting round with status quo q_0.[6] Suppose that at the beginning of the decisive voting round, the beliefs of the different types are described by the map $\pi : Z \to \Delta^n$. Suppose that citizen z with belief $\pi(z)$ becomes agenda–setter in a decisive voting round with status quo q_0. Clearly, if he proposes q_0, the outcome of the decisive voting round can only be q_0. If he proposes some $q \in Q \setminus \{q_0\}$, then it follows from sincere voting that the outcome is q if $q \leq q^*$ and the outcome is q_0 otherwise. More formally, if an agenda–setter

[6]A similar analysis could be carried out for the case where the status quo is some $q \in Q \setminus \{q_0\}$, but only the case with the status quo q_0 is important for the analyis in the rest of the paper.

of type z with belief $\pi(z)$ proposes $q \in Q$ in the decisive voting round with status quo q_0, then the expected public good level to result from this decisive voting round can be written as

$$v\left(q, z, \pi(z)\right) = \begin{cases} \left(\sum_{j=k}^{n} \pi_j(z)\right) q_k & \text{if } q = q_k,\ k \in N, \\ 0 & \text{if } q = q_0. \end{cases}$$

Due to the tax–exemption the agenda–setter of the decisive voting round seeks to maximize the expected public good level. That is, he makes a proposal which belongs to

$$\omega(z, \pi(z)) = \arg\max_{q \in Q} v(q, z, \pi(z)).$$

Definition 4.3. *Suppose that beliefs prior to a decisive voting round are given by $\pi : Z \to \Delta^n$. We say that the state k^* is* revealed *by π if $\omega(z, \pi(z)) = \{q^*\}$ for almost all $z \in Z$.*

Our notion of the revelation of the state does not require that all types believe with probability one in the correct state. Instead, we consider the state to be revealed if almost all types have beliefs which make it optimal for them to propose the Condorcet winner. This notion of revelation is appropriate for the purpose of the analysis because we are ultimately not interested in the exact beliefs but in whether the outcome of the mechanism is the Condorcet winner. If the state is not revealed by the beliefs π, then we say that the state is *hidden* under those beliefs.

Lemma 4.4. *Consider a decisive voting round with status quo q_0. If citizens beliefs are such that the state is hidden, then the outcome of this decisive voting round is a public good level strictly smaller than q^*.*

Proof. Indeed, consider beliefs π under which the state is hidden. Let

$$\hat{Z} = \{z \in Z |\ \omega(q_0, z, \pi(z)) \supset \{q\} \text{ for some } \ q \in Q \setminus \{q^*\}\}.$$

By the supposition that the state is hidden, it follows that $\hat{Z}$ is of strictly positive mass. Then, the probability that the type of the agenda–setter in the decisive voting round is some $z \in \hat{Z}$ is strictly positive. Consequently, there is strictly positive probability that the proposal made in the decisive voting round is some $q \in Q \setminus \{q^*\}$. If $q < q^*$, then q is accepted by the majority and the resulting public good level is indeed strictly smaller than q^*, as claimed by the lemma. If, on the contrary, $q > q^*$, then q is rejected by the majority. Then, the outcome of the decisive voting round is $q_0 < q^*$, and the proof of the lemma is complete.

□

As pointed out before, we are interested in democratic mechanisms consisting of several voting rounds. So far, we have focused on the decisive voting round. In the remainder of the paper, we will be interested in the *preliminary voting rounds.* Prior to these preliminary voting rounds, each citizen z has belief $\beta(z)$ about the state of nature, that is, he has updated the prior p with private information about the type z. Our interest in the remainder of this paper is how the preliminary voting rounds may allow citizens to update their beliefs $\beta(z)$ in such a way that after the preliminary rounds, they have some beliefs π which reveal the state of nature in the sense defined above. Similar to the literature on mechanism design, the crucial issue is whether the citizens have incentives to reveal information about their types truthfully. Contrary to the mechanism design literature, types cannot be reported explicitly but any information must be revealed indirectly by voting Yes or No. For the following analysis, two groups of citizens are particularly important. One group (which we will call Z_+) consists of the citizens who benefit most from the public good, and the other group (to be called Z_-) consists of the citizens who benefit least. To be more precise, we define the sets Z_+ and Z_- as

$$\begin{aligned} Z_+ &= \left\{ z \in Z | z > \frac{c(q_n) - c(q_{n-1})}{q_n - q_{n-1}} \right\}, \\ Z_- &= \{ z \in Z | z < c(q_1)/q_1 \}. \end{aligned}$$

From the convexity of $c(q)$, it follows that a citizen $z \in Z_+$ strictly prefers q_{k+1} over q_k for all $k \in N \setminus \{n\}$, while a citizen $z \in Z_-$ strictly prefers q_k over q_{k+1} for all $k \in N \setminus \{n\}$. In other words, a citizen $z \in Z_+$ prefers "more to less" while a citizen $z \in Z_-$ prefers "less to more" of the public good. The set inclusion (4) says that the sets Z_+ and Z_- have non–empty interiors and thus are of strictly positive mass. The preference for "more to less" and the preference for "less to more" are present in non–degenerate subsets of the society.

Given that $Z \subset \mathbb{R}_{++}$, it is straight–forward that the preference for "more to less" can be emulated in a citizen of any type by exempting him from taxation. Emulating the preference for "more to less" in a citizen of arbitrary type will be crucial for the implementation results in Theorem 6.2 and Theorem 8.1. The significance of the set Z_- lies in the fact that any member of Z_- who is not tax–exempt prefers the state to be hidden rather than revealed.

In this section, we have seen that citizens vote sincerely in the decisive voting round. Therefore, the agenda–setter faces a simple decision problem. Given that the agenda–setter of the decisive voting round is tax–exempt, this decision problem consists of maximizing the expected level of public good provision. The only strategic interaction in the decisive voting round takes place when all citizens decide simultaneously whether or not to apply for the role of the agenda–setter. No public good provision would take place if no citizen had applied for agenda–setting. In order to conclude the analysis, we argue that such a behavior would not be consistent with Nash equilibrium.

Lemma 4.5. *In any Nash equilibrium of the decisive voting round, some citizen applies for agenda–setting.*

Proof. The proof can be found in the appendix. □

In conclusion, we have seen in this section that the Condorcet winner is implemented under a condition on the beliefs which citizens hold prior to the decisive voting round. These beliefs arise from updating the prior p with the private information about the type z and with the preliminary voting rounds. We have shown that citizens of a sufficiently high type are interested in the revelation of the state, while citizens of a sufficiently low type want to block the revelation of the state. The latter group has an incentive to "manipulate" any mechanism which serves to elicit information in the preliminary voting rounds. In the remainder of this paper, we assess the robustness of different democratic mechanisms to such manipulations. This is similar in spirit to a standard mechanism design problem in which one is interested in the "incentive–compatibility" or "truthfulness" of some mechanism. In the next section, we present a negative result: Democratic mechanisms which consist of repeated voting alone – in the absence of experimentation – cannot generally reveal the state and implement the Condorcet winner.

5 Repeated voting mechanisms

5.1 Definition and equilibrium analysis

In this section, we introduce a class of mechanisms which we call *two–stage voting mechanisms.* Such a mechanism consists of one preliminary round and a decisive voting round. We will see that two–stage voting mechanisms can implement the Condorcet winner if there are only two states and if a certain kind of inconsequential vote (to be defined later) is ruled out. However, we show that this positive result does not extend to the general class of public good problems which we consider. Thus, we have an impossibility result: The state cannot generally be revealed by a two–stage voting mechanism.

A two–stage voting mechanism can be described in extensive form as follows.

Two–Stage Voting Mechanism. *An agenda–setter for the preliminary round is randomly chosen from the population. The agenda–setter announces a preliminary proposal $q \in Q$. All citizens vote simultaneously to accept or reject the preliminary proposal. Let the share of Yes–votes be δ. If $\delta \leq \frac{1}{2}$, then a decisive voting round with status quo q_0 follows. If $\delta > \frac{1}{2}$, and if the preliminary proposal was q_0, then a decisive voting round with status quo q_0 follows. If, however, $\delta > \frac{1}{2}$, and the preliminary proposal was some $q \in Q \setminus \{q_0\}$, then* *either q or q_0 may be taken as the status quo of the decisive voting round. (Both possibilities will be explored below.) The agenda–setter of the decisive voting round is tax–exempt, while all other citizens are taxed uniformly.*

The above definition of a two–stage voting mechanism does not specify the status quo of the decisive voting round for $\delta > \frac{1}{2}$. Therefore, it defines a class of mechanisms rather than one specific mechanism in this class. Two alternative specifications of the status quo of the decisive voting round seem particularly relevant. First, one could specify that a proposal q becomes the status quo of the decisive voting round if it is accepted in the preliminary round. We will call this alternative a mechanism with *evolving status quo.* Second, one could specify the status quo in the decisive voting round to be q_0, irrespective of the result of the preliminary voting round. In that case, we will say that the status quo is *unresponsive.*

In this section, we will derive an impossibility result. In particular, we show that a two–stage voting mechanism does not generally implement the Condorcet winner in our public good problem. This result is true for all two–stage voting mechanisms under the above definition, regardless of whether their status quo evolves or is unresponsive. We note that even under a two–stage voting mechanism with an unresponsive status quo, the preliminary round may be used for information aggregation. The reason is that even a seemingly inconsequential vote in the preliminary round may affect citizens' beliefs and their beliefs may in turn influence the choice of the agenda–setter in the decisive voting round and ultimately the outcome of the mechanism. In what follows, we are going to refer to a vote in the preliminary round as a *mock vote* if this vote does not directly affect the status quo of the decisive voting round. If the status quo of a two–stage voting mechanism is unresponsive, then every vote in the preliminary round is a mock vote. If the status quo of a two–stage voting mechanism can evolve, then the vote in the preliminary round is still a mock vote if the proposal and the status quo of the preliminary round are identical, that is, if the preliminary proposal is q_0.

Two–stage voting mechanisms extend important properties of a democratic mechanism to the preliminary round. More precisely, all citizens (except the agenda–setter) can only send binary messages anonymously and simultaneously. Moreover, every citizen has the same chance to be the agenda–setter in the preliminary round.

Specifying a particular two–stage voting mechanism together with a public good problem $P \in \mathcal{P}$, we obtain a *two–stage voting game.* In a two–stage voting game, a *proposal strategy* is a map $\rho : Z \to Q$, where $\rho(z)$ is the preliminary proposal when citizen z is the agenda–setter. Moreover, a *voting strategy* is a map $\sigma : Z \times Q \to \{Yes, No\}$ which describes how every type reacts to every possible proposal. It gives rise to a function $\delta : N \times Q \to [0, 1]$ which indicates for each possible proposal the share of Yes–votes it will receive in each state. Finally, a *belief function* $\pi : Z \times Q \times [0,1] \to \Delta^n$ indicates the probabilities that citizen z assigns to the states of nature after observing which proposal was made and how many citizens approved it. The beliefs $\mu : Z \times Q \times [0,1] \to \Delta^{n+1}$ indicate the probabilities which citizen z assigns to the public good levels, given the preliminary proposal and a certain share of Yes–votes. Since we have shown before that the decisive voting round reduces to a decision–problem of the agenda–setter, we do not include it in the definition of the strategies and beliefs. To define an equilibrium of a two–stage voting game, let $X \subset Z$ stand for an arbitrary non–empty set of types such that all citizens $z \in X$ have the same preferences over Q. Let σ and $\hat{\sigma}$ be two voting strategies with $\hat{\sigma}(z) = \sigma(z)$ for all $z \in Z \setminus X$, but $\hat{\sigma}(z) \neq \sigma(z)$ for some or all $z \in X$. Then, we say that $\hat{\sigma}$ is a *joint deviation* from σ by the voters in X. A joint deviation by the members of X is *profitable* if all members of X are strictly better off under $\hat{\sigma}$ than under σ. We say that there is *no profitable joint deviation* if there does not exist $X \subset Z$ such that the members of X have a profitable joint deviation. We assume that in any decisive voting round which is a subgame of a two–stage voting game, all players behave optimally. Since we have analyzed the optimal behavior in the decisive voting round before, we do not specify the actions in these subgames. We are now ready to define the equilibrium concept. A *Bayesian equilibrium* in a two–stage voting game is a profile of strategies and beliefs $(\rho^*, \sigma^*, \pi^*, \mu^*)$ which satisfies the following conditions.

1. Given (σ^*, π^*, μ^*), there is no $z \in Z$ so that an agenda–setter of type z could benefit from a unilateral deviation to a different proposal than $\rho^*(z)$.

2. Given (ρ^*, π^*, μ^*), there is no profitable joint deviation from σ^*.

3. Given (ρ^*, σ^*), the beliefs (π^*, μ^*) are consistent.

The first and third requirements enumerated above are standard. The rationale for the second requirement is as follows. With regard to deviations from a strategy profile at the voting stage of the game, one cannot easily adopt the usual notion of a profitable deviation which is "unilateral" as well as "one–shot." At any rate, a unilateral deviation by a single voter would be pointless because an individual voter has zero mass and does not influence the outcome of the vote. Therefore, we require that an equilibrium should be robust against a coordinated deviation by all players whose preference ranking over the possible alternatives is the same. This requirement reduces to the robustness against a deviation by an arbitrarily small group if one has sufficiently many alternatives. Another possible equilibrium concept would require the absence of a profitable deviation by a subset of players with at most arbitrarily small but strictly positive mass. Such a concept of equilibrium would not greatly restrict the players' behavior in the game at hand, however.

Because of the discreteness of Q, an arbitrarily small set of citizens with positive measure cannot affect voting outcomes and thus almost all conceivable voting patterns could be rationalized as equilibria. Such circumstances are typical of voting games and are avoided if deviations by larger groups are allowed.

5.2 The impossibility result

Suppose that $(\rho^*, \sigma^*, \pi^*, \mu^*)$ is a Bayesian equilibrium of a two–stage voting game. Furthermore, suppose that no matter which state of nature has realized, playing the two–stage voting game according to $(\rho^*, \sigma^*, \pi^*, \mu^*)$ results with probability one in the outcome which corresponds to the Condorcet winner. In that case, we say that the Bayesian equilibrium at hand *implements the Condorcet winner.* If for every public good problem $P \in \mathcal{P}$, the two–stage voting game consisting of the two–stage voting mechanism and the public good problem P admits a Bayesian equilibrium which implements the Condorcet winner, then we say that the two–stage voting mechanism *implements the Condorcet winner.*

The purpose of this subsection is to demonstrate that no two–stage voting mechanism implements the Condorcet winner. In order to proof this result, we first derive Lemma 5.1. We begin by defining a *difference function* which captures the vertical distance between two consecutive cumulative distribution functions as

$$d_k(z) = F_k(z) - F_{k+1}(z), \; k \in N \setminus \{n\}.$$

Since cumulative distribution functions are twice continuously differentiable, the difference function d_k is also twice continuously differentiable. We have assumed the monotonicity of likelihood ratios in the family of probability distributions associated with $(F_k)_{k \in N}$. A simple geometric argument can be used to show that this assumption implies the single–crossing property of the density functions f_k and f_{k+1} for every $k \in N \setminus \{n\}$. As a consequence, for every $k \in N \setminus \{n\}$, there is a unique maximizer $z_k^* \in Z$ of the difference function d_k. For every $z \in Z$ such that $z < z_k^*$, the function d_k is monotonically increasing, while it is monotonically decreasing for all $z \in Z$ such that $z > z_k^*$. We say that the state of nature in a public good problem is *concealable* if there is a state $k \in N \setminus \{n\}$ such that $z_k^* \in Z_-$. From now on, we denote $d_k^* = d_k(z_k^*)$ for every $k \in N \setminus \{n\}$. We will show that no Bayesian equilibrium implements the Condorcet winner in a public good problem where the state is concealable. Because we want to find a mechanism which implements the Condorcet winner for the whole set $\mathcal{P}$ of public good problems, we will end up with an impossibility result.

Lemma 5.1. *Consider a two–stage voting game involving a public good problem in which the state is concealable. In particular, consider a subgame of this game which starts with a* *mock vote on some preliminary proposal* $q \in Q$. *If the profile* $(\rho^*, \sigma^*, \pi^*, \mu^*)$ *is a Bayesian equilibrium of the two–stage voting game, then the state is not revealed in the restriction of* $(\rho^*, \sigma^*, \pi^*, \mu^*)$ *to the subgame under consideration.*

Proof. Suppose by way of contradiction that $(\rho^*, \sigma^*, \pi^*, \mu^*)$ is a Bayesian equilibrium of a two–stage voting game, and its restriction to some subgame starting with a mock vote on $q \in Q$ reveals the state.

In the subgame, voting strategy σ^* assigns a vote in favor or against q to every type $z \in Z$. (Notice that the vote in the preliminary round is a mock vote; its outcome affects the sequel of the game only through the beliefs. Since no citizen knows the state, the vote can only be conditional on the type.) The voting strategy σ^* thus gives rise to the sets

$$\begin{aligned} Y' &= \{z \in Z \setminus Z_- | \sigma^*(z, q) = Yes\}, \\ Y'' &= \{z \in Z_- | \sigma^*(z, q) = Yes\}. \end{aligned}$$

of types who vote Yes. The masses of the sets Y' and Y'' depend on the state. Indeed, for any $k \in N$, let the mass of Y' be η_k and the mass of Y'' be χ_k. Moreover, let $\hat{z} = \frac{c(q_1)}{q_1}$; that is, all citizens $z < \hat{z}$ belong to Z_-, while all citizens $z > \hat{z}$ belong to $Z \setminus Z_-$. By the supposition that the state is concealable in the public good problem at hand, there is a state $i \in N \setminus \{n\}$ such that the unique maximizer of $d_i(z)$ belongs to Z_-. Let us suppose that $\eta_{i+1} \geq \eta_i$. Since d_i attains a unique maximum at $z_i^* < \hat{z}$ and is decreasing on $Z \setminus Z_-$, we have that

$$0 \leq \eta_{i+1} - \eta_i \leq F_i(\hat{z}) - F_{i+1}(\hat{z}).$$

By construction, the set Z_- has mass $F_{i+1}(\hat{z})$ when the state is $i + 1$, hence

$$0 \leq \chi_{i+1} \leq F_{i+1}(\hat{z}).$$

Adding up the two above inequalities, we obtain

$$0 \leq \eta_{i+1} - \eta_i + \chi_{i+1} \leq F_i(\hat{z}).$$

Since $F_i(z)$ is a continuous function, it can attain any value in the interval $[0, F_i(\hat{z})]$ for appropriately chosen $z \in Z_-$. In particular, there is a $\bar{z} \in Z_-$ such that

$$F_i(\bar{z}) + \eta_i = \eta_{i+1} + \chi_{i+1}.$$

Now consider a joint deviation by the members of Z_- from σ^*, under which citizens $z \in Z_-$ with $z < \bar{z}$ vote Yes, and citizens $z \in Z_-$ with $z > \bar{z}$ vote No. Under this deviation, the share of Yes votes among all citizens in state i is equal to the share of Yes votes under voting strategy σ^* in state $i + 1$. Recall the supposition that the state is revealed in the subgame under consideration. This supposition implies that if the share $F_i(\bar{z}) + \eta_i = \eta_{i+1} + \chi_{i+1}$ of Yes–votes is observed, then with probability one, the proposal in the decisive voting round will be q_{i+1}. Thus, under the deviation considered here, the proposal q_{i+1} will be made in the decisive voting round with probability one if the true state is i. Since citizens vote sincerely in the decisive voting round, this proposal will be rejected, and the resulting public good level will be q_0. By construction, members of Z_- prefer q_0 over q_i. Since the prior probability p_i of state i is strictly positive, it follows that in the supposed Bayesian equilibrium, the state cannot be revealed in a preliminary round with a mock vote.

We have now completed the proof of the lemma for the case where $\eta_{i+1} \geq \eta_i$. Recall that the vote under consideration here is a mock vote, its outcome changes only the beliefs. We can repeat the argument by alternatively defining η_k and χ_k as the share of No– rather than Yes–votes in state $k \in N \setminus \{n\}$, and by constructing a joint deviation for members of Z_- such that citizens $z < \bar{z}$ vote No, and citizens $z \in Z_-$ with $z > \bar{z}$ vote Yes. In this sense, the earlier supposition that $\eta_{i+1} \geq \eta_i$ is without loss of generality, and thus the proof of the lemma is complete. □

The above lemma has two main implications. First, a two–stage voting mechanism with an unresponsive status quo cannot implement the Condorcet winner – in such a mechanism, every vote in a preliminary round is a mock vote. Second, if a two–stage voting mechanism with an evolving status quo does implement the Condorcet winner, then the preliminary proposal must be different from q_0 – otherwise the vote in the preliminary round would again be a mock vote.

Lemma 5.2. *As in the previous lemma, consider a two–stage voting game involving a public good problem in which a state is concealable. Then, there is no Bayesian equilibrium* $(\rho^*, \sigma^*, \pi^*, \mu^*)$ *of this two–stage voting game which implements the Condorcet winner.*

Proof. Suppose by way of contradiction that $(\rho^*, \sigma^*, \pi^*, \mu^*)$ is a Bayesian equilibrium of this two–stage voting game which does implement the Condorcet winner. According to Lemma 5.1 above, the Condorcet winner cannot be implemented in a Bayesian equilibrium of a subgame following the preliminary proposal q_0. Now the supposition that $(\rho^*, \sigma^*, \pi^*, \mu^*)$ is a Bayesian equilibrium implies that this profile involves a preliminary proposal $\hat{q} \in Q \setminus \{q_0\}$. Since $(\rho^*, \sigma^*, \pi^*, \mu^*)$ implements the Condorcet winner, the agenda–setter of the preliminary round obtains the payoff $u(z, q^*)$ if he is of type z. Consider the possible deviation by this agenda–setter to proposing q_0 instead of $\hat{q}$. According to Lemma 5.1 above, the state will not be revealed after such a deviation. By Lemma 4.4, an expected public good level $\check{q} < q^*$ results. This deviation is profitable if $z \in Z_-$. The preliminary agenda–setter belongs to Z_- with strictly positive probability. We conclude that $(\rho^*, \sigma^*, \pi^*, \mu^*)$ cannot in fact be a Bayesian equilibrium which implements the Condorcet winner.

□

We have now shown that a two–stage voting mechanism fails to implement the Condorcet winner when the state is concealable. Example 5.4 below shows that the set of public good problems in which the state is concealable is a non–empty, non–degenerate subset of the set $\mathcal{P}$ of public good problems under consideration in this paper. Thus, we know that the statement of the previous lemma is not vacuous. Indeed, we have the following impossibility result:

Theorem 5.3. *No two–stage voting mechanism implements the Condorcet winner.*

One potential objection against this impossibility result is that the state might be concealable only in rather "complex" problems featuring many states or some specific functional form for the cost function. In that case, implementation of the Condorcet winner might still be possible in simpler public good problems. However, we can use Example 5.4 below to counter this objection. In fact, the example is minimal in the sense that there are only two states. Moreover, the cost function in the example is of a very simple quadratic form.

5.3 Example

Example 5.4. *Let the type space be given by $Z = [\theta, \theta + 1]$ for some $\theta > 0$, and let the cost function be*

$$c(q) = q^2 + \alpha q, \ \alpha \in \left(0, \frac{1}{2}\right).$$

There are two states $N = \{1, 2\}$, and the corresponding cumulative distribution functions are given by

$$\begin{aligned} F_1(z) &= z - \theta, \\ F_2(z) &= (z - \theta)^\gamma, \ \gamma \in (1, 2). \end{aligned}$$

We can easily compute the Condorcet winners in each state as

$$\begin{aligned} q_1 &= \frac{1}{4} + \frac{1}{2}(\theta - \alpha), \\ q_2 &= \frac{1}{2^{(\frac{1}{\gamma}+1)}} + \frac{1}{2}(\theta - \alpha). \end{aligned}$$

In order to verify whether the state is concealable, we need to maximize $d_1(z) = F_1(z) - F_2(z) = z - \theta - (z - \theta)^\gamma$. Indeed, taking the relevant first order condition with respect to z yields that

$$z^* = \theta + \left(\frac{1}{\gamma}\right)^{\frac{1}{\gamma-1}}$$

is the maximizer of d_1. The condition that $\gamma \in (0,1)$ ensures that z^* belongs to the interval Z. By definition, the state is concealable if $z^* < c(q_1)/q_1$. Substituting for q_1 and z^*, this condition can be expressed as

$$\left(\frac{1}{\gamma}\right)^{\frac{1}{\gamma-1}} < \frac{1}{4} + \frac{\alpha}{2} - \frac{\theta}{2}.$$

Many configurations of the parameters (α, γ, θ) satisfy this condition. To give a specific example, suppose that $\alpha = \frac{4}{10}$ and $\gamma = \frac{12}{10}$. Then, the state is concealable when $\theta < 0.096244$.

5.4 Generalization

The proof of the impossibility result is driven by two main considerations. First, the members of Z_- find it in their interest to conceal the state of nature, and for some public good problems they have the ability to do so when the vote in the preliminary round is a mock vote. Second, one of the members of Z_- is the agenda–setter with positive probability. Such citizens find it in their interest to turn the preliminary round into a mock vote.

In fact, the impossibility result is easily extended to voting mechanisms with several preliminary rounds. To be more specific, consider an m–stage voting mechanism in which there are $m-1$ preliminary rounds of the same kind as in the two–stage voting mechanism. For every natural number m, the joint probability that the agenda–setters of all $m-1$ preliminary rounds belong to Z_- is strictly positive. As a result, the implementation of the Condorcet winner fails with strictly positive probability. This leads to the following corollary.

Corollary 5.5. *No m–stage voting mechanism implements the Condorcet winner.*

5.5 Implementation in the absence of mock votes

We now turn to a special case which will be important for the result in Section 6. The proof of the impossibility result given by Theorem 5.3 is based on the insight that a mock vote blocks the revelation of information about the state and the implementation of the Condorcet winner. Next, one may wonder if the Condorcet winner can be implemented when the vote in the preliminary round is not a mock vote. We claim that this is indeed the case when the public good problem includes only two states of nature so that $Q = \{q_0, q_1, q_2\}$, and when the two–stage voting mechanism is such that the outcome of the preliminary round becomes the status quo of the decisive voting round. In such a two–stage voting game, the claim is that the Condorcet winner can be implemented in an equilibrium of a subgame in which the agenda–setter of the preliminary round has made a preliminary proposal $q' \in \{q_1, q_2\}$. Indeed, we construct the Bayesian equilibrium for such a subgame as follows.[7] Suppose that citizen z votes in favor of the preliminary proposal q' if and only if $z \in Z_+$. More formally, consider the voting strategy

$$\bar{\sigma}(z, q') = \begin{cases} Yes & \text{if } z \in Z_+, \\ No & \text{otherwise.} \end{cases}$$

Denote by e_k the n–vector with k^{th} component equal to one and all other components equal to zero. For the remainder of the paper, we will say that a belief $\pi \in \Delta^n$ is *deterministic* if $\pi = e_k$ for some $k \in N$. Moreover, if $\pi = e_k$ for some $k \in N \setminus \{k^*\}$, then we say that the belief π is *deterministic but erroneous.* It is straight–forward that a belief $\pi = e_{k^*}$ reveals the state. In the rest of the paper, we will consider several implementation results. All of these results will involve the revelation of the state through a deterministic belief.

Now let us define the belief

$$\bar{\pi}(z, q', \delta) = \begin{cases} e_2 & \text{if } \delta > \frac{1}{2}, \\ e_1 & \text{otherwise.} \end{cases}$$

Notice that neither the voting strategy $\bar{\sigma}$ above nor the belief $\bar{\pi}$ depend on whether the preliminary proposal is q_1 or q_2. One more important feature of the profile $(\bar{\pi}, \bar{\sigma})$ is that the belief depends only on whether a preliminary proposal was accepted or rejected, that is, whether δ did or did not exceed one half. The exact share of favorable votes does not influence the resulting belief.

Theorem 5.6. *Consider a two–stage voting game which consists of a public good problem with* $n = 2$, *and of the two–stage voting mechanism in which the outcome of the preliminary round becomes the status quo of the decisive voting round. In a subgame where an agenda–setter has made a preliminary proposal* $q' \in \{q_1, q_2\}$, *the voting strategy* $\bar{\sigma}$ *and the belief* $\bar{\pi}$ *as stated above define a revealing Bayesian equilibrium.*

Proof. The proof can be found in the appendix. □

This result is driven by the fact that manipulating the beliefs which result from the preliminary round is impossible without changing the decisive voting round. More precisely, when the share of favorable votes exceeds one half, this changes the (deterministic) beliefs and the status quo of the ensuing decisive voting round. While Theorem 5.6 does not readily extend to the case with $n \geq 3$, we will later make use of a similar logic in order to show how the impossibility result can be overcome using experimentation. This will be crucial in the derivation of Theorem 8.1.

[7] In the subgame under consideration, a Bayesian equilibrium is defined only by a voting strategy and the resulting beliefs.

6 An existence result based on signaling

6.1 The signaling mechanism

In this section we will introduce a democratic mechanism which allows for the revelation of the state of nature and the choice of the Condorcet winner. This mechanism involves experimentation in the preliminary round and the default outcome q_0 as status quo for the decisive voting round. In this context, experimentation means that some subset of the society is randomly chosen to receive a different tax treatment. The idea of the tax treatment is to change the incentives of that group in such a way that its members are interested in the revelation of the state of nature whatever their types may be, that is, even if they belong to Z_-. The voting behavior of the experimentation group can be commonly observed separately from the voting behavior of the rest of the society. In the preliminary round of the mechanism, citizens can use their votes to signal information about their types. That is, the votes only affect the outcome of the mechanism by aggregating information. Therefore, we refer to this mechanism as the *signaling mechanism.* More formally, the signaling mechanism is defined as follows:

Signaling Mechanism. *A subset of mass $\lambda > 0$ is randomly drawn from the population, we call it the* experimentation group. *Each member of the experimentation group takes a binary decision to send or not to send a signal. The share of experimentation group members who have sent a signal becomes common knowledge. Then a decisive voting round with status quo q_0 is played. Experimentation group members as well as the agenda–setter in the decisive voting round are tax–exempt, while all other citizens are taxed uniformly.*

We emphasize that it only matters whether an experimentation group member does or does not send a signal. The content of the signal is irrelevant. Since we consider voting mechanisms, it is natural to think of "sending a signal" as saying "Yes." However, the signal may also consist of some arbitrary message, so long as the message space from which the citizen can choose remains binary. Contrary to the analysis in the previous section, under the signaling mechanism, the public good is no longer financed by uniform taxation. From the point of view of citizens outside the experimentation group, the provision of a public good quantity $q \in Q$ is no longer associated with a tax burden of $c(q)$, but of $\left(\frac{1}{1-\lambda}\right) c(q)$. Consequently, the utility of citizen z outside the experimentation group from the public good level $q \in Q$ is given by

$$\hat{u}(z, q) = zq - \left(\frac{1}{1-\lambda}\right) c(q). \tag{8}$$

According to the above definition of the signaling mechanism, the decisive voting round always has status quo q_0. Recalling that voting is sincere in the decisive voting round, we see that citizen z votes in favor of a proposal $q \in Q$ in the decisive voting round of the signaling mechanism if and only if

$$(1-\lambda)z > \frac{c(q)}{q}.$$

Lemma 6.1. *A proposal $q \in Q$ which would be accepted in state $k \in N$ in a decisive voting round with status quo q_0 under uniform taxation, will also be accepted in the decisive voting round of the signaling mechanism, provided that $\lambda > 0$ is sufficiently small.*

Proof. The proof can be found in the appendix. □

While the introduction of the tax–exemption for the experimentation group does distort the voting behavior in the decisive voting round, this distortion has no effect on the acceptance or rejection of a proposal when $\lambda > 0$ is chosen sufficiently small. In particular, it is still true that in state $k^* \in N$ and with status quo q_0, a proposal q_k with $k \leq k^*$ will be accepted, and a proposal q_k with $k > k^*$ will be rejected.

6.2 The existence result

The signaling mechanism combined with a public good problem $P \in \mathcal{P}$ constitutes the *signaling game.* The strategies, beliefs, and the equilibrium concept for the signaling game are defined as follows. A strategy for the experimentation group is a map $\sigma : Z \to \{Yes, No\}$ which indicates for each type of an experimentation group member whether he does or does not send a signal. A belief for the agenda–setter is a map $\pi : [0,1] \times Z \to \Delta^n$ which assigns to each possible share of signals and types of the agenda–setter a probability distribution on the states of nature. For a subset of experimentation group members $Y \subset Z$ of strictly positive mass, a joint deviation from a pair (σ, π) is some $\tilde{\sigma}$ such that $\tilde{\sigma}(z) = \sigma(z)$ for all $z \in Z \setminus Y$ and $\tilde{\sigma}(z) \neq \sigma(z)$ for some $z \in Y$. The joint deviation is profitable if each citizen $z \in Y$ obtains a strictly greater payoff under $(\tilde{\sigma}, \pi)$ than under (σ, π).

A pair (σ^*, π^*) is a Bayesian equilibrium of the signaling game if there is no profitable joint deviation from σ^* given π^* and, moreover, π^* is consistent with σ^*.

The signaling game is a strategic game between the experimentation group members and the agenda–setter. The strategic interaction between them is "trivial" in the sense that they all share the same preferences over the possible outcomes. In order to advance their common interest, they need to accomplish coordination on the meaning of the signals, and thereby allow the dissemination of information. In order to show that the experimentation group members and the agenda–setter can indeed coordinate their actions successfully, we construct a Bayesian equilibrium which implements the Condorcet winner.

Define the set

$$\bar{Z} = \{z \in int(Z) | d_k(z) > 0 \ \forall k \in N \setminus \{n\}\}.$$

Our assumptions on the cumulative distribution functions $(F_k)_{k \in N}$ imply that $\bar{Z}$ is non–empty. Take any $\bar{z} \in \bar{Z}$, and define the strategy $\sigma^{\bar{z}}$ as follows,

$$\sigma^{\bar{z}}(z) = \begin{cases} 1 & \text{if } \ z \geq \bar{z}, \\ 0 & \text{otherwise.} \end{cases}$$

We associate to the strategy $\sigma^{\bar{z}}$ a belief $\pi^{\bar{z}}$ which is defined as follows,

$$\pi^{\bar{z}}(\delta, z) = \begin{cases} e_k & \text{if } \delta = 1 - F_k(\bar{z}) \ k \in N, \\ \beta(z) & \text{otherwise.} \end{cases}$$

Theorem 6.2. *For all $\bar{z} \in \bar{Z}$, the profile $(\sigma^{\bar{z}}, \pi^{\bar{z}})$ is a Bayesian equilibrium of the signaling game.*

Proof. The proof can be found in the appendix. □

6.3 Discussion

Theorem 6.2 shows that the signaling mechanism implements the Condorcet winner since the true state k^* is revealed and any agenda–setter in the decisive voting round proposes q_{k^*}. The signaling mechanism consists of one preliminary round (where signaling takes place), and a decisive voting round. Important democratic properties of the decisive voting round extend to the preliminary round. In particular, every experimentation group member makes a binary and anonymous decision, and the decisions of all experimentation group members have the same weight because it only matters whether they do or do not send a signal. Moreover, ex ante all members of the society have the same probability of being selected for experimentation group membership and the concomitant tax–exemption.

However, contrary to tax–exemption for the agenda–setter of a decisive voting round, this tax–exemption has an adverse effect on everybody outside the experimentation group. One drawback of the signaling mechanism is therefore that it creates a group within the society which enjoys a privilege at the significant expense of everyone else. However, the consequences for the rest of the electorate in terms of an additional tax burden can be made arbitrarily small by choosing the mass λ of the experimentation group arbitrarily small. Moreover, in the next section, we are going to discuss the possibility of a revealing mechanism under which the tax–exemption for an experimentation group is only needed off the path of equilibrium play, but does not occur in the revealing equilibrium itself. On the equilibrium path, no experimentation group is formed and everyone is taxed equally.

A game based on the signaling mechanism admits a multitude of Bayesian equilibria which implement the Condorcet winner. In fact, if the cumulative distribution functions are such that $F_1(z), \ldots, F_n(z)$ are n distinct numerical values for every $z \in int(Z)$, then every interior type can serve as the threshold type $\bar{z}$ used in Theorem 6.2 and is thus associated with one revealing Bayesian equilibrium. This multiplicity of equilibria can be viewed as a somewhat problematic feature of the mechanism. After all, the implementation of the Condorcet winner hinges on the ability of the players to coordinate on one particular threshold type $\bar{z}$. The mechanism includes no "communication device" to accomplish this coordination.

The signaling mechanism, however, can be modified and extended to ease the coordination of experimentation group members. Since all experimentation group members are interested in the revelation of the state, there is no inherent obstacle to coordination. The following extension simplifies this task: Before the decisions to send or not to send the signals are made, one (randomly appointed) citizen announces a particular public good level $q \in Q$. An experimentation group member of type z thereupon sends a signal if and only if he prefers q over any lower public good level,that is, if and only if $zq - c(q) > zq' - c(q')$ for every $q' < q$. This generates a uniquely defined threshold type $\bar{z}$ and the Bayesian equilibrium associated with $\bar{z}$ can then be played. We note that the existence theorem does not depend on the prior belief of citizens about the probability distribution p on the states. Hence, the mechanism with experimentation is *prior–free,* which is an important and desirable robustness property of the mechanism.[8]

7 Conditional experimentation

We have found that two–stage voting mechanisms cannot generally reveal the state of nature and implement the Condorcet–winning alternative in our model. These objectives can be achieved by using a signaling mechanism with an experimentation group. The effectiveness of a signaling mechanism does not depend on the number of states of nature in the model. One major drawback of the signaling mechanism, however, is that a small subset of the society is tax–exempt. As a result, citizens are not treated equally ex post. Some citizens do not contribute to the financing of the public good. The defense of this taxation rule is that every citizen has equal probability of being an experimentation group member and, therefore, one can argue that the signaling mechanism treats all citizens

[8]For the general theory of robust mechanisms in the standard framework, we refer to Bergemann and Morris (2005).

equally ex ante, that is, at a stage where it not yet known who will be an experimentation group member.

In this section, we introduce a further extension which combines elements of the two–stage voting and signaling mechanisms. With this extension, all citizens are treated equally ex ante and, moreover, all citizens are treated equally ex post on the equilibrium path of play. A tax–exemption for a subset of citizens (of positive mass) occurs only off the equilibrium path. In this section, we establish the result for the case with two states. In the next section, we show a similar result for public good problems with an arbitrary number of states.

Conditional experimentation mechanism. *One agenda–setter is randomly drawn from the whole population. This agenda–setter announces a preliminary proposal $q' \in Q$. If $q' = q_0$, then a subset of mass $\lambda > 0$ is randomly drawn from the population, we call it the* experimentation group. *Each experimentation group member simultaneously votes Yes or No on the preliminary proposal. The share of experimentation group members who have voted Yes is publicly observable. A decisive voting round with status quo q_0 follows, and the experimentation group members as well as the agenda–setter of the decisive voting round are tax–exempt. If, however, the preliminary proposal q' is different from q_0, then all citizens simultaneously vote in favor or against q' in the preliminary round. If the majority votes in favor of q', then a decisive voting round with status quo q' follows. Otherwise, a decisive voting round with status quo q_0 follows. Only the agenda–setter of the decisive voting round is tax–exempt.*

This conditional experimentation mechanism is a hybrid of the signaling mechanism and a two–stage voting mechanism. We obtain:

Theorem 7.1. *Consider any public good problem with $n = 2$. The game consisting of this public good problem and the conditional experimentation mechanism admits a Bayesian equilibrium which implements the Condorcet winner, and does not involve mock votes nor experimentation on the equilibrium path.*

Proof. The proof can be found in the appendix. □

One implication of Theorem 7.1 is that all citizens except the agenda–setter of the decisive voting round are treated equally at the equilibrium which implements the Condorcet winner. Two remarks are in order.

First, the agenda–setter in the preliminary round is indifferent between making a proposal $q' \in \{q_1, q_2\}$ and calling a mock vote. Indeed, there also exists a Bayesian equilibrium in which q_0 is proposed, experimentation occurs, and experimentation group members are tax–exempt in equilibrium. One might argue here that the equilibrium would not exist if the agenda–setter in the preliminary round cares slightly about equal treatment or is marginally risk–averse. In addition, if the equilibrium exists, it can be avoided by excluding the agenda–setter from becoming an experimentation group member if he calls for a mock vote with q_0.

Second, there are two other interesting variants of democratic mechanisms with conditional experimentation which we discuss next. One possible variant of the conditional experimentation mechanism is a mechanism where the experimentation group members are drawn in the beginning of the procedure. In the preliminary round of such a mechanism, the agenda–setter would be strictly better off by calling a mock vote if he is an experimentation group member, and strictly better off by proposing a non–zero quantity if he is not an experimentation group member. The drawback of such a mechanism would be that ex post unequal treatment of the citizens occurs with a strictly positive (yet arbitrarily small) probability λ. This problem can be avoided in a further variant in which the first agenda–setter is excluded from being an experimentation group member. Under that variant, the selection of the agenda–setter and the experimentation group members are combined in a hierarchical selection procedure. Such a *hierarchical selection procedure* fulfills the requirements of democratic mechanisms and has the following properties:

1. At the beginning, every citizen has the same probability λ of being pre–selected for experimentation group membership.

2. Each pre–selected citizen has the same probability of becoming the agenda–setter of the preliminary round.

3. All pre–selected citizens except the agenda–setter form the experimentation group.

We obtain:

Corollary 7.2. *Consider any public good problem with $n = 2$. Then, the game consisting of the conditional experimentation mechanism with a hierarchical selection procedure admits a Bayesian equilibrium which implements the Condorcet winner. The agenda–setter has a strict preference for making a proposal $q' \in Q \setminus \{q_0\}$ in the preliminary round.*

The proof of this corollary follows from the same considerations as the proof of Theorem 7.1, complemented by the observation that the expected utility of the agenda–setter from calling the mock vote in the preliminary round is $zq^* - \frac{c(q^*)}{1-\lambda}$ which is strictly smaller than the expected utility $zq^* - c(q^*)$ from making a proposal $q' \in Q \setminus \{q_0\}$.

To sum up, conditional experimentation leads to an equilibrium with equal treatment of all citizens ex ante as well as equal treatment of all citizens except the agenda–setter ex post on the equilibrium path.[9] Equal treatment of citizens (with the same income) has been a prominent theme and desideratum in public finance and its constitutional foundations.[10] While we cannot avoid different tax treatments off the equilibrium path, conditional experimentation can ensure equal treatment in equilibrium.

8 Conditional experimentation with many alternatives

As a final result, we show that a conditional experimentation mechanism can implement the Condorcet winner of a public good problem with an arbitrary number n of public good levels while avoiding experimentation on the constructed equilibrium path of play.

Conditional experimentation mechanism in $n + 1$ rounds. *The mechanism consists of (up to) n preliminary rounds followed by a decisive voting round. Each preliminary round is of the following form. First, one citizen is drawn at random from the entire population to be the agenda–setter. This citizen puts forward a preliminary proposal, say q'. If $q' = q_0$ or if q' has been proposed in a previous preliminary round, then we say that a* mock vote *has been called. Indeed, if a mock vote has been called, then the rest of the mechanism corresponds to the signaling mechanism. That is, an experimentation group is drawn, signals are sent, and a decisive voting round with status quo q_0 follows. If the vote in a preliminary round is not a mock vote, then all citizens simultaneously vote Yes or No to the proposal q'. If q' is accepted, then we say that q' has* prevailed *and that q' becomes the new* default. *If q' is rejected, then we say that the previous default has* prevailed, *and it remains in place. In the first preliminary round, the initial default is q_0. If n preliminary rounds have passed without a mock vote, a decisive voting round follows. The status quo of this decisive voting round is the highest quantity against which no other quantity has prevailed. In particular, if no quantity has ever prevailed against the initial default q_0, then q_0 is the status quo of the decisive voting round. The agenda–setter in the decisive voting round is tax–exempt.*

Analogously to the signaling game and the two–stage voting game, we define a *conditional experimentation game* as consisting of the conditional experimentation mechanism in $n + 1$ rounds and a public good problem $P \in \mathcal{P}$. In a conditional experimentation game, a strategy profile must, among other things, specify a preliminary proposal to be

[9] If one wants to treat also the agenda–setter of the decisive voting round equally, one could levy an ex ante fee for agenda–setting equal to the expected tax burden of citizens.

[10] See for instance Gersbach, Hahn and Imhof (2013) for a discussion.

made in each preliminary round. The notion of a Bayesian equilibrium with up to n preliminary rounds is a straightforward extension of the equilibrium notion in Section 5 for two–stage voting games, consisting of proposal strategies, voting strategies, and beliefs. Joint deviations are defined analogously as in Section 5.

For the analysis to follow, it is useful to formally introduce some particular strategies and beliefs. We define the *descending proposal strategy* as the strategy under which the agenda–setter in every preliminary round makes the highest proposal which has not been made before. We denote this strategy by ρ^* in what follows. Moreover, we define a belief $\gamma^* \in \Delta_+^n$ which is held by all citizens regardless of their types and which depends only on which element of Q is the default after all preliminary rounds. In particular, if the default after n preliminary rounds is $q_k \in Q \setminus \{q_0\}$, then we define $\gamma_k^* = 1$. If the default after n preliminary rounds is q_0, then we define $\gamma_n^* = 1$. Finally, we denote by σ^* the sincere voting strategy; that is, under σ^*, citizen z votes in favor of the preliminary proposal q' against the default $\bar{q}$ if and only if $u(z, q') > u(z, \bar{q})$. In what follows, we are going to show that the profile of strategies and beliefs $(\rho^*, \sigma^*, \gamma^*)$ is a Bayesian equilibrium of the conditional experimentation game. Since this should implement the Condorcet winner, this is tantamount to a constructive proof of the following theorem.

Theorem 8.1. *The conditional experimentation mechanism implements the Condorcet winner.*

Proof. Step 1. In order to verify the consistency of the beliefs γ^*, let us first describe the path of play induced by the profile $(\rho^*, \sigma^*, \gamma^*)$. In the n preliminary rounds, all possible quantities are proposed in descending order. Each proposal which is higher than the Condorcet winner is rejected, so that q_0 remains the default. When the Condorcet winner is proposed, it prevails and thus becomes the new default. Subsequently, all quantities smaller than the Condorcet winner will be proposed and rejected so that the Condorcet winner remains the default until the end of the last preliminary round. We conclude that the belief γ^* is consistent with the path of play induced by ρ^* and σ^*.

In the next steps, we show that the profile under consideration does implement the Condorcet winner.

Step 2. As a next step, consider the preliminary proposals. With positive probability, one (or even all) of the preliminary agenda–setters belong to Z_-, and may therefore have incentives to obstruct the revelation of the Condorcet winner. The question is whether the preliminary agenda–setters can manipulate the conditional experimentation mechanism given that all citizens vote sincerely and given that the beliefs are as specified by γ^*. By construction, if one preliminary agenda–setter calls a mock vote, the remainder of the mechanism amounts to the signaling mechanism. Hence, the Condorcet winner will be implemented following a mock vote. Consequently, it cannot be a profitable deviation for any preliminary agenda–setter to call a mock vote. Suppose that the preliminary agenda–setters deviate from ρ^* in some way which does not trigger a mock vote. By the definition of a mock vote, this implies that each quantity $q \in Q \setminus \{q_0\}$ must be the preliminary proposal in exactly one preliminary round.

Step 3. In particular, the Condorcet winner is proposed in some preliminary round. Because voting is sincere, the Condorcet winner prevails in that round and becomes the new default. Again, due to sincere voting, the Condorcet winner then remains the default until the end of the n preliminary rounds. It follows that after the preliminary rounds, the belief γ^* assigns probability one to the Condorcet winner. The agenda–setter of the decisive voting round will therefore propose the Condorcet winner, which will then become the outcome of the mechanism. Indeed, the "conditional experimentation" prevents manipulation by the preliminary agenda–setters given that all citizens vote sincerely.

Step 4. In order to complete the proof of Theorem 8.1, we therefore have to show that sincere voting is optimal from the citizens' point of view. This is the claim of the next lemma.

□

Lemma 8.2. *The sincere voting strategy* σ^* *is optimal given the preliminary proposals prescribed by* ρ^* *and given the belief* γ^*.

Proof. The proof can be found in the appendix. □

In order to implement the Condorcet winner through the conditional experimentation mechanism, we have made use of the descending proposal strategy in the preliminary rounds. This strategy closely resembles the *amendment procedure,* which is one well–known way of organizing a legislature's agenda in practice. Under the amendment procedure, the legislators vote in several rounds; in each round, two alternatives are opposed to each other, and the loser of the vote is subsequently discarded. Under this procedure, the alternatives are typically voted upon in a natural order from the most extreme to the least extreme. Indeed, the descending proposal strategy emulates these features.

9 Discussion and conclusion

The main insight of the current paper is that democratic decision–making procedures can be used to identify and implement socially desirable policies even in the presence of profound uncertainty. The resolution of profound uncertainty and the implementation of the most socially desirable policy hinge on the use of experimentation as part of the decision–making process. A democratic mechanism based solely on repeated voting does not guarantee the revelation of the distribution of the types of citizens. We have established these findings in the context of a choice problem from a discrete set of possible public good levels. We stress that the introduced mechanisms with experimentation are prior–free, that is, they do not depend on the ex ante beliefs of citizens about the states of nature. This is a particularly desirable robustness property of democratic mechanisms, since they should be applicable to a variety of situations and their rules should not depend on citizens' current beliefs.

There is a variety of extensions and further applications which can be considered in future research. For instance, one could examine to what extent our results carry over to choices from different sets of possible policies, such as continuous policy spaces, or multi–dimensional public good problems in which several public goods can be combined in a bundle of public goods. Moreover, one might consider an electorate with different income levels and the possibility to differentiate the tax burden as a function of income. In such a model, one could investigate the effect of a policy chosen by a democratic mechanism on the degree of inequality among citizens.

While such extensions will considerably further the scope of democratic mechanisms in a polity, it is unlikely that optimal mechanisms will not involve experimentation. On the contrary, in the presence of profound uncertainty we expect experimentation to be an essential ingredient of democratic mechanisms.

Appendix

Proof of Lemma 4.5

Suppose by way of contradiction that in some Nash equilibrium of the decisive voting round no citizen applies for agenda–setting. Consider a deviation by some citizen z to the following strategy. Citizen z applies for agenda–setting and, if chosen as the agenda–setter, we assume he makes the proposal $\hat{q}$, where $\hat{q} = \bar{q}$ if $\bar{q} > q_0$ and $\hat{q} = q_1$ if $\bar{q} = q_0$. Sincere voting together with the assumption that $F_1(c(q_1)/q_1) < 1/2$ and the first–order stochastic dominance of F_k over F_1 for every $k \in N \setminus \{1\}$ jointly imply that $\hat{q}$ will be the outcome of the decisive voting round. Thus, the deviation leads to a payoff of $z\hat{q} > 0$ for citizen z. But in the supposed Nash equilibrium, citizen z would obtain a zero payoff, the desired contradiction.

Proof of Lemma 5.6

It is straight–forward that the profile $(\bar{\sigma}, \bar{\pi})$ leads to the revelation of the state and the implementation of the Condorcet winner. It is also easy to see that the belief $\bar{\pi}$ is consistent with the strategy $\bar{\sigma}$. We need to show that the strategy $\bar{\sigma}$ is optimal given the belief $\bar{\pi}$. For this purpose, we verify whether deviations by some citizens $z \in Z_+$ or $z \in Z \setminus Z_+$ can be profitable. Suppose first that some subset of $X \subset Z_+$ deviates from $\bar{\sigma}$ by voting No. If the true state is the first state, then this deviation is inconsequential. If the true state is the second state and the deviation is not inconsequential, then the deviation leads to a decisive voting round with status quo q_0 and to the deterministic but erroneous belief e_1. Clearly, q_1 will be the outcome of the mechanism, whereas without the deviation the outcome would have been q_2. But all the deviating players prefer q_2 over q_1. We see that the deviation is not profitable. Now suppose that some subset $X \subset Z \setminus Z_+$ deviates from $\bar{\sigma}$ by voting Yes. If the true state is the second state, this is inconsequential. If the true state is the first state and the deviation is not inconsequential, then the deviation leads to a decisive voting round with status quo q' and to the deterministic but erroneous belief e_2. If $q' = q_1$, then the decisive voting round will have status quo q_1 and the proposal will be q_2. Since we are in the first state, the outcome will be q_1, as it would have been without the deviation. If $q' = q_2$, then the decisive voting round will lead to the outcome q_2 instead of q_1. But all deviating players prefer q_1 over q_2, hence the deviation is not profitable.

Proof of Lemma 6.1

To demonstrate the lemma, it suffices to show that the following statements hold when $\lambda > 0$ is sufficiently small.

$$\begin{aligned} \left\{ \frac{c(q_1)}{q_1(1-\lambda)}; \frac{c(q_n) - c(q_{n-1})}{(q_n - q_{n-1})(1-\lambda)} \right\} &\subset \quad int(Z), \\ F_k\left(\frac{c(q_k) - c(q_j)}{(q_k - q_j)(1-\lambda)} \right) &< \quad \frac{1}{2},\ k \in N,\ j \in \{0, 1, \ldots, k-1\}, \\ F_k\left(\frac{c(q_{k+1})}{q_{k+1}(1-\lambda)} \right) &> \quad \frac{1}{2},\ k \in N \setminus \{n\}. \end{aligned}$$

These statements follow from our assumptions on Q and mirror (4), (5), and (6). The first set inclusion comes from the fact that the two fractions change continuously with λ and from the fact that the set inclusion (4) places the fractions in the interior of the interval Z when $\lambda = 0$. The inequality in the second line above follows from the facts that inequality (5) is strict, that the fraction in the argument of F_k changes continuously with λ, and, moreover, that the function F_k itself is continuous. The inequality in the third line above follows from the same considerations. It holds for any $\lambda > 0$ such that $\frac{c(q_{k+1})}{q_{k+1}(1-\lambda)} \in Z$.

Proof of Theorem 6.2

It is straight–forward that the belief $\pi^{\bar{z}}$ is consistent with the strategy $\sigma^{\bar{z}}$. We show that $\sigma^{\bar{z}}$ is optimal given $\pi^{\bar{z}}$. Suppose that the true state is k^*. Suppose that citizens in some $Y \subset Z$ jointly deviate from $\sigma^{\bar{z}}$. Denote the resulting share of favorable votes by δ'. We first consider the case where there is some $k \in N$ such that $\delta' = 1 - F_k(\bar{z})$. Then, we have $\pi^{\bar{z}}(\delta', z) = e_k$ for all $z \in Z$. Consequently, the proposal in the decisive voting round will be q_k. If $k = k^*$, then the deviation under consideration is inconsequential, and therefore not profitable. Suppose now that the true state differs from k ($k^* \neq k$). If $k^* > k$, then q_k will be approved in the decisive voting round, and the quantity $q_k < q_{k^*}$ will be implemented. But without the deviation by Y, the quantity q_{k^*} would have been implemented. Because all experimentation group members (and, in particular, all members of Y) are tax–exempt, they prefer more to less of the public good, so the deviation is not profitable. If $k^* < k$, then the quantity q_k will not be approved in the decisive voting round. Hence, no public good will be provided. Again, the deviation is not profitable for the members of Y. Now consider the case where there is no $k \in N$ such that $\delta' = 1 - F_k(\bar{z})$. Then, $\pi^{\bar{z}}(\delta', z) = \beta(z)$; the state remains hidden. By Lemma 4.4, the expected public good level is strictly lower when the state is hidden than when it is revealed.Since all experimentation group members are tax–exempt, they prefer a higher expected public good level– again, the deviation by Y is not profitable.

Proof of Theorem 7.1

This result follows from Theorems 5.6 and 6.2. To be more specific, suppose that the agenda–setter in the preliminary round has proposed q_0 and thereby called a mock vote. Then, the remainder of the decision–making procedure amounts to the signaling mechanism. A favorable vote by an experimentation group member is tantamount to "sending the signal," while voting against the preliminary proposal amounts to "not sending the signal." It follows from Theorem 6.2 that the state of nature can be revealed and the Condorcet winner implemented. If the agenda–setter decides to call a mock vote in the preliminary round by putting forward q_0 as the preliminary proposal, he expects to become an experimentation group member with probability λ. If $q^* \in Q \backslash \{q_0\}$ is the Condorcet winner and $z \in Z$ is the type of the agenda–setter in the preliminary round, then his expected utility from calling the mock vote is $zq^* - (1-\lambda)\frac{c(q^*)}{1-\lambda} - \lambda 0 = zq^* - c(q^*)$. Now suppose that the agenda–setter of the preliminary round has proposed some $q' \in \{q_1, q_2\}$. In that case, it follows from Theorem 5.6 that the state can be revealed and the Condorcet winner implemented in an equilibrium of the ensuing subgame, in which case the agenda–setter of the preliminary voting round receives the same utility $zq^* - c(q^*)$ which he would also receive if he called for a mock vote. A deviation from a strategy profile where he proposes some $q' \in Q \setminus \{q_0\}$ to calling the mock vote would not be profitable. Consequently, the game consisting of the conditional experimentation mechanism and a public good problem with $n = 2$ admits an equilibrium in which the Condorcet winner is implemented although no mock vote is called and thus no experimentation takes place.

Proof of Lemma 8.2

Suppose that q_i is the Condorcet winner. Consider a joint deviation from σ^* by a subset $X \subset Z$ of the citizens. Suppose that following this deviation, the outcome of the mechanism is $q_j \in Q \setminus \{q_i\}$. We need to show that there is some $z \in X$ for whom this deviation is not profitable. We distinguish the following two cases.

Case 1 ($q_j > q_i$). Given ρ^*, the supposition that q_j is the outcome of the mechanism implies that q_j has prevailed against q_{j-1}. Let $M' = \{z \in Z | u(z, q_{i+1}) < 0\}$. Since q_i is the Condorcet winner, it follows from the inequality (6) that M' contains a majority of citizens. By the convexity of $c(q)$, we have that $u(z, q_k) < u(z, q_{k-1})$ for all $z \in M'$ and $k \in N$ such that $k > i$. In particular, this implies the inequality $u(z, q_j) < u(z, q_i)$. From the inequality $u(z, q_k) < u(z, q_{k-1})$ for all $z \in M'$ with $k = j$ and the fact that q_j prevailed against q_{j-1}, it follows that some members of M' must have voted insincerely so that $M' \cap X \neq \emptyset$. Indeed, let $z' \in M' \cap X$. Then, it follows from the above that $u(z', q_j) < u(z', q_i)$. We have now found a citizen $z' \in X$ who is worse off after the deviation by X, as desired.

Case 2 ($q_j < q_i$). Given ρ^*, the supposition that q_j is the outcome of the mechanism implies that q_{i-1} must have prevailed against q_i in some preliminary round. Let $M'' = \{z \in Z | u(z, q_i) > u(z, q_{i-1})\}$. Since q_i is the Condorcet winner, the set M'' contains a majority of citizens. Hence, some members of M'' have not voted sincerely, and thus $M'' \cap X \neq \emptyset$. Now let $z'' \in M''$. By the convexity of $c(q)$, it follows that $u(z'', q_i) > u(z'', q_j)$. Hence, citizen z'' is worse off following the deviation from σ^*, which completes the proof of the lemma.

References

Aghion, P. and P. Bolton (2003) "Incomplete social contracts," *Journal of the European Economic Association, 1,* 38-67.

Bednar, J. (2011) "Nudging federalism towards productive experimentation," *Regional and Federal Studies, 4-5,* 503-521.

Bendor, J. and D.Mookherjee (1987) "Institutional structure and the logic of ongoing collective action," *American Political Science Review, 81,* 129-154.

Bergemann, D. and S.Morris (2005) "Robust mechanism design," *Econometrica, 73,* 1771-1813.

Bergemann, D. and J.Välimäki (2006) "Bandit problems," Cowles Foundation Discussion Paper No.1551.

Bernhardt, D., J. Duggan and F.Squintani (2008) "A survey on polling in elections," in Aragones, E., Bevia, C., Llavador, H., Schofield, N., Eds., The Political Economy of Democracy, Fundacion BBVA.

Buchanan, J.M. and G.Tullock (1962) "The Calculus of Consent: Logical Foundations of Constitutional Democracy." Ann Arbor: University of Michigan Press.

Buera, F.J., A.Monge-Naranjo, and G.E.Primiceri (2011) "Learning the wealth of nations," *Econometrica, 79,* 1-45.

Butenko, V. (2013) Democratic Decision–Making and Experimentation. Dissertation No. 21208, ETH Zurich, Switzerland.

Cai, H. and Treisman, D. (2009) "Political decentralization and policy experimentation," *Quarterly Journal of Political Science, 4(1),* 35-58.

Callander, S. and B.Harstad (2013) "Experimentation in federal systems," NBER Working Paper 19601.

Gersbach, H. (2009) "Democratic Mechanisms," *Journal of the European Economic Association, 7,* 1436-1469.

Gersbach, H., Hahn, V., and S.Imhof (2013) "Tax Rules," *Social Choice and Welfare, 41,* 19-42.

Kollman, K. (2003) "The rotating presidency of the European Council as a search for good policies," *European Union Politics, 4,* 51-74.

Kollman, K., J.H.Miller, and S.E.Page (2000) "Decentralization and the search for policy solutions," *The Journal of Law, Economics, and Organization, 16,* 102-128.

Ledyard, J. (2008) "Voting and Efficient Public Good Mechanisms," in Wittman, D.A. and B.R. Weingast (Eds.), The Oxford Handbook of Political Economy, doi. 10.1093/oxfordhb/9780199548477.001.0001

Ledyard, J. and T.Palfrey (1994) "Voting and Lottery Drafts as Efficient Public Good Mechanisms," *Review of Economic Studies, 61,* 327-355.

Ledyard, J. and T.Palfrey (2002) "The approximation of efficient public good mechanisms by simple voting schemes," *Journal of Public Economics, 83,* 153-171.

Rose-Ackerman, S. (1980) "Risk-taking and reflection: Does federalism promote innovations," Rothschild, M. (1974) "A two-armed bandit theory of market pricing," *Journal of Economic Theory 9,* 185-202.

Shipan, C.R. and C.Volden (2006) "Bottom–up federalism: The diffusion of anti–smoking policies from U.S. cities to states," *American Journal of Political Science, 50,* 825-843.

Strumpf, K.S. (2002) "Does government decentralization increase policy innovation?" *Journal of Public Economic Theory, 4,* 207-241.

Volden, C. (2006) "States as policy laboratories: Emulating success in the children's health insurance program," *American Journal of Political Science, 50,* 294-312.

Volden, C., M.M.Ting, and D.P.Carpenter (2008) "A formal model of learning and policy diffusion," *American Political Science Review, 102,* 319-332.