An Optimal Voting System When Voting is Costly*

Katalin Bognar[†], Tilman Börgers[‡], Moritz Meyer-ter-Vehn[§]

May 10, 2012

Abstract

We propose a formulation for the problem of designing an optimal dynamic voting system when voting is costly. The possibility to coordinate participation by voting sequentially requires explicit maximization over a set of dynamic mechanisms which involves some arguments that are new to the theory of mechanism design. For a highly stylized specification of our model with private values and two alternatives we show the optimality of a voting system that combines two main elements: (i) there is an arbitrarily chosen default decision and non-participation is interpreted as a vote in favor of the default; (ii) voting is sequential and is terminated when a supermajority requirement, which declines over time, is met. We show the optimality of such a voting system by arguing that it is first best, that is, it maximizes welfare when incentive compatibility constraints are ignored, and by showing that individual incentives and social welfare can be sufficiently aligned to make a first best system incentive compatible.

*This paper supersedes a previous draft entitled "Dynamic Mechanism Design with Costly Participation." We are grateful to Arthur Lupia, Romans Pancs, Alessandro Pavan, Marek Pycia, Ilya Segal, Ricky Vohra, and two anonymous referees for helpful comments and to Yan-Min Choo for research assistance. Tilman Börgers thanks the Center for Mathematical Studies in Economics and Management Science at Northwestern University for the hospitality that it offered while he was working on this paper. Moritz Meyer-ter-Vehn thanks the Hausdorff Research Institute for Mathematics in Bonn for its hospitality, which he enjoyed while working on this paper.

[†]kbognar@umich.edu.

[‡]Department of Economics, University of Michigan, tborgers@umich.edu.

[§]Department of Economics, UCLA, mtv@econ.ucla.edu.

1 Introduction

We consider the design of expected welfare maximizing voting systems when participation in votes is costly. The design of an expected welfare-maximizing voting system must trade off the potential expected welfare loss that results from sub-optimal decisions against the expected cost of voting. Sequential voting procedures are often better solutions than static procedures. It may for example happen that a margin for one candidate becomes established by the votes of the first voters that cannot be overturned by the remaining voters. In this case it is clearly optimally to terminate voting, and thus save the voting cost of the remaining voters.

In this paper we propose a formulation of the problem of the design of a welfare maximizing voting system with costly voting that allows for sequential mechanisms. The designer of the voting scheme maximizes over a set of extensive-form games, anticipating that agents will play a sequential equilibrium of this game. We determine an optimal extensive-form game with a corresponding sequential equilibrium.

We consider the following setting. A set of voters has to choose one candidate from a set of two possible candidates for a public office. Each voter has a strict preference over candidates, and knows his or her own preference, but not other voters' preferences. Voters' preferences are independent, and each voter is equally likely to prefer either candidate. Participation in the voting procedure has a known positive cost that is common to all voters.

Our main result is that there is a simple optimal voting system: voters vote sequentially, one after the other. No voter is forced to participate. Non-participation is interpreted as a vote for a default candidate. A candidate is elected if a sufficiently large majority of voters who have expressed their preference, either by voting or by not voting, favors the candidate. In that case, voting is terminated. It may be that at that stage the remaining voters could still overturn the majority. The majority requirement decreases as the number of remaining voters decreases.

It is intuitive that the optimal procedure makes use of the fact that non-participation is free, and that non-participation is interpreted as if it was a vote. Arguably, non-participation is in practice sometimes observed as an expression of a preference. If, for example, a department chair informs department members that she will take one particular course of action unless a majority of members objects, she establishes through this announcement that non-participation is interpreted as a vote.

We interpret in our model non-participation as a vote for a default candidate. In the baseline model the choice of default candidate is arbitrary, and could, in principle, change during the voting procedure. Small perturbations of the baseline model make the candidate who is more likely to be preferred by a majority of voters the uniquely optimal default candidate.[1]

It is not surprising that not all voters are necessarily consulted. The cost of voting may outweigh the potential welfare gain from a more accurate elicitation of the will of the majority. In practice, the number of voters who is consulted is, unlike in our model, often determined ex ante, for example by forming a committee. In our model, by contrast, voters vote sequentially, and the decision whether to consult further voters depends on the size of the majority established so far. The optimality of this procedure in our model is related to a strong assumption that we make: observing the progress of voting is free. Only actual participation carries cost. In practice, it may be easier to implement, say, a two stage procedure: first a small group of voters cast votes. If they favor some decision with a large margin, then this decision is final. Otherwise, a larger group of voters, perhaps all voters, cast votes, and the majority preference is implemented. This procedure, which may be worth considering in practice, is built on the assumption that it is less costly to observe the result of a committee meeting than to attend a meeting oneself.

Our findings are for a very special model. We shall explain the difficulties that are raised by generalizations of the model. A central motivation for studying this very special model is methodological. The main feature of our work that is of methodological interest is that we are studying a mechanism design problem in which it would not be sufficient to consider static games only. Rather, we have to consider extensive game forms with sequences of moves and non-trivial information sets. Formally, the choice set of the designer of the voting system in this paper is the set of all finite extensive game forms with endogenous participation choices. Choosing optimally from this set is a much more complex problem than choosing the optimal mechanism from static games, which is traditionally considered in mechanism design, where in addition typically one can assume that the agents' message spaces are identical to their type spaces. In this paper we present one instance of this problem that we have found to be analytically tractable. The more general problem is well-defined, but currently a solution is outside of our reach.

A mechanism designer who chooses from the set of all extensive game forms has also been considered in a recent paper by Gershkov and Szentes (2009). They, too, study the optimal design of a voting scheme. However, whereas we consider a private value setting, they consider a setting with common values. While in our setting participation is costly, in their set-

[1] This is explained in Section 4.3 below.

ting information acquisition is costly. As is the case in our paper, they cannot restrict attention to static revelation games but instead need to consider all extensive game forms. Their paper therefore deals with a similar methodological issue as ours.

It is interesting to compare Gershkov and Szentes' analytical approach to ours. We proceed by first identifying a canonical class of mechanisms to which we can restrict attention. We then show that within this class the mechanisms that maximize ex ante expected welfare when incentive constraints are relaxed, i.e. "first best" mechanisms, are also incentive compatible. Gershkov and Szentes, too, find a class of canonical mechanisms to which they can restrict attention, and then optimize within this class. However, first best mechanisms are not incentive compatible in their setting.

The reason why the difference between private and social interests has more severe consequences in Gershkov and Szentes' setting than in ours is that in our set-up, by interpreting non-participation as a signal, the mechanism designer can effectively choose which action is costly, and which action is free. So, in particular, he can always make whatever action has positive expected externalities freely available to agents. Therefore, if in the first best mechanism the mechanism designer asks an agent to take a costly action, we can conclude that asking the agent to take this action is not based on the positive externalities of this action, but on the benefits that the mechanism designer expects to accrue to this particular agent. The agent will therefore also find that it is in his interest to take the costly action, and the mechanism designer's request will be incentive compatible.

In Gershkov and Szentes' set-up, by contrast, the costs of actions are intrinsic to those actions. They cannot be chosen by the mechanism designer. If information acquisition has positive externalities, then the mechanism designer does not have any way of making information acquisition free. If the mechanism designer in the first best mechanism asks an agent to acquire information, this request may well be based on the positive externalities from information acquisition rather than on that agent's own interests. The request may therefore not be incentive compatible.

Gershkov and Szentes' canonical mechanisms provide minimal information to voters to relax voters' incentive constraints. Even with this construction, they do not obtain incentive compatibility of the first best mechanism. By contrast, we consider mechanisms in which all information about previous votes is revealed to agents and show that even though such a mechanism maximizes the opportunities for deviations, the optimal canonical mechanism is incentive compatible. We view this maximal informativeness as an attractive feature of our mechanism because in practice it may be difficult to conceal information from voters.

Some authors have considered welfare properties of particular voting schemes either in a private or in a common value setting, without examining the general mechanism design problem. For example, Börgers (2004) showed in a model similar to ours the superiority of voluntary voting participation over required voting participation when voting is costly. Ghosal and Lockwood (2009) and Krasa and Polborn (2009) describe models in which the opposite conclusion can be reached. Gershkov and Szentes (2009) reference a number of papers that study particular voting institutions, and in which the emphasis is not on participation cost but on information acquisition cost.

Issues analogous to the ones investigated in this paper arise in other mechanism design problems whenever participation is costly and it is possible that participants observe other players' actions without incurring all participation cost. Procurement auctions seem a realistic example. Participants in procurement auctions might first observe the participation decision of other bidders before incurring the cost of preparing their own bidding material. The existing literature has studied auction design problems in the case in which all participation decisions have to be made simultaneously. Examples of relevant papers are Stegeman (1996) who considers welfare maximizing auction design in this environment, and Celik and Yilankaya (2009), who consider profit maximizing auction design. Our paper indicates how one can set up the auction designer's problem if one wants to consider the possibility of reducing participation cost, or perhaps of encouraging participation, through sequential participation decisions.[2]

Considering information acquisition cost rather than participation cost, Bergemann and Välimäki (2002) find for a general environment with transferable utility that Vickrey-Clarke-Groves mechanisms provide socially optimal incentives for information acquisition when agents have independent private values. Our main result is similar in that it also proves the incentive compatibility of a first best solution, but the underlying intuition is different: the transfer payments in a Vickrey-Clarke-Groves mechanism induce perfectly aligned individual and social incentives, whereas in our setting individual and social incentives potentially diverge, yet equilibrium participation decisions are socially optimal. Further results in the literature on information acquisition in environments with transferable utility are surveyed in Bergemann and Välimäki (2006). A recent paper by Pancs (2011) considers a setting with two potential buyers of a single, indivisible object, and finds a sequential mechanism with partial information disclo-

[2] Celik and Yilankaya (2009, p. 3) note that dynamic mechanisms might be of advantage in their setting.

sure to be maximizing the expected value of a weighted average of welfare and revenue. This paper is related to our work because it also considers an optimization problem where the objects among which the mechanism designer chooses are extensive game forms.

A different type of cost, communication cost, are considered in Fadel and Segal (2009). The communication cost of a decision rule equals the minimal number of bits of information agents must transmit in an incentive compatible mechanism that computes this decision rule. In the model specification with two equally likely types studied in this paper, our participation costs are equivalent to Fadel and Segal's communication costs. One difference between the papers is that they fix a decision rule and study the minimal communication costs to implement it, while we also study the tradeoff between suboptimal decisions and participation costs. In a more general specification of our model with three or more types, our participation cost correspond to a fixed cost of communication, whereas Fadel and Segal's communication costs are proportional to the number of bits communicated.

We present our framework in Section 2. In Section 3 we state and prove the main results. Section 4 contains some concluding remarks.

2 Set-Up

There are $n \in \mathbb{N}$ voters. We denote the voters by: $i \in N \equiv \{1, 2, \ldots, n\}$. The voters have to pick one of two candidates: $y \in \{A, B\}$. Each voter $i \in N$ has a type $\theta_i \in \{A, B\}$. The type of voter i indicates which candidate voter i prefers. Types are random, and different voters' types are independent. Each voter is equally likely to prefer candidate A and candidate B. Each voter observes his or her own type, but not the other voters' types. Voter i's Bernoulli utility equals $1 - c$ if i's preferred candidate wins and i participated in the decision process, $-c$ if i's preferred candidate loses and i participated in the decision process, 1 if i's preferred candidate wins and i did not participate in the decision process, and 0 if i's preferred candidate loses and i did not participate in the decision process. Here, $c > 0$ is a constant, the cost of participation in the decision making process. The distribution of types, what each voter observes, the Bernoulli utility function, and the value of c are common knowledge among voters and the mechanism designer. The mechanism designer does not observe any voter's type. The mechanism designer seeks to maximize the sum of all voters' ex ante expected utilities.

The setting described so far is obviously very special. We study this simple setting to focus on methodological issues without distraction and because elementary modifications of the set-up make the question that we examine analytically significantly harder, as we explain in the last section of the paper. As anticipated in the Introduction, despite its simplicity, the setting that we study yields results that are of some real world plausibility.

The mechanism designer chooses firstly an extensive game form that the voters then use to pick one of the two candidates, and secondly a sequential equilibrium of the game implied by the game form, the voters' Bernoulli utilities, and the voters' information structure. It is assumed that the voters play the sequential equilibrium that the mechanism designer chooses.

We begin by describing the set of extensive game forms that the mechanism designer can choose from. We have adapted the following definition of *extensive game forms* from Osborne and Rubinstein (1994, pp. 200-201).

Definition 1. *An* extensive game form *consists of:*

1. *The set N of* players *(identical to the set of voters).*
2. *A finite set H of finite sequences with the following properties:*
   (a) *The empty sequence $\emptyset$ is an element of H;*
   (b) *If $\left(a^k\right)_{k=1,\dots,K} \in H$ and $L < K$ then $\left(a^k\right)_{k=1,\dots,L} \in H$.*

   (Each element of H is a history*, or, equivalently, a* node*. Each component of a history is an* action *taken by a player, or a chance move. A history $\left(a^k\right)_{k=1,\dots,K} \in H$ is* terminal *if there is no a^{K+1} such that $(a^k)_{k=1,\dots,K+1} \in H$. The set of terminal histories is denoted by Z. The set of actions available after a nonterminal history $h \in H \setminus Z$ is denoted by $\mathcal{A}(h) \equiv \{a : (h; a) \in H\}$.)*
3. *A function P that assigns to each nonterminal history $h \in H \setminus Z$ an element of $N \cup \{C\}$. (P is the* player function*, $P(h)$ being the player who takes an action after history h. If $P(h) = C$ then chance determines the action taken after history h.)*
4. *A function f_C that associates with every history h for which $P(h) = C$ a probability measure $f_C(\cdot|h)$ on $\mathcal{A}(h)$. ($f_C(a|h)$ is the probability that a occurs after history h.)*
5. *For each player $i \in N$ a partition $\mathcal{I}_i$ of $\{h \in H : P(h) = i\}$ with the property that $\mathcal{A}(h) = \mathcal{A}(h')$ whenever h and h' are in the same element of the partition. ($\mathcal{I}_i$ is the* information partition *of player i. Any set $I_i \in \mathcal{I}_i$ is an* information set *of player i.)*

6. *A function D that assigns to each terminal history $h \in Z$ a decision $D(h) \in [0,1]$. ($D(h)$ as the probability that A is chosen after history h. Candidate B is chosen with the remaining probability $1 - D(h)$. We also write $D(h) = A$ if $D(h) = 1$, and $D(h) = B$ if $D(h) = 0$.)*

We next single out particular extensive game forms, namely those that can be interpreted as decision making procedures that include voters' decisions about participation in the procedure. We call such extensive game forms *mechanisms with participation decisions.*

Definition 2. *An extensive game form is a* mechanism with participation decisions *if:*

1. *For every nonterminal history $h \in H \setminus Z$ we either have:*

   (a) $P(h) \in N$ and $\mathcal{A}(h) \subseteq \{0,1\}$, or

   (b) $\mathcal{A}(h) \cap \{0,1\} = \emptyset$.

   (Histories that satisfy (a) will be called participation nodes *of player $P(h)$. For such histories we interpret the action 1 as participation by voter $P(h)$, and the action 0 as non-participation by voter $P(h)$.)*

2. *For every terminal history $\left(a^k\right)_{k=1,\dots,K} \in Z$ and every $i \in N$ we either have:*

   (a) if $\ell < K$ and $P(a^1, \dots, a^\ell) = i$ then $a^{\ell+1} = 0$

   or there is a unique $L < K$ such that:

   (b) if $\ell < L$ and $P(a^1, \dots, a^\ell) = i$ then $a^{\ell+1} = 0$;

   (c) $P(a^1, \dots, a^L) = i$ and $a^{L+1} = 1$;

   (d) if $\ell > L$ and $P(a^1, \dots, a^\ell) = i$ then $a^{\ell+1} \notin \{0,1\}$.

   (Either player i does not participate in this terminal history, or he participates at the $L+1$st action node of the history. In the latter case, all of i's earlier action nodes are participation nodes at which i does not participate, but none of the later action nodes are participation nodes. This definition implies that before doing anything else players have to choose to participate.)

For every mechanism with participation decisions there is an associated game of incomplete information in which voters first learn privately their types, and then play the mechanism. Voters evaluate outcomes according to their Bernoulli utility function. The cost c is incurred by a player i if and only if a participation history of that player i is reached and the player chooses to participate, that is, chooses action 1 at that point. We denote strategies of player i by σ_i. The solution concept that we shall use is sequential equilibrium. The mechanism designer's problem is then as follows:

The Mechanism Designer's Problem. *Choose a mechanism with participation decisions $\mathcal{M}$ and a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ for the incomplete information game associated with $\mathcal{M}$ to maximize the sum of ex ante expected utilities of all voters, subject to the constraint that σ is a sequential equilibrium of the incomplete information game associated with $\mathcal{M}$.*

An important assumption of our construction is that voters can observe the actions of other voters without incurring participation cost. Put differently, the cost to observing the outcome of an earlier vote are assumed to be negligible in comparison to the cost of participation in the voting mechanism. Observation may only require watching a TV or computer screen, which is essentially costless, whereas participation requires actual presence in a room. We have discussed this assumption in the Introduction. As an extreme case, it may be useful to have in mind members of the US congress who watch roll-call votes on a screen in their offices, and therefore can let their decision whether to leave their offices and to enter the chamber to cast their own votes depend on the current margin of votes.

3 Results

We solve the mechanism designer's problem in two steps. In Proposition 1, we exhibit basic properties of a solution. In Proposition 2 we use the insights from Proposition 1 to develop a more detailed characterization of an optimal mechanism.

Proposition 1. *A solution $(\mathcal{M}, \sigma)$ to the mechanism designer's problem exists. Moreover, there is at least one solution with the following properties:*

(i) There is no $h \in H \setminus Z$ such that $P(h) = C$. (There are no chance moves.)

(ii) *For every* $i \in N$ *and every* $I_i \in \mathcal{I}_i$*:* $\#I_i = 1$*. (The extensive game form is of perfect information.)*

(iii) *For every* $h = (a^1, \ldots, a^K) \in Z$ *and every* $i \in N$ *there is at most one* $\ell < K$ *such that* $P(a^1, \ldots, a^\ell) = i$*. Moreover then:* $A(a^1, \ldots, a^\ell) = \{0, 1\}$*. (Every voter makes at most one decision. This decision is a participation decision. Participation is voluntary.)*

(iv) *There is an* $x \in \{A, B\}$ *such that for every* $h = (a^1, \ldots, a^K) \in Z$*:* $D(h) = x$ *if* $\#\{k|a^k = 1\} < K/2$*, and* $D(h) \neq x$ *if* $\#\{k|a^k = 1\} > K/2$*. (Candidate* x *is the* default *candidate. If a majority of voters who make a choice decide not to participate, then* x *is chosen. If a majority of voters who make a choice decide to participate, then the default is overturned.)*

(v) *If* $P(h) = i$ *and* $\mathcal{A}(h) = \{0, 1\}$*, then* $\sigma(h, \theta_i) = 1$ *if and only if* $\theta_i \neq x$*. (A voter participates if and only if she opposes the default.)*

(vi) *If* $h = (a^1, \ldots, a^K)$ *is such that* $\#\{k|a^k = 0\} \geq n/2$ *or* $\#\{k|a^k = 1\} \geq n/2$ *then* $h \in Z$*. (The decision process ends at the latest when a majority of potential voters has indicated a preference for one candidate.)*

Proof. The proof has two parts. In the first part we consider a relaxed version of the mechanism designer's problem in which the constraint that σ has to be a sequential equilibrium is dropped. We show that a pair $(\mathcal{M}^*, \sigma^*)$ exists that solves the relaxed problem, and that has the properties described in Proposition 1. In the second part we show for any $(\mathcal{M}^*, \sigma^*)$ that solves the relaxed problem, and that has the properties described in Proposition 1, that σ^* is a sequential equilibrium of the extensive game with incomplete information corresponding to $\mathcal{M}^*$. Therefore such a solution of the relaxed maximization problem is also a solution of the mechanism designer's original problem.

Part 1. We consider the relaxed maximization problem. Let $\mathcal{M}$ be any mechanism with participation decisions, and let σ be a strategy combination for the corresponding extensive game with incomplete information. We shall explain how one can transform $\mathcal{M}$ into a mechanism with participation decisions $\hat{\mathcal{M}}$, and σ into a strategy combination $\hat{\sigma}$ for the extensive game with incomplete information corresponding to $\hat{\mathcal{M}}$, such that $(\hat{\mathcal{M}}, \hat{\sigma})$ have the properties described in Proposition 1, and expected welfare resulting from $(\hat{\mathcal{M}}, \hat{\sigma})$ is at least as large as expected welfare from $(\mathcal{M}, \sigma)$.

This implies that when solving the relaxed maximization problem, we can restrict attention to pairs $(\hat{\mathcal{M}}, \hat{\sigma})$ that have the properties described in Proposition 1. Existence of an optimal solution with the properties described in Proposition 1 then follows immediately, as there are only finitely many such pairs.

We now describe the transformation. If $\mathcal{M}$ has information sets with multiple elements, we begin by modifying $\mathcal{M}$ so that all information sets are singletons. We adjust σ so that each player's strategy assigns to every non-terminal history at which a player moves in the new extensive game form the same action that σ assigned previously to that history. In this step we thus provide players with additional information, but assume that they do not make use of this additional information. This step leaves expected welfare unchanged.

Next, we remove all randomization from the extensive game form and from the strategies. If the strategy of some player i involves randomization, we can find the equivalent mixed strategy, and then pick from the support of that mixed strategy that yields the largest expected welfare. Thus we remove all randomization from the players' strategies without lowering expected welfare. We can remove all chance moves in the extensive game form, or in the decision rule, by the same argument, treating "chance" as one further player. We can then remove all nodes at which the chance player chooses, replacing the branch from the chance move's predecessor to the chance move, and from the chance move to the chance move's successor, by a single branch that connects the predecessor to the successor.

We now pick one of the two candidates as the "default candidate." Let x be that candidate, and let y be the other candidate. It does not matter which candidate we pick to be the default candidate. Proceeding in an arbitrary order of players, we successively for each player i make the following changes to the extensive game form: We first consider the earliest nodes h at which player i moves. These must be participation nodes. If one type of player i participates at h, and another one does not participate, then we label the choice that type $\theta_i = x$ makes as "0" (i.e. don't participate), and the choice that type $\theta_i = y$ makes as "1" (i.e. participate). If the game tree that follows the choice 0 contains further choices of player i, then we remove all choices except the one that type $\theta_i = x$ would make at those nodes. We can then just as well remove these nodes from the game tree, because there is only one feasible choice at these nodes. Similarly, if the game tree that follows the choice 1 contains further choices of player i, then we remove all choices except the choice that type $\theta_i = y$ would have make at those nodes, and then we remove these nodes themselves from the game tree. If both types of player i participate at h, then we remove h from the game tree, assuming that player i participates at this node, and label the next decision nodes of player i as participation nodes. If both types of player i don't participate at h, we remove h from the game tree, assuming that player i does not participate at this node. We iterate this operation for player i until there are no further nodes left to consider. Then we proceed to the next player until there are no further players left to consider.

The step just described either leaves expected welfare unchanged or it increases it. This is because we don't change for any vector of realized types of all voters the candidate chosen. Moreover, in each step, if at the node that we are considering exactly one type of a player participates, then we don't change the expected participation cost or we reduce them. We may have switched which type participates, but because both types are equally likely, and both types have the same participation cost, this is inconsequential. Moreover, for the type that does not participate, we have fixed that this type does not participate at future nodes either. This potentially reduces expected participation cost. If both types participate then we have clearly either left participation cost unchanged, or reduced them. Moreover, if both types did not participate at the node that we were considering then we have left participation cost unchanged.

Our next step is to modify the candidate chosen after each history to be the candidate who maximizes expected welfare conditional on the information revealed by the voters' choices. This obviously requires that the candidate is chosen who is preferred by the majority of players who have made a choice, where we count non-participation as an expression of a preference for x, and participation as an expression of a preference for y.

The final step is to remove all decision nodes that don't affect the final decision, starting at the end of the game tree and moving iteratively to the beginning. This step leaves the collective decision unchanged, and either leaves expected participation cost unchanged or reduces them. After this step we have obtained a mechanism with participation decisions, and a strategy vector for the extensive game with incomplete information corresponding to this mechanism, which have all the properties described in Proposition 1.

Part 2. Let $(\mathcal{M}^*, \sigma^*)$ be a solution to the relaxed problem that has the properties listed in Proposition 1. We now show that σ^* is a sequential equilibrium of the incomplete information game corresponding to $\mathcal{M}^*$. We first note that all information sets in the extensive game with incomplete information corresponding to $\mathcal{M}^*$ are reached with positive probability. Therefore, beliefs are given by Bayesian updating. The strategies form a sequential equilibrium if and only if they are sequentially rational given these beliefs.

To prove sequential rationality we consider, in an intermediate step, a modification $(\tilde{\mathcal{M}}, \tilde{\sigma})$ of the solution $(\mathcal{M}^*, \sigma^*)$ of the relaxed problem. The modification differs from $(\mathcal{M}^*, \sigma^*)$ in that at every node the action preferred by voters other than the voter who chooses at this node is available for free, while the action that others do not prefer requires costly participation.

To construct $\tilde{\mathcal{M}}$, we first consider any non-terminal history h of $\mathcal{M}^*$. Denote by i the voter who moves at h, and, for $z \in \{A, B\}$, denote by $w_{-i}(z, h)$ the expected welfare of voters other than i, conditional on reaching h, if voter i "votes for z." By "voting for z" we mean "not participating" if z is the default candidate, and "participating" if z is not the default candidate. Observe that the expected values $w_{-i}(z, h)$ do not depend on player i's type because different players' types are independent. Let $z^*(h)$ be the candidate for whom $w_{-i}(z, h)$ is largest, and let $z_*(h)$ be the other candidate.[3] If $z^*(h)$ is the default candidate, we don't change $\mathcal{M}^*$. But if $z^*(h)$ is not the default candidate, we switch labels: we label the action that voter i's type θ_i takes as "0," i.e. non-participation, if $\theta_i = z^*(h)$, and as "1," i.e. "participation," if $\theta_i = z_*(h)$. Doing this for any non-terminal history h of $\mathcal{M}^*$, we obtain the new mechanism $\tilde{\mathcal{M}}$. The strategy vector $\tilde{\sigma}$ is defined by assuming that at any non-terminal history h, denoting by i the player moving after that history, voter i chooses "0," i.e. non-participation, if $\theta_i = z^*(h)$, and "1," i.e. "participation," if $\theta_i = z_*(h)$. Observe that $(\tilde{\mathcal{M}}, \tilde{\sigma})$ is another solution to the planner's relaxed problem. This is because $(\tilde{\mathcal{M}}, \tilde{\sigma})$ and $(\mathcal{M}^*, \sigma^*)$ lead to the same outcomes, the same expected aggregate participation costs, and therefore the same expected welfare.[4]

Because $(\tilde{\mathcal{M}}, \tilde{\sigma})$ solves the planner's relaxed problem, conditional on reaching any node h, the conditional expected welfare if players choose according to $\tilde{\sigma}$ is at least as large as the conditional expected welfare if both types of i abstain at h, thus expressing a preference for the candidate $z^*(h)$ preferred by the others, and all later players follow $\tilde{\sigma}$. Letting $\pi(0, h)$ be the probability that candidate $z^*(h)$ is chosen if i abstains, and $\pi(1, h)$

[3] It does not matter which candidate we pick if $w_{-i}(z, h)$ is the same for both candidates.

[4] Note that $(\tilde{\mathcal{M}}, \tilde{\sigma})$ need not have the properties listed in Proposition 1 because different voters who have the same candidate preference need not make the same participation decision. We shall resolve this issue later in the proof when we infer properties of $(\mathcal{M}^*, \sigma^*)$ from properties of $(\tilde{\mathcal{M}}, \tilde{\sigma})$

be the probability that $z^*(h)$ is chosen if i participates,[5] this implies:

$$\Big(w_{-i}(z^*(h), h) + \pi(0,h)\Big)\Big/2 + \Big(w_{-i}(z_*(h), h) + 1 - \pi(1,h) - c\Big)\Big/2 \geq \Big(w_{-i}(z^*(h), h) + \pi(0,h)\Big)\Big/2 + \Big(w_{-i}(z^*(h), h) + 1 - \pi(0,h)\Big)\Big/2. \quad (1)$$

Simplifying, we obtain:

$$w_{-i}(z_*(h), h) + \pi(0,h) - \pi(1,h) \geq w_{-i}(z^*(h), h) + c. \quad (2)$$

By the definition of $z^*(h)$ we have $w_{-i}(z_*(h), h) \leq w_{-i}(z^*(h), h)$, and therefore the above inequality implies:

$$\pi(0,h) - \pi(1,h) \geq c. \quad (3)$$

Thus, agent i's participation reduces the probability that $z^*(h)$ is chosen by at least c. This shows that participation is sequentially rational in $\tilde{\mathcal{M}}$ if player i is of type $\theta_i = z_*(h)$. That non-participation is sequentially rational in $\tilde{\mathcal{M}}$ if player i is of type $\theta_i = z^*(h)$ is obvious because non-participation increases at no cost the probability of $z^*(h)$ by at least c. This concludes the proof of sequential rationality of $\tilde{\sigma}$ in $\tilde{\mathcal{M}}$.

Intuitively, sequential rationality holds at every non-terminal history h because, when it is optimal for the mechanism designer to ask i's type $\theta_i = z_*(h)$ to participate and cause a negative externality to others, it is certainly optimal for player i, too, to participate, because player i ignores the negative externality for others in her own optimization whereas the mechanism designer includes it.

Now we return to the original solution $(\mathcal{M}^*, \sigma^*)$ of the planner's relaxed maximization problem. Note that the action in $\mathcal{M}^*$ to which we referred earlier as "voting for $z^*(h)$," although we might have relabeled it when moving from $\mathcal{M}^*$ to $\tilde{\mathcal{M}}$, increases the probability of $z^*(h)$ by the same amount in $\mathcal{M}^*$ as in $\tilde{\mathcal{M}}$, that is by $\pi(0,h) - \pi(1,h)$. This change in probability is not affected by the labeling. Similarly, the action in $\mathcal{M}^*$ to which we referred earlier as "voting for $z_*(h)$" reduces the probability of $z^*(h)$, and thus increases the probability of $z_*(h)$, by the same amount in $\mathcal{M}^*$ as in $\tilde{\mathcal{M}}$, that is by $\pi(0,h) - \pi(1,h)$. Thus, whichever type of player i is asked in $\mathcal{M}^*$ to vote for her preferred alternative by participating increases the probability of her favorite alternative by $\pi(0,h) - \pi(1,h)$ and thus has, by inequality (3), an incentive to participate. The other type obviously has no incentive to participate. This completes the proof of the sequential rationality of σ^*. □

[5] These probabilities do not depend on player i's type because different players' types are independent.

Our next result describes in more detail at least one of the solutions to the mechanism designer's problem. Consider any mechanism of the type described in Proposition 1. If $h = (a^1, \ldots, a^K) \in H$, we define:

$$r(h) = \#\{i \in I : P(a^1, \ldots, a^k) \neq i \text{ for all } k = 1, 2, \ldots, K\}, \tag{4}$$

and:

$$m(h) = \#\{k \leq K : a^k = 0\} - \#\{k \leq K : a^k = 1\}. \tag{5}$$

Thus, $r(h)$ is the number of remaining voters, that is, voters who have not yet been offered an opportunity to participate, and $m(h)$ is the current, positive or negative, margin of the default candidate.

Proposition 2. *Consider a given $c > 0$. There is a weakly increasing function $\mu : \mathbb{N}_0 \to \mathbb{N}_0$ such that for every $n \in \mathbb{N}$ there is a solution to the planner's problem with n voters that has the properties listed in Proposition 1, and in which the planner terminates voting, i.e. $h \in Z$, if and only if $|m(h)| \geq \mu(r(h))$. The function μ is independent of the number of voters, n.*

Proposition 2 shows that for given cost of voting c the optimal stopping decision depends on the number of remaining voters, and on the current margin for the default candidate, but not on the total number of voters. The larger the number of remaining voters is, the larger must the current margin be if the planner is to stop voting. Voting is thus terminated only if a supermajority requirement is met. As voting progresses, the required supermajority declines. An optimal stopping rule is illustrated in Figure 1. When there are 10 remaining voters, voting is terminated only if the current majority for one candidate is at least 3, whereas if there are 4 remaining voters, voting is terminated if the current majority for one candidate is at least 2. Note that in each case it is possible that the remaining voters overturn the existing majority if voting were continued.

Proof. Throughout the proof we assume that c is given and fixed. We examine the planner's first best problem, assuming that the planner proceeds as described in the first part of the proof of Proposition 1. Voters are ordered in some arbitrary way. Starting with the initial voter, the planner has to choose whether to consult one further voter, or to stop the voting process. Voters are assumed to indicate their preference truthfully, either through participation or through abstention. Once the planner stops, he picks the candidate favored by the majority, breaking ties arbitrarily.

Step 1: We consider a given and fixed number $n \in \mathbb{N}$ of voters. We take the vector (n, r, m) as the planner's state variable, where $r \in \mathbb{N}_0$ is the number of remaining voters who have not yet been consulted, and $m \in \mathbb{Z}$

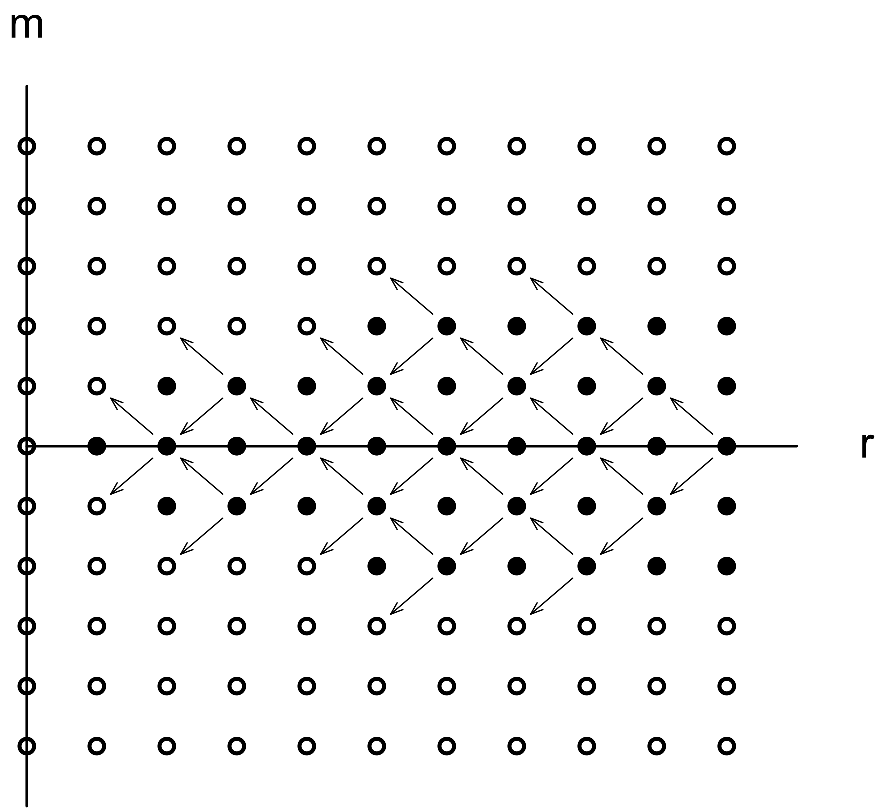

Figure 1: The optimal stopping rule for $c = 0.15$. The number of remaining voters r is on the horizontal axis. The current margin m is on the vertical axis. The axes intersect in the point (0,0). Empty circles represent states where the planner terminates voting, while solid circles represent states where the planner consults a further voter. As voting progresses, we move from the right to the left in the coordinate system. From each full circle there are two equally likely transitions: to the circle in the column to the left that is one row higher, and to the circle in the column to the left that is one row lower. Which transition occurs depends on how the consulted voter votes. We indicate the possible transitions for the case that there are 10 voters, in which case voting starts in the point $(r, m) = (10, 0)$.

is the difference between votes cast for the default candidate and votes cast for the alternative candidate among the voters consulted so far. Only states in which $|m| + r \leq n$ and $n - r + m$ is even are feasible.[6]

We define $D(n, r, m)$ to be the optimal decision in state (n, r, m), where $D(n, r, m) = 1$ if the mechanism designer continues the voting process in state (n, r, m), and $D(n, r, m) = 0$ otherwise. Obviously, D is only well defined if $r \geq 1$. We define $W(n, r, m)$ to be expected welfare if the planner makes optimal decisions starting in state (n, r, m), and we set:

$$V(n, r, m) = W(n, r, m) - (n - r - m)c/2 - n/2. \quad (6)$$

V differs from W only in that we have subtracted the participation cost that are sunk in stage (n, r, m), that is, $(n - r - m)c/2$ (see footnote 6), and also the expected welfare that the mechanism designer would achieve by making a random decision without consulting any voter: $n/2$. We refer to $V(n, r, m)$ as the "value of the planner's problem in state (n, r, m)." In this step we develop recursive formulas for D and V where the recursion is over r, the number of remaining voters.

We first consider the case $r = 0$, that is, all voters have been consulted. If $m > 0$, the planner chooses the default candidate x, and thus:

$$V(n, 0, m) = \#\{i \in N : \theta_i = x\} - n/2 = (n + m/2) - n/2 = m/2. \quad (7)$$

Here, we have used the fact that the number of voters favoring the default is $(n + m)/2$ (see footnote 6). The case $m < 0$ is symmetric, and if $m = 0$ welfare obviously equals $n/2$. Thus, we conclude for all n and m:

$$V(n, 0, m) = |m|/2. \quad (8)$$

Now consider states where $r > 0$. If the planner terminates voting V is again $|m|/2$. To see this, suppose that $m > 0$ and the planner decides on the default candidate. Then his payoff equals:

$$(n - r + m)/2 + r/2 - n/2 = m/2, \quad (9)$$

where we have used the fact that among the voters already consulted the number of voters favoring the default is $(n - r + m)/2$ (see footnote 6). The case $m \leq 0$ is analogous.

When the planner does not terminate the process but consults an additional voter he incurs expected cost $c/2$, the number of remaining voters

[6] Letting n_x denote the voters consulted so far who voted for the default, and n_y denote the voters consulted so far who voted against the default, then we have to have: $n_x + n_y = n - r$, and $n_x - n_y = m$, so that $n_x = (n - r + m)/2$ and $n_y = (n - r - m)/2$, which are integers only if $n - r + m$ (equivalently: $n - r - m$) is even.

decreases to $r-1$, and the margin is equally likely to increase or decrease by 1. The planner's optimal decision rule is thus:

End the voting process, that is, $D(n,r,m)=0$*, if*

$$(V(n,r-1,m+1)+V(n,r-1,m-1)-c)/2 \leq |m|/2 \tag{10}$$

and continue, that is, $D(n,r,m)=1$*, if*

$$(V(n,r-1,m+1)+V(n,r-1,m-1)-c)/2 > |m|/2. \tag{11}$$

where we have assumed that the planner terminates the voting process if indifferent. For the expected value, we obtain the recursive formula:

$$\begin{aligned} V(n,r,m) &= \max\{(V(n,r-1,m+1)+V(n,r-1,m-1)-c)/2, \\ & \quad |m|/2\}. \end{aligned} \tag{12}$$

Note that V does not depend on n. One can prove this by induction over r. Our calculations show that it is true for $r=0$, and for $r>0$ the above recursive formula shows that V does not depend on n for $r>0$ if it does not depend on n for $r-1$. Because V does not depend on n, D does not depend on n either.

Step 2: We now re-define the domains of the functions D and V as the set of all pairs $(r,m)\in\mathbb{N}_0\times\mathbb{Z}$ by setting $D(r,m)=1$ if $D(n,r,m)=1$ for some n,and $D(r,m)=0$ if $D(n,r,m)=0$ for some n. Step 1 established that it does not matter which n we choose. We set $D(0,m)=0$ for all m. Similarly, if $v\in\mathbb{R}$, then $V(r,m)=v$ if and only if $V(n,r,m)=v$ for some n.

The functions V and D are symmetric: $V(r,m)=V(r,-m)$ and $D(r,m)=D(r,-m)$. This is obvious from the symmetry of our model. We also note that $V(r,m)=|m|/2$ and $D(r,m)=0$ whenever $r\leq|m|$, because $r\leq|m|$ implies that the remaining voters cannot overturn the existing majority, and therefore consulting them will not change the outcome.

Step 3: We prove that $D(r,m)$ is weakly increasing in r:

$$D(r,m)=1 \Rightarrow D(r+1,m)=1 \text{ for all } r\in\mathbb{N}_0 \text{ and all } m\in\mathbb{Z}. \tag{13}$$

The value of stopping[7] is the same in state (r,m) and in state $(r+1,m)$: $|m|/2$. In state $(r,m+1)$ the mechanism designer could also adopt the

[7]We define the "value" of a decision in a state as in Step 1: the expected welfare if the mechanism designer continued with optimal decisions, minus what expected welfare would be if the planner had made a random decision without consulting any voters, and ignoring participation cost incurred in the past.

strategy that is optimal in state (r, m) until the last voter is reached, and then always stop, that is, ignore the last voter. This strategy yields the same value in state $(r+1, m)$ as the optimal strategy does in state (r, m). Denote this value $\bar{v}$. Because this strategy is optimal in state (r, m) we have: $\bar{v} \geq |m|/2$. The optimal strategy among all strategies that continue in the voting problem with $(r+1, m)$ voters must have a value that is at least as large as $\bar{v}$. Therefore, $\bar{v} \geq |m|/2$, and it is optimal to continue in state $(r+1, m)$.[8] Hence $D(r+1, m) = 1$.

Step 4: We prove that D is weakly decreasing in m if $m \geq 0$, and weakly decreasing in $|m|$ when $m \leq 0$. Because of the symmetry of D, it is sufficient to consider the case that $m \geq 0$. Thus, we show:

$$D(r, m) = 0 \Rightarrow D(r, m+1) = 0 \text{ for all } r, m \in \mathbb{N}_0. \tag{14}$$

The value of stopping is $m/2$ in state (r, m), and it is $(m+1)/2$ in state $(r, m+1)$. Thus, stopping in state $(r, m+1)$ is optimal if the difference in expected value between the optimal strategy among all strategies that continue voting in state $(r, m+1)$ and the optimal strategy among all strategies that continue voting in state (r, m) is at most $1/2$. Suppose the decision maker used the optimal strategy that continues voting in state $(r, m+1)$ also in state (r, m). Then the two expected values are identical up to the expected utility of the one added voter in state $(r, m+1)$, whose preference is known to be for the default candidate. Therefore, the expected value difference is the probability of choosing the default candidate multiplied by $1/2$, and thus cannot be more than $1/2$. If instead the decision maker uses in state (r, m) the optimal strategy that continues in that state, the difference between the two expected values cannot increase, and therefore must be no more than $1/2$, as we had to show.

Step 5: We now define for every $r \in \mathbb{N}_0$:

$$\mu(r) = \min\{|m| \in \mathbb{N}_0 : D(r, m) = 0\}. \tag{15}$$

Because $D(r, m) = 0$ whenever $r \leq |m|$, a fact that we mentioned in Step 2 above, the set over which we take the minimum is non-empty, and therefore the minimum exists. The monotonicity of D in m from Step 4 implies that for every r $D(r, m) = 0$ whenever $m \geq \mu(r)$ and $D(r, m) = 1$ otherwise. We can complete the proof by showing that $\mu(r)$ is weakly increasing in r. This follows from the fact that the set of which $\mu(r)$ is the minimum, decreases in terms of set inclusion as r increases because $D(r, m)$ is weakly increasing in r, as we showed in Step 3. □

[8] Recall that we assume that, if indifferent, the planner continues voting.

4 Discussion

4.1 Sub-optimality of Static Mechanisms

Propositions 1 and 2 describe features that *at least one* optimal pair of mechanism and its equilibrium has. The results do not assert that *all* optimal pairs of mechanisms and their corresponding equilibria have these features. We argue in this subsection that when n is at least 3 and the voting cost c is sufficiently small *no* static mechanism will be optimal. Here, we call a mechanism with participation decisions *static* if all moves, including participation decisions, are made simultaneously.

We begin with the observation that a static mechanism with participation decisions and its equilibrium cannot be optimal if in equilibrium there is a positive probability that at least $n/2+1$ (if n is even) or $n/2+1.5$ (if n is odd) voters participate. This is an implication of the fact that the procedure described in part 1 of the proof of Proposition 1, when applied to such a static mechanism and its equilibrium, strictly increases expected welfare. It strictly increases expected welfare because it allows voting to be terminated when sufficiently many agents have participated to establish a majority which cannot be overturned. This economizes on voting cost.

Now suppose the mechanism designer only considered static mechanisms. When c is sufficiently close to zero the optimal static mechanism and strategy have all voters participate (if they reject the default) if n is odd, and all except one voter participate if n is even.[9] Therefore, with positive probability, as least as many voters as given by the above thresholds participate. This implies that when n is at least three, and the voting cost c are sufficiently small, no optimal mechanism is static.

Quantitatively we can use the characterization in Proposition 2 to calculate welfare in the planner's solution and in the optimal static mechanisms. Neither problem gives rise to a closed-form value function but both are easy to evaluate numerically. For $n = 10$ voters and voting costs of $c = 0.15$, we find that the planner's solution is the mechanism depicted in Figure 1 while the optimal static mechanism asks seven voters to participate. Expected welfare equals 5.72 in the planner's solution, 5.57 in the optimal static mechanism, and 5.00 in the mechanism that does not consult any voter. Taking the latter mechanism as a benchmark, the planner's solution increases welfare by 0.72, that is by 26% more than the optimal static mechanism.

[9] We omit the formal proof of this assertion.

4.2 The Revelation Principle

We have not used the revelation principle in our analysis, although this principle plays a central role in many studies of mechanism design problems. The revelation principle, applied to our model, says that for any sequential equilibrium of a mechanism with participation decisions one can construct a direct mechanism in which voters reveal their type to the mechanism designer who then implements the outcome that would have resulted in the original mechanism if the voters had played their equilibrium strategies. Moreover, in this direct mechanism, all voters reporting their types truthfully is a Bayesian Nash equilibrium, that is, the direct mechanism is incentive compatible. In our context, the outcome that the mechanism designer implements once types have been revealed consists of a selection of one of the two candidates, and also participation decisions for each voter. Reporting one's type in the direct mechanism is costless. A voter incurs participation cost only when the direct mechanism's outcome specifies that the voter participates. With this notion of a direct mechanism, the standard proof of the revelation principle can be used to prove the revelation principle also for our model.

Note that the direct mechanisms described in the previous paragraph are *not* mechanisms with participation decisions in the sense of Definition 2. According to Definition 2 an agent has to incur a participation cost before taking any other action. In particular, in Definition 2, reporting one's type is not possible without incurring a participation cost. When solving the mechanism designer's problem, we cannot therefore proceed in the usual way and maximize expected welfare among all incentive compatible direct mechanisms, and then conclude that the optimal direct mechanism is also a solution to the mechanism designer's original problem. Direct mechanisms are not in the feasible set of the mechanism designer's original problem.

One might try instead to first maximize expected welfare among all incentive compatible direct mechanisms that are equivalent[10] to some mechanism with participation decisions and a corresponding sequential equilibrium, re-constructing the underlying mechanism with participation decisions and its sequential equilibrium only in a second step. However, there is no obvious characterization of the incentive compatible direct mechanisms that the designer can choose from. Clearly, not all incentive compatible direct mechanisms can be allowed. For example, any direct mechanism

[10]We call a direct mechanism "equivalent" to a mechanism with participation decisions and a corresponding sequential equilibrium if in this sequential equilibrium for every vector of voters' types, the same candidate is chosen, and the same voters participate, as in the direct mechanism if voters reveal their types truthfully in that mechanism.

where no voter ever participates, yet the collective decision depends on voters' types, must be ruled out. Which restrictions exactly describe the admissible set of incentive compatible direct mechanisms seems difficult to determine.

A possibly more useful version of the revelation principle in our model considers mechanisms with participation decisions where any voter who has decided to participate reports her type subsequently, and then does not get to move again.[11] One might call such mechanisms *direct mechanisms with participation decisions*. The standard logic of the revelation principle implies that for every mechanism with participation decisions, and corresponding sequential equilibrium, there is an equivalent direct mechanism with participation decisions and a sequential equilibrium such that any voter who is asked to reveal his type reveals that type truthfully. This result is true in more general models than ours, and in such more general models may simplify the search for optimal mechanisms. In our model there is no need for participating voters to reveal their types because the types can be inferred from the participation decision. Therefore, our analysis in this paper seems best conducted without explicit appeal to this version of the revelation principle.

4.3 Correlated Types

A natural extension of our model is a model in which agents' types are correlated. For concreteness suppose there are two equally likely possible states of the world, a and b, where conditional on state a individual voters' preferences are i.i.d. with the probability that a voter prefers candidate A being $p > 0.5$ and conditional on state b individual voters' preferences are i.i.d. with the probability that a voter prefers candidate B being p. Voters observe the state and their own preferences. This structure is common knowledge.

If the mechanism designer does not observe the state of the world before designing the mechanism, then, unlike in our setting, it may be that no first best solution is incentive compatible. To see this suppose for simplicity that preferences are perfectly correlated, i.e. $p = 1$. Then any first best mechanism will invite at most one voter to participate. Non-participation will be interpreted as a preference for one candidate, and participation will

[11]This revelation principle is related to the revelation principle invoked in Myerson (1986) where players in a multi-stage mechanism report in each period their privately observed information to the mechanism designer. However, participation decisions play no role in Myerson's set-up, and the rationale for reporting private information not initially but later is in Myerson's paper not that a player did not participate earlier, but that the information was not available to the player at an earlier stage.

be interpreted as a preference for the other candidate. This one voter's preference will determine the collective choice. However, this voter's individual incentives to participate do not reflect the positive externality that he exerts on other voters, and he may individually prefer not to participate even if the first best mechanism requires him to participate.[12] [13]

If the mechanism designer does observe the state of the world before designing the mechanism, then the first best can be analyzed as in the proof of Proposition 1 except that the choice of the default candidate is no longer arbitrary, but the candidate who is more likely to be preferred by voters must be made the default candidate. A difficulty is that the argument in the proof of Proposition 1 for incentive compatibility of the first best no longer applies. This is because our proof of the sequential rationality of the first best $(\tilde{\mathcal{M}}, \tilde{\sigma})$ is built on agent i's participation imposing a non-positive externality on other voters. In general, this need not be the case.

We illustrate this with a simple example. For this example we generalize our framework by allowing for known, voter-specific preference intensities u_i; that is voter i's cardinal utility when his favorite candidate is elected equals u_i rather than 1. Suppose the state is a. Consider a history in the first best solution where all voters but i have cast their votes and other voters slightly favor candidate B, that is sum of the utilities from B for voters in favor of B minus the sum of the utilities from A for voters who voted in favor of A is some $\varepsilon > 0$. Assume that the state is such that it is more likely that a voter favors A. If ε is sufficiently small, voter i's expected preference for candidate A will imply that the planner would optimally choose candidate A if he were to terminate the process without learning i's type. If additionally $u_i \in (c - \varepsilon, c)$ the planner prefers i's type B to vote because the expected benefit of participation $(1-p)(u_i + \varepsilon)$ exceeds expected participation cost $(1-p)c$. However, this social calculation incorporates the positive externality of i's participation on others, while in the private calculation of i's type B participation costs c exceed the benefit u_i. Thus, the participation of i's type B to participate at this node is socially optimal but not sequentially rational.[14]

[12]If preferences are specified as in this paper, the single voter consulted will be unwilling to participate only if $c > 1$. In this case, with independent preferences, the planner would not consult any voter. We have not constructed an example with correlated preferences where $c < 1$ implies that first best is not incentive compatible.

[13]Note the analogy with the analysis of Gershkov and Szentes (2009) that we cited in the Introduction. In their paper, too, first best mechanisms are not incentive compatible because information acquisition by one voter exerts a positive externality on other voters, and therefore individual incentives for information acquisition need not be as large as the social benefit from information acquisition.

[14]This example is unfortunately not tight in the sense that it violates not just the symmetry assumption, but additionally assumes that voters' cardinal utility u_i from

Despite the failure of Propositions 1 and 2 for general asymmetric types $p \in [0, 1]$, these results do not depend on the knife-edge assumption that p is exactly equal to 0.5. Consider the case in which p is commonly known by the planner and by the voters, and is slightly different from 0.5. Part 1 of the proof of Proposition 1 remains valid if we choose the candidate who is more likely to be preferred as the default candidate. Part 2 remains valid for generic c because, for generic c, inequality (1) is strict and all terms in this inequality depend continuously on p. In addition, one may verify that for generic c the inequalities that define the optimal stopping rule in Proposition 2 are strict, and that all terms in these inequalities depend continuously on p, and therefore the optimal stopping rule remains unchanged. Note one major difference between the result for p slightly different from 0.5, and the result for p exactly equal to 0.5. It is that when p is different from 0.5, the default candidate is uniquely determined whereas with p exactly equal to 0.5, the choice of default candidate is arbitrary, and can even be changed as voting progresses.

4.4 Privately Observed Voting Costs

Another modification of our model assumes that not only voters' preferences but also their voting cost c are privately observed random variables. The reason why the analysis of Section 3 does not apply to this modified model is as follows. When analyzing the incentive issues raised in the second part of the proof of Proposition 1 we may find that voters with large voting costs who oppose the default have an incentive to participate even if the mechanism designer prefers them not to participate. For voters who oppose the default the argument in the second part of our proof shows that their individual incentives to participate are at least as large as c whenever the designer wants them to participate. But the argument allows the possibility that the individual incentives are strictly larger than c. Thus, when the planner's first best mechanism requests a voter to abstain when her cost realization is high, the voter may find it privately beneficial to mimic a low cost type and participate anyway. If c is privately observed

electing their favorite candidate is voter specific and may fall short of participation costs c for some voters. It would be desirable to find an example where Proposition 1 fails, that violates only the symmetry assumption. Such a tight example would construct a first best where at some history h type B's participation increases the probability of electing B by only $p(B, h) - p(A, h) < c$, but the planner would like type B to participate anyway because of the positive externality on others. Unfortunately we have not been able to construct such an example. We view this shortcoming of our example as acceptable, because it still highlights where the proof of Proposition 1 fails for asymmetric types. Also, Proposition 1 extends to the case of known voter specific preference intensities, so its failure must be due to the assumption of asymmetrically distributed types.

the planner cannot prevent her from doing so. Therefore, this modification of our model, too, requires a separate analysis.

References

Dirk Bergemann and Juuso Välimäki (2002), "Information Acquisition and Efficient Mechanism Design," *Econometrica* 70, 1007-1033.

Dirk Bergemann and Juuso Välimäki (2006), "Information in Mechanism Design," in: Richard Blundell, Whitney Newey and Torsten Persson (eds.), *Proceedings of the 9th World Congress of the Econometric Society*, Cambridge: Cambridge University Press, 186-221.

Tilman Börgers (2004), "Costly Voting," *American Economic Review* 94, 57-66.

Gorkem Celik and Okan Yilankaya (2009), "Optimal Auctions With Simultaneous and Costly Participation," *The B. E. Journal of Theoretical Economics (Advances)* 9, Article 24.

Ronald Fadel and Ilya Segal (2009), "The Communication Costs of Selfishness," *Journal of Economic Theory* 144, 1895-1920.

Alex Gershkov and Balázs Szentes (2009), "Optimal Voting Schemes With Costly Information Acquisition," *Journal of Economic Theory* 144, 36-68.

Sayantan Ghosal and Ben Lockwood (2009), "Costly Voting When Both Information And Preferences Differ: Is Turnout too High or too Low?," *Social Choice and Welfare* 33, 25-50.

Stefan Krasa and Mattias Polborn (2009), "Is Mandatory Voting Better Than Voluntary Voting?," *Games and Economic Behavior* 66, 275-291.

Roger B. Myerson (1986), "Multistage Games With Communication," *Econometrica* 56, 323-358.

Martin Osborne and Ariel Rubinstein (1994), *A Course in Game Theory*, Cambridge (MA): The MIT Press.

Roman Pancs (2011), "Sequential Negotiations with Costly Information Acquisition," University of Rochester.

Mark Stegeman (1996), "Participation Costs and Efficient Auctions," *Journal of Economic Theory* 71, 228-259.