Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 144 (2009) 36–68
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Optimal voting schemes with costly information acquisition

Alex Gershkov [a], Balázs Szentes [b,*]

[a] *Department of Economics, University of Bonn, Adenauerallee 24-42, Bonn 53113, Germany*
[b] *Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60657, USA*

Received 26 October 2006; final version received 31 January 2008; accepted 1 February 2008

Available online 23 May 2008

Abstract

A group of individuals with identical preferences must make a decision under uncertainty about which decision is best. Before the decision is made, each agent can privately acquire a costly and imperfect signal. We discuss how to design a mechanism for eliciting and aggregating the collected information so as to maximize ex-ante social welfare.

We first show that, of all mechanisms, a sequential one is optimal and works as follows. At random, one agent at a time is selected to acquire information and report the resulting signal. Agents are informed of neither their position in the sequence nor of other reports. Acquiring information when called upon and reporting truthfully is an equilibrium.

We next characterize the ex-ante optimal scheme among all ex-post efficient mechanisms. In this mechanism, a decision is made when the precision of the posterior exceeds a cut-off that decreases with each additional report. The restriction to ex-post efficiency is shown to be without loss when the available signals are sufficiently imprecise. On the other hand, ex-post efficient mechanisms are shown to be suboptimal when the cost of information acquisition is sufficiently small.

JEL classification: D71; D72; D81

Keywords: Information aggregation; Voting

* Corresponding author.
E-mail addresses: alex.gershkov@uni-bonn.de (A. Gershkov), szentes@uchicago.edu (B. Szentes).

0022-0531/$ – see front matter
doi:10.1016/j.jet.2008.02.004

1. Introduction

Consider a setting in which a group of individuals with a common goal must make a single choice among a number of alternatives, where the alternative that is best depends on an unknown state. Further, suppose that each group member can privately choose to invest in a costly and imperfectly informative signal about the state. Such situations are not uncommon: recruiting committees in academic departments, boards of directors evaluating a potential merger, teams of doctors deciding upon the best course of treatment for a patient; each of these group decision-making situations, and many more like them, fits reasonably well into the basic framework of this paper.

In such environments, there is typically a free-rider problem at work; if sufficiently many individuals are expected to invest in and report their information, then any one of them, recognizing that his information is unlikely to be pivotal, will rationally choose not to invest, thereby saving his investment cost. This creates a social inefficiency because the free-rider's decision does not take into account the reduction in the others' payoffs. Note that punishment severe enough to deter free-riding when information acquisition is observable might not be effective here, where it is private, because the free-rider can always pretend to have invested by providing a report that some signal might have generated. What then is the most efficient way to deal with this free-rider problem? The purpose of the present paper is to bring the tools of mechanism design to bear on this fundamental question of group decision-making.

An interesting preliminary observation is that the most commonly employed version of the revelation principle is insufficient for our purposes. To understand this, recall that the standard revelation principle asserts that it is without loss of generality to ask each agent to simultaneously report his type. But this is not possible in our environment because, when first approached by the mechanism-designer, our agents possess no information and hence have no type to report. The version of the revelation principle that we require must apply when agents optimally choose *whether* to obtain private information, i.e. when private information is endogenous. Our first result is to show that it is without loss of generality to restrict attention to mechanisms within the following class. First, the agents are randomly ordered. Next, the agents are sequentially asked to acquire and report their signal. An agent receives no information about either her position in the order or the signals reported by previous agents in the sequence. After each agent's report, either a final decision is made or the next agent in the sequence is asked to acquire and report a signal. Finally, each such mechanism must be incentive compatible. That is, it must be an equilibrium for each agent to acquire and report truthfully their signal when asked to do so.

Let us call a mechanism *ex-post efficient* if the resulting decision is efficient, given the information acquired by the agents. We view the class of ex-post efficient mechanisms as particularly important and natural for at least two reasons. First, committing to ex-post inefficient decisions is often very difficult in practice. For example, it is unlikely that a group of doctors can commit to recommending a treatment they all believe is less likely than another to cure a patient. Second, there are legal considerations that make ex-post inefficient decisions potentially very costly and, if they are costly enough, the fully optimal mechanism will be ex-post efficient.[1] For these and

[1] We are grateful to Richard Epstein and Richard Posner for pointing this out. In general, the law treats failure to act on a known risk more severely than failure to act on an unknown risk. To illustrate the relevance of this issue to our model consider the following two scenarios. Case 1: A hospital provides a patient with treatment X. The patient dies and it turns out that treatment Y could have saved his life. The hospital could have known this if it had acquired information about the patient. But the hospital did not know. Case 2: The hospital examines the patient, conducts several tests. The

related reasons, the agents may well be able to commit to a mechanism that sometimes makes a decision without asking all group members to gather information, but may be unable to commit to deliberately making a decision that is suboptimal given the collected information.

If the precision of the signal is sufficiently low relative to the number of agents, the restriction to ex-post efficient mechanisms is without any loss; the ex-ante optimal mechanism, among all possible mechanisms, is ex-post efficient. On the other hand, when the cost of information is sufficiently small relative to the precision of the signal, ex-ante social surplus is maximized by making ex-post inefficient decisions with positive probability. In this case, ex-post inefficient decisions serve as a cost-efficient threat to induce agents to acquire information.

Our main result is a characterization of the optimal ex-post efficient mechanism. It has the following simple form: if after receiving reports from n agents, the posterior distribution over states of the world has precision above some cut-off $f(n)$, information acquisition stops and the efficient decision given the reports is made. Otherwise, an additional agent is asked to acquire information and to report her resulting signal. A key property is that f is decreasing. Consequently, less precise information is required to induce a decision as the number of solicited agents increases.

The intuition behind this result is the following. An agent is willing to acquire information only if her probability of being pivotal is sufficiently large. To see how the mechanism provides this probability, suppose that an agent is pivotal, i.e., that her report changes the decision. Given the ex-post efficient decision-rule, the resulting decision must then be taken with an imprecise posterior, otherwise her single report would unlikely to change it. Consequently, the mechanism can provide an agent with opportunities to be pivotal only by sufficiently often making a decision with an imprecise posterior. Thus, providing incentives to invest runs counter to making informed decisions. We should therefore expect the optimal mechanism to make imprecise decisions that have a large impact on the agents' incentives to invest, but a small impact on the potential inefficiency of the outcome. And indeed, this is what our optimal mechanism achieves. Because the precision cut-off function $f(n)$ is decreasing, when a decision is made, it is made with an imprecise posterior only if sufficiently many agents have been asked to acquire information. Consequently, from an ex-ante point of view, decisions with imprecise priors are unlikely (because a decision is unlikely to require asking many agents), and their efficiency cost is therefore small. But here is the critical point. Despite the fact that the event in which a decision is made with an imprecise posterior has an *ex-ante* small probability, it has a sufficiently large *conditional* probability in the eyes of an agent who is asked to acquire information that gives her the correct incentives. This is because the random ordering of agents implies that it is unlikely that any particular agent is near the beginning of the order. Hence, conditional upon being asked to acquire information, it is unlikely that only a small number of agents have previously been asked. Consequently, the conditional probability that a relatively large number of agents have already been asked, increases compared with the ex-ante situation, thereby increasing the conditional probability that (a) the precision cut-off f is low (because f is decreasing), (b) that the decision is made with an imprecise posterior, and finally (c) that the agent is pivotal.

---

tests indicate that treatment X is the wrong treatment, and treatment Y is the right one. Nevertheless, the hospital applies treatment X and the patient dies. The law would describe Case 1 as a case of palpable negligence. The hospital would be liable for compensatory damages. In Case 2, the law would consider the hospital's conduct reckless, and the hospital would likely be forced to pay punitive as well as compensatory damages. Further, if it can be proven that the hospital acted with deliberate disregard, it can be convicted for murder as well as malpractice.

Most of the existing literature has focused on voting models with an *exogenously* given information structure. See [1,5,9–11,14,17]. For example, [8] compares sequential and simultaneous voting and find an equivalence between different voting schemes in terms of the equilibrium outcome. In particular, the authors of [8] find that for any choice rule, there is a Pareto best equilibrium outcome that is the same whether voting is sequential or simultaneous. In contrast, this paper shows that sequential schemes always dominate simultaneous ones when information acquisition is costly.

There is a growing literature on costly information acquisition in mechanism design. Most of this literature deals with auction and public goods models where utilities are perfectly transferable, unlike in voting models, where monetary transfers are not feasible. See [3,18,22], and the references therein. These papers focus on simultaneous information acquisition. Their goal is to analyze the incentives to acquire information in different classical mechanisms. They do not seek the optimal mechanism, as we do here.

[23] examines two-stage voting games of the following form. In the first stage, all committee members simultaneously decide whether to acquire a noisy signal about the state of the world or to remain uninformed. In the second stage, they vote. The author analyzes the optimal voting scheme among threshold voting rules and the optimal size of the committee. [12] analyzes environments similar to those in [23]. The authors enrich the model by introducing a communication stage before voting and demonstrate that this cheap talk stage can make the mechanism more efficient. [21] analyzes the effect of committee size on the accuracy of the final decision. He shows that in symmetric mixed-strategy equilibria, increasing the committee size may lead to a less accurate decision. In asymmetric pure-strategy equilibria, however, changing the committee size does not affect the accuracy of the final decision. All of these papers employ specific ex-post efficient voting procedures and embrace the idea that it is important to improve outcomes through careful design. We agree and take this view a step further by optimizing over all ex-post efficient mechanisms.

All of the papers described in the previous paragraph analyze settings with homogeneous preferences. [4] allows for heterogeneous preferences with both non-verifiable information acquisition and costly participation in a committee. In this model, members are supposed to report their information to a principal. The principal makes the final decision based on the reports. [4] characterizes the optimal committee size when the signals as well as the decision are continuous variables and the principal uses the mean decision rule to determine the final decision. [13] examines a committee with a fixed number of members. Each member acquires a noisy signal. The signals become public after acquisition. He analyzes the properties of the optimal decision rule for the case of simultaneous information acquisition. To provide incentives for acquiring information, it is optimal to distort the decision rule away from the ex-post efficient one. [4] treats the decision rule as given and focuses on committee size. [13] does the opposite. [2] also considers a model where the voters have heterogeneous preferences over two alternatives. In addition, the voters are allowed to communicate prior to voting. The authors show that unanimous voting makes it impossible to reveal all private information.

[15] discusses environments in which voters have common preferences but are uncertain which of two alternatives is better for them. Each voter can acquire a costly signal; the precision of this signal depends on the amount of investment. The author proves that if the cost and the marginal cost of precision are zero at zero, then the majority decision induced by simultaneous voting is asymptotically efficient.

Several key features of our version of the revelation principle are discussed in [20]. For example, the author notes that when agents can take actions that affect their private information, they need only report their new information to the mechanism and the mechanism need only communicate to the agents the actions they should take. In all cases, communication is private between the mechanism and each agent. However, [20] focuses on the communication between the mechanism and the agents, taking the order of play and actions of the agents as given.

[24] analyzes the problem of free riding in multi-agent computations. Each agent can retrieve a binary input at a cost. The goal of the group is to compute the value of a given function, G, mapping from input vectors into $\{0, 1\}$. The agents face a free-rider problem similar to ours: If the probability that a single input affects the value of G is small, an individual prefers not to incur the cost and provides an input at random. The authors independently derive a version of the revelation principle that is similar to the one we employ here. In particular, their canonical mechanism is also sequential and each agent is told only whether she should or should not retrieve the input. The goal of the authors is to characterize the set of those functions that can be computed with probability one. This mechanism design problem is conceptually different from ours. Therefore, the results in [24], apart from the derivation of the canonical mechanisms, are not comparable to ours.

The paper is organized as follows. Section 2 describes the model and characterizes the first-best voting scheme. We preview our main results in Section 3. Section 4 characterizes the canonical mechanisms, and explicitly derives the incentive compatibility constraint. Section 5 introduces the notion of continuation mechanism and proves some basic properties of the optimal mechanisms. The main results are in Section 6. Section 7 concludes. Some of the proofs are posted on the JET Supplementary Materials website and on the homepages of the authors.

2. The model

There is a society consisting of $K\,(\in \mathbb{N} \cup \{\infty\})$ agents. The state of the world can take one of two values, A or B. Each state occurs with probability one-half. The society must take an action, either α or β. An agent's utility is

$$u(\alpha|A) = u(\beta|B) = 1,$$
$$u(\alpha|B) = u(\beta|A) = 0.$$

Every agent can draw a signal at most once at a cost c. The signal can take one of two values, a or b and is distributed as follows:

$$P(a|A) = P(b|B) = p > 1/2,$$
$$P(a|B) = P(b|A) = 1 - p = q.$$

The signals are independently distributed across agents conditional on the state of the world. Information acquisition is unobservable. (We also implicitly assume that information acquisition takes no time.) An agent's payoff is u if she does not invest into information, and $u - c$ otherwise.

There is a Social Planner (SP) who wants to maximize expected sum of the agents' utilities, net of the expected total cost of information acquisition. However, we allow the SP to weight costs and benefits differently from the agents. The objective of the SP is to maximize

$$K_0 Eu - c\overline{L}, \tag{1}$$

where $\overline{L}$ is the expected number of agents who collect information, and K_0 (> 1) is an integer. If $K < \infty$, one can assume that $K_0 = K$, in which case (1) is the standard social welfare. The SP cannot use a transfer scheme to induce the agents to acquire information.

It is useful to introduce the following notation. Let $\#_a(s)$ ($\#_b(s)$) denote the number of signals a (b) in the finite signal sequence s. The following lemma states that the posteriors about the state of the world and about the next signal after observing a sequence of signals depend only on the difference between the numbers of signals a and b observed so far.

Lemma 1. *Let s be a finite sequence of signals such that $\#_a(s) - \#_b(s) = d$. Then*

$$P(A|s) = \frac{p^d}{p^d + q^d} \quad \text{and} \quad P(a|s) = \frac{p^{d+1} + q^{d+1}}{p^d + q^d}.$$

Proof. See Appendix A. □

We introduce the following notation: $P(d) = p^d/(p^d + q^d)$, $Q(d) = 1 - P(d)$, $p(d) = (p^{d+1} + q^{d+1})/(p^d + q^d)$, and $q(d) = 1 - p(d)$. Notice, $P(d)$ denotes the probability that the action α (β) is the efficient one, given that $\#_a(s) - \#_b(s) = d$ ($\#_b(s) - \#_a(s) = d$). $Q(d)$ denotes the probability that the action α (β) is the efficient one, given that $\#_a(s) - \#_b(s) = -d$ ($\#_b(s) - \#_a(s) = -d$).

In order to make our problem interesting, we assume $P(1) - 1/2 > c$. This means that the benefit to an individual of acquiring a signal exceeds its cost. If this assumption does not hold, no agent would ever invest in information, and hence the optimal mechanism would be to take an action randomly without information acquisition.

2.1. The first-best mechanism

Next, we characterize the first-best voting scheme. That is, the optimal mechanism in which the agents do not behave strategically, and hence the scheme does not have to be incentive compatible. Since information acquisition takes no time, the first-best mechanism can be assumed to be sequential. The SP sequentially asks the agents to draw signals and report them. After each agent, the SP has to decide whether to ask an additional agent or to take a final action. Clearly, whenever the SP takes an action, it corresponds to the majority of the reports. That is, the final action is α if and only if more signals a were reported than signals b. The problem of the SP is a standard stochastic dynamic programming problem. From Lemma 1, it follows that one of the state variables is the difference between the number of signals received of each type. In fact, if $K = \infty$ this is obviously the only state variable. On the other hand, if $K < \infty$ there is an additional state variable: the number of agents who have previously been asked. The reason the number of agents already asked is a state variable can be shown by the following example. Suppose there are ten agents. Compare the following two cases. In case one, only three agents have been asked and each of them has reported signal a. In case two, nine agents have been asked; six of them reported signal a, and three of them reported signal b. By Lemma 1, the posteriors of the SP are identical in the two cases. However, unlike in case one, in case two the value of asking an additional agent is zero, since even if she reports signal b, taking action α will remain optimal. In general, the fewer agents that remain to possibly ask, the smaller the value of asking an additional agent given a certain posterior. Hence, conditional on the same posterior, the larger is the number of agents who have already been asked, the less likely the SP will ask an additional agent.

The following proposition characterizes the first-best mechanism.

Proposition 1. *There exists a weakly decreasing function $g:\mathbb{N}\to\mathbb{N}$ such that if, after asking l agents, the reported signal sequence is s and $|\#_a(s)-\#_b(s)|\geqslant g(l)$, $|\#_a(s)-\#_b(s)|\geqslant g(l)$ the SP makes the majority decision. Otherwise, the SP asks an additional agent.*

(i) *If $K=\infty$ then there exists $k\in\mathbb{N}$ such that $g\equiv k$. In addition, $k\to\infty$ as $K_0\to\infty$.*
(ii) *If $K<\infty$ then for all $l\in\mathbb{N}$, $g(l+1)=g(l)$ or $g(l+1)=g(l)-1$, and $g(K-1)=1$.*

Part (i) is the standard Wald Theorem. Part (ii) can be deduced from [7, Theorems 2 and 3 in Chapter 12.5], and hence the proofs are omitted.

Observe that $g(K-1)=1$ does not mean that, in the first-best mechanism, all the agents are potentially asked to acquire information. In general, the value of g is one for a number smaller than $K-1$.

3. Preview of the results

This section is devoted to an illustration of our main theorem by a numerical example. Suppose that $K=K_0=9$, $p=2/3$, and $c=0.04$. Below, we describe the optimal mechanism and the strategies of the agents. We note that the optimal mechanism is ex-post efficient in this example (see Theorem 4).[2] We shall compare this mechanism with the optimal simultaneous voting scheme, in which some of the agents have to acquire information at the same time.

The optimal mechanism works as follows. The SP asks, at random, one agent at a time to invest in information and to report the resulting signal. Agents are informed of neither their position in the sequence nor the reports of previous agents. After receiving a report, the SP has to decide whether to take a final action, or to ask an additional agent. In our example, this decision can roughly be characterized by the function

$$f(l)=\begin{cases}4 & \text{if } l\in\{1,2,3,4\},\\ 3 & \text{if } l\in\{5,6,7\},\\ 2 & \text{if } l=8,\\ 1 & \text{if } l=9.\end{cases}$$

If, after receiving l signals, $d=|\#_a(s)-\#_b(s)|\geqslant f(l)$ the SP takes the optimal action given the collected signals. If $d=|\#_a(s)-\#_b(s)|<f(l)$ then the SP asks an additional agent to acquire information. There is only one exception: if $l=5$ and $|\#_a(s)-\#_b(s)|=3$, the SP randomizes between taking an action and asking an additional agent. This mechanism is represented by squares in Fig. 1.

The strategy of each agent is to invest in information whenever she is asked to do so, and to report the resulting signal truthfully.

Consider now the simultaneous voting game which is the one most commonly analyzed in the literature. Suppose that five out of nine agents are asked to acquire information and report the resulting signals. In addition, the action α is taken if and only if at least three agents report signal a. This is the optimal mechanism among the simultaneous ones. Although agents have to take actions at the same time, one can still represent this scheme as a sequential mechanism, in

[2] Since the cost of information is low, the first-best mechanism specifies asking an agent to acquire information whenever the reports of the remaining agents affect the final decision with positive probability.

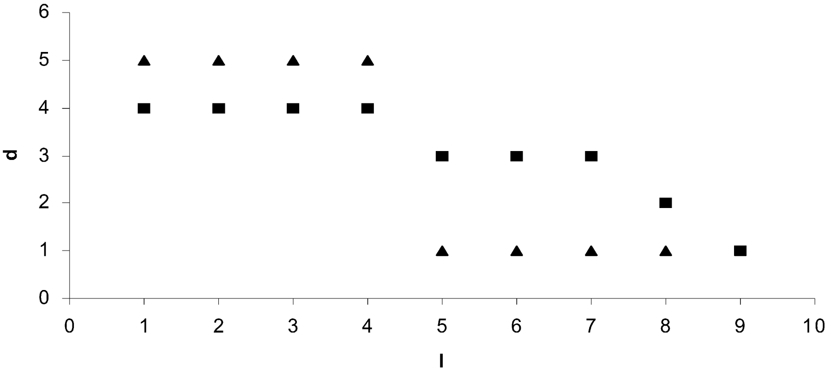

Fig. 1. The optimal and the simultaneous mechanisms.

which no matter what the reports are, exactly five agents are asked to invest.[3] This scheme is described by the function

$$h = \begin{cases} 5 & \text{if } l \leqslant 4, \\ 1 & \text{if } l \geqslant 5, \end{cases}$$

and is represented by triangles in Fig. 1.

The main difference between the optimal scheme and the simultaneous one is that in the simultaneous scheme, the history of reports does not affect whether the SP asks an additional agent. Indeed, no matter what the reports are, the SP always asks exactly five agents. In contrast, in the optimal mechanism, the SP stops asking agents if his posterior given a history of reports exceeds a cutoff level. In addition, given that the function f is gradually decreasing, the greater is the number of agents who have already acquired information, the smaller is this cutoff level. There are two reasons explaining this feature: (i) it makes the mechanism cost efficient, and (ii) it provides the agents with proper incentives to acquire information. Next, we elaborate on these reasons.

3.1. Cost efficiency

The reason the decreasing function f saves costs is the same as the reason the function g characterizing the first-best mechanism decreases when $K < \infty$ (see part (ii) of Proposition 1). That is, as more agents invest in information, fewer agents remain to acquire signals. It is therefore less likely that the reports of the remaining agents change the posterior of the SP. To illustrate, consider the points $(3,3)$ and $(5,3)$ in Fig. 1. These points correspond, for example, to the report sequences $s_1 = (a,a,a)$ and $s_2 = (a,a,a,b,a)$, respectively. The optimal mechanism solicits information from an additional agent after s_1, but recommends action α after s_2 with positive probability. The reason is that, after the sequence s_1, the SP still has available six more agents. Therefore, there is a relatively high probability that the reports of the remaining six agents will induce the SP to take an action different from α. In contrast, after the sequence s_2, the SP has only four available agents. The SP would take an action different from α only if all four agents

[3] This is because the simultaneous and the sequential mechanisms are strategically equivalent as long as the agents receive no information when they are asked to invest in signals.

report the signal b. Since this event has a low probability, it is optimal to take the action α now, avoiding additional costs to gathering information.

Now, consider the simultaneous mechanism and the point $(3, 3)$ in Fig. 1. This point corresponds, for example, to the report sequence $s = (a, a, a)$. The simultaneous mechanism specifies asking two more agents after the sequence s. No matter what these agents report, the action α is taken because the function h jumps down from five to one at $l = 4$. Therefore, after the sequence s, asking two more agents is costly but yields no benefit. In fact, any large downward jump of the stopping rule is associated to acquiring irrelevant information. Our optimal mechanism saves the cost of this type of redundant information for the following reason: since the function f never jumps down more than one step at a time, whenever an agent is asked to acquire information his report influences the final decision with positive probability.

In general, unlike the simultaneous scheme, the optimal mechanism stops soliciting information if the cost of information acquisition outweighs the benefit of making a more efficient decision. As expected, the social welfare generated by our optimal mechanism is higher than the social welfare generated by the simultaneous scheme.

3.2. Efficient incentive provision

The second reason f is gradually decreasing is that it provides the agents with better incentives to acquire information. An agent is willing to acquire information only if her probability of being decisive is sufficiently large. The only way to provide opportunities for being pivotal is to take actions even if the posterior is imprecise,—thatis, when a single report can potentially determine the final action. Since f is decreasing, final actions when the posterior is imprecise are made after long sequences of reports rather than after short ones. The key insight is that conditional on being asked, an agent believes that she is more likely to have been asked along a long sequence than along a short one. Next, we explain why the probabilities of long sequences are *exaggerated,* conditional on being asked.

Suppose that there are one hundred agents, and consider the following mechanism. With probability .9 only a single agent is asked at random to acquire information, and with probability .1 all agents are asked. When an agent is asked, she receives no additional information. In particular, she does not know if she is the only one who has been asked to acquire information, or if there are 99 other agents who have also been asked. We emphasize that this mechanism has little to do with the optimal one, and we merely use it to illustrate the difference between the conditional probabilities of short and long sequences. Using Bayes' rule, the probability that the SP asks each agent conditional on being asked is $.1/(.1 + .9/100) > .9$. That is, although the probability of a hundred-long sequence is only .1, the probability of the same event conditional on being asked is larger than .9. Since the actual probability of a long sequence is small, the effect of such a sequence on social welfare is small. However, since the conditional probability of a long sequence is large, the effect on the incentive compatibility constraint is also large. That is why the cut-off function f is decreasing; after a short sequence of reports the SP takes an action only if the posterior is precise, but after a long sequence a decision is made even if the posterior is fuzzy.

The main result of our paper, see Theorem 1, is that, no matter what the values of the parameters are, the optimal ex-post efficient mechanism has similar attributes to those of the mechanism described above. In particular, the SPs decision about whether to ask an additional agent, or to take a final action can always be characterized by a function f, which has the following two properties:

(i) for all $l \in \mathbb{N}$, $f(l+1) = f(l)$ or $f(l+1) = f(l) - 1$, and
(ii) there exists an $N \in \mathbb{N}$, $N \leqslant K$, such that $f(N) = 1$.

Property (i) implies that the function f gradually decreases, but never jumps down more than one step. Property (ii) implies that there is a bound, N, above which the value of the function f is one. Notice that after asking either N or $N+1$ agents, the difference between the number of different signals is at least one. Hence, even if $K = \infty$, the SP never asks more than $N+1$ agents to acquire information.

The contrast between the first-best and second-best mechanisms is especially sharp when $K = \infty$. Recall from part (i) of Proposition 1, that the first-best mechanism is then characterized by a constant. The second-best mechanism, on the other hand, is still characterized by a decreasing step function. In addition, unlike in the first-best mechanism, the maximum number of agents acquiring information is bounded in the second-best mechanism. When $K = \infty$, the function f is decreasing only to provide incentives for the agents to incur the cost of information acquisition. Indeed, since $K = \infty$, the number of agents already asked has no impact on how likely it is that the remaining agents change the posterior of the SP.

4. The revelation principle and incentive compatibility

4.1. Canonical mechanisms

A voting mechanism is an extensive-form game (with perfect recall) where the players are the agents. At each information set, either chance or an agent moves.[4] Since information acquisition is unobservable, at any information set where a certain agent has to move, she can also acquire information if she has not done so at a previous information set.[5] Furthermore, each terminal node corresponds to taking either action α or action β. Let G be such a mechanism and NE be a Bayesian Nash equilibrium in the game. The profile NE can be a mixed-strategy equilibrium.[6] The pair (G, NE) is said to implement $k \in \mathbb{R}$ if the value of (1) is k if players follow the equilibrium strategies specified by NE. A mechanism is said to be *sequential* if the agents never have to take actions simultaneously.

The next lemma states that we can restrict attention to a manageable class of mechanisms. We call the elements of this class *canonical mechanisms*. A mechanism is canonical if it satisfies the following three properties: (i) it is sequential, (ii) for each agent, there is at most one information set where the agent has to take an action, and (iii) the possible actions of the agents are reporting signal a and reporting signal b after either acquiring a signal or not. That is, in a canonical mechanism, the agents are sequentially asked to acquire and report information according to a possibly random order. After an agent is asked, either a final decision is made or an additional agent is asked to acquire information. Property (ii) says that when an agent is asked to invest

[4] First, chance determines the state of the world, and also the realizations of the signals. However, a move of the chance does not necessarily correspond to a randomization by nature. It may mean, for example, asking more agents to collect information, revealing some information to certain agents, or making a final decision. It can be useful to think of the chance as a mediator or a machine as in [20].

[5] Remember, each agent can acquire a signal once at most.

[6] Chance may also randomize when taking an action. In fact, we will show that the optimal mechanism generically involves such randomizations.

in a signal, she receives information neither about her position in the sequence, nor about other agents' reports.[7]

Lemma 2. *Suppose that* (G, NE) *implements* k $(\in \mathbb{R})$. *Then there exists a canonical mechanism,* G', *and an equilibrium* NE' *in* G' *such that* (G', NE') *also implements* k. *Furthermore,* NE' *is a pure-strategy equilibrium and specifies that whenever an agent has to take an action she acquires information and reports it truthfully.*

Proof. Since information acquisition takes no time, one can assume that G is sequential.[8] Let us modify G in the following ways.

(a) Suppose that at a certain information set, denoted by I, a certain agent takes an action without acquiring information with probability γ (> 0). Then introduce a new information set I' into the game. The information set I' is similar to I except that chance moves instead of the agent. Furthermore, from any node preceding I, the information set I is reached with probability $1 - \gamma$ and I' is reached with probability γ. At I', chance moves and takes the same (possibly random) action as the one taken by the agent after not acquiring information at I in G.

Let G'' denote the game after applying (a) whenever possible. Let NE'' denote the strategy profile, where each agent, at each information set, takes the same action as in G at the corresponding information set after acquiring information. Since NE is an equilibrium in G, the profile NE'' is an equilibrium in G''. In addition, each agent collects information before taking an action in NE''. Next, we modify G'' and NE'' in the following ways.

(b) Suppose that, at information set I, some agent can take an action other than reporting a or reporting b after acquiring information. Modify the game such that the agent's available actions at I are reporting a and reporting b. In addition, if the agent reports signal s $(\in \{a, b\})$ then chance moves, and takes the same action as the one taken by the agent with signal s in NE''. Finally, specify the strategy of an agent for reporting the acquired signal.

(c) Notice that, after the modifications above, it is still possible that an agent, after acquiring information, reports her signal multiple times along a path of the game. Modify the game such that whenever an agent has to move for at least the second time along a certain path of the play, then chance moves instead of the agent and takes the same action as the agent took at the first information set along this path.

(d) After the changes above, an agent always acquires information and reports her signal truthfully at each information set. Unify all of the information sets where a certain agent has to take an action, so each agent moves at most at one information set.

After applying (b)–(d), we defined a new game G' and a strategy profile NE'. The game G' satisfies properties (i)–(iii). The agents indeed collect information and report it truthfully in NE'. Obviously (G', NE') implements k, too. It remains to show that in G', the strategy profile NE' is indeed an equilibrium. But if there were a profitable deviation, then by (b), there would have been a profitable deviation in NE too. □

[7] [19] and [20] derive similar results in general principal-agent models. These papers show that there is no loss of generality in assuming that the agents always report their new information to the principal, and the principal only recommends actions to the agents. [24] also arrives at similar conclusions.

[8] If after some history, several agents have to take actions simultaneously, one can separate out these simultaneous moves by modifying the game as follows. After this history, make these agents move sequentially without receiving information about the actions taken by the other agents. This change does not affect the equilibrium strategies.

Property (ii) of the canonical mechanisms, which says that each agent has at most one information set where she can take an action, deserves some discussion. One might think, for example, that the optimal mechanism has the following feature. If after a long sequence of reports the posterior is imprecise, the SP reveals some information to the next agent in order to give her more incentive to invest. But in this case, if an agent is asked to acquire a signal and does not receive additional information, she concludes that she was not asked after a long sequence where the posterior of the SP is imprecise. As a result, she might find it optimal not to acquire a signal. In general, if a mechanism provides information to the agents about previous reports, there are several information sets corresponding to each agent, and there is an incentive compatibility constraint corresponding to each of these information sets. If this mechanism is modified such that no information is revealed to the agents, each agent will have only one information set and a single incentive constraint. This constraint is the average of the constraints corresponding to the agent's information sets in the original mechanisms. The reason one can restrict attention to mechanisms where no information is revealed to the agents is that it is easier to satisfy the average of several constraints than each of them individually.

One can also assume that the optimal mechanism is ex-ante symmetric with respect to $a(A)$ and $b(B)$. The reason is the following. Let us assume that the optimal mechanism is asymmetric. Then let us consider another mechanism where the role of $a(A)$ and $b(B)$ are switched. This is clearly also an optimal mechanism. Now consider the mechanism in which the SP uses the previous two mechanisms, each with probability one-half. This mechanism is also optimal, and it is ex-ante symmetric with respect to $a(A)$ and $b(B)$. From now on, we restrict attention to canonical mechanisms that are symmetric with respect to $a(A)$ and $b(B)$.

If the maximum number of agents asked in a mechanism is bounded, say by N ($\in \mathbb{N}$), then one can also assume that N agents have to acquire information according to a uniform order, and the rest of them are never asked. If $K < \infty$, the maximum number of agents asked in a mechanism is obviously bounded by K. If $K = \infty$ and the maximum number of agents asked in a mechanism is unbounded, one cannot assume that agents are ordered uniformly, because there is no uniform distribution over the integers. This implies that agents cannot be treated identically by such a mechanism, and different agents will have different incentive constraints. In order to avoid this complication, in the rest of this section, and throughout Section 5, we restrict attention to the model where there are finitely many agents, that is, $K < \infty$.[9] In addition, we restrict attention to canonical mechanisms in which the SP uniformly orders the agents and asks them accordingly.

4.2. Incentive compatibility

The goal of this subsection is to explicitly characterize the incentive compatibility constraint, that is, the constraint that guarantees an agent indeed has an incentive to collect information when asked. (Another incentive compatibility constraint guarantees that an agent, upon acquiring information, will report her signal truthfully. As we will show, that constraint is trivially satisfied.) The derivation of this constraint is somewhat different from the usual incentive constraint in Bayesian mechanism design, since the agents do not have private information to start with. The idea of deriving this constraint is to compare an agent's expected payoff if she acquires information and reports it truthfully with her payoff if she just randomizes between reporting signal

[9] We shall show in the proof of Theorem 1 that, even if $K = \infty$, the maximum number of agents asked in the optimal ex-post efficient mechanism is bounded.

a and reporting signal b without actually drawing a signal. Surprisingly, the difference between these two payoffs can be expressed in a fairly simple form.

Let $S(k)$ denote the set of signal sequences that are weakly shorter than k. Let $l(s)$ denote the length of the signal sequence s. Let us fix a voting mechanism G. The mechanism may involve some randomization by the SP. Let $r(M:s)$ ($r(m:s)$) denote the probability that the SP stops the mechanism after the reported sequence s (but not before) and makes the majority (minority) decision, conditional on s being reported. (A majority (minority) decision after a sequence s means taking action α if and only if there are more (fewer) signals a in s than signal b.) Notice that the randomization of the SP is independent of the realization of the sequence. (All of these randomizations can be made ex-ante.) Let $p(s)$ denote the unconditional probability of the sequence s. Let $d(s)$ denote $\#_a(s) - \#_b(s)$. Finally, let $p_c(s, M)$ ($p_c(s, m)$) denote the probability that the reported signal sequence is s, and decision M (m) is made after s (but not before) conditional on being asked, and given that each agent acquires a signal and reports it truthfully. Hence the payoff to an agent who has been asked to collect information is

$$\sum_{s \in S(K)} \left[p_c(s, M)P\left(\left|d(s)\right|\right) + p_c(s, m)Q\left(\left|d(s)\right|\right)\right] - c. \tag{2}$$

Let $\overline{L}$ denote the expected length of a sequence of signals after which the SP stops and makes a decision, conditional on each agent acquiring information upon being asked. That is, $\overline{L}$ is the expected number of agents asked in G and

$$\overline{L} = \sum_{s \in S(K)} p(s)\left(r(M:s) + r(m:s)\right)l(s).$$

The next proposition characterizes the incentive compatibility constraint of an agent.

Proposition 2. *The agents have incentive to acquire information if and only if*

$$\frac{p-q}{4pq} \sum_{s \in S(K)} r(M:s)p(s)P\left(\left|d(s)\right|\right)\left[\left|d(s)\right| - (p-q)l(s)\right]$$
$$+ \frac{p-q}{4pq} \sum_{s \in S(K)} r(m:s)p(s)Q\left(\left|d(s)\right|\right)\left[-\left|d(s)\right| - (p-q)l(s)\right] \geqslant c\overline{L}. \tag{3}$$

Proof. Suppose that an agent is contemplating a deviation. We shall refer to this agent as Agent 1. Since the mechanism and the information structure are symmetric with respect to a and b one can assume that if Agent 1 deviates, she reports signal a without acquiring it. First, we rewrite the payoff of Agent 1 if she acquires information conditional on being asked and given that all other agents invest in information. Then, we compare it to her deviation payoff.

Fix a sequence of signals, s. Suppose that, in this sequence, there are $i + d$ signals a and i signals b (where $d \geqslant 0$). Furthermore, suppose that the action of the SP is α after the sequence s, and Agent 1 acquires information if she is asked to do so. Next, we compute $p_c(s, M)$ using Bayes' rule. Let X, Y, Z, and $\overline{Z}$ denote the following events:

- X: if the first $l(s)$ agents acquire information, then the reported sequence is s.

- Y: the SP decides that, if the reported sequence was s, he does not stop the mechanism after any subsequence of s and makes decision M after s.[10]
- Z: Agent 1 is asked to collect information.
- $\overline{Z}$: Agent 1 is ordered among the first $l(s)$ individuals.

Using these notations:

$$p_c(s,M) = P(X \cap Y : Z) = \frac{P(X \cap Y \cap Z)}{P(Z)}.$$

The unconditional probability that Agent 1 is asked, $P(Z)$, is clearly $\overline{L}/K$. Notice that

$$P(X \cap Y \cap Z) = P(X \cap Y \cap \overline{Z}).$$

The reason is that if the decision is made after an $l(s)$-long report sequence, then Agent 1 is asked if and only if she is ordered among the first $l(s)$ agents. Since X, Y, and $\overline{Z}$ are independent, $P(X \cap Y \cap \overline{Z}) = P(X)P(Y)P(\overline{Z})$. The probability that Agent 1 is ordered among the first $l(s)$ agents, $P(\overline{Z})$, is $l(s)/K = (2i + d)/K$. In addition, $P(X) = p(s)$ because each agent, including Agent 1, collects information upon being asked. Furthermore, $P(Y) = r(M : s)$ by definition. Hence,

$$p_c(s,M) = \frac{P(X)P(Y)P(\overline{Z})}{P(Z)} = \frac{2i+d}{\overline{L}} p(s)r(M:s). \tag{4}$$

Next, we compute the difference between Agent 1's payoff when she collects information and her payoff if she does not, conditional on her being asked. As we mentioned earlier, one can assume that the deviator reports a. Let $p'_c(s,M)$ ($p'_c(s,m)$) denote the probability of making a decision M (m) after a reported sequence s conditional on Agent 1 is being asked, and given that she reports a and all other agents report their acquired information. Since Agent 1 reports a without acquiring information the probabilities of the events X, Y, and Z might be different from the previous ones. Let P' denote the probabilities instead of P. We need one more piece of notation: let $\overline{Z}'$ denote the event where Agent 1's position in the order corresponds to those places where there is a signal a in the sequence s.[11] Then

$$p'_c(s,M) = P'(X \cap Y : Z) = \frac{P'(X \cap Y \cap Z)}{P'(Z)} = \frac{P'(X \cap Y \cap \overline{Z}')}{P'(Z)}.$$

The second equality follows from $X \cap Y \cap Z = X \cap Y \cap \overline{Z}'$. This is because when decision M is made after the reported sequence s, Agent 1 is asked if and only if her position in the order corresponds to one of the signals a in the sequence s. Furthermore, $P'(X \cap Y \cap \overline{Z}') = P'(Y)P'(X \cap \overline{Z}')$ because Y is independent of X and $\overline{Z}'$. By the Bayes' rule $P'(X \cap \overline{Z}') = P'(X : \overline{Z}')P'(\overline{Z}')$. Hence,

$$p'_c(s,M) = \frac{P'(Y)P'(X : \overline{Z}')P'(\overline{Z}')}{P'(Z)}.$$

[10] We emphasize that the SP makes this decision ex-ante, before any signal is reported. Also notice that the event Y does not imply the realization of s. This is because, for example, the SP might decide to take action α if the first reported signal is a, but the first reported signal can turn out to be b.

[11] For example if $s = (b,a,b,a)$, then $\overline{Z}$ denotes the event where Agent 1 is ordered either as the second or as the fourth agent.

Notice that $P'(Z) = P(Z)$ and $P'(Y) = P(Y)$ because Z and Y do not depend on the deviation. When one computes the probability that the reported sequence is s, given that Agent 1 reported a, one must take away a signal a and compute the likelihood of the remainder of the sequence. This is because one of the signals a in sequence s is due to the deviation and has nothing to do with the distribution of the signals. Hence,

$$P'(X : \overline{Z}') = \frac{1}{2}(pq)^i \left(p^{d-1} + q^{d-1}\right) = p(s)\frac{p^{d-1} + q^{d-1}}{p^d + q^d}.$$

Furthermore, $P'(\overline{Z}') = (i + d)/K$, because there are $i + d$ signals a in the sequence s. Therefore,

$$p'_c(s, M) = \frac{i + d}{\overline{L}} p(s)\frac{p^{d-1} + q^{d-1}}{p^d + q^d} r(M : s) = p_c(s, M)\frac{i + d}{2i + d}\frac{p^{d-1} + q^{d-1}}{p^d + q^d}. \tag{5}$$

It turns out to be useful to execute the same computation for the complement of s, denoted by $\bar{s}$. (‘Complement’ means that a and b are reversed. For example, if $s = (a, a, b)$, then $\bar{s} = (b, b, a)$.) The same argument as the one above[12] yields that

$$p'_c(\bar{s}, M) = \frac{i}{\overline{L}} p(\bar{s})\frac{p^{-d-1} + q^{-d-1}}{p^{-d} + q^{-d}} r(M : \bar{s}) = p_c(\bar{s}, M)\frac{i}{2i + d}\frac{p^{-d-1} + q^{-d-1}}{p^{-d} + q^{-d}}. \tag{6}$$

If s is the reported sequence of signals, decision M is made after s and Agent 1 reports a, then her expected utility from the majority decision is $P(d - 1)$. Again, this is because Agent 1 computes the posterior as if there were one fewer signal a in s. Similarly if $\bar{s}$ is the reported sequence of signals, her expected utility is $P(d + 1)$. Notice that by Lemma 1,

$$P(d - 1)\frac{p^{d-1} + q^{d-1}}{p^d + q^d} = \frac{P(d)}{p} \quad \text{and} \quad P(d + 1)\frac{p^{-d-1} + q^{-d-1}}{p^{-d} + q^{-d}} = \frac{P(d)}{q}. \tag{7}$$

Hence, the difference between the expected benefit of Agent 1 multiplied by the probabilities of making decision M after s and $\bar{s}$ conditional on being asked when she collects information and when she does not is:

$$\begin{aligned}
&p_c(s, M)P(d) + p_c(\bar{s}, M)P(d) - \left[p'_c(s, M)P(d - 1) + p'_c(\bar{s}, M)P(d + 1)\right] \\
&\quad = p_c(s, M)\left(P(d) - \frac{i + d}{2i + d}\frac{p^{d-1} + q^{d-1}}{p^d + q^d}P(d - 1)\right) \\
&\qquad + p_c(\bar{s}, M)\left(P(d) - \frac{i}{2i + d}\frac{p^{-d-1} + q^{-d-1}}{p^{-d} + q^{-d}}P(d + 1)\right) \\
&\quad = p_c(s, M)\left(P(d) - \frac{i + d}{2i + d}\frac{P(d)}{p}\right) + p_c(\bar{s}, M)\left(P(d) - \frac{i}{2i + d}\frac{P(d)}{q}\right) \\
&\quad = \frac{p_c(s, M)P(d)}{2i + d}\left(2(2i + d) - \frac{i + d}{p} - \frac{i}{q}\right) \\
&\quad = \frac{p_c(s, M)P(d)}{l(s)}\left(2l(s) - \frac{l(s) + d}{2p} - \frac{l(s) - d}{2q}\right)
\end{aligned}$$

[12] In this computation we use the formula $p(s) = (1/2)(pq)^{i+d}(p^{-d} + q^{-d})$. The probability that the deviator’s position in the ordering corresponds to a signal a in $\bar{s}$ is i/K. In addition, the probability that the sequence of reports is $\bar{s}$ is $(1/2)(pq)^{i+d}(p^{-d-1} + q^{-d-1})$. Notice that $d(\bar{s}) \leqslant 0$ if $d(s) \geqslant 0$.

$$= \frac{p_c(s,M)P(d)}{l(s)}\frac{p-q}{2pq}\big[d-(p-q)l(s)\big]$$
$$= \frac{p(s)r(M:s)P(d)}{\bar{L}}\frac{p-q}{2pq}\big[d-(p-q)l(s)\big]. \tag{8}$$

The first equality follows from (5) and (6). The second equality follows from (7). The third one follows from the assumption that the mechanism is symmetric with respect to a and b. The last equality follows from (4).

What happens if, after the realization of s and $\bar{s}$, the SP makes the minority decision? The same arguments leading to (5) and (6) now yield

$$p'_c(s,m) = p_c(s,m)\frac{i+d}{2i+d}\frac{p^{d-1}+q^{d-1}}{p^d+q^d} \quad \text{and}$$
$$p'_c(\bar{s},m) = p_c(\bar{s},M)\frac{i}{2i+d}\frac{p^{-d-1}+q^{-d-1}}{p^{-d}+q^{-d}}.$$

If s is the reported sequence of signals, decision m is made and Agent 1 reported a, then her expected utility from the minority decision is $Q(d-1)$. Again, this is because the deviator computes the posterior as if there were one fewer signal a in s. Similarly if $\bar{s}$ is the reported sequence of signals, her expected utility is $Q(d+1)$. From Lemma 1,

$$Q(d-1)\frac{p^{d-1}+q^{d-1}}{p^d+q^d} = \frac{Q(d)}{q} \quad \text{and} \quad Q(d+1)\frac{p^{-d-1}+q^{-d-1}}{p^{-d}+q^{-d}} = \frac{Q(d)}{p}.$$

Hence, similarly to (8), we get

$$p_c(s,m)Q(d) + p_c(\bar{s},m)Q(d) - \big[p'_c(s,m)Q(d-1) + p'_c(\bar{s},m)Q(d+1)\big]$$
$$= \frac{p_c(s,m)Q(d)}{l(s)}\frac{p-q}{2pq}\big[-d-(p-q)l(s)\big]$$
$$= \frac{p(s)r(m:s)Q(d)}{\bar{L}}\frac{p-q}{2pq}\big[-d-(p-q)l(s)\big]. \tag{9}$$

To compute the difference between the expected benefit of Agent 1 when she acquires information and when she just reports signal a, one has to execute the following computation. Sum up for all possible sequences of signals the last expressions in (8) and in (9). Then, divide this sum by two because in each term of the summation we have counted a sequence and its complement. A mechanism is incentive compatible if and only if this difference exceeds the cost of information acquisition. Hence, the mechanism is incentive compatible if and only if (3) is satisfied. □

So far, we have ignored the constraint that guarantees that once an agent acquires information, she indeed reports the true signal. We show, however, that this constraint is satisfied whenever (3) is satisfied.

Remark 1. Suppose that (3) is satisfied. Then the agents have an incentive to report the true signals.

Proof. Let $\bar{u}$ denote the probability of making the correct final decision, conditional on an agent being asked and reporting a signal corresponding to the true state of the world and the rest of the agents acquiring information and reporting it truthfully. Let $\underline{u}$ denote the probability of making the correct final decision, conditional on an agent being asked and reporting a signal not corresponding to the true state of the world and the rest of the agents acquiring information and reporting it truthfully. Since the probability that a signal corresponds to the true state of the world is p, the incentive compatibility constraint, (3), can be rewritten as

$$p\bar{u} + (1-p)\underline{u} - c \geqslant \frac{1}{2}(\bar{u} + \underline{u}).$$

Since $c > 0$ and $p > 1/2$ the previous inequality implies $\bar{u} > \underline{u}$, and in turn

$$p\bar{u} + (1-p)\underline{u} \geqslant p\underline{u} + (1-p)\bar{u}.$$

This is exactly the constraint that guarantees that an agent reports the true signal if she acquires one. □

4.3. Markovian mechanisms

Next, we show that we can restrict our search for the optimal mechanisms to the class of Markovian mechanisms. To this end, let us define $V(l,d)$ as follows:

$$\{s \mid l(s) = l,\ d(s) = d\}.$$

That is, $V(l,d)$ is the set of signal sequences of length l in which the difference between the number of signals a and the number of signals b is d. We refer to $V(l,d)$ as a *state*. Call a state $V(l,d)$ *feasible* if $l \leqslant K$, $|d| \leqslant l$ and $l-d$ is divisible by two. Let $\mathcal{V}(K)$ denote the set of feasible states. We can extend the definition of the functions d and l to these states. That is, if $V \in \mathcal{V}(K)$, then $d(V) = d$ and $l(V) = l$ if and only if $V = V(l,d)$.

Three decisions can be made at any state: the majority decision (M), the minority decision (m), and the decision to ask one more agent (C). The mechanism determines the probabilities over these decisions at each state. What the mechanism specifies after reaching $V(l,d)$ may depend on the history, that is, on how $V(l,d)$ is actually reached.[13] However, for each mechanism, there exists another one that operates as follows: The decision at $V(l,d)$ depends only on the probabilities determined by the original mechanism, not on the actual sequence.[14] Notice that the unconditional probability of any sequence in a certain state is the same. Hence, from (3), it is clear that if the original mechanism was incentive compatible, the new one is also incentive compatible. The value of the objective function is also the same, so one can restrict attention to those mechanisms where the decision at each state is independent of the history.

Definition 1. A mechanism is said to be Markovian if for all sequence $s \in S(K)$, the decision $D(\in \{M, m, C\})$ after s depends only on $l(s)$ and $d(s)$.

From now on, we restrict attention to Markovian mechanisms. Let G denote a mechanism. Let $p(D : V, G)$ (where $D \in \{M, m, C\}$, $V \in \mathcal{V}(K)$) denote the probability of making decision D at

[13] Suppose, for example, that after the sequence (b, a, a) the SP asks an agent, but after the sequence (a, b, a) she makes the majority decision. Notice, that $(b, a, a), (a, b, a) \in V(3, 1)$.

[14] That is, the probability of making the majority decision at state $V(l,d)$ is

$$\sum_{s \in V(l,d)} p(s) r(M : s) \Big/ \sum_{s \in V(l,d)} p(s).$$

state V, conditional on G and being at state V. A Markovian mechanism G can be identified by the collection of these probabilities, that is, by

$$\{p(D:V,G) \mid D \in \{M,m,C\},\ V \in \mathcal{V}(K)\}. \tag{10}$$

Notice that the cardinality of $\mathcal{V}(K)$ is less than K^2, and $\sum_{D\in\{M,m,C\}} p(D:V(l,d),G) = 1$. Therefore, a mechanism can be defined as a point in a compact subset of $[0,1]^{3K^2}$. Next, we show that both (1) and (3) can be rewritten in terms of $p(D:V,G)$'s and they are continuous in these probabilities. Hence, a solution to our optimal mechanism design problem exists.

Define $\rho(V,D:G)$ ($V \in \mathcal{V}(K)$, $D \in \{M,m,C\}$) as the probability of reaching state V and making a decision D given mechanism G. Furthermore, let $p(V:G)$ denote the probability of reaching V conditional on G. Then clearly $p(D:V,G) = \rho(V,D:G)/p(V:G)$.

With this notation, we can rewrite the maximization problem of the SP and (3) as follows:

$$\max_G K_0 \sum_{V\in\mathcal{V}(K)} \big[\rho(V,M:G)P\big(|d(V)|\big) + \rho(V,m:G)Q\big(|d(V)|\big)\big] - c\overline{L}(G), \tag{11}$$

subject to the incentive compatibility constraint

$$\begin{aligned}&\frac{p-q}{4pq}\sum_{V\in\mathcal{V}(K)} \rho(V,M:G)P\big(|d(V)|\big)\big[|d(V)| - (p-q)l(V)\big]\\ &\quad + \frac{p-q}{4pq}\sum_{V\in\mathcal{V}(K)} \rho(V,m:G)Q\big(|d(V)|\big)\big[-|d(V)| - (p-q)l(V)\big]\\ &\geqslant c\overline{L}(G),\end{aligned} \tag{12}$$

where $\overline{L}(G)$ is the expected number of agents asked in mechanism G. The left side of the last inequality is basically the probability of some agent being pivotal times $\overline{L}(G)$.

Suppose that a mechanism G is a probability mixture of some incentive compatible mechanisms $G_1,\dots,G_n$. That is, the SP uses mechanism G_i with probability p_i. (The agents only know p_i's, but not the realization of the randomization.) Then

$$\rho(V,D:G) = \sum_{i=1}^{n} p_i \rho(V,D:G_i),$$

and clearly

$$\overline{L}(G) = \sum_{i=1}^{n} p_i \overline{L}(G_i).$$

Notice that both (11) and (12) are linear in $\rho(V,D:G)$. Hence we can claim the following:

Remark 2. The probability mixture of incentive compatible mechanisms is also incentive compatible. Furthermore, the value of the objective function is the probability mixture of the values of the objective functions corresponding to the mechanisms used in the mixture.

Next, we show that in optimal mechanisms, the agents are indifferent between collecting information and making a random report.

Lemma 3. *Suppose that the first-best mechanism is not incentive compatible and G^* is either the ex-ante optimal mechanism or the ex-ante optimal mechanism among the ex-post efficient ones. Then* (12) *holds with equality.*

Proof. Suppose, on the contrary, that (12) holds with strict inequality. Consider the following mechanism: employ the first-best mechanism with probability ε (> 0) and mechanism G^* with probability $1-\varepsilon$. Clearly, this new mechanism increases the value of social welfare. Furthermore, (12) is continuous in the probabilities $\rho(V, D : G)$ ($D \in \{M, m\}$), and these probabilities are continuous in ε. Hence, if ε is small enough, the new mechanism is incentive compatible. This contradicts to the optimality of G^*. □

5. Continuation mechanisms

The arguments of most proofs regarding the optimality of mechanisms involve modifying the mechanism at some states. We introduce the notion of a continuation mechanism, which is the part of a mechanism that follows after reaching a certain state. In the first-best mechanism, when the SP had to decide whether to ask an additional agent or to make a decision, he only had to compare the cost of a continuation mechanism from a particular state with the expected increase in precision due to the continuation mechanism. It did not matter how the mechanism would operate when starting from another state. Such a simple argument cannot be used when the mechanism is subject to the incentive compatibility constraint, because there is an interaction between continuation mechanisms through this constraint. Intuitively, if at a certain state, the SP decides to ask more agents to invest in order to make a more informed final decision, each agent becomes less likely to be pivotal. Therefore, the SP might be forced to stop asking agents in another state in order to satisfy the incentive constraints of the agents. This means that employing a continuation mechanism at a certain state can make it impossible to employ another continuation mechanism at a different state. Hence, finding an optimal incentive compatible voting scheme is a stochastic dynamic problem with a non-recursive structure.

Recall, the reason the function defining the first-best scheme was decreasing (see part (ii) of Proposition 1) is that the larger the number of agents have already been asked, the less valuable it is to ask an additional agent given a certain posterior. In what follows, we establish a similar monotonicity property of continuation mechanisms, even when the schemes are subject to the incentive compatibility constraint, as long as we also require that final decisions are ex-post efficient. To that end, we define the efficiency of continuation mechanisms. Then, we show that the efficiency of a given continuation mechanism decreases in the number of agents already asked. In addition, we prove that in an optimal mechanism, if a certain continuation mechanism is employed with positive probability, any other continuation mechanism that is more efficient and feasible must also be employed.

A continuation mechanism at state $V(l, d)$ is just a mechanism G defined by (10) specifying the decisions of the SP after reaching $V(l, d)$. Let $G(l, d)$ denote the continuation mechanism G at $V(l, d)$. Similarly to $\rho(V, D : G)$, one can define $\rho^d(V, D : G(l, d))$ as the probability of reaching $V(l + l(V), d + d(V))$ from $V(l, d)$ and making a decision D given the continuation mechanism G. However, a majority (minority) decision, M (m), means taking action α if and only if $d(V) + d > (<) 0$. A continuation mechanism G at $V(l, d)$ is said to be ex-post efficient if

$$\rho^d\big(V, m : G(l, d)\big) = 0 \quad \text{for all } V \in \mathcal{V}(K).$$

That is, the minority decision is never made. Notice that whether a continuation mechanism is ex-post efficient depends on the state V at which it is employed. However, it depends only on $d(V)$ and not on $l(V)$. Hence, an ex-post efficient continuation mechanism at $V(l,d)$ is necessarily asymmetric with respect to $a(A)$ and $b(B)$ unless $d=0$.

For example of a continuation mechanism, consider G defined by

$$p\big(M : V(1,1), G\big) = p\big(M : V(1,-1), G\big) = 1.$$

That is, G specifies action α after signal a and action β after signal b. Suppose that G is used at $V(3,3)$. Then,

$$\rho^3\big(V(1,1), M : G(3,3)\big) = p(3), \qquad \rho^3\big(V(1,-1), m : G(3,3)\big) = q(3),$$

and all the other probabilities are zero. Observe that the continuation mechanism $G(3,3)$ is not ex-post efficient.

Definition 2. A continuation mechanism G is feasible at $V(l,d)$ with respect to a mechanism G' if (i) $p(V(l,d) : G') > 0$, and (ii) if $p(V' : G) > 0$ then $l(V') \leqslant K - l$.

The previous definition states that given a mechanism G', a continuation mechanism is feasible at $V(l,d)$ if and only if $V(l,d)$ is reached with positive probability according to G'. Furthermore, the continuation mechanism G never specifies that the SP asks more agents than the number available after reaching $V(l,d)$, that is $K-l$.

Let $\Delta W(G(l,d))$ denote the change in the objective function of the SP if, instead of making the majority decision at $V(l,d)$, the mechanism continues according to G conditional on being at $V(l,d)$. That is,

$$\begin{aligned}\Delta W\big(G(l,d)\big) = K_0 &\sum_{V\in\mathcal{V}(K)} \rho^d(V, M : G) P\big(\big|d(V)+d\big|\big)\\ &+ K_0 \sum_{V\in\mathcal{V}(K)} \rho^d(V, m : G) Q\big(\big|d(V)+d\big|\big)\\ &- c\overline{L}_d(G) - K_0 P\big(|d|\big),\end{aligned} \tag{13}$$

where $\overline{L}_d(G)$ denotes the expected number of agents asked according to G if G is employed at $V(l,d)$. Furthermore, define $\Delta IC(G(l,d))$ as

$$\begin{aligned}&\frac{p-q}{4pq} \sum_{V\in\mathcal{V}(K)} \rho^d(V, M : G) P\big(\big|d(V)+d\big|\big)\big[\big|d(V)+d\big| - (p-q)\big(l(V)+l\big)\big]\\ &\quad+ \frac{p-q}{4pq} \sum_{V\in\mathcal{V}(K)} \rho^d(V, m : G) Q\big(\big|d(V)+d\big|\big)\big[-\big|d(V)+d\big| - (p-q)\big(l(V)+l\big)\big]\\ &\quad- c\overline{L}_d(G) - \frac{p-q}{4pq} P\big(|d|\big)\big[|d| - (p-q)l\big].\end{aligned} \tag{14}$$

Roughly, $\Delta IC(G(l,d))$ is the change in the incentive compatibility constraint if the SP continues asking agents according to G instead of making the majority decision at $V(l,d)$ conditional on being at $V(l,d)$.

Suppose that for a mechanism G', $\rho(V(l,d), M : G') > 0$. Then, if G' is modified such that the SP continues asking additional agents according to G at $V(l,d)$ instead of making the majority decision, the value of the SPs objective function increases by $\rho(V(l,d),$

$M : G')\Delta W(G(l,d))$. Furthermore, the new incentive compatibility constraint is almost the same as the one corresponding to G', but $\rho(V(l,d), M : G')\Delta IC(G(l,d))$ must be added to the left side of (12). Notice that neither $\Delta W(G(l,d))$ nor $\Delta IC(G(l,d))$ depends on G'. This is why one can evaluate the efficiency of a continuation mechanism independent of G'.

Definition 3. Let G be a mechanism. The efficiency of G at $V(l,d)$ is defined as

$$e\big(G(l,d)\big) = \begin{cases} \frac{\Delta W(G(l,d))}{|\Delta IC(G(l,d))|} & \text{if } \Delta W(G(l,d))\Delta IC(G(l,d)) < 0, \\ \infty & \text{if } \Delta W(G(l,d)), \Delta IC(G(l,d)) \geqslant 0, \\ -\infty & \text{if } \Delta W(G(l,d)), \Delta IC(G(l,d)) \leqslant 0. \end{cases}$$

The efficiency of a continuation mechanism essentially specifies how much the objective function of the SP increases if the incentive compatibility constraint changes by one unit. There are continuation mechanisms for which $\Delta W(G(l,d))$ as well as $\Delta IC(G(l,d))$ are negative, but those continuation mechanisms are obviously never employed in optimal mechanisms. Their efficiencies are defined to be $-\infty$. There may also exist continuation mechanisms for which $\Delta W(G(l,d))$ and $\Delta IC(G(l,d))$ are positive, but then, at $V(l,d)$ the SP always continues to ask more agents. The efficiency of such a continuation mechanism is set to be ∞. The subjects of our interest are those continuation mechanisms for which $\Delta W(G(l,d))$ is positive but $\Delta IC(G(l,d))$ is negative, or vice versa. That is, employing G at $V(l,d)$ increases the value of the SPs objective function at the cost of incentive compatibility, or vice versa.

Lemma 4. *Let G be a feasible ex-post efficient continuation mechanism at $V(l,d)$. Then*

$$\sum_{V \in \mathcal{V}(K)} \rho^d\big(V, M : G(l,d)\big) P\big(\big|d(V)+d\big|\big) \geqslant P\big(|d|\big).$$

The right side of the inequality is the SPs posterior at $V(l,d)$ about the true state of the world. The left side is the expected posterior if continuation mechanism G is used. The lemma states that the expected posterior at V if an ex-post efficient continuation mechanism is employed at V is at least as precise as the posterior at V. Clearly, more information cannot result in a less accurate expected posterior.

Proof. See the Online Appendix. □

The next lemma shows that $G(l',d)$ is more efficient than $G(l,d)$ whenever $l' < l$.

Lemma 5. *Let G be a mechanism such that G is ex-post efficient at V if $d(V) = d$. Suppose that G is feasible at $V(l,d)$ as well as at $V(l',d)$. If $l > l'$, then*

$$e\big(G(l,d)\big) \leqslant e\big(G(l',d)\big),$$

and the inequality is strict whenever $|e(G(l',d))| \neq \infty$.

This lemma is the key to characterizing the optimal ex-post efficient mechanism. It establishes a monotonicity property of the optimal mechanism; given a certain posterior, the longer the sequence, the less efficient it is to ask more agents instead of making the final decision. The intuition behind this lemma is the following. If the SP continues to ask more agents at a certain state, the agents have less incentive to collect information because it becomes more likely that other agents correct the noise introduced by a deviation. The lemma says that given a certain posterior, this effect on the incentive compatibility is less severe if the SP asks additional agents after a short sequence and more severe after a long sequence. The reason is the following. When an agent computes the probability of a sequence s conditional on being asked, she multiplies the unconditional probability of the sequence by $l(s)/\overline{L}(G)$. Therefore, long sequences enter into the incentive compatibility constraints with larger weights relative to the unconditional probabilities and the short sequences enter with smaller weights. Since the actual probability of a long sequence is small the effect of such a sequence on the objective function is small.

Proof. Because of Lemma 1, the probability of realizing a certain sequence conditional on being at $V(l,d)$ is the same as that conditional on being at $V(l',d)$. Hence, $\rho^d(V,D:G(l,d)) = \rho^d(V,D:G(l',d))$ for all $V \in \mathcal{V}(K)$ and $D \in \{M,m,C\}$. From (13), it follows that

$$\Delta W\big(G(l,d)\big) = \Delta W\big(G(l',d)\big).$$

Furthermore, from (14),

$$\Delta IC\big(G(l,d)\big) - \Delta IC\big(G(l',d)\big) = -(l-l')\frac{(p-q)^2}{4pq}\left[\sum_{V\in\mathcal{V}(K)} \rho^d\big(V,M:G(l,d)\big)P\big(|d(V)+d|\big) - P\big(|d(V)|\big)\right].$$

From the previous lemma, we know that the term in square brackets is positive. Since $l > l'$ this expression is negative. Hence $e(G(l,d)) \leqslant e(G(l',d))$. Furthermore, whenever $|e(G(l',d))| \neq \infty$ the inequality is obviously strict. □

Given a mechanism G' and a feasible continuation mechanism G at $V(l,d)$ with respect to G', let $P(G(l,d) \mid G')$ denote the probability of using continuation mechanism G according to G' conditional on reaching $V(l,d)$.

Lemma 6. *Let G^* be an optimal mechanism among the ex-post efficient ones. Suppose that $G(l,d)$ and $G'(l',d')$ are feasible ex-post efficient continuation mechanisms with respect to G^*. Suppose that $P(G'(l',d') \mid G^*) > 0$.*

(i) *If $e(G(l,d))e(G'(l',d')) > 0$ and $e(G(l,d)) > e(G'(l',d'))$, then $p(V(l,d),C,G^*) = 1$.*
(ii) *If $e(G(l,d))e(G'(l',d')) < 0$, $P(G(l,d)|G^*) \in (0,1)$, and $P(G'(l',d') \mid G^*) < 1$, then $|e(G(l,d))| = |e(G'(l',d'))|$.*

Part (i) of this lemma essentially states that if G^* is an optimal ex-post efficient mechanism, then it never happens that the mechanism stops at $V(l,d)$ with positive probability if there exists a feasible continuation mechanism at $V(l,d)$ that is more efficient than another continuation mechanism that is used in G^* with positive probability.

Proof. We only prove part (i) of the lemma for the case when

$$\Delta IC\big(G(l,d)\big) < 0 \quad \text{and} \quad \Delta W\big(G(l,d)\big) > 0.$$

The proofs of the other cases and part (ii) are similar to this one.

Suppose that the hypothesis of the lemma is satisfied but $p(V(l,d),C,G^*) < 1$. Let us consider the following mechanism, denoted by G^{**}. G^{**} is almost identical to G^*, only it is modified at $V(l,d)$ and $V(l',d')$. At $V(l,d)$, it uses the continuation mechanism G with probability ε (>0) instead of making the majority decision, and at $V(l',d')$ it stops with probability

$$\varepsilon \frac{\Delta IC(G(l,d))p(V(l,d),G^*)}{\Delta IC(G'(l',d'))p(V(l',d'),G^*)}$$

and makes the majority decision instead of using G' as a continuation mechanism. If

$$\varepsilon < \min\left\{p\big(M:V(l,d),G^*\big),\, P\big(G'(l',d') \mid G^*\big)\frac{\Delta IC(G(l',d'))p(V(l',d'):G^*)}{\Delta IC(G(l,d))p(V(l,d):G^*)}\right\},$$

these changes are indeed feasible and the new mechanism is well defined. (Since $p(V(l,d), C, G^*) < 1$, ε can be chosen to be bigger than zero.) The change in the incentive compatibility constraint is

$$\begin{aligned} &p\big(V(l,d):G^*\big)\varepsilon\Delta IC\big(G(l,d)\big) \\ &\quad - p\big(V(l',d'):G^*\big)\varepsilon\frac{\Delta IC(G(l,d))p(V(l,d):G^*)}{\Delta IC(G'(l',d'))p(V(l',d'):G^*)}\Delta IC\big(G'(l',d')\big) \\ &= 0. \end{aligned}$$

Hence, since the mechanism G^* is incentive compatible, the new mechanism is also incentive compatible. Furthermore, the change in the objective function of the SP is

$$\begin{aligned} &p\big(V(l,d):G^*\big)\varepsilon\Delta W\big(G(l,d)\big) \\ &\quad - p\big(V(l',d'):G^*\big)\varepsilon\frac{\Delta IC(G(l,d))p(V(l,d):G^*)}{\Delta IC(G'(l',d'))p(V(l',d'):G^*)}\Delta W\big(G'(l',d')\big) \\ &= p\big(V(l,d):G^*\big)\varepsilon e\big(G(l,d)\big)\big|IC\big(G(l,d)\big)\big| \\ &\quad - p\big(V(l',d'):G^*\big)\varepsilon\frac{|\Delta IC(G(l,d))|p(V(l,d):G^*)}{p(V(l',d'):G^*)}e\big(G'(l',d')\big) \\ &= p\big(V(l,d):G^*\big)\varepsilon\big|IC\big(G(l,d)\big)\big|\big(e\big(G(l,d)\big) - e\big(G'(l',d')\big)\big). \end{aligned}$$

Since $e(G(l,d)) > e(G'(l',d'))$, the new mechanism increases the value of the objective function, contradicting the optimality of G^*. □

Notice that the only reason $G(l,d)$ and $G'(l',d')$ were required to be ex-post efficient is that G^* was the optimal mechanism among the ex-post efficient ones. The proof of the lemma has not used the ex-post efficiency property of the continuation mechanisms. Hence, we can claim a similar result for the optimal (not necessarily ex-post efficient) mechanism.

Remark 3. Let G^* be an optimal mechanism. Suppose that $G(l,d)$ and $G'(l',d')$ are feasible continuation mechanisms with respect to G^*. Suppose that $e(G(l,d))e(G'(l',d')) > 0$, $e(G(l,d)) > e(G'(l',d'))$ and $P(G'(l',d') \mid G^*) > 0$. Then $p(C:V(l,d),G^*) = 1$.

Lemma 7. *Suppose that G^* is an optimal mechanism among the ex-post efficient ones, $P(G(l,d)|G^*) > 0$, and $d \neq 0$. Then there exists a state $V' \in \mathcal{V}(K-l)$ such that $d(V')d < 0$ and*

$$\rho^d\big(V', M: G(l,d)\big) > 0.$$

The condition $d(V')d < 0$ means that the majority decision at V' is different from the majority decision at $V(l,d)$. Hence, the previous lemma says that if a continuation mechanism in an optimal mechanism is employed at state $V(l,d)$, then it reaches and specifies the majority decision in at least one state where the decision differs from the majority decision at $V(l,d)$. This means that an agent is never asked to collect information if her report will not affect the final decision.

Proof. See the Online Appendix. □

Lemma 8. *There exists an optimal mechanism G^* among the ex-post efficient ones such that the SP either does not randomize or there is a pair $(\widehat{l},\widehat{d})$ such that he randomizes only at states $V(\widehat{l},\widehat{d})$ and $V(\widehat{l},-\widehat{d})$. Furthermore, generically, if the first-best mechanism is not incentive compatible, the optimal mechanism G^* among the ex-post efficient ones must involve randomization.*

Generically means that the Lebesgue measure of those (p,c) pairs for which the statement of the lemma is false is zero.

Proof. See the Online Appendix. □

6. Optimal mechanisms

In this section, we first characterize the optimal ex-post efficient mechanism. Then we discuss some attributes of this mechanism. Finally, we show that the ex-ante optimal mechanism sometimes does and sometimes does not involve ex-post inefficient decisions.

6.1. Optimal ex-post efficient mechanism

We are ready to characterize the optimal mechanism in the class of ex-post efficient mechanisms, that is, the class of mechanisms where the SP always makes a majority decision.

Theorem 1. *Suppose that the first-best mechanism does not satisfy the incentive compatibility constraint,* (12). *Let $G^\star$ be an ex-ante optimal mechanism among the ex-post efficient ones. Then, there exists a decreasing step function $f:\mathbb{N}\to\mathbb{N}$, and $N\in\mathbb{N}$, such that*

$$\text{for all } l\in\mathbb{N},\ f(l+1)=f(l)\quad\text{or}\quad f(l+1)=f(l)-1,\quad\text{and}\quad f(N)=1.$$

Let

$$T=\bigl\{V: f\bigl(l(V)\bigr)=\bigl|d(V)\bigr|,\ f\bigl(l(V)-1\bigr)=f\bigl(l(V)\bigr)+1\bigr\}.$$

$G^\star$ is defined by the following three conditions:

(i) *if $V\notin T$ and $f(l(V))\leqslant |d(V)|$ then $p(M:V,G^\star)=1$,*
(ii) *if $V\notin T$ and $f(l(V))>|d(V)|$ then $p(C:V,G^\star)=1$,*
(iii) *if $V\in T$ then $p(M:V,G^\star)\geqslant 0$, $p(C:V,G^\star)>0$.*

Furthermore, generically, there exists an optimal ex-post efficient mechanism for which there are only two states, $V_1,V_2\in T$, such that $p(M:V_i,G^\star)>0$ for $i=1,2$. In addition, if $K<\infty$, then $N=K$.

Theorem 1 essentially claims that an optimal ex-post efficient mechanism can be described by the decreasing step function f. The SP keeps asking the agents sequentially to collect information and report it to him. Once a state $V(l,d)$ is reached where $f(l)=|d|$, the SP makes the majority decision. Since the function f is decreasing, the more agents the SP has already asked, the less precise is the required posterior to induce the SP to stop asking agents and take an action. Furthermore, the function f never jumps down by more than one steps. Since, generically, in such a deterministic mechanism the incentive compatibility constraint holds with inequality, the SP randomizes at certain states. Property (iii) says randomization can happen at a state $V(l,d)$ only if $f(l)=|d|$ and the function f jumps down at $l-1$.

We emphasize that Theorem 1 characterizes the optimal mechanism even when $K=\infty$. In this case, the contrast between the first-best and second-best mechanisms is particularly sharp. While the first-best mechanism is characterized by a constant, the function f characterizing the second-best mechanism is a decreasing step function. If $K=\infty$, the feasibility of a continuation mechanism at a state V does not depend on the number of agents who have already acquired information, $l(V)$, as long as V is reached with positive probability. Hence, the reason the function f is decreasing cannot be that the number of available continuation mechanisms decreases as the number of agents asked increases. It is only the incentive compatibility constraint that forces f to be decreasing.

Indeed, the way to guarantee a large probability of being pivotal to the agents is to make decisions after signal sequences where the difference between the numbers of different signals is small. That is, the SP makes a decision after sequences that generate imprecise posteriors. We claim that providing a large probability of being pivotal is more efficient after long sequences than after short sequences. The intuition is as follows. Conditional on being asked, a voter believes that she is more likely to have been asked along a long sequence rather than a short one. This is because when a voter computes the probability of a sequence s conditional on being asked, she multiplies the unconditional probability of s by $l(s)/\overline{L}(G)$. The actual probability of a long sequence is small, and hence the effect of such a sequence on the social welfare is small. However, the conditional probability of such a sequence is large, because $l(s)$ is large. Hence, the effect on the incentive compatibility constraint is also large. For this reason the function f is decreasing. That is, after a short sequence the Social Planner makes a decision only if his posterior is precise, but after a long sequence he makes a decision even if his posterior is imprecise.

Sketch of the proof. The proof of Theorem 1 is relegated to Appendix A. Here, however, we present the main argument for $K<\infty$. Let $\overline{d}$ denote the largest d for which there exists an l such that $p(V(l,d):G^*)>0$. For each d $(\in\{1,\ldots,\overline{d}-1\})$, let $\phi(d)$ denote the largest possible l for which $p(C:V(l,d),G^*)$ is positive. (It can be shown that the function ϕ is well defined.) We claim that, conditional on reaching $V(l,d)$, the mechanism continues at $V(l,d)$ if and only if $l<\phi(d)$. If $l<\phi(d)$, then the continuation mechanism at $V(\phi(d),d)$ is feasible at $V(l,d)$, too. By Lemma 5, the continuation mechanism used at $V(\phi(d),d)$ is more efficient at $V(l,d)$. Hence, from Lemma 6, it is impossible that the mechanism stops at $V(l,d)$ conditional on reaching $V(l,d)$. If $\phi(d)<l$, then the mechanism stops at $V(l,d)$ by the definition of ϕ. We show that ϕ is a strictly decreasing function. Suppose not, and $d>d'$ and $\phi(d)\geqslant\phi(d')$. This means that the mechanism continues at $V(\phi(d),d)$ with positive probability. Notice, however, that after continuing from $V(\phi(d),d)$, the mechanism must stop whenever the difference between the number of different signals is d'. (This is because $\phi(d')<\phi(d)$.) Therefore, the decision of the SP after reaching $V(\phi(d),d)$ is always the majority decision at this state, no matter what the agents report in the continuation mechanism. This contradicts Lemma 7. It can be shown that the mechanism continues at state $V(l,d)$ if and only if $l \leqslant \phi(d)$. Furthermore, since ϕ is decreasing, $p(M : V(l,d), G^*)$ is positive if and only if $\phi(d-1) \geqslant l-1$ and $\phi(d) < l$. The first inequality guarantees that the mechanism continues at $V(l-1,d-1)$; hence $V(l,d)$ can indeed be reached. The second inequality guarantees that the mechanism does not continue at $V(l,d)$. The function f in Theorem 1 can be defined as follows:

$$f(l) = \min_{d \in \{1,\dots,\bar{d}-1\}} \{ d \mid \phi(d) \leqslant l \}.$$

Notice that $f(\phi(d)) = d$, so one can think of the function f as the inverse of ϕ. Since ϕ is strictly decreasing, the function f is also decreasing. It is easy to show that the rest of the properties of f described in the theorem hold. As we pointed out earlier, from Lemma 3, it follows that the optimal mechanism generically involves randomization. From Lemmas 6 and 7, it follows that randomization can happen only at states in T. From Lemma 8, it follows that there generically exists an optimal mechanism in which randomization happens at two states. □

Next, we claim that the optimal voting scheme described in Theorem 1 is essentially unique.

Theorem 2. *The optimal ex-post efficient mechanism described in Theorem* 1 *is generically unique and involves randomization at two states.*

Proof. See the Online Appendix. □

The argument of the proof of Theorem 2 is as follows. First, we show that if, for a certain pair (p,c), there are at least two different optimal mechanisms, then there exists an optimal mechanism that involves randomizations in at least two different states, V_1 and V_2, such that $l(V_1) \neq l(V_2)$. By Lemma 6, it follows that there are at least two deterministic continuation mechanisms that have the same efficiency. Since there are only finitely many deterministic continuation mechanisms, if the optimal mechanism were not generically unique, there would exist two continuation mechanisms with the same efficiency for a positive measure of (p,c). In the proof of Theorem 2, we show that this is impossible.

6.2. Ex-ante optimal mechanisms

In this section, we prove two results regarding the ex-post efficiency of the ex-ante optimal mechanisms. First, we show that the optimal mechanism sometimes involves ex-post inefficient decisions. Second, we characterize a set of parameter values for which the ex-ante optimal mechanism never makes ex-post inefficient decisions.

The first result is proved by contradiction. It will be shown that if the cost of information acquisition is small enough, then the optimal ex-post efficient mechanism can be improved upon by replacing a continuation mechanism with an ex-post inefficient continuation mechanism.

Next, we define two continuation mechanisms, $G_1(N,1)$ and $G_2(N,1)$. Employing G_1 at $V(N,1)$ means the following. The SP asks an additional agent. If the agent confirms the posterior of the SP, he makes the majority decision. If not, the SP asks an additional agent and makes the decision corresponding to the report of the last agent. In contrast, employing G_2 at $V(N,1)$ means the following. At $V(N,1)$, the SP asks an agent, and if she confirms his posterior, he makes the majority decision. If not, the SP asks two more agents. If their reports are the same, he again makes the majority decision. However, if the two agents report different signals, the SP asks an additional agent and makes the minority decision. In the following lemma we show that if c is small enough, then G_2 is more efficient than G_1 at $V(N,1)$.

Lemma 9. *There exist a $\widetilde{c}$ (>0) and $\overline{N} \in \mathbb{N}$ such that for all $N \geqslant \overline{N}$ and $c \in (0, \widetilde{c})$,*

$$e\big(G_2(N,1)\big) > e\big(G_1(N,1)\big).$$

Proof. See the Online Appendix. □

We have to show that for some values of c, the optimal ex-post efficient mechanism indeed uses $G_1(N,1)$ with probability strictly between zero and one for some N.

Lemma 10. *For all $\widehat{c}$ (>0) and $\overline{N} \in \mathbb{N}$ there exist $K, N \in \mathbb{N}$, $c < \widehat{c}$ such that $P(G_1(N,1) \mid G^*) \in (0,1)$ and $N \in \{\overline{N}, \ldots, K-4\}$.*

Proof. See the Online Appendix. □

Now we are ready to state our third theorem.

Theorem 3. *For all $\widehat{c}$ (>0) there exist $K \in \mathbb{N}$, $c < \widehat{c}$ such that the ex-ante optimal mechanism involves ex-post inefficient decisions.*

Proof. By Lemmas 9 and 10, for all $\widehat{c}$ (>0) and $\overline{N} \in \mathbb{N}$ there exist $K, N \in \mathbb{N}$, $c < \widehat{c}$ such that $P(G_1(N,1) \mid G^*) \in (0,1)$, $e(G_2(N,1)) > e(G_1(N,1))$ and $N \in \{\overline{N}, \ldots, K-4\}$. In addition, G_2 is feasible at $V(N,1)$ because $N \leqslant K-4$. Hence, by Remark 3, the optimal mechanism among the ex-post efficient ones, G^*, is not optimal. Therefore, the optimal mechanism involves ex-post inefficient decisions. □

Next, a set of parameter values is characterized for which the ex-ante optimal mechanism is always ex-post efficient. The proof of this theorem is based on the following observation. For a fixed K, if the precision of the signal, p, is small enough then in any mechanism switching a minority decision to a majority decision makes the mechanism more incentive compatible. Since these switches also increase social welfare, the ex-ante optimal mechanism cannot involve minority decisions.

Theorem 4. *If $(2p-1) \leqslant 1/\sqrt{K}$, then in the ex-ante optimal mechanism only ex-post efficient decisions are made.*

Proof. Suppose that at state V the optimal mechanism G^* specifies the ex-post inefficient decision with positive probability. This implies that on the left-hand side of the incentive compatibility constraint, (12), there is a term in the following form:

$$\rho(V, m : G^*) Q\big(\left|d(V)\right|\big)\big[-\left|d(V)\right| - (p-q)l(V)\big],$$

where $\rho(V, m : G^*)$ is strictly positive. We show that turning the ex-post inefficient decision into an ex-post efficient one increases the incentives of the agents to collect information. (This change obviously increases social welfare.) If we do so, then instead of the previous term the following expression appears in the incentive compatibility constraint:

$$\rho(V, m : G^*) P\big(\left|d(V)\right|\big)\big[\left|d(V)\right| - (p-q)l(V)\big].$$

We have to show that

$$Q\big(\big|d(V)\big|\big)\big[-\big|d(V)\big|-(p-q)l(V)\big]\leqslant P\big(\big|d(V)\big|\big)\big[\big|d(V)\big|-(p-q)l(V)\big].$$

Recall that $Q(|d(V)|)=1-P(|d(V)|)$. Hence, the previous inequality can be rewritten as

$$\big(2P\big(\big|d(V)\big|\big)-1\big)(p-q)l(V)\leqslant\big|d(V)\big|.$$

Equivalently

$$l(V)\leqslant\frac{|d(V)|}{(2P(|d(V)|)-1)(p-q)}. \tag{15}$$

Notice however that

$$\frac{1}{(p-q)^2}\leqslant\frac{|d(V)|}{(2P(|d(V)|)-1)(p-q)}.$$

Since $l(V)$ must be smaller than K, and $K\leqslant 1/(p-q)^2$, the inequality (15) is satisfied. □

7. Discussion

This paper analyzed optimal voting schemes in environments where information acquisition is costly and unobservable. Theorem 1 characterized the optimal ex-post efficient voting scheme. We showed that it can essentially be defined by a decreasing step function. For each number of agents already asked, this function assigns a value of the posterior. The Social Planner stops asking agents only if his posterior is more precise than the value of this function corresponding to the number of agents already asked. Since the function is decreasing, this decision rule means that when more agents have already been asked, a less precise posterior will induce the Social Planner to take a final action instead of acquiring more information. On the one hand, the Social Planner wants to act only if his posterior is sufficiently precise. On the other hand, to give agents enough incentive to acquire information, he must make sure each agent is sufficiently likely to be decisive. An agent is more likely to be pivotal if the Social Planner's posterior is imprecise when the Social Planner makes a decision. Theorem 1 says that the Social Planner should make the agents pivotal along relatively long sequences. This is because an agent, conditional on being asked, assigns a high probability to long sequences relative to the actual probability of those sequences. This intuition does not involve any assumption about the distribution of states of the world and signals. Hence, we conjecture that the statement of Theorem 1 is valid in one form or another under virtually any information structure.

What was essential in our computations is that the posterior of the SP was a function only of the number of reported signals and the difference between the numbers of different signals. It does not seem to be important that the posterior actually does not depend on the number of signals. This is because these two variables define the state space. Suppose, for example, that the range of signals is $\{a,b,\emptyset\}$ instead of $\{a,b\}$, $p(a|A)=p(b|B)$, and $p(\emptyset|A)=p(\emptyset|B)$. That is, an agent may observe a signal $\emptyset$ and be unable to update her prior about the state of the world. The posterior of the SP remains a function only of the difference between the number of signals a and b. We conjecture that the claim of Theorem 1 is still valid. The only difference is that all states of the form $V(l,d)$, $|d|\leqslant l$, must be considered and not only those where $l-d$ is an even number.

What happens if the prior about the state of the world is asymmetrically distributed? Notice, that if for example $p(A)>1/2$, then there is an initial bias toward action α. In this case, one cannot restrict attention to mechanisms that are symmetric with respect to the reported signals. Therefore, one cannot assume that if an agent decides not to acquire information, she is indifferent between reporting either of the signals. There will be two incentive compatibility constraints, corresponding to each signal. It can be shown that our main result is robust to local perturbation in the prior. That is, for all (p, c, K) there exists an $\varepsilon > 0$, such that if $|p(A) - 1/2| < \varepsilon$, then the statement of Theorem 1 is essentially the same. Furthermore, our approach turns out to be useful even if in the optimal mechanism an agent who has not acquired information strictly prefers to report one of the signals, say signal a. Then, it can be shown that the decision α can again be characterized by a decreasing step function.[15]

Nonetheless, we view our model's assumptions about the information structure as restrictive. Only the particular assumed distributions enabled us to explicitly characterize the incentive compatibility constraint. Recall that having the explicit form of the incentive compatibility constraint made it possible to compute the efficiency of continuation mechanisms. These computations played a major role in the proof of Theorem 3, which states that if the cost of information acquisition is small enough, the ex-ante optimal voting mechanism sometimes necessarily involves ex-post inefficient decisions. The fact that ex-post inefficient decisions can increase ex-ante efficiency in voting models was also reported in [5,12,13]. We find this result surprising. It says that the Social Planner can threaten the agents by committing to make ex-post inefficient decisions. This threat induces the agents to acquire information to avoid inefficient decisions.

We have also characterized a set of parameters where in the ex-ante optimal mechanism ex-post inefficient decisions are never made (see Theorem 4). However, if the ex-ante optimal mechanism involves ex-post inefficient decisions, hardly anything is known about the optimal scheme. Nonetheless, we view the problem of identifying optimal mechanisms that are ex-post inefficient as rather theoretical. We believe that in most economic and political situations of interest, commitment to inefficient actions is not feasible.

How large is the efficiency loss due to incentive problems? It can be shown that as the number of agents goes to infinity the probability of making the right decision goes to one in the first-best scheme. And so does the expected payoff of an agent. In contrast, in the optimal incentive compatible mechanism the probability of making the right decision converges to a number that is strictly less than one. The expected payoff of an agent also converges to this number. This implies that the ratio of the expected payoff of an agent in the optimal incentive compatible mechanism and in the first best scheme converges to a number strictly less than one as K goes to infinity. Hence, the total welfare loss goes to infinity.

A feature of the optimal mechanism is that agents cannot abstain. What happens if for some reason the SP must allow agents to abstain? We claim that the SP can achieve the same value of his objective function as without this restriction. The mechanism described in Theorem 1 can be modified as follows. The possible actions of an agent are reporting signal a or b or abstaining. If an agent abstains but no agent abstained before, the SP randomizes between signals a and b and continues to operate the mechanism as if the agent had reported the outcome of this randomization. If an agent abstains and there was already an agent who abstained in the sequence, the SP stops asking agents and makes the ex-post efficient decision. From the agent's point of view, abstaining is the same as reporting a signal randomly, given that the other agents acquire information. Ex-post efficiency is not violated, either, since the SP makes a final decision only if

[15] All these results can be found at http://home.uchicago.edu/~szentes/asymmetric.pdf.

the difference between the numbers of each type of signal is not zero. Hence the deviation of a single agent cannot make the final decision ex-post inefficient.

We assumed that the cost of acquiring information is the same for all agents. What happens if the agents' costs are heterogeneous? The derivation of the canonical mechanism (Lemma 2) is still valid. However, one cannot restrict attention to the uniform ordering of agents anymore. Clearly, the SP prefers to place the low-cost agents at the beginning of the sequences. On the other hand, an agent has little incentive to acquire information if she knows that she is at the beginning of a sequence. Hence, we conjecture that in the optimal mechanism, the low-cost agents are placed at the beginning of the ordering relatively more frequently but not always. Also, the difficulty of analyzing that problem is that each agent who is ordered differently will have a different incentive compatibility constraint.

Finally, throughout the paper we maintained the assumption that the Social Planner cannot use a transfer scheme to induce the agents to acquire information. If he could do so, he could implement the first-best mechanism. Such a mechanism would specify a transfer scheme that rewards an agent if her report matches the majority of other agents' reports and punishes the agent otherwise. (For further details of such mechanisms, see [6] and [16].) We believe that many environments where small groups make decisions should be modeled with nontransferable utilities.

Appendix A. The proof of Theorem 1 for $K < \infty$

Let $\bar{d}$ denote the largest d for which there exists an l such that $p(V(l,d), G^*) > 0$. For each d, ($d \in \{0, \dots, \bar{d}-1\}$), let $\phi(d)$ denote the largest possible l for which $\rho(V(l,d), C : G^*)$ (and by symmetry $\rho(V(l,-d), C : G^*)$) is positive. Since there exists a state $V(l,\bar{d})$ such that $p(V(l,\bar{d}), G^*) > 0$ for all d ($|d| \in \{0, \dots, \bar{d}-1\}$), there must exist an l such that $\rho(V(l,d), C : G^*) > 0$. (Otherwise the state $V(l,\bar{d})$ could not have been reached with positive probability.) Hence the function ϕ is well defined. Next, we show that conditional on reaching $V(l,d)$, the mechanism continues at $V(l,d)$ if and only if $l < \phi(d)$.

Lemma 11. *Suppose $p(V(l,d) : G^*) > 0$. If $\phi(|d|) \geqslant l$, then $p(C : V(l,d), G^*) > 0$. If $\phi(|d|) < l$, then $p(C : V(l,d), G^*) = 0$.*

Proof. Without loss of generality, assume $d \geqslant 0$. First, suppose that $p(V(l,d) : G^*) > 0$ and $\phi(d) > l$. Assume, by contradiction, that $p(C : V(l,d), G^*) \neq 1$. (That is, $p(M : V(l,d), G^*) > 0$.) By the definition of the function ϕ, a continuation mechanism is employed at the state $V(\phi(d), d)$ with positive probability. Since $p(V(l,d) : G^*) > 0$ and $\phi(d) < l$, this continuation mechanism is also feasible at the state $V(l,d)$. Furthermore, by Lemma 5, the continuation mechanism is more efficient at $V(l,d)$ than at $V(\phi(d), d)$. But then, by Lemma 6, $p(M : V(l,d), G^*) > 0$ is impossible. If $\phi(d) = l$, then by the definition of the function ϕ, $p(C : V(l,d), G^*) > 0$.

If $\phi(d) < l$, then the mechanism stops at $V(l,d)$ by the definition of ϕ. □

Lemma 12. *The function ϕ is strictly decreasing.*

Proof. Suppose, by contradiction, that $d > d'$ ($\geqslant 0$) and $\phi(d) \geqslant \phi(d')$. Notice that the mechanism continues at the state $V(\phi(d), d)$. However, after continuing from $V(\phi(d), d)$, whenever the difference between the numbers of different signals is d', the mechanism must stop. (This is because $\phi(d') < \phi(d)$.) Hence the decision of the SP after reaching $V(\phi(d), d)$ is always the majority decision at $V(\phi(d), d)$, no matter what the agents report in the continuation mechanism. This contradicts the statement of Lemma 7. $\square$

Lemma 11 only characterized those states where the mechanism continues conditional on reaching those states in terms of the function ϕ. The next lemma essentially says that the states characterized in Lemma 11 are actually reached.

Lemma 13. *Let $V(l, d) \in \mathcal{V}(K)$. Then*

$$\rho\big(V(l, d), C : G^*\big) > 0 \quad \Leftrightarrow \quad l \leqslant \phi\big(|d|\big).$$

Proof. If $\rho(V(l, d), C : G^*) > 0$, then by the definition of the function ϕ, $l \leqslant \phi(|d|)$.

Suppose that there exists a state $V(l, d) \in \mathcal{V}(K)$ such that $l \leqslant \phi(|d|)$ but $\rho(V(l, d), C : G^*) = 0$. That is, the set

$$\Omega = \big\{V(l, d)\colon\ V(l, d) \in \mathcal{V}(K),\ l \leqslant \phi\big(|d|\big),\ \rho\big(V(l, d), C : G^*\big) = 0\big\}$$

is non-empty. Among such states, consider the set $\widetilde{\Omega}$ of those where $l + |d|$ is minimal. That is,

$$\widetilde{\Omega} = \arg\min_{V \in \Omega} \big\{l(V) + \big|d(V)\big|\big\}.$$

Among these states, consider one where l is minimal. Let $V(l', d')$ be such a state. That is,

$$V(l', d') \in \arg\min_{V \in \widetilde{\Omega}} \big\{l(V)\big\}.$$

Since the mechanism is symmetric with respect to a and b, $d' \geqslant 0$ can be assumed. There can be two reasons why $\rho(V(l', d'), C : G^*) = 0$: either the state $V(l', d')$ is never reached, or although it is reached, the mechanism stops there. First, suppose that $p(V(l', d') : G^*) > 0$. Then, by Lemma 11, $\rho(V(l', d'), C : G^*) > 0$, a contradiction. Hence, $p(V(l', d') : G^*) = 0$. Then $\rho(V(l' - 1, d' - 1), C : G^*) = 0$, for otherwise the state $V(l', d')$ would be reached with positive probability. We consider two different cases.

Case 1. $d' > 0$. Then, since $l' \leqslant \phi(d')$, from Lemma 12 it follows that $l' - 1 \leqslant \phi(d' - 1)$. That is, $V(l' - 1, d' - 1) \in \Omega$. Since $l' - 1 + d' - 1 < l' + d'$, $V(l', d') \notin \widetilde{\Omega}$, a contradiction.

Case 2. $d' = 0$. Then $p(V(l' - 1, -1) : G^*)$ must be zero. (If not, then since $V(l', 0) \in \widetilde{\Omega}$, $V(l' - 1, -1) \in \widetilde{\Omega}$ also. But then

$$V(l', 0) \notin \arg\min_{V \in \widetilde{\Omega}} \big\{l(V)\big\},$$

a contradiction.) Since $p(V(l' - 1, -1) : G^*) = 0$, it follows that the mechanism cannot continue at $V(l' - 2, 0)$. That is, $\rho(V(l' - 2, 0), C : G^*) = 0$. This implies $V(l', 0) \notin \widetilde{\Omega}$, again a contradiction. $\square$

Let $\overline{T}$ denote the set of states where the mechanism involves randomization:

$$\overline{T} = \big\{V\colon\ \rho(V, C : G^*), \rho(V, M : G^*) > 0\big\}.$$

Lemma 14. *Let $V(\widehat{l}, \widehat{d}) \in \overline{T}$. Then $\phi(|\widehat{d}|) = \widehat{l}$.*

Proof. Since $\rho(V(\widehat{l},\widehat{d}), C : G^*) > 0$, $\widehat{l} \leqslant \phi(|\widehat{d}|)$ by Lemma 13. Suppose that $\widehat{l} < \phi(|\widehat{d}|)$. By the definition of the function ϕ, $\rho(V(\phi(|\widehat{d}|), \widehat{d}), C : G^*) > 0$. Notice that the continuation mechanism induced by G^* at $V(\phi(|\widehat{d}|), \widehat{d})$ is feasible at $V(\widehat{l},\widehat{d})$ and more efficient by Lemma 5. Since $\rho(V(\widehat{l},\widehat{d}), M : G^*) > 0$, this contradicts Lemma 6, and therefore $\phi(|\widehat{d}|) = \widehat{l}$. □

The function f in Theorem 1 can be defined as follows:

$$f(l) = \min\{d \mid \bar{d} > d \geqslant 1,\ \phi(d) \leqslant l\}. \tag{A.1}$$

We are ready to prove Theorem 1 for $K < \infty$.

Proof of Theorem 1 for $K < \infty$. We have to show that the function f defined by (A.1) satisfies the claim of the theorem. Since ϕ is decreasing (by Lemma 12) the function f is also decreasing. Next, we show that

$$f(l+1) = f(l) \quad \text{or} \quad f(l+1) = f(l) - 1.$$

If $l + 1 = \phi(d)$ then, by (A.1), $f(l+1) = d$. Since $\phi(d) > l$ and ϕ is decreasing, it follows from (A.1) that $f(l) > d$. But $\phi(d+1) < \phi(d) = l + 1$. Therefore,

$$f(l) = d + 1 = f(l+1) + 1.$$

If $f(l+1) = d$ and $l + 1 < \phi(d)$, then clearly $l \leqslant \phi(d)$. Hence $f(l) = d = f(l+1)$. From this argument, it follows that the set T in the claim of Theorem 1 can be defined as

$$T = \{V\colon \phi(|d(V)|) = l(V)\}. \tag{A.2}$$

$f(K)$ is one for the following reason. From (A.1), it follows that $f(K) \geqslant 1$. Suppose, by contradiction, that $f(K) > 1$. Then at a state $V(K-1, d)$ where $d \geqslant 1$, an additional agent is asked to collect information even though she cannot change the posterior of the SP. This contradicts to the statement of Lemma 7.

It remains to show that conditions (i)–(iii) are also satisfied.

(i) Suppose that $f(l) \leqslant d$ and $V(l,d) \notin T$. Also assume that $p(V(l,d) : G^*) > 0$. Suppose, by contradiction, that the mechanism continues with positive probability, that is, $p(C : V(l,d), G^*) > 0$. But then $\phi(|d|) \geqslant l$ by the definition of ϕ. If $\phi(|d|) = l$, then $V(l,d) \in T$ by (A.2), a contradiction. If $\phi(d) > l$, then $f(l) = d$ is impossible by the definition of f.

(ii) Suppose that $f(l) < d$, $V(l,d) \notin T$, and $p(V(l,d) : G^*) > 0$. Since $f(l) > d$ and $V(l,d) \notin T$, it follows that $\phi(d) > l$. From Lemma 13 $p(C : V(l,d), G^*) > 0$ follows. By Lemma 14, $p(M : V(l,d), G^*) = 0$. Hence $p(C : V(l,d), G^*) = 1$.

(iii) This follows immediately from (A.2) and Lemma 14.

From Lemma 8, it follows that there exists an optimal ex-post efficient mechanism that involves randomization at only a single state. Hence, for this mechanism, there is a single state $V \in T$ such that $p(M : V, G^\star) > 0$. □

Acknowledgments

We are grateful for helpful discussions to Eddie Dekel, Jeff Ely, Tim Feddersen, Dino Gerardi, Ali Hortacsu, Motty Perry, Rob Shimer, and Hugo Sonnenschein, Leeat Yariv, and especially to Phil Reny. We thank seminar participants at UCLA, University of Chicago, Iowa State University, Tel Aviv University, Hebrew University, Boston University, Northwestern, Queen's University, Harvard, MIT, Stanford, Penn, NYU, and Wisconsin Madison for comments.

Supplementary material

The online version of this article contains additional supplementary material.
Please visit DOI: 10.1016/j.jet.2008.02.004.

References

[1] D. Austen-Smith, J.S. Banks, Information aggregation, rationality, and the Condorcet Jury Theorem, Amer. Polit. Sci. Rev. 90 (1996) 34–45.
[2] D. Austen-Smith, T. Feddersen, The inferiority of deliberation under unanimity rule, Amer. Polit. Sci. Rev., in press.
[3] D. Bergemann, J. Valimaki, Information acquisition and efficient mechanism design, Econometrica 70 (2002) 1007–1033.
[4] H. Cai, Costly participation and heterogeneous preferences in informational committees, Discussion Paper 257, UCLA Department of Economics, 2003.
[5] M. Chwe, Minority voting rights can maximize majority welfare, Amer. Polit. Sci. Rev. 93 (1999) 85–97.
[6] J. Cremer J., R. McLean, Full extraction of the surplus in Bayesian and dominant strategy auctions, Econometrica 56 (1988) 1247–1258.
[7] M.H. DeGroot, Optimal Statistical Decisions, McGraw–Hill, New York, 1970.
[8] E. Dekel, M. Piccione, Sequential voting procedures in symmetric binary elections, J. Polit. Economy 108 (2000) 34–55.
[9] T.J. Feddersen, W. Pesendorfer, The swing voter's curse, Amer. Econ. Rev. 86 (1996) 408–424.
[10] T.J. Feddersen, W. Pesendorfer, Voting behavior and information aggregation in elections with private information, Econometrica 65 (1997) 1029–1058.
[11] T.J. Feddersen, W. Pesendorfer, Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting, Amer. Polit. Sci. Rev. 92 (1998) 23–35.
[12] D. Gerardi, L. Yariv, Information acquisition in committees, Games Econ. Behav. 62 (2) (2008) 436–459.
[13] H. Li, A theory of conservatism, J. Polit. Economy 109 (2001) 617–636.
[14] H. Li, S. Rosen, W. Suen, Conflicts and common interests in committees, Amer. Econ. Rev. 91 (2001) 1478–1497.
[15] C. Martinelli, Would rational voters acquire costly information? J. Econ. Theory 129 (2006) 225–251.
[16] R.P. McAfee, P.J. Reny, Correlated information and mechanism design, Econometrica 60 (1992) 395–421.
[17] A. McLennan, Consequences of the Condorcet Jury Theorem for beneficial information aggregation by rational agents, Amer. Polit. Sci. Rev. 92 (1998) 413–418.
[18] P. Milgrom, Rational expectations, information acquisition, and competitive bidding, Econometrica 49 (1981) 921–943.
[19] R.B. Myerson, Optimal coordination mechanisms in generalized principal-agent problems, J. Math. Econ. 10 (1982) 67–81.
[20] R.B. Myerson, Multistage games with communication, Econometrica 56 (1986) 323–358.
[21] K. Mukhopadhaya, Jury size and the free rider problem, J. Law, Econ., Organ. 19 (2003) 24–44.
[22] N. Persico, Information acquisition in auctions, Econometrica 68 (2000) 135–148.
[23] N. Persico, Committee Design with Endogenous Information, Rev. Econ. Stud. 71 (2004) 165–194.
[24] R. Smorodinsky, M. Tennenholtz, Overcoming free riding in multi-party computations—The anonymous case, Games Econ. Behav. 55 (2003) 385–406.