Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 159 (2015) 1056–1073

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

An optimal voting procedure when voting is costly ☆

Katalin Bognar [a], Tilman Börgers [b], Moritz Meyer-ter-Vehn [c,*]

[a] *Precision Health Economics, United States*
[b] *Department of Economics, University of Michigan, United States*
[c] *Department of Economics, UCLA, United States*

Received 30 September 2013; final version received 4 November 2014; accepted 10 November 2014

Available online 18 November 2014

Abstract

We study optimal dynamic voting procedures when voting is costly. For a highly stylized specification of our model with private values, two alternatives, and binary, equally likely types we show the optimality of a voting procedure that combines two main elements: (i) there is an arbitrarily chosen default decision and abstention 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 procedure by arguing that it is first best, that is, it maximizes welfare when equilibrium constraints are ignored, and by showing that individual incentives and social welfare are sufficiently aligned to make a first best procedure incentive compatible. We also provide counterexamples where no first best procedure is incentive compatible when voters' binary types are not equally likely.

JEL classification: D72; D82

Keywords: Dynamic mechanism design; Costly voting; Participation costs

☆ Thanks to Arthur Lupia, Romans Pancs, Alessandro Pavan, Marek Pycia, Ilya Segal, Ricky Vohra, the editors and four anonymous referees for helpful comments and to Yan-Min Choo and Siwei Kwok 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.
* Corresponding author.
E-mail addresses: katalin.bognar@gmail.com (K. Bognar), tborgers@umich.edu (T. Börgers), mtv@econ.ucla.edu (M. Meyer-ter-Vehn).

http://dx.doi.org/10.1016/j.jet.2014.11.008
0022-0531/

1. Introduction

The mechanism design literature usually imposes no restrictions on the outcomes that a planner can assign to a terminal history of a game form. This assumes that agents have no intrinsic preferences over the actions that they take in a mechanism. An exception is the literature on endogenous participation and mechanism design. In this literature positive participation costs arise at every terminal history along which an agent takes an action other than the action "abstention." Papers in this line of research include Stegeman (1996) and Celik and Yilankaya (2009) who consider welfare/profit maximizing auction design with endogenous participation.

The papers listed above only allow the planner to choose mechanisms in which agents make their participation decisions simultaneously. But if the objective of the planner includes participation costs, then it may be advantageous to allow for sequential participation. This may save participation costs because some agents may recognize that it is not worthwhile for them to participate, given their knowledge of previous agents' actions. The planner thus needs to optimize over a large class of extensive game forms and their sequential equilibria.

Here we solve this problem for a particular application, voting with costly participation, but we hope that our formulation will be useful also for other applications, such as auction design. Participation costs in voting have been studied before,[1] but the problem of the optimal design of voting rules in the presence of participation costs has not been addressed. We propose a formulation of this problem that allows for sequential mechanisms, i.e. we allow some voters to condition their participation decision on previous voters' choices. In practice, it may be that a small committee votes first, and a larger population is only asked to vote if no sufficiently large majority can be established in the smaller committee.[2] This resembles referendum procedures in California or Switzerland where a referendum must first attract a quorum of signatures before being put in front of the general electorate. In the first stage the default is that the referendum fails; this default needs to be overcome by sufficient participation.

In the model, a set of voters has to choose one of two candidates A, B. Each voter has a strict preference over candidates, and knows his or her own preference, but not other voters' preferences. Participation in the voting procedure has known positive costs that may differ across voters. The game form has to satisfy some restrictions that make the interpretation of certain branches as "abstention", or "participation" meaningful.

Our main result is that if voters are equally likely to prefer A or B there is a simple optimal voting procedure: voters vote sequentially in the order of increasing participation costs. Participation is voluntary and abstention is interpreted as a vote for a default candidate. A candidate is elected if a sufficiently large supermajority of voters who have expressed their preference — either by voting or by not voting — favors the candidate. This procedure exploits the assumption that abstention is free and can be interpreted as if it was a vote. Arguably, abstention 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, this establishes that abstention is interpreted as a vote. It may be that voting terminates at a stage when remaining voters could still overturn the majority. Continued voting until the will of majority is established without any possibility of mistake is costly and requires many voters to be polled, particularly when the alternative to the default is preferred.

[1] See, e.g. Börgers (2004), Ghosal and Lockwood (2009), Krasa and Polborn (2009).

[2] The economics department at the University of Michigan operates such a rule for decisions about certain types of offers of faculty positions.

In practice, when not all voters participate, the number of voters who is consulted is 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 relies on our assumption that observing the progress of voting is free. Only actual participation carries costs. Observation may only require watching a computer screen, whereas participation requires actual presence in a room. As an extreme case, one may have in mind members of the US congress who watch roll-call votes on a screen in their offices and enter the chamber to cast their vote only if the current margin of votes makes it optimal.

Our efficiency result relies on the assumption that voters prefer candidates A and B with equal probability.[3] In this case the choice of the default candidate is arbitrary which is important when dealing with the externalities that one voter's participation imposes on other voters. In Section 5 we provide two counterexamples that show how first best procedures are not incentive compatible when A, say, is more likely preferred than B and must thus be the default candidate in a first best procedure.

Our work contributes to the literature on dynamic mechanism design. Important papers in this literature, e.g. Pavan, et al. (2014), consider problems where the underlying problem itself is dynamic in that additional actions or information become available over time. In our work the basic problem is static, that is, all actions and all private information are available at the outset, yet the optimal mechanism is dynamic.

The problem of maximizing welfare over the set of all dynamic voting processes and their sequential equilibria has also been considered by Gershkov and Szentes (2009).[4] However, whereas we consider a private value setting, they consider a setting with common values. While in our setting participation is costly, in their setting information acquisition is costly.[5] As is the case in our paper, they cannot restrict attention to static revelation games but instead need to consider all extensive game forms.

It is interesting to compare Gershkov and Szentes' analytical approach to ours. Both papers first identify a canonical class of mechanisms to which one can restrict attention. We then show that the *first best mechanisms* in this class, i.e. those that maximize ex ante expected welfare when incentive constraints are relaxed, are incentive compatible when voters are equally likely to prefer either candidate. The key idea is that, by interpreting abstention as a vote for some candidate, the planner can effectively choose which action is costly and which action is free. In particular, he can always make the action with positive expected externalities freely available to agents. Therefore, if the first best mechanism asks an agent to take a costly action, this request is not based on some positive externalities of this action, but on its benefits to this particular agent. Therefore, the request is incentive compatible.

In contrast, first best mechanisms are not incentive compatible in Gershkov and Szentes' model. In their set-up, the costs of actions are intrinsic to those actions. If information acquisition has positive externalities, then the planner has no way of making information acquisition free. If a 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.

---

[3] However, we point out in footnote 10 that our result is not knife-edge and carries through when A and B are preferred with almost equal probability.

[4] Fadel and Segal (2009) also maximize over a set of extensive games.

[5] However this difference is not essential. Section 7.1 shows that our model is equivalent to a model with information acquisition.

The request may therefore not be incentive compatible. Similarly, when voters are more likely to prefer candidate A in our model, then a first best mechanism must declare A the default to economize on voting costs. But then it may be the case that participation and voting for B imposes a positive externality on others and incentive compatibility of the first best mechanism may fail. We illustrate this with the two counterexamples in Section 5 mentioned before.

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.[6] We view this maximal informativeness as an attractive feature of our mechanism because in practice it may be difficult to conceal information from voters.

Considering information acquisition costs rather than participation costs, Bergemann and Välimäki (2002) find for a general environment with independent private values and transferable utility that Vickrey–Clarke–Groves mechanisms provide socially optimal incentives for information acquisition. Our main result is related 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).

Some authors have considered welfare properties of particular voting schemes in the presence of voting costs, either in a private or in a common value setting, without examining the general mechanism design problem, and assuming that participation and voting choices are made at the same time. For example, Börgers (2004) showed in a model similar to ours the superiority of voluntary voting over mandatory voting 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, but in which the emphasis is not on participation costs but on information acquisition costs.

We present our framework in Section 2. In Section 3 we consider the planner's first best problem where equilibrium constraints are ignored. In Section 4 we show that certain first best solutions are incentive compatible when types are equally likely. Section 5 shows how the result of Section 4 fails when voters' types are not equally likely. Section 6 then discusses difficulties with solving for the second best mechanism when first best is not attainable. Section 7 establishes that our model is equivalent to a model with information acquisition; it also discusses model extensions that further illustrate the limits of our efficiency result.

2. Set-up

There are $n \in \mathbb{N}$ voters $i \in N \equiv \{1, 2, \ldots, n\}$. The voters have to pick one of two candidates $z \in \{A, B\}$. Each voter $i \in N$ has a type $\theta_i \in \{A, B\}$, indicating which candidate i prefers. Types are random, privately observed, and independent across voters. Each voter prefers candidate A with probability $p \geq 1/2$ and candidate B with probability $1 - p$. Voter i's Bernoulli utility

[6] However, our counterexamples in Section 5 are robust to withholding information.

is additive across the outcome of the voting process, 1 if i's preferred candidate wins and -1 otherwise, and a participation utility, $-c_i$ if i participated in the voting process and 0 otherwise. The distribution of types, what each voter observes, and the Bernoulli utility functions are common knowledge among voters and the planner. The planner seeks to maximize the sum of all voters' ex ante expected utilities.

Our setting is obviously very special. We study this setting to focus on methodological issues without distraction and because even in this simple setting we can solve the planner's problem only for some parameter values. Despite its simplicity, this setting yields results that are of some real world plausibility, as argued in the Introduction.

The planner chooses an extensive game form that the voters use to pick one of the two candidates, and a strategy profile in this game form. We first describe the set of extensive game forms that the planner can choose from. The following definition of *extensive game forms* is adapted 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 h with the following properties:
   (a) the empty sequence $\emptyset$ is an element of H;
   (b) if $(a^k)_{k=1,\dots,K} \in H$ and $L < K$ then $(a^k)_{k=1,\dots,L} \in H$.
   (Each h is a *history/node*. Each component of a history is an *action*, or a chance move. A history $(a^k)_{k=1,\dots,K} \in H$ 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 Z. The set of actions available after a nonterminal history $h \in H \setminus Z$ is $\mathcal{A}(h) \equiv \{a : (h; a) \in H\}$.)
3. A function P that assigns to each nonterminal history $h \in H \setminus Z$ an element of $N \cup \{C\}$. ($P(h)$ is the player who takes an action after history h. If $P(h) = C$ then h is a chance node.)
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 action 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 $D(h) \in [0, 1]$. ($D(h)$ is the probability that A is chosen after history h. Candidate B is chosen with 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.

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) are called *participation nodes* of player $P(h)$. Action 1 is interpreted as participation by voter $P(h)$, and action 0 as abstention by $P(h)$.)

2. For every terminal history $(a^k)_{k=1,\dots,K} \in Z$ 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. Thus, before doing anything else players have to choose to participate, although this "choice" needn't be voluntary.)

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. Voter i incurs the cost c_i if one of i's participation nodes is reached and i chooses to participate. We denote strategies of player i by σ_i. Stegeman (1996) we refer to a pair of a mechanism with participation decisions and a strategy profile $(\mathcal{M}, \sigma)$ as a *procedure*. A procedure $(\mathcal{M}, \sigma)$ is *incentive compatible* if σ is a sequential equilibrium in the incomplete information game associated with $\mathcal{M}$.

The Planner's problem. *Choose an incentive compatible procedure $(\mathcal{M}, \sigma)$ to maximize the sum of ex ante expected utilities of all voters.*

3. The first-best problem

We first relax the planner's constraints by ignoring voters' incentives.

The first best problem. *Choose any procedure $(\mathcal{M}, \sigma)$ to maximize the sum of ex ante expected utilities of all voters.*

We solve the first best problem in two steps. In Proposition 1, which holds for all $p \geq 1/2$, we exhibit basic properties of a first best procedure. In Proposition 2 we refine this characterization for the case $p = 1/2$.

Definition 3. A procedure $(\mathcal{M}, \sigma)$ is *canonical* if:

(i) For each $i \in N$ and $I_i \in \mathcal{I}_i$: $\#I_i = 1$. (Perfect information.)
(ii) There is no $h \in H \setminus Z$ such that $P(h) = C$. (No chance moves.)
(iii) There are $x, y \in \{A, B\}$ with $x \neq y$ such that whenever $P(h) = i$ and $\mathcal{A}(h) = \{0,1\}$, then $\sigma_{i,x}(h) = 0$ and $\sigma_{i,y}(h) = 1$. If $p > 1/2$ then $x = A$ and $y = B$. (Candidate x is the *default* candidate. At a participation node, voter i participates if and only if she opposes the default. If voters are more likely to prefer candidate A then A is the default candidate.)
(iv) For every $h = (a^1,\dots,a^K) \in Z$ and $i \in N$ there is at most one $\ell < K$ such that $P(a^1,\dots,a^\ell) = i$. Moreover then $A(a^1,\dots,a^\ell) = \{0,1\}$. (A voter makes at most one decision in any history. This decision is a participation decision. Participation is voluntary.)
(v) For every $h = (a^1,\dots,a^K) \in Z$ define $n(x) = \#\{k \mid a^k = 0\} + (n-K)p$ and $n(y) = \#\{k \mid a^k = 1\} + (n-K)(1-p)$: Then $D(h) = x$ if $n(x) > n(y)$, and $D(h) = y$ if $n(x) < n(y)$. ($n(z)$ is the expected number of voters preferring candidate z. To calculate $n(z)$, note that voters k with $a^k = 0$ prefer A, voters with $a^k = 1$ prefer B and the $n - K$ voters who

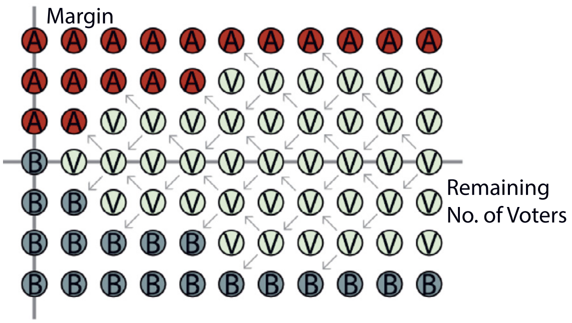

Fig. 1. The optimal stopping rule for $p = 1/2$ and voting costs $c_i = 0.3$ for all $i = 1, 2, \ldots, n$. The number of remaining voters r is on the horizontal axis. The current margin m is on the vertical axis. The axes intersect in the point $(0, 0)$. Circles labeled with A/B represent states where the planner terminates voting and selects the respective candidate, while circles labeled with V represent states where the planner consults a further voter.

have not been consulted prefer A with probability p and B with probability $1 - p$. The mechanism chooses the candidate who is preferred by the greater expected number of voters, given all available information.)

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

Proposition 1. *A solution $(\mathcal{M}, \sigma)$ to the first best problem exists. Moreover, there is at least one solution that is a canonical procedure.*

Proposition 1 is proved in Appendix A. Our second result describes one canonical first best procedure $(\mathcal{M}, \sigma)$ in more detail. For history $h = (a^1, \ldots, a^K)$, we define the number of remaining voters, i.e. voters who have not had an opportunity to participate:

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

and the current — positive or negative — margin of the default candidate:

$$m(h) = \#\big\{k \leq K : a^k = 0\big\} - \#\big\{k \leq K : a^k = 1\big\}. \tag{2}$$

Proposition 2. *Suppose $p = 1/2$ and fix an infinite decreasing sequence of participation costs $c_1 \geq c_2 \geq \cdots$. 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 first best procedure in the problem with n voters and costs $c_1, \cdots, c_n$ that has the properties listed in Proposition 1, in which the planner consults the voters in order of increasing voting costs c_i, and in which the planner terminates voting when $|m(h)| \geq \mu(r(h))$. The value $\mu(r)$ depends only on the costs of the remaining voters $c_1, \cdots, c_r$.*

Proposition 2 is proved in Appendix A. Proposition 2 shows that the optimal stopping decision depends on the number of remaining voters, their voting costs, and the current margin for the default candidate, but not on the total number of voters. The larger the number of remaining voters, the larger must the current margin be for the planner 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 Fig. 1. With five remaining voters, voting is terminated

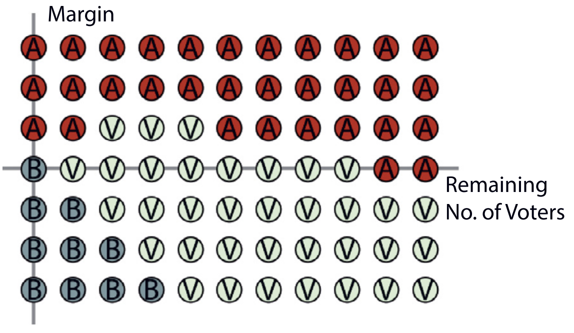

Fig. 2. The optimal stopping rule for $p = 0.68$ and voting costs $c_i = 0.3$ for all $i = 1, 2, \ldots, n$. If voting starts in the point $(r, m) = (10, 0)$, candidate A is chosen without vote.

only if the current majority for one candidate is at least $\mu(5) = 3$; with four remaining voters, voting is terminated if the current majority for one candidate is at least $\mu(4) = 2$. In each case the remaining voters could overturn the existing majority if voting were continued.

The assumption $p = 1/2$ is important for Proposition 2. With $p > 1/2$ and large n the law of large numbers implies that most likely the realized majority in the population is in favor of A. Therefore only a small majority, if any at all, is required to terminate the vote. Of course, the required majority for A is also small if the remaining number of voters is small. The function μ for candidate A is then single-peaked; for candidate B, the function μ is similar to the function described in Proposition 2 as illustrated in Fig. 2.

Propositions 1 and 2 describe features of one first best procedure, but there may be other first best procedures. We conclude this section by arguing that when $n \geq 3$ and voting costs c_i are sufficiently small no *static procedure* — in which all moves, including participation decisions, are simultaneous — can be optimal. We assume for simplicity that costs are constant, $c_i = c$ for all $i = 1, 2, \ldots, n$ and that n is odd.

When c is sufficiently close to zero, all voters who reject the default participate in the optimal static procedure. Thus, with positive probability all voters participate. Applying the transformation described in the proof of Proposition 1 to this procedure then strictly increases expected welfare by terminating voting when the current majority cannot be overturned anymore. The optimal static procedure is therefore not first best.

We can use Proposition 2 to calculate numerically the welfare in the planner's solution and in the optimal static procedure. For $n = 10$, $c = 0.3$, and $p = 1/2$ the planner's solution is the mechanism illustrated in Fig. 1 while the optimal static mechanism asks seven voters to participate.[7] Expected welfare is 1.14 in the optimal static mechanism, and 1.44 in the optimal dynamic mechanism. Thus, allowing for dynamic participation increases welfare by 26%.

4. Incentive compatibility of first best procedures

In Propositions 1 and 2 we neglected voters' incentive constraints. As any voter only moves once in a canonical procedure, a strategy profile σ is an equilibrium if a voter who is supposed to participate and bear the costs c_i finds it optimal to do so. Now the utilitarian planner consults

[7] Thus, unlike in the above argument where c is assumed to be sufficiently small, this optimal static mechanism may choose the ex post "wrong" candidate.

voter i to base the choice of the candidate on i's preferences. This suggests that incentive compatibility may be easy to achieve. This intuition is borne out by the following proposition for the case $p = 1/2$, but if $p > 1/2$ it may be wrong, as we shall see in Section 5.

Call a canonical first best procedure $(\mathcal{M}, \sigma)$ *minimal* if it has the smallest number of histories among canonical first best procedures.

Proposition 3. *Suppose $p = 1/2$. Any canonical first best procedure $(\mathcal{M}, \sigma)$ is incentive compatible. Moreover, if $(\mathcal{M}, \sigma)$ is minimal then σ is the unique sequential equilibrium of the incomplete information game induced by $\mathcal{M}$.*

Proof. We begin by proving that $(\mathcal{M}, \sigma)$ is incentive compatible. First, all information sets of the mechanism are reached with positive probability. Therefore, beliefs are given by Bayesian updating and 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 $(\mathcal{M}, \sigma)$. The modification differs from $(\mathcal{M}, \sigma)$ in that at every node h the action preferred by voters other than $P(h)$ is available for free, while the action that others do not prefer requires costly participation. We then show that $(\tilde{\mathcal{M}}, \tilde{\sigma})$ is incentive compatible and finally argue that therefore $(\mathcal{M}, \sigma)$ must also be incentive compatible.

To formally construct $(\tilde{\mathcal{M}}, \tilde{\sigma})$ fix a non-terminal history h of $\mathcal{M}$ with player $\mathcal{P}(h) = i$ and, for $z \in \{A, B\}$, let $w_{-i}(h, z)$ be 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. Assume wlog that $w_{-i}(h, A) \geq w_{-i}(h, B)$. If A is the default candidate, we don't change $(\mathcal{M}, \sigma)$ at h. If B is the default candidate, we switch labels: we label the action that voter i's type A takes as "0", i.e. "abstention", and the action that voter i's type B takes as "1", i.e. "participation". Doing this for any non-terminal history h of $\mathcal{M}$, we obtain the modification $(\tilde{\mathcal{M}}, \tilde{\sigma})$.

Clearly, $(\tilde{\mathcal{M}}, \tilde{\sigma})$ is another solution to the first best problem as $(\tilde{\mathcal{M}}, \tilde{\sigma})$ and $(\mathcal{M}, \sigma)$ lead to the same outcomes, the same expected aggregate participation costs (as $p = 1/2$), and therefore the same expected welfare.

As $(\tilde{\mathcal{M}}, \tilde{\sigma})$ solves the first best 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 A preferred by the others, and all later players follow $\tilde{\sigma}$. Letting $\pi(h, 0)$ be the probability that candidate A is chosen if i abstains, and $\pi(h, 1)$ be the probability that A is chosen if i participates,[8] this implies:

$$\begin{aligned}&\big(w_{-i}(h, A) + 2\pi(h, 0) - 1\big)/2 + \big(w_{-i}(h, B) + \big(1 - 2\pi(h, 1)\big) - c_i\big)/2 \\ &\quad\geq \big(w_{-i}(h, A) + 2\pi(h, 0) - 1\big)/2 + \big(w_{-i}(h, A) + \big(1 - 2\pi(h, 0)\big)\big)/2. \qquad (3)\end{aligned}$$

Simplifying, we obtain:

$$w_{-i}(h, B) + 2\big(\pi(h, 0) - \pi(h, 1)\big) \geq w_{-i}(h, A) + c_i. \qquad (4)$$

[8] These probabilities, as well as the expected values $w_{-i}(h, z)$, do not depend on player i's type because types are independent across players.

By assumption we have $w_{-i}(h, B) \leq w_{-i}(h, A)$, so that (4) implies:

$$2\big(\pi(h,0) - \pi(h,1)\big) \geq c_i. \tag{5}$$

Thus, voter i's participation reduces the probability that A is chosen by at least $c_i/2$. This is exactly the condition for participation of voter i to be sequentially rational in $\tilde{M}$ if i prefers candidate B. That abstention is sequentially rational in $\tilde{M}$ if i prefers candidate A is obvious because abstention increases at no cost A's election probability by at least $c_i/2$. This concludes the proof of sequential rationality of $\tilde{\sigma}$.

Now we return to the original solution $(\mathcal{M}, \sigma)$ of the first best problem. At history h, "voting for candidate $z \in \{A, B\}$", that is, taking type z's action prescribed by σ increases the probability of electing candidate z by the same amount as in mechanism $\tilde{\mathcal{M}}$, namely $\pi(h,0) - \pi(h,1)$, which exceeds $c_i/2$. Thus, $(\mathcal{M}, \sigma)$ is incentive compatible.

To prove the last sentence of the proposition assume $(\mathcal{M}, \sigma)$ minimal. Then inequality (3) must be strict at every history h: Otherwise, pruning history h and the subtree following participation at h from the modified procedure $(\tilde{\mathcal{M}}, \tilde{\sigma})$ keeps welfare constant; thus the canonical version of this pruned procedure is a first best procedure with fewer nodes, contradicting the minimality of $(\mathcal{M}, \sigma)$. Inequality (3) being strict implies that inequality (5) is also strict. Therefore, by backward induction, the sequential equilibria $\tilde{\sigma}, \sigma$ are unique. □

5. Deviation incentives in first best procedures

If $p > 1/2$, the canonical first best voting procedure need not be incentive compatible. In this case, A must always be the default candidate in a first best procedure to economize on participation cost. But then, if voter i's participation and vote for B creates positive externalities, i's participation may be desirable to the planner while i does not have sufficient incentives to participate.[9,10]

When $p > 1/2$, participation at a given history can impose externalities on three classes of voters: previous voters, potential future voters, and non-voters. The latter can be thought of as being excluded from the decision making committee; as these voters are more likely to prefer A, participation imposes a negative externality on them. The externality on past voters is positive iff a majority voted for B. Externalities for future voters are more complex. Participation may make it more or less likely that they are pivotal and may increase or decrease their expected participation costs. Therefore, the externality's sign is ambiguous. Below we give two examples illustrating these externalities.

We have chosen the two examples such that the positive externalities are sufficiently large to make participation socially desirable, but individually undesirable, so that the first best is not incentive compatible. In a large number of examples that we have numerically investigated, we have found many parameter values for which the first best is incentive compatible.[11] Curiously,

[9] In contrast, if $p = 1/2$, the planner can avoid positive externalities by labeling the choice with a positive externality "non participation" and the other "participation."

[10] While Propositions 2 and 3 may fail for general $p \in (0.5, 1]$, they do not depend on the knife-edge assumption $p = 1/2$: For generic c_i, inequalities (3)–(5) — as well as the inequalities that define the optimal stopping rule in Proposition 2 — are strict and all terms depend continuously on p. Thus, they also hold for p close to 1/2.

[11] Drawing a uniform random sample of size 10^6 with p between 1/2 and 1, n between 3 and 10, and c_i independent and between 0.05 and 2.1, we found that the first best satisfied incentive compatibility (at all histories) in approximately 65% of all cases.

Fig. 3. Canonical first best procedure for Examples 1 and 2. The figure is to be read like Figs. 1 and 2, except that the horizontal axis shows the number of voters already consulted rather than the number of voters remaining. In Example 1 (left panel) voter 2's type B prefers not to participate at the double circled node. In Example 2 (right panel) voter 1's type B prefers not to participate at the double circled node.

numerical analysis suggests that first best is incentive compatible whenever all voters have identical voting costs, regardless of the value of p. We have, however, no proof that this is indeed the case. For each example, we present the canonical first best procedure and then discuss private incentives.

Example 1. Assume $n = 3$, $p = 0.75$, $c_1 = 0.3$, $c_2 = 1.6$ and $c_3 = 1.8$. The canonical first best voting procedure is shown in the left panel of Fig. 3.

In this example, in the canonical fist best procedure, voter two, when asked to participate after voter one voted for B, will not find it in his interest to participate, and the first best is not incentive compatible: By participating, voter two secures that the outcome is B. If he does not participate, he passes the decision to voter three, and there is a probability of 0.75 that the outcome will not be B. This yields for voter two an expected utility gain of 1.5 from participating, but the costs of participation are 1.6. Therefore, voter two does not want to participate.

The reason why the planner wants voter two to participate is that his participation has a net positive externality for the other voters. Specifically, there are a positive and a negative externality from his participation. The positive externality is on voter one. Voter two's participation decision only arises when voter one favors B, and thus by participating voter two clinches the vote for B, voter one's preferred outcome. This raises voter one's expected utility by 1.5. The negative externality from voter two's participation concerns voter three. If voter two participates, candidate B is chosen, and voter three has no say. If voter two does not participate, the choice of candidate is left to voter three. Voter three's expected utility from the candidate chosen is lowered by 1.5 if B is chosen regardless of his preferences. From this we have to subtract voter three's expected participation, $0.25{*}1.8 = 0.45$. Thus, the total externality of voter two's participation on other voters' welfare is 0.45, which outweighs voter two's expected loss of 0.1 and makes participation welfare increasing. Note the subtlety of the argument explaining the net positive externality of voter two's participation: the key observation is that voter one's participation costs are sunk when voter two has to decide whether to participate, whereas voter three's participation costs are not yet sunk.

We now modify our parameters to create an example in which the externalities of the problematic voter's decision concern only future voters.

Example 2. Assume $n = 3$, $p = 0.75$, $c_1 = 0.55$, $c_2 = 0.6$ and $c_3 = 3$. The canonical first best voting procedure is shown in Fig. 3.

In this example, the first best voting procedure is not incentive compatible for voter one. His participation raises the probability of choosing B by 0.25, implying a utility gain of 0.5, less than the participation costs. The planner wants voter one to participate because of the positive externality on voter two. By participating, voter one hands the decision to voter two, and thus raises voter two's expected utility. If voter one did not participate, candidate A would be chosen without voter two being consulted. Voter two's utility increases from 0.5 if he gets A for sure to $1 - (1 - p)c_2 = 0.85$. On the other hand, the externality on voter three is negative. His utility decreases from 0.5 if he gets A for sure to 0.25 if he gets A with probability p. Thus, the aggregate positive externality $0.35 - 0.25 = 0.1$ outweighs one's expected loss of $0.55 - 0.5 = 0.05$ if he participates, and makes one's participation welfare optimal.

In the first best procedure in Example 2 the decision is delegated to a committee consisting of voters one and two. However, voter three's existence is important, for it is his existence that makes the optimal committee size two. Without him the optimal committee size drops to one.

When the canonical first best procedure is not incentive compatible, one might wonder whether there are perhaps non-canonical procedures that achieve first best and are incentive compatible. We now show that this is not the case in Examples 1 and 2.

Proposition 4. *Suppose $n, c_1, \ldots, c_n$, and p are such that among all canonical procedures there is a unique procedure that achieves first best welfare. If in this procedure some voter i has only one action node, and if participation at this node is not incentive compatible for voter i, then no first best procedure $(\mathcal{M}, \sigma)$ is incentive compatible.*

This proposition applies to Examples 1 and 2 because in both examples every voter, including the deviating voter, has at most one action node, so that the optimal canonical procedure is unique.

Proof. Fix a first best procedure $(\mathcal{M}, \sigma)$. We shall prove that the player i referred to in the proposition can improve his payoff by deviating from σ_i. In general, σ can involve mixed strategies and thus constitute a probability distribution over pure strategy profiles. Let $s = (s_i, s_{-i})$ be a pure strategy profile that arises with positive probability under σ. Clearly, $(\mathcal{M}, s)$ is also a first best procedure.[12]

When we apply the transformation in the proof of Proposition 1 to $(\mathcal{M}, s)$ we must obtain the unique first best canonical procedure, and no step of the transformation can strictly increase expected welfare. Consider the participation nodes of player i in $(\mathcal{M}, s)$ that are reached with positive probability if the other players play s_{-i} and at which different types of player i separate, i.e. type $\theta_i = A$ participates and type $\theta_i = B$ does not. It must be that at each such participation node player i's choice affects the outcome distribution. Otherwise the node could be removed, and participation costs would be saved. But if player i's choice affects the outcome, then the procedure described in the proof of Proposition 1 does not remove the participation node. Because in the unique first best canonical procedure player i moves only once, it must be the case that there is only one participation node of player i in $\mathcal{M}$ that is reached with positive probability given s_{-i}, and at which different types of player i make different choices. We denote this participation node by $h(s)$.

In $(\mathcal{M}, s)$, player i knows when $h(s)$ has been reached. Also, when i abstains at $h(s)$ the probability of electing A must equal the probability of electing A when i abstains in the first best canonical mechanism; the respective probabilities must also coincide when i participates:

[12] If there are chance nodes, where player "Chance" mixes, this step selects one possible action out of the support of σ at these nodes.

No step in the transformation in the proof of Proposition 1 changes these probabilities as $(\mathcal{M}, s)$ is non-random and chooses a welfare maximizing candidate at every history. Thus, as i's type $\theta_i = B$ prefers to abstain at his unique participation node in the canonical procedure, he also has an incentive to deviate from s and abstain at the information set that contains $h(s)$.

Now we return to the original, mixed strategy profile σ. To show that i has an incentive to deviate from σ, it suffices to show that there is an s_i in the support of σ_i such that player i has an incentive to deviate from (s_i, σ_{-i}). Consider any information set of player i that contains only participation nodes, and that is reached with positive probability under σ_{-i}. At such an information set, player i's beliefs attach positive probability only to nodes $h(s_i, s_{-i})$ for s_{-i} in the support of σ_{-i}. But by the previous paragraph player i's type $\theta_i = B$ prefers to deviate from s at any of these nodes. Therefore, player i's type B does not want to participate for any beliefs over these nodes. Thus, i can gain by deviating from σ. □

6. The second best problem

In Examples 1 and 2 the planner's problem with incentive compatibility constraints involves a welfare loss compared to the first best. How does one solve the planner's problem in this case? We are unable to answer this question. This Section offers a conjecture, though, that might be useful.

A tool that is often used to find second best mechanisms is the revelation principle. The revelation principle, applied to our model says that for any sequential equilibrium of a mechanism with participation decisions one can construct an equivalent direct mechanism in which voters reveal their types to the planner who then implements the outcome that would have resulted in the original mechanism if the voters had played their equilibrium strategies. A key point in our context is the correct definition of "outcomes". An outcome must specify the choice of one of the two candidates, as well as participation decisions for each voter. Reporting one's type in the direct mechanism is costless. A voter incurs participation costs 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 shows that truth telling is a Bayesian Nash equilibrium in the direct mechanism.

However, the direct mechanisms described in the previous paragraph are *not* mechanisms with participation decisions in the sense of Definition 2 which requires an agent to incur participation costs before taking any other action, such as reporting his type. When solving the planner's problem, we cannot therefore proceed in the usual way and maximize expected welfare among all incentive compatible direct mechanisms.

We conjecture a more useful version of the revelation principle for our model. Define "dynamic direct mechanisms with participation" as mechanisms with participation decisions in which any voter who participates subsequently reports her type without receiving further information, and then does not get to move again. We conjecture that for every mechanism with participation decisions, and corresponding sequential equilibrium, there is an equivalent dynamic direct mechanism with participation and a sequential equilibrium where voters truthfully report their types.[13,14]

[13] This is related to the revelation principle in Myerson (1986) where players in a multi-stage mechanism report in each stage their private information to the planner.

[14] In our model with binary types, only one type of each player participates, obviating the need to communicate the type explicitly after the participation decision.

While this version of the revelation principle reduces the class of mechanisms for the second best problem, it remains to be determined which players potentially participate, in which order players make their decisions, and which information they possess when making their decisions. We hope that future research will further explore the second best problem.

7. Discussion

7.1. Information acquisition

Our model assumes that voters know their types for free, but that participation is costly. By contrast, the literature on mechanism design with information acquisition, e.g. Gershkov and Szentes (2009), assumes that it is costly to know one's type, but that participation is free. We now argue that in the confines of our binary type model this distinction does not matter. Specifically, assume that voters do not know their types and have to learn them at cost c_i', but that participation is free. We maintain all other assumptions of our setup, in particular private values and independent types.[15] A mechanism with information acquisition is then an extensive game form with special "information acquisition" nodes where voter i can learn at cost c_i' whether she prefers candidate A or B.

To solve the mechanism design problem of a utilitarian planner, we can proceed as in the proof of Proposition 1 by first ignoring the voters' incentive constraints. The proof of this proposition then shows that any mechanism with information acquisition is weakly welfare-dominated by a canonical mechanism, where along a game path voter i moves at most twice, first learning his type and then, immediately afterwards, announcing it truthfully to the planner. Setting information acquisition costs equal to expected participation costs, $c_i' \equiv (1-p)c_i$, the planner's first best problems for information acquisition and costly participation coincide.

Now consider the incentives of voter i who is asked to learn his type at history h. Learning his type costs c_i' but allows his types to separate. Writing $\pi(h,z)$ for the probability of electing candidate A when announcing type z, his expected utility equals $p(2\pi(h,A)-1)+(1-p)(1-2\pi(h,B))-c_i'$ when he learns his type and announces it truthfully, and $p(2\pi(h,A)-1)+(1-p)(1-2\pi(h,A))$ when he does not learn his type and announces type A. Information acquisition is optimal if $2(1-p)(\pi(h,A)-\pi(h,B)) \geq c_i' = (1-p)c_i$. This is equivalent to inequality (5), so that information acquisition is incentive compatible if and only if at the corresponding history h of the mechanism with costly participation voter i finds it optimal to participate if he opposes the default. Moreover, voter i's truth-telling constraint $\pi(h,A)-\pi(h,B) \geq 0$ is implied by the information acquisition constraint. Hence, with independent binary types and private values, our results for dynamic voting with costly participation also apply to dynamic voting with costly information acquisition.

7.2. Correlated types

A natural extension of our model is a model with correlated types. For concreteness suppose there are two equally likely possible states of the world, a and b, where conditional on state a (resp. b) voters' preferences are i.i.d. with the probability of preferring A (resp. B) being $p > 1/2$.

[15] With common values and conditionally independent types we would recover the model by Gershkov and Szentes (2009).

In this model, it may be that no first best procedure 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. Abstention is interpreted as a preference for one candidate, and participation is interpreted as a preference for the other candidate. This one voter's preference determines the collective choice. However, this voter's individual incentives to participate do not reflect his positive externality on other voters, and he may individually prefer not to participate even if the first best procedure requires him to participate.

7.3. *Privately observed voting costs*

Another modification of our model assumes that not only voters' preferences but also their voting costs c_i are privately observed random variables. In this case, there is a new reason why first best procedures may not be incentive compatible. In our model, if he so wishes, the planner can force a type B voter to abstain by not offering him a participation node. In a model with privately observed voting costs the planner may want some type t of voter i to participate, and some other t' not. To allow for t's participation, a participation node must be available to voter i. But then, as social and private incentives are not aligned, type t' may also choose to participate. Therefore, this modification requires a separate analysis.

Appendix A

Proof of Proposition 1. We will transform any procedure $(\mathcal{M}, \sigma)$ into a canonical procedure with weakly larger welfare. Then we can restrict attention to canonical procedures in the first best problem. Existence of a first best canonical procedure then follows from the finiteness of canonical procedures.

We now describe the transformation. If $\mathcal{M}$ has information sets I_i with multiple elements, we begin by modifying $\mathcal{M}$ so that all information sets are singletons. We adjust σ so that i's strategy assigns to every $h \in I_i$ the same action that it previously assigned to I_i. This step provides players with additional information, but then they do not make use of this additional information. This 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 pick from the support of that mixed strategy the pure strategy with the largest 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. This step weakly increases expected welfare, and does so strictly if expected welfare differs across strategies in the support.

We now pick one of the two candidates as the "default candidate." Let x be that candidate, and let y be the other candidate. If $p = 1/2$ it does not matter which candidate we pick to be the default candidate; if $p > 1/2$ we pick $x = A$. 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 i, we remove all choices except $\theta_i = y$'s choice, and then remove these nodes themselves. If both types of i participate at h, then we remove h from the game tree, assuming that i participates at this node. If i's types separate at some node h' after h we relabel type x's action as 0 and type y's action as 1 at that node and remove all other actions. We also remove all nodes between h and h' where both types take the same action and the nodes after the types have separated at node h'. If both types of i don't participate at h, we remove h from the game tree, assuming that i does not participate at this node. We iterate this operation for 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 weakly increases expected welfare. Decision costs are unchanged because we didn't change the chosen candidate for any type realization. Moreover, in each step, if at node h exactly one type of a player participates, the expected participation costs weakly decrease. We may have switched which type participates, but if $p = 1/2$ this is inconsequential and if $p > 1/2$ we have switched participation to the less likely type $y = B$. 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 costs. If both types participate then we have reduced participation costs. Moreover, if both types did not participate then we have left participation costs unchanged.

Next, at every terminal history, we modify the chosen candidate to maximize expected welfare conditional on all revealed information.[16]

Finally, we 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 leaves the collective decision unchanged, and weakly decreases expected participation costs. After this step, we obtain a canonical procedure as defined in Definition 3. □

Proof of Proposition 2. Throughout the proof we assume that n and c_i are fixed. We examine the planner's stopping rule in a canonical first best procedure. First, as c_i decreases, the planner approaches voters in order of decreasing index i; clearly, if voters were not ordered according to increasing costs, doing so would yield a mechanism with weakly higher welfare. Then, starting with voter n, 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 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 $m \in \mathbb{Z}$ 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.[17]

We define $\Delta(n, r, m)$ to be the optimal decision in state (n, r, m), where $\Delta(n, r, m) = 1$ if the planner continues the voting process in state (n, r, m), and $\Delta(n, r, m) = 0$ otherwise. Obviously, Δ is only well defined if $r \geq 1$. We define $V(n, r, m)$ to be the expected welfare *net of past voting*

[16] If $p = 1/2$ this is the candidate preferred by the majority of polled players.

[17] 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.

costs if the planner makes optimal decisions starting in state (n, r, m). 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 Δ and V where the recursion is over r, the number of remaining voters.

When in state (n, r, m) the planner terminates the voting process, $(n - r + m)/2$ voters have expressed a preference for the default, $(n - r - m)/2$ have expressed a preference against the default, and the r remaining voters are equally likely to prefer the default or not. The expected welfare of selecting the default and the alternative candidate are then m and $-m$, so the optimal choice when terminating has an expected welfare of $|m|$.

When the planner does not terminate the process but consults voter r he incurs expected costs $c_r/2$, the number of remaining voters 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, $\Delta(n, r, m) = 0$*, if*

$$\big(V(n, r-1, m+1) + V(n, r-1, m-1) - c_r\big)/2 \le |m| \tag{6}$$

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

$$\big(V(n, r-1, m+1) + V(n, r-1, m-1) - c_r\big)/2 > |m|, \tag{7}$$

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

$$V(n, r, m) = \max\big\{\big(V(n, r-1, m+1) + V(n, r-1, m-1) - c_r\big)/2, |m|\big\}. \tag{8}$$

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, Δ does not depend on n either.

Step 2: As Δ and V do not depend on n we now re-define their domains as the set of all pairs $(r, m) \in \mathbb{N}_0 \times \mathbb{Z}$ by setting $\Delta(r, m) = 1$ if $\Delta(n, r, m) = 1$ for some n, and $\Delta(r, m) = 0$ if $\Delta(n, r, m) = 0$ for some n. Similarly, $V(r, m) = v$ if and only if $V(n, r, m) = v$ for some n.

The functions V and Δ are symmetric: $V(r, m) = V(r, -m)$ and $\Delta(r, m) = \Delta(r, -m)$. This is obvious from our symmetry assumption $p = 1/2$. We also note that $V(r, m) = |m|$ and $\Delta(r, m) = 0$ whenever $r \le |m|$, because $r \le |m|$ implies that the remaining voters cannot overturn the existing majority, and therefore consulting them will not change the outcome.

Step 3: We prove that $\Delta(r, m)$ weakly increases in r:

$$\Delta(r, m) = 1 \quad \Rightarrow \quad \Delta(r+1, m) = 1 \quad \text{for all } r \in \mathbb{N}_0 \text{ and all } m \in \mathbb{Z}. \tag{9}$$

The value of stopping[18] is the same in state (r, m) and in state $(r + 1, m)$: $|m|$. In state $(r + 1, m)$ the planner could also adopt the strategy that is optimal in state (r, m) until the last voter is reached, and then always stop, that is, ignore the last voter. As costs c_i decrease in i, this strategy has weakly lower participation costs and therefore a weakly higher value in state $(r + 1, m)$ than the optimal strategy does in state (r, m). The optimal strategy in state $(r + 1, m)$ yields a weakly higher value, so that $V(r + 1, m) \ge V(r, m)$. The latter strictly exceeds $|m|$ because $\Delta(r, m) = 1$, and so $\Delta(r + 1, m) = 1$.

[18] We define the "value" of a decision in a state as in Step 1: the expected welfare if the planner continued with optimal decisions, net of participation costs incurred in the past.

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

$$\Delta(r,m) = 0 \quad \Rightarrow \quad \Delta(r,m+1) = 0 \quad \text{for all } r, m \in \mathbb{N}_0. \tag{10}$$

Let $\Delta(r,m)$ and $\Delta(r,m+1)$ be the optimal stopping rules starting at states (r,m) and $(r,m+1)$. Also, let $W(r,m)$ be the value of using rule $\Delta(r,m+1)$ at state (r,m). As $V(r,m+1)$ and $W(r,m)$ are calculated based on the same stopping rule $\Delta(r,m+1)$, and the payoff when stopping $V(r',m') = |m'|$ increases at rate at most 1, we have $V(r,m+1) - W(r,m) \leq 1$. By definition of the continuation value at state (r,m) we have $V(r,m) \geq W(r,m)$ and so $V(r,m+1) - V(r,m) \leq 1$. Thus, if $\Delta(r,m) = 0$ and so $V(r,m) = m$, then $V(r,m+1) \leq m+1$ and so $\Delta(r,m+1) = 0$.

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

$$\mu(r) = \min\left\{|m| \in \mathbb{N}_0 : \Delta(r,m) = 0\right\}. \tag{11}$$

Because $\Delta(r,m) = 0$ 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 Δ in m from Step 4 implies that for every r, $\Delta(r,m) = 0$ whenever $m \geq \mu(r)$ and $\Delta(r,m) = 1$ 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 $\Delta(r,m)$ weakly increases in r, as we showed in Step 3. □

References

Bergemann, Dirk, Välimäki, Juuso, 2002. Information Acquisition and Efficient Mechanism Design. Econometrica 70, 1007–1033.

Bergemann, Dirk, Välimäki, Juuso, 2006. Information in mechanism design. In: Blundell, Richard, Newey, Whitney, Persson, Torsten (Eds.), Proceedings of the 9th World Congress of the Econometric Society. Cambridge University Press, Cambridge, pp. 186–221.

Börgers, Tilman, 2004. Costly voting. Amer. Econ. Rev. 94, 57–66.

Celik, Gorkem, Yilankaya, Okan, 2009. Optimal auctions with simultaneous and costly participation. J. Theoretical Econ. (Advances) 9. Article 24.

Fadel, Ronald, Segal, Ilya, 2009. The communication costs of selfishness. J. Econ. Theory 144, 1895–1920.

Gershkov, Alex, Szentes, Balázs, 2009. Optimal voting schemes with costly information acquisition. J. Econ. Theory 144, 36–68.

Ghosal, Sayantan, Lockwood, Ben, 2009. Costly voting when both information and preferences differ: is turnout too high or too low? Soc. Choice Welfare 33, 25–50.

Krasa, Stefan, Polborn, Mattias, 2009. Is mandatory voting better than voluntary voting? Games Econ. Behav. 66, 275–291.

Myerson, Roger B., 1986. Multistage games with communication. Econometrica 56, 323–358.

Osborne, Martin, Rubinstein, Ariel, 1994. A Course in Game Theory. The MIT Press, Cambridge, MA.

Pavan, Alessandro, Segal, Ilya, Toikka, Juuso, 2014. Dynamic mechanism design: a Myersonian approach. Econometrica 82, 601–653.

Stegeman, Mark, 1996. Participation costs and efficient auctions. J. Econ. Theory 71, 228–259.