

Persuading Voters ^{*}

RICARDO ALONSO[†]

ODILON CÂMARA[‡]

London School of Economics

University of Southern California

November 10, 2015

Abstract

In a symmetric information voting model, an individual (politician) can influence voters' choices by strategically designing a policy experiment (public signal). We characterize the politician's optimal experiment. With a non-unanimous voting rule, she exploits voters' heterogeneity by designing an experiment with realizations targeting different winning coalitions. Consequently, under a simple-majority rule, a majority of voters might be *strictly* worse off due to the politician's influence. We characterize voters' preferences over electoral rules and provide conditions for a majority of voters to prefer a supermajority (or unanimity) voting rule, in order to induce the politician to supply a more informative experiment.

JEL classification: D72, D83.

Keywords: Strategic experimentation, persuasion, voting.

^{*}We thank the following audiences for their comments: Caltech, FGV-EESP, University of Chicago — Harris School of Public Policy, University of Iowa, University of Southern California, University of Warwick, 2015 AEA Meeting, 2014 SWET Conference, 2014 MPSA Conference, Political Economy Workshop at the University of Bristol, and the Princeton Conference on Political Economy, in particular Dan Bernhardt and the discussants Tiberiu Dragu, Navin Kartik, Stephan Lauermaann and Daniel Seidmann. We are also grateful for the suggestions of the editor and four anonymous referees.

[†]LSE, Houghton Street, London WC2A 2AE, United Kingdom. R.Alonso@lse.ac.uk

[‡]USC FBE Dept, 3670 Trousdale Parkway Ste. 308, BRI-308 MC-0804, Los Angeles, CA 90089-0804. ocamara@marshall.usc.edu.

1 Introduction

Uncertainty gives rise to persuasion.

— Anthony Downs (1957)

Information is the cornerstone of democracy, as it allows voters to make better choices. In many important cases, however, uninformed voters are not free to launch their own investigations and must rely on the inquiries of others. For example, in most trials, a juror may not choose which tests are performed during the investigation or which questions are asked of a witness — jurors must rely on the prosecutor’s investigation and questions. In politics, the Legislative branch often must rely on the information in investigative reports produced by the Executive. In firms, shareholders and the Board of Directors typically depend on reports commissioned by the CEO. If the individual choosing the questions and the voters have different preferences, then that individual might strategically design her investigation to persuade voters to choose her preferred alternative. Our main goal in this paper is to study how different voting rules affect this strategic provision of information and the equilibrium payoff of voters.

The main features of our model are: (i) a group of uninformed voters must choose whether to keep the status quo (default) policy or to implement a proposed new policy; and (ii) an individual can influence this collective decision by strategically designing an experiment that reveals information about a payoff-relevant state — as in Kamenica and Gentzkow (2011), KG henceforth. To simplify presentation, we interpret our model as one in which a *politician* tries to persuade a group of *voters*. The politician is the head of the executive branch (or the head of a government agency), and the voters are members of the legislative branch (or general voters in the case of a ballot proposal). Voters must approve or reject a proposed policy, and the politician’s objective is to maximize the probability of approval.¹ The politician can sway voters’ decision by strategically designing a *policy experiment* (a pilot test). After observing the results of the experiment, voters apply Bayes’ rule and reach a common posterior belief. They then choose an action (vote), and the proposal is implemented if and only if it receives the approval of at least k voters, where k is the established voting rule.

We first characterize the politician’s optimal experiment. We then ask: Given a k -voting

¹We first consider a politician with a state-independent payoff who has no private information. In online Appendix B, we consider a politician with a state-dependent payoff who may also privately learn the state.

rule, do voters benefit from the politician’s experiment? We show that under a simple-majority rule, the politician’s influence always makes a majority of voters weakly worse off. A majority of voters is *strictly* worse off whenever the politician’s experiment targets different winning coalitions — that is, whenever the politician exploits voters’ preference disagreement to increase the probability of approving the proposal. The next example illustrates this point.

Example 1: Suppose that there are three voters, A , B and C , three equally likely states, a , b and c , and a simple-majority voting rule. The status quo yields a payoff of zero to each voter, while the proposal’s payoff for each voter, conditional on each state, is described in Table 1.

		states		
		a	b	c
voters	A	+1.1	-1	-1
	B	-1	+1.1	-1
	C	-1	-1	+1.1

Table 1: Payoffs from Approving the Proposal.

For each state, the proposal is better than the status quo for one “winner” and worse for two “losers.” Without a policy experiment, all voters reject the proposal, as it yields a negative expected payoff. With a fully informative experiment, a majority of voters (the two losers) reject the proposal. However, the politician can design the following experiment with three possible realizations. One realization reveals that the state is not c , thus revealing voter C to be a loser. This information targets coalition $\{A,B\}$ to approve the proposal since, now, the proposal’s expected payoff is positive for both A and B . Another realization reveals that the state is not b ; hence, coalition $\{A,C\}$ approves the proposal. The last realization shows that the state is not a ; hence, coalition $\{B,C\}$ approves the proposal. Since the politician’s optimal experiment guarantees that the proposal will be implemented, all voters are strictly worse off because of the politician’s influence. \square

It is important to note that this negative influence can happen even when voters’ preferences are very aligned. We say that *(i) voters rank states in the same order* if they agree on

how to rank states according to the net payoff from approving the proposal;² and *(ii) voters agree under full information* if they would always agree on the approval/rejection decision if they knew the true state. Even when both conditions hold, we show that under a simple-majority rule, a majority of voters may still be strictly worse off because of the politician’s influence (see Example 2 in Section 3.2).

Anticipating the politician’s influence, which k -voting rule do voters prefer? Voting rules affect outcomes not only by the consensus required to approve the proposal, but also by the equilibrium amount of information that the politician’s experiment provides about the proposal’s relative merits. In particular, while requiring a higher consensus (higher k) may lead to excessive rejection of the proposal, it may also induce the politician to provide a more informative experiment. Voters then face a trade-off between control and information. We first show that if voters rank states in the same order, then each voter has single-peaked preferences over k -voting rules. This implies that a majority of voters prefer *any* supermajority voting rule over a simple-majority rule.³ Our last result shows that, if voters rank states in the same order and agree under full information, then *every* voter prefers unanimity over any other k -voting rule. That is, even heterogeneous voters may agree on the optimal electoral rule.

Our paper is related to the recent literature on strategic experimentation⁴ — KG in particular. KG develop the fundamental methodology to solve a broad class of strategic experimentation problems when players have common priors. Alonso and Câmara (2015) study strategic experimentation when players have different prior beliefs. As in our paper, Michaeli (2014), Taneva (2014) and Wang (2013) focus on strategic experimentation when there are multiple receivers. Moreover, in our paper, the incumbent politician strategically generates information about the payoff consequences of a new policy through a small scale policy experiment. In other papers — e.g., Callander (2011) — the incumbent strategically generates information by fully implementing policies.

Our paper also relates to the broad literature on how institutional rules endogenously affect

²For example, voters rank states in the same order if the unknown state represents the “quality” of the proposal, and each voter’s net payoff from approving the proposal increases in its quality.

³Here, we assume that a majority would reject the proposal in the absence of the policy experiment.

⁴E.g., Brocas and Carrillo (2007), Gill and Sgroi (2008), Duggan and Martinelli (2011), and Rayo and Segal (2010).

the information available to voters. Following the work of Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996), a large literature has studied how voting rules affect information aggregation via the strategic behavior of privately-informed voters. Recent papers focus on how voting rules affect information provision by privately-informed experts. Jackson and Tan (2013) consider experts who can reveal verifiable information, while Schnakenberg (2015a,b) considers cheap talk. In these papers, experts are endowed with private information about the state, while in our paper, the uninformed politician chooses the information content of a policy experiment (public signal).

2 The Model

Voters: A finite group of $n \geq 1$ voters must choose one alternative from a binary policy set $X = \{x_0, x_1\}$, where x_0 is the status quo (or default) policy, and x_1 is the proposal. Each voter $i \in I = \{1, \dots, n\}$ has preferences over policies that are characterized by a von Neumann-Morgenstern utility function $u_i(x, \theta)$, $u_i : X \times \Theta \rightarrow \mathbb{R}$, with Θ a finite state space. To simplify presentation, let $\Theta = \{\theta_1, \dots, \theta_T\}$, with $\theta_1 < \dots < \theta_T$ and $T \geq 2$. All players share a common prior belief $p = (p_\theta)_{\theta \in \Theta}$, which has full support on Θ .

Preference parameter $\delta_\theta^i \equiv u_i(x_1, \theta) - u_i(x_0, \theta)$ captures the net payoff from approving the proposal. We use vector $\delta^i \equiv (\delta_\theta^i)_{\theta \in \Theta}$ to represent the voter’s “type” (his preference profile). To ease exposition, throughout the paper, we assume that $\theta \neq \theta' \Rightarrow \delta_\theta^i \neq \delta_{\theta'}^i$.

Politician: An incumbent politician (she), who is not a member of the group of voters, has preferences over policies characterized by a von Neumann-Morgenstern utility function. We first consider the case of pure persuasion in which the politician’s preferences are state-independent, and she has no private information. Without loss of generality, suppose that the politician receives a payoff of one if the proposal is approved and zero otherwise. Therefore, her expected payoff is simply the probability of the proposal being approved.

Policy Experiment: The politician can influence voters’ decision by designing a policy experiment that is correlated with the state. Before the group selects a policy, the politician chooses an experiment π , consisting of a finite realization space S and a family of likelihood functions over S , $\{\pi(\cdot|\theta)\}_{\theta \in \Theta}$, with $\pi(\cdot|\theta) \in \Delta(S)$. Experiment π is “commonly understood”:

π is observed by all players who agree on the likelihood functions $\pi(\cdot|\theta), \theta \in \Theta$. Players process information according to Bayes' rule. Let $q(s|\pi, p)$ be the updated posterior belief of voters after experiment π generates a realization s . To simplify notation, we use $q(s)$ or q as shorthand for $q(s|\pi, p)$.

k-voting rule: After observing the experiment's result, each voter chooses one policy $x \in X$ — we abstract from abstention. Proposal x_1 is selected if and only if it receives at least k votes, where $k \in \{1, \dots, n\}$ is the established electoral rule.

Equilibrium selection: We apply the following two equilibrium selection criteria :

1. If policy x yields voter i a strictly higher expected payoff than x' does, then he votes for x .
2. If the two policies yield voter i the same expected payoff, then he votes for x_1 .

The first criterion rules out uninteresting equilibria — for example, when $k < n$, all voters vote for the status quo independently of expected payoffs. Note that, in our model, voters have no private information about the state, so there is no information aggregation problem. Hence, the strategic voting considerations related to the probability of being pivotal are not relevant in our setup. From the set of equilibria satisfying the first criterion, the second criterion selects the subset of sender-preferred equilibria, which guarantees that the politician's expected payoff is an upper semicontinuous function of posterior beliefs (see KG).

Electoral Outcome: Consider a voter with type δ and belief q . His expected net payoff from implementing the proposal is $\langle q, \delta \rangle \equiv \sum_{\theta \in \Theta} q_\theta \delta_\theta$. Therefore, he votes for the proposal if and only if $\langle q, \delta \rangle \geq 0$. We write optimal voting strategies $a : \Delta(\Theta) \times \mathbb{R}^T \rightarrow \{0, 1\}$ as follows: $a(q, \delta) = 1$ if $\langle q, \delta \rangle \geq 0$, and $a(q, \delta) = 0$ if $\langle q, \delta \rangle < 0$. Given an electorate $\{\delta^1, \dots, \delta^n\}$, a belief q and a k -voting rule, it is useful to define the *win set* $W_k = \{q \in \Delta(\Theta) | \sum_{i=1}^n a(q, \delta^i) \geq k\}$. That is, voters implement proposal x_1 if and only if $q \in W_k$.

Politician's Problem: For any experiment π and realization $s \in S$ that yields posterior q , the politician's payoff v is: $v(q) = 1$ if $q \in W_k$, and $v(q) = 0$ if $q \notin W_k$. The politician selects an experiment π that maximizes $E_\pi[v(q)]$. Upper-semicontinuity of v ensures that an optimal experiment π^* exists (see KG for details).

Definitions: We say that voters δ^i and δ^j *agree under full information* if for every state

$\theta \in \Theta$, we have $\delta_\theta^i \geq 0 \iff \delta_\theta^j \geq 0$. We say that voters δ^i and δ^j *rank states in the same order* if for every pair of states $\theta, \theta' \in \Theta$, we have $\delta_\theta^i > \delta_{\theta'}^i \iff \delta_\theta^j > \delta_{\theta'}^j$. For instance, if δ_θ^i strictly increases in θ for all voters, then voters agree that a higher θ means a “better” proposal.

For a given voter δ , define the *set of approval states* $D(\delta) = \{\theta \in \Theta | \delta_\theta \geq 0\}$ and the *set of approval beliefs* $A(\delta) = \{q \in \Delta(\Theta) | \langle q, \delta \rangle \geq 0\}$. Under full information, voter δ approves x_1 if and only if $\theta \in D(\delta)$, while under uncertainty, he approves x_1 if and only if $q \in A(\delta)$. Define the *set of strong rejection beliefs* $R(\delta) = \{q \in \Delta(\Theta) | \theta \in D(\delta) \Rightarrow q_\theta = 0\}$ — that is, the set of beliefs that assign probability zero to every approval state.

Given the electorate, define \mathcal{B} as the collection of all coalitions of at least $n - k + 1$ voters, with typical element $b \in \mathcal{B}$. Define the *set of strong rejection beliefs* $R_k = \cup_{b \in \mathcal{B}} (\cap_{\delta \in b} R(\delta))$. That is, R_k is the set of beliefs such that there exists a “blocking” coalition b , with voters $\delta \in b$ assigning probability zero to every approval state.

3 Strategic Experimentation

3.1 Dictator

How does the politician optimally design a policy experiment when the approval decision is delegated to a single voter (a “dictator”) δ ? If $p \in A(\delta)$, then the politician does not run a policy experiment (or runs a completely uninformative experiment), as the voter approves the proposal in the absence of additional information. Now suppose that $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. The politician can always construct an optimal experiment π^* with only two realizations: one leads to approval, the other to rejection. We now provide a geometric construction of π^* . Approval realization s^+ must induce posterior $q^+ \in A(\delta)$. After observing rejection realization s^- , voter δ must assign zero probability to every approval state $\theta \in D(\delta)$; otherwise, the politician would benefit from further disclosing information (see, also, Proposition 4 in KG). Thus, s^- must induce posterior belief $q^- \in R(\delta)$. Holding q^- constant, realization s^+ becomes more likely as the posterior q^+ moves closer to the prior p . Conversely, holding q^+ constant, s^- becomes less likely as q^- moves further away from p . Consequently, the politician would

like to resort to both an approval belief $q^+ \in A(\delta)$ that is closest to the prior, and a strong rejection belief $q^- \in R(\delta)$ that is farthest from the prior — Figure 1(a) illustrates this point. The martingale property of Bayesian updating requires, however, that q^+ , q^- and p must all be collinear. The following lemma shows that an optimal experiment balances these two goals. It corresponds to a line through p that maximizes the ratio of the distances from p to $A(\delta)$ and $R(\delta)$, as described by (1) below — see, also, Figure 1(b).

Lemma 1 *Consider a dictator δ , with $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. Let π^* be any politician's optimal experiment supported on $\{s^-, s^+\}$, where dictator δ approves the proposal if and only if $s = s^+$, with $l^* = q(s^+) - p$. Let $d_l(p, A(\delta))$ and $d_l(p, R(\delta))$ be the (Euclidean) distances from the prior belief to the sets $A(\delta)$ and $R(\delta)$ along the line l . Then,*

$$\frac{d_{l^*}(p, R(\delta))}{d_{l^*}(p, A(\delta))} = \max_l \frac{d_l(p, R(\delta))}{d_l(p, A(\delta))}. \quad (1)$$

The proof of Lemma 1 in Appendix A further shows that the equilibrium probability of approval is $\Pr[\text{Approval}] = \frac{d_{l^*}(p, R(\delta))}{d_{l^*}(p, R(\delta)) + d_{l^*}(p, A(\delta))}$. The next proposition shows that the solution to (1) can be understood as the optimal choice of a cutoff state.

Proposition 1 *Consider a dictator δ , with $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. Let π^* be any politician's optimal experiment supported on $\{s^-, s^+\}$, where dictator δ approves the proposal if and only if $s = s^+$. Letting $\alpha_\theta = \Pr[s^+|\theta]$, there exists a unique $\theta^* \in \Theta$ such that*

$$\alpha_\theta = \begin{cases} 0 & \text{if } \delta_\theta < \delta_{\theta^*} \\ 1 & \text{if } \delta_\theta > \delta_{\theta^*} \end{cases}, \text{ and } \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0. \quad (2)$$

The politician's expected utility (probability of approval) under π^ is $\sum_{\theta \in \Theta} \alpha_\theta p_\theta$.*

To understand (2), first consider the voter's ideal experiment. Voter δ would like to know whether an approval state occurred; thus, his preferred experiment induces s^+ for each θ such that $\delta_\theta \geq 0$, and s^- for each θ such that $\delta_\theta < 0$. The approval probability is, then, $\sum_{\{\theta: \delta_\theta \geq 0\}} p_\theta$, and the voter's net value from approval is $\sum_{\{\theta: \delta_\theta \geq 0\}} p_\theta \delta_\theta$. If $\delta_\theta > 0$ for at least one θ , then the politician can increase the probability of approval by distorting this experiment in such a way that rejection states with a small incremental loss (i.e., small $|\delta_\theta|$) still induce approval. The politician can do so until the voter's net value from approval is identically zero, as indicated by (2). This also implies that the voter gains nothing from making decisions with π^* , as he is indifferent between approval and rejection after observing s^+ .




Figure 1: Simplex Representing the Belief of Dictator δ , with $\delta_{\theta_3} > 0 > \delta_{\theta_2} > \delta_{\theta_1}$

3.2 k -voting rule

A basic insight from persuading a dictator is the existence of a binary optimal experiment. This is possible, as the set of approval beliefs $A(\delta)$ of any voter is convex. However, with a k -voting rule, the win set W_k is, in general, not convex. Nevertheless, persuading voters with win set W_k is, for the politician, payoff-equivalent to persuading voters with a win set equal to the convex hull of W_k . To see this, note that any belief in $co(W_k)$ can be expressed as a convex combination of posterior beliefs that ensure approval. Therefore, if $q \in co(W_k)$, then there is an experiment that ensures approval with certainty.

Consider an electorate $\{\delta^1, \dots, \delta^n\}$ and a k -voting rule, with $p \notin co(W_k)$ and $W_k \neq \emptyset$. We now construct an optimal experiment that is the composition of two experiments, π_1^* and π_2^* , defined as follows. Let π_1^* be a binary experiment supported on $S_1 = \{s^-, s^+\}$, with $q(s^+) \in co(W_k)$, $q(s^-) \in R_k$ and $l^* = q(s^+) - p$. Let $d_l(p, co(W_k))$ and $d_l(p, R_k)$ be the (Euclidean) distances from the prior to the sets $co(W_k)$ and R_k along the line l . We say that π_1^* is an *optimal collective experiment* if

$$\frac{d_{l^*}(p, R_k)}{d_{l^*}(p, co(W_k))} = \max_l \frac{d_l(p, R_k)}{d_l(p, co(W_k))}, \text{ and } d(p, q(s^+)) = d_{l^*}(p, co(W_k)). \quad (3)$$

As we move beliefs along the line l^* from prior p to posterior $q(s^+)$, we get closer to the win set; hence, π_1^* captures collective persuasion. The politician then runs experiment π_2^* only after observing s^+ . We say that experiment π_2^* with realization space S_2 is an *optimal targeted experiment* if, for every $s_2 \in S_2$, we have $q(s^+, s_2) \in W_k$. That is, after observing s^+ and any re-

alization s_2 of π_2^* , at least k voters would approve the proposal. Note that π_2^* captures targeted persuasion as different realizations of π_2^* convince different coalitions of voters. That is, experiment π_2^* exploits voter disagreement to increase the chance of the proposal being approved.

Proposition 2 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$ and a k -voting rule, with $p \notin \text{co}(W_k)$ and $W_k \neq \emptyset$. Experiment π is optimal if it is the composition of an optimal collective experiment π_1^* and an optimal targeted experiment π_2^* .*

The proof of Proposition 2 follows immediately by replacing $A(\delta)$ with $\text{co}(W_k)$, and $R(\delta)$ with R_k , in Lemma 1 and then applying the same reasoning as in the proof of Lemma 1.

Note that the proposal is approved with certainty following realization s^+ of π_1^* , and it is rejected with certainty following realization s^- . Hence, knowledge of the binary experiment π_1^* suffices to compute the expected payoff of all players. If voters rank states in the same order, then we can construct a *weak representative voter* $\delta^*(k)$ such that, for all players, persuading dictator $\delta^*(k)$ is payoff-equivalent to persuading the electorate under the k -voting rule. Furthermore, we can choose $\delta^*(k)$ such that it ranks states in the same order as voters in the electorate (see Proposition B.2 in online Appendix B). Example 2 illustrates these insights.

Example 2: Consider states $\Theta = \{\theta_1, \theta_2, \theta_3\}$; voters $\{\delta^A, \delta^B, \delta^C\}$; and simple majority $k = 2$. Let $\delta_{\theta_3}^i > 0 > \delta_{\theta_2}^i > \delta_{\theta_1}^i$ for $i \in \{A, B, C\}$, so that all voters (i) rank states in the same order and (ii) agree under full information. The prior belief and the win set are depicted in Figure 2(a). Figure 2(b) depicts the posterior beliefs $\{q^-, q_A^+, q_B^+\}$ induced by an optimal experiment π^* supported on $S = \{s^-, s_A^+, s_B^+\}$. Realization s_A^+ induces posterior q_A^+ and coalition $\{A, C\}$ approves the proposal; realization s_B^+ induces posterior q_B^+ and coalition $\{B, C\}$ approves the proposal; realization s^- induces posterior q^- and rejection.

To understand this optimal experiment, we first analyze the weak representative voter δ^* . Voter δ^* is represented in Figure 2(c) by the red dotted line, which delineates the convex hull of the win set W_2 . The optimal experiment with delegation to δ^* coincides with π_1^* in Proposition 2 and induces posteriors q_*^+ and q^- . The line connecting q^- and q_*^+ in Figure 2(c) is a direction of *common interest*: all voters agree that moving beliefs from q^- in the direction of q_*^+ represents “good news” about the proposal. Thus, π_1^* captures *collective persuasion*. Although q_*^+ is good news about the proposal, it is not enough to convince voters δ^A and

δ^B . Therefore, the politician relies on *targeted persuasion*. Starting from q_*^+ , experiment π_2^* moves the belief to either q_A^+ or q_B^+ . The straight line connecting q_A^+ and q_B^+ in Figure 2(b) is a direction of *opposing interest*: moving beliefs from q_A^+ in the direction of q_B^+ represents “good news” about the proposal to voter δ^B , but “bad news” to voter δ^A . It is important to note that the weak representative voter corresponds precisely to this direction of opposing interest — as Figures 2(b) and (c) illustrate. From belief q_*^+ , the politician ensures approval by exploiting the voters’ opposing interests. The fact that, at posterior q_*^+ , the politician ensures approval, but voters δ^A and δ^B strictly prefer to reject, implies that the politician’s influence *strictly reduces* the expected payoff of a majority of voters — voters δ^A and δ^B . \square




Figure 2: Optimal Experiment for Example 2

4 Institutional Design

4.1 Do voters benefit from the politician’s experiment?

The next corollary summarizes some welfare consequences of the politician’s influence by studying two scenarios: 1) voters choose a policy on the basis of their prior beliefs; and 2) voters choose a policy after observing the outcome of the experiment π^* .

Corollary 1 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$. Compare voters’ ex ante expected payoff under the politician’s optimal experiment π^* and under no experimentation.*

(i) *If $k = n$, then all voters are weakly better off under the politician’s influence; and*

(ii) if $k < n$, then, at most, $k - 1$ voters are strictly better off under the politician's influence. Thus, at least $n - k + 1$ voters are weakly worse off under the politician's influence. These voters are strictly worse off if there is no optimal experiment with a binary realization space.

In particular, with a simple-majority voting rule, a majority of voters are weakly worse off because of the politician's influence.

Part (i) follows from the veto power of voters: under unanimity, the politician must simultaneously convince all voters to approve the proposal. However, for any non-unanimous voting rule, the politician can exploit preference disagreement by choosing realizations that target different winning coalitions. Part (ii) highlights that it cannot be the case that k voters are strictly better off by the politician's influence. Otherwise, the politician could strictly increase the probability of approval by choosing a less informative experiment that leaves the same k voters weakly better off, with at least one of them indifferent. Moreover, whenever the posterior belief $q^+ \in co(W_k)$, obtained from combining all approval realizations of π^* , is such that $q^+ \notin W_k$, $n - k + 1$ voters are strictly worse off. This is the case if there is no optimal experiment with only two realizations, which implies that the politician must be targeting different winning coalitions.

4.2 Voter preferences over k -voting rules

When facing a dictator, the politician provides just enough information to leave the dictator indifferent between the proposal and the status quo when he approves the proposal. As a result, the dictator may prefer to delegate the approval decision to someone with different preferences than his own, but who will elicit more information about the benefits of the proposal.

Voters face a similar trade-off between control and information when evaluating different k -voting rules: a higher consensus (i.e., higher k) may lead to excessive rejection of the proposal, but it may induce the politician to provide a more informative experiment. So how does each voter rank different voting rules? The following result follows from Corollary 1.

Corollary 2 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$ with an odd number $n \geq 3$ of voters. If $p \notin W_{\frac{n+1}{2}}$, then a majority of voters weakly prefer unanimity over simple majority.⁵*

⁵Corollary 1(i) implies that all voters weakly prefer unanimity with the politician's influence to rejecting

The previous result applies to any preference heterogeneity across voters. However, in many important cases, voters' preferences are partially aligned, in the sense that voters agree on the ranking of the state — for example, when the state captures the overall “quality” of the proposal. We next derive sharper results in these cases.

Lemma 2 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$. If voters rank states in the same order, then each voter δ^i has single peaked preferences over k : there exists $k^*(\delta^i)$ such that the voter's expected utility is non-decreasing in k for $k < k^*(\delta^i)$, and it is non-increasing in k for $k > k^*(\delta^i)$.*

To prove Lemma 2, suppose that voters rank states in the same order.⁶ Without loss of generality, suppose that, for each voter, the net payoff δ_θ^i increases with θ . This implies that all voters view a higher θ as a “higher-quality” proposal. Proposition B.2 in online Appendix B implies that, for all players, the k -voting rule is payoff-equivalent to delegating the decision to a weak representative voter $\delta^*(k)$ who also ranks states in the same order. Proposition 1 implies that, for each k , there exists a unique cutoff state θ_k^* such that only proposals with $\theta \geq \theta_k^*$ are approved. A higher k implies that the decision is delegated to a “tougher” weak representative voter, in the sense that the approval set “shrinks,” $A(\delta^*(k+1)) \subseteq A(\delta^*(k))$. In order to convince the tougher weak representative voter, the politician has to design a more informative experiment, which implies that the cutoff state θ_k^* weakly increases with k . Consequently, a higher k results in an experiment that discriminates better between states of higher net value and states of lower net value.

For any given voter δ^i , the marginal value of increasing θ_k^* is positive if θ_k^* is a rejection state, $\delta_{\theta_k^*}^i < 0$. In this case, voter δ^i views the weak representative voter $\delta^*(k)$ as too easy to persuade (approves the proposal too often) and prefers the tougher $\delta^*(k+1)$. Conversely, for voter δ^i , the marginal value of increasing θ_k^* is negative if θ_k^* is an approval state, $\delta_{\theta_k^*}^i > 0$. In this case, voter δ^i views the tougher weak representative voter $\delta^*(k+1)$ as too hard to

the proposal without further information; assumption $p \notin W_{\frac{n+1}{2}}$ and Corollary 1(ii) imply that a majority of voters prefer to reject the proposal without further information to have a simple-majority rule under the politician's influence. Corollary 2 then follows immediately. However, if $p \in W_{\frac{n+1}{2}}$, then a majority of voters prefer a simple-majority rule whenever unanimity makes approval too unlikely — e.g., if the win set W_n is empty.

⁶If voters do not agree on the ranking of the states, then preferences might not be single-peaked, even when voters agree under full information. See Example B.3 in online Appendix B.

persuade (rejects the proposal too often) and prefers to stay with $\delta^*(k)$. Therefore, preferences over k are single-peaked.

An important implication of Lemma 2 is that a majority of voters prefer a supermajority voting rule over a simple-majority voting rule.

Corollary 3 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$ with an odd number $n \geq 3$ of voters and $p \notin W_{\frac{n+1}{2}}$. If voters rank states in the same order, then a majority of voters weakly prefer any supermajority voting rule $k' > \frac{n+1}{2}$ over simple majority $k = \frac{n+1}{2}$.*

The result follows since a majority of voters weakly prefer unanimity over simple majority (Corollary 2), and voters have single-peaked preferences over k -voting rules (Lemma 2). Also note that a majority of voters strictly prefer supermajority k' over simple majority if it leads to a lower (but positive) equilibrium probability of approval.

The next proposition provides sufficient conditions for all voters to have the same preferences over k -voting rules.

Proposition 3 *Suppose that all voters (i) rank states in the same order and (ii) agree under full information. Then, every voter weakly prefers a $(k+1)$ -voting rule to a k -voting rule, for $k \in \{1, \dots, n-1\}$. Consequently, every voter weakly prefers unanimity over any other k -voting rule.*

When (i) holds, we know from the proof of Lemma 2 that a higher k -voting rule implies a tougher weak representative voter $\delta^*(k)$ and, hence, a higher cutoff θ_k^* associated with the optimal experiment. Recall that voter δ^i prefers a higher θ_k^* whenever it is a rejection state, $\delta_{\theta_k^*}^i < 0$. When we have condition (ii) in addition to condition (i), it implies that, for every k , the equilibrium cutoff state θ_k^* is always a rejection state for all voters (if it were an approval state for all voters, then the politician could increase the approval probability by decreasing the cutoff state). It then follows that all voters view the weak representative voter as too easy to persuade and, thus, prefer a higher k rule. This proposition implies that even heterogeneous voters may have the same preferences over electoral rules. Essentially, sufficient alignment among voters can induce perfect agreement over electoral rules if information is endogenous to the electoral rule. Voters prefer rules that require more consensus only because they induce the politician to supply a more informative experiment.

5 Discussion

Extensions: In online Appendix B, we consider extensions of the model. In particular, we show that if the politician (i) privately learns the state before choosing a policy experiment and/or (ii) has a state-dependent payoff function so that she ranks states in the same order as voters, then the results from Lemma 2 and Proposition 3 continue to hold.

Commitment: Consider a group of voters who rank states in the same order, agree under full information and currently operate under a simple-majority rule. Instead of changing the voting rule to unanimity, they would be better off committing to (i) keep the status quo if the politician does not implement a fully informative policy experiment, and (ii) vote according to a simple majority if the experiment is fully informative. However, suppose that the politician implements a very informative (but not perfectly informative) experiment. After observing its result, a majority of voters strictly prefer to implement the proposal, while a minority strictly prefer the status quo. Without external commitment devices, it is optimal for the majority to renege on their promise and approve the proposal. Even if voters tried to write a contract based on the implicit informativeness of policy experiments, it would typically be hard for the minority to prove in a court of law that the experiment was not sufficiently informative. In contrast, it is easier to contract on a k -voting rule.

Optimal k -voting rule: Our results establish voters' preferences over voting rules for a given preference profile of voters. However, voters' preferences vary across different proposals, and, in practice, voters cannot choose a different voting rule for each new proposal. Nevertheless, voters can (and do) set issue-specific voting rules for recurring topics, especially if the likely direction of voters' and politicians' preferences is known. For instance, if prosecutors likely have a higher preference for conviction than jurors, then a higher k -voting rule would illicit more informative experiments from them.

Optimal Endorsement: In online Appendix B, we consider an alternative interpretation of the model, in which we substitute the politician's choice of a policy experiment for the choice of an optimal endorser (intermediary). In that model, the politician has access to a diverse set of potential endorsers: established individuals (other politicians, legislators or bureaucrats) whose policy preferences are publicly known. The politician can privately show the state to

the endorser, who then strategically sends a cheap-talk message to voters. If the politician is unconstrained in her choice of an endorser, then she can replicate her ex ante payoff from the optimal experiment π^* . The optimal endorser is someone less biased towards the proposal than the politician, but more biased than voters. Note that the equilibrium cheap-talk message is, in general, not a simple “support” or “not support” statement. Isomorphic to the optimal policy experiment in our benchmark model, the equilibrium message is a “targeted endorsement” in which the endorser specifies which coalition of voters should approve the proposal.

6 Conclusion

In important cases, acquiring information is infeasible or prohibitively expensive for individual voters. Voters must then rely on the information generated by certain individuals, who control the design of a public experiment (e.g., jurors and prosecutor, voters and media, shareholders and CEO). In our main application, a politician designs a policy experiment whose outcome is observed by voters. Obviously, if the politician and voters share the same preferences, then the politician’s experiment always benefits voters, as it allows them to make better decisions. However, this is not true if there is a conflict of interest between the politician and voters. We show that, with a simple-majority rule, a majority of voters are *always* weakly worse off by observing the experiment’s outcome. In fact, all voters can be *strictly* worse off, even when they would agree on their choice if they knew the true state. This is so because the politician strategically designs an experiment with realizations targeting different winning coalitions. To prevent this negative impact, voters may switch to a supermajority voting rule that induces the politician to supply a more informative experiment. We also provide conditions for unanimity to be the rule that all voters prefer.

Two interesting extensions of our model are to allow for voters to privately acquire information and then deliberate prior to voting, and to allow voters to choose among multiple policy options. We see these extensions as promising and leave them for future work.

A Appendix

Proof of Lemma 1: Let π' be an arbitrary binary experiment that induces posteriors $\{q^-(\pi'), q^+(\pi')\}$ with $q^-(\pi') \in R(\delta)$, $q^+(\pi') \in A(\delta)$ and $\sum_{\theta \in \Theta} q_\theta^+(\pi') \delta_\theta = 0$. Define $l = q^+(\pi') - p$. Bayesian rationality implies that average posteriors must equal the prior so that

$$\Pr[\text{Approval}] \langle q^+(\pi') - p, l \rangle + (1 - \Pr[\text{Approval}]) \langle q^-(\pi') - p, l \rangle = 0,$$

and $q^-(\pi') - p$ and $q^+(\pi') - p$ are collinear, so

$$\langle q^-(\pi') - p, q^+(\pi') - p \rangle = - \| (q^+(\pi') - p) \| \| (q^-(\pi') - p) \|.$$

Therefore,

$$\Pr[\text{Approval}] = \frac{\langle p - q^-(\pi'), l \rangle}{\langle q^+(\pi') - p, l \rangle + \langle p - q^-(\pi'), l \rangle} = \frac{\| (q^-(\pi') - p) \|}{\| (q^+(\pi') - p) \| + \| (q^-(\pi') - p) \|},$$

where, by construction, $\| (q^+(\pi') - p) \| = d_l(p, A(\delta))$ and $\| (q^-(\pi') - p) \| = d_l(p, R(\delta))$. As π^* maximizes $\Pr[\text{Approval}]$, then $l^* = q^+(\pi^*) - p$ must satisfy (1). ■

Proof of Proposition 1: The existence of an optimal binary experiment is established in KG (Proposition 1, pg. 2595). Consider δ , π^* and α_θ as in Proposition 1. Dictator δ will approve after observing s^+ if and only if

$$E[\delta | s^+] = \sum_{\theta \in \Theta} q_\theta(s^+) \delta_\theta = \sum_{\theta \in \Theta} \frac{\alpha_\theta p_\theta \delta_\theta}{\Pr[\text{Approval}]} \geq 0,$$

with $\Pr[\text{Approval}] = \sum_{\theta \in \Theta} \alpha_\theta p_\theta$. Therefore, $(\alpha_\theta)_{\theta \in \Theta}$ must solve the following linear program:

$$\sum_{\theta \in \Theta} \alpha_\theta p_\theta = \max \sum_{\theta \in \Theta} \alpha'_\theta p_\theta, \quad s.t. \quad 0 \leq \alpha'_\theta \leq 1, \sum_{\theta \in \Theta} \alpha'_\theta p_\theta \delta_\theta \geq 0. \quad (4)$$

For any θ' , if $\delta_{\theta'} \geq 0$, then $\alpha_{\theta'} = 1$, as increasing $\alpha_{\theta'} < 1$ relaxes the approval constraint and increases the approval probability. Suppose that $\delta_\theta < \delta_{\theta'} < 0$ for $\theta, \theta' \in \Theta$. If $\alpha_\theta > 0$, but $\alpha_{\theta'} < 1$, then increasing $\alpha_{\theta'}$ by $\varepsilon (|\delta_\theta| p_\theta / |\delta_{\theta'}| p_{\theta'})$ while reducing α_θ by ε leaves the approval constraint unchanged, but increases the probability of approval by $\varepsilon p_\theta (|\delta_\theta| / |\delta_{\theta'}|) - \varepsilon p_\theta > 0$, thus leading to a contradiction. Therefore, if $\alpha_\theta > 0$, then $\alpha_{\theta'} = 1$ for any $\delta_{\theta'} > \delta_\theta$. Given our assumption $\theta \neq \theta' \Rightarrow \delta_\theta \neq \delta_{\theta'}$ and a binding approval constraint, there exists a unique optimal binary experiment and unique cutoff θ^* . ■

The proofs of Lemma 2 and Propositions 2 and 3 are in the text.

References

- [1] ALONSO, R., AND O. CÂMARA (2015): “Bayesian Persuasion with Heterogeneous Priors,” *mimeo*.
- [2] AUSTEN-SMITH, D., AND J. S. BANKS (1996): “Information Aggregation, Rationality, and the Condorcet Jury Theorem,” *American Political Science Review*, 90(1), 34-45.
- [3] BROCAS, I., AND J. CARRILLO (2007): “Influence through Ignorance,” *Rand Journal of Economics*, 38(4), 931-947.
- [4] CALLANDER, S. (2011): “Searching for good policies,” *American Political Science Review*, 105(04), pp. 643-662.
- [5] DOWNS, A. (1957): *An Economic Theory of Democracy*, Harper and Row, New York.
- [6] DUGGAN, J., AND C. MARTINELLI (2011): “A Spatial Theory of Media Slant and Voter Choice,” *Review of Economic Studies*, 78(2), 640-666.
- [7] FEDDERSEN, T. J., AND W. PESENDORFER (1996): “The Swing Voter’s Curse,” *American Economic Review*, 86(3), 408-424.
- [8] GILL, D., AND D. SGROI (2008): “Sequential Decisions with Tests,” *Games and Economic Behavior*, 63(2), 663-678.
- [9] JACKSON, M. O., AND X. TAN (2013): “Deliberation, Disclosure of Information, and Voting,” *Journal of Economic Theory*, 148(1), 2-30.
- [10] KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590-2615.
- [11] MICHAELI, B. (2014): “Divide and Inform: Rationing Information to Facilitate Persuasion,” *mimeo*.
- [12] RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118(5), 949-987.
- [13] SCHNAKENBERG, K. E. (2015a): “Expert Advice to a Voting Body,” *Journal of Economic Theory*, 160, 102-113.
- [14] SCHNAKENBERG, K. E. (2015b): “Informational Lobbying and Legislative Voting,” *mimeo*.
- [15] TANEVA, I. (2014): “Information Design,” *mimeo*.
- [16] WANG, Y. (2013): “Bayesian Persuasion with Multiple Receivers,” *mimeo*.