

Political Disagreement and Information in Elections*

RICARDO ALONSO[†]

ODILON CÂMARA[‡]

London School of Economics and CEPR

University of Southern California

October 13, 2016

Abstract

We study the role of re-election concerns in incumbent parties' incentives to shape the information that reaches voters. In a probabilistic voting model, candidates representing two groups of voters compete for office. In equilibrium, the candidate representing the majority wins with a probability that increases in the degree of political disagreement—the difference in expected payoffs from the candidates' policies. Prior to the election, the office-motivated incumbent party (IP) can influence the degree of disagreement through policy experimentation—a public signal about a payoff-relevant state. We show that if the IP supports the majority candidate, then it strategically designs this experiment to increase disagreement and, hence, the candidate's victory probability. We define conditions such that the IP chooses an upper-censoring experiment and the experiment's informativeness decreases with the majority candidate's competence. The IP uses the experiment to increase disagreement even when political disagreement is due solely to belief disagreement.

JEL classification: D72, D83.

Keywords: Disagreement, Bayesian persuasion, strategic experimentation, voting.

*For their suggestions, we thank Dan Bernhardt, Alessandra Casella, Navin Kartik, the anonymous referees, and the discussants Patrick Le Bihan, Ben Golub and Galina Zudenkova, as well as the audiences at the 2016 Cowles Foundation Summer Conference in Economic Theory, 2015 MPSA Conference, 2015 Princeton-Warwick Conference on Political Economy, 2015 ESSET Conference, 2015 SAET Conference, Columbia University, Michigan State University, Princeton University, University of Illinois at Urbana-Champaign, and University of Nottingham.

[†]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

Voters and politicians are often uncertain about the possible repercussions of different policies. When candidates advocate different policies, this uncertainty plays an important role in defining electoral outcomes. Learning about the payoff consequences of policies can then change voters’ preferences over politicians and affect electoral outcomes. The incumbent party (IP), through its control over the government, is in a privileged position to affect what voters can learn. In this paper, we study the effects of re-election concerns on the incumbents’ incentives to shape the information that reaches voters.

There are many ways in which the incumbent can affect voters’ learning. For example, the IP can run a small-scale pilot test of a novel policy or design an experiment to evaluate unobserved effects of existing policies.¹ Moreover, when designing the rollout of a complex new law, the IP can determine which aspects of the law will be enforced before and after the next election and which preliminary information will be released during the early stages of the reform.² Similarly, the IP can establish disclosure rules for government agencies.³ In all these cases, government’s control of information affects what voters can learn and, consequently, electoral outcomes.⁴

Although our model fits all these interpretations, to simplify presentation, we will say that the IP engages in strategic policy experimentation — i.e., the IP can design a public signal that generates information about the expected payoffs from different policies. We model this strategic supply of information as a persuasion game (see Kamenica and Gentzkow (2011), KG henceforth).

Our probabilistic voting model has the following ingredients: (i) Electorate: Uninformed voters are divided into two groups, majority A and minority B, with differing preferences over

¹See Greenberg and Shroder (2004) for many examples of social experiments.

²For instance, the Affordable Care Act (commonly called “ObamaCare”) was signed into law in March 2010, but many of its payoff-relevant features were implemented only after the 2012 presidential election — e.g., the program’s website HealthCare.org was launched in October 2013.

³For example, the IP can regulate what government agencies can investigate regarding current trends in gun violence, including what information can be collected. The information (or lack of information) generated by the government can influence voters’ beliefs about the most appropriate gun control laws.

⁴See Bernecker, Boyer and Gathmann (2015) for an empirical study of how re-election concerns shape the incumbent’s incentives to experiment.

policies. (ii) Parties and Candidates: Two parties compete for office. Parties' candidates are differentiated in two dimensions — a position issue (policy) and a valence issue (competence). With regard to policy, the candidate from Party \mathcal{A} will implement the preferred policy of group A if elected, while the candidate from party \mathcal{B} defends the preferred policy of group B . Candidates also differ in their competence. (iii) Policy Experiment: Party \mathcal{A} currently controls the government and, hence, has the authority to carry out a policy experiment that reveals information about voters' policy payoffs. Party leaders (or bureaucrats) are purely office-motivated; thus, Party \mathcal{A} chooses an experiment that maximizes its re-election probability. (iv) Election: After observing the experiment's outcomes, candidates revise their beliefs, and, therefore, the policies that they will implement if elected, while voters update their evaluation of the candidates' policies. Voters already know the valence of the incumbent from party \mathcal{A} . During the electoral campaign, voters also observe a noisy signal about the valence of the untried candidate from party \mathcal{B} . Each voter then chooses candidate \mathcal{A} if she is expected to deliver a higher total payoff (valence + policy) than \mathcal{B} .

We first note that, in equilibrium, the candidate representing the majority group wins with a probability that increases in the degree of political disagreement — defined as the difference in expected payoffs from the policies supported by the candidates. Therefore, the IP designs the experiment with the sole purpose of *increasing* political disagreement, which benefits its candidate.⁵

We start our analysis by studying, in Section 3, the effect that the incumbent's valence v^A has on the informativeness of the IP's optimal policy experiment. We first consider the case in which the valence distribution of the untried candidate has a log-concave probability density function (p.d.f.), such as a Normal Distribution. Then, regardless of the preferences of majority and minority voters, the following single-crossing property holds: If an experiment does not increase the incumbent's probability of victory when her competence is v^A , then this experiment does not increase her victory probability if her competence is higher than v^A (Lemma 1). This result implies that there are two cutoffs in the extended real line, such that the IP prefers to be fully transparent about policy payoffs and, thus, favors fully informative experiments when the majority candidate is sufficiently incompetent; prefers to be partially

⁵Stokes (1963) highlights the strategic use of information to shift the salience of issues. See Iyengar and Simon (2000) for a survey.

transparent for intermediate levels of competence; and prefers to be completely opaque — thus providing a completely uninformative experiment — when the majority candidate is sufficiently competent (Proposition 1 and Corollary 1).

The single-crossing property in Lemma 1 holds for *every* specification of the preferences of voters in the majority and minority groups. To characterize the optimal experiment, in Section 4, we focus on cases in which political disagreement *endogenously* increases in the voters' expectation over an unknown state. Experimental outcomes that lead to an upward revision of the average state would then magnify political disagreement, which benefits the IP, and outcomes that produce a downward revision of the average state would reduce disagreement. We show that, under the assumption of a log-concave p.d.f., it is optimal for the IP to use an upper-censoring experiment (Proposition 3). Such experiments define a cutoff state, and voters learn the true state when it falls below this cutoff; otherwise, voters learn only that the state is above the cutoff. That is, an upper-censoring experiment fully reveals low-disagreement states and pools high-disagreement states. An important implication is that, as the incumbent's competence improves, the IP monotonically provides less information to the electorate.

All of our results derive from the curvature properties of the incumbent's re-election probability. Under the log-concave assumption, the re-election probability is locally convex when the incumbent's valence and disagreement are low and locally concave when they are high. Intuitively, convexity gives the incumbent incentives to gamble on information — i.e., to generate an experiment that might increase or decrease disagreement — while concavity gives incentives to avoid such a gamble. Our results are reversed if the p.d.f. of the challenger's valence is log-convex. In the log-convex case, the incumbent's re-election probability is convex when her valence and disagreement are low and concave when they are high. Therefore, the single-crossing property goes in the opposite direction: lower values of the incumbent's competence induce less experimentation, while higher competence induces more experimentation. Moreover, if political disagreement increases with the expected state, then the IP would favor lower-censoring experiments in the log-convex case.

Our results highlight that one should view the idea of *gambling for resurrection* (e.g., Downs and Rocke (1994)) with caution. In a simple version of this story, an incumbent politician with a bad reputation and a low probability of re-election is willing to engage in

war, hoping that a favorable outcome in the battlefield will increase her re-election chances. Our model complements these papers, by explicitly showing how the gambling for resurrection result relies on the shape of the probability of re-election — in particular, the opposite result holds in the log-convex case. Moreover, engaging in war can be viewed as implementing a “full scale,” very informative experiment. As we show, in many cases, the optimal experiment is only partially informative (upper-censoring in the log-concave case).⁶

In Section 5, we present alternative interpretations of our model to emphasize that our results apply to a wide set of economic models. We consider the case in which the incumbent’s public signal generates information about her own competence, the case in which the incumbent generates information about multiple policy dimensions, and the case in which the IP generates information during the rollout of a major policy reform.

In Section 6, we present the main extension of our model: we allow for heterogeneous prior beliefs. As Callander (2011, pg. 657) notes, “[M]uch political disagreement is over beliefs rather than outcomes” — that is, much disagreement is rooted in members of the electorate holding different views of the likely effects of various policies. To focus on the role of belief disagreement, we restrict attention to cases in which voters share the same payoff function, so that political disagreement stems solely from belief disagreement. That is, in the absence of uncertainty, all voters would agree on the optimal policy, and candidates would be judged solely on their valence. In this case, one may conjecture that public information creates consensus among voters; hence, the IP will seldom benefit from persuasion, and belief disagreement will foster opaqueness. However, we show that this view is flawed. For example, if there are more than four possible states, and political disagreement is increasing in the distance between each group’s expectation of the state, then the IP can generically design an experiment that increases political disagreement with probability one (Proposition 4).

Section 7 presents additional extensions of the model (the IP supports the minority candidate; parties are policy motivated; and the case of competition in information provision). Section 8 concludes. All proofs are in Appendix A, and additional results are available online in Appendix B. We next discuss the related literature.

⁶In our model, the incumbent can choose any signal that is correlated with the state, while other models consider restrictions on the choice of signals — e.g., the gamble in Carrillo and Mariotti (2001) is constrained to be normally distributed, while the gamble in Duggan and Martinelli (2011) must be a “slant.”

Related Literature: Our paper is related to, and borrows from, various literatures.

Policy experimentation and electoral outcomes: A number of papers explore how policy experimentation (learning how different policies map into payoffs) can influence future policies and electoral outcomes, as well as how re-election concerns by office-motivated politicians guide the choice of policy experiments. One strand of the literature focuses on the role that experimentation plays in uninformed voters’ learning about the incumbent’s or the challenger’s characteristics (Biglaiser and Mezzeti, 1997; Majumdar and Mukand, 2004; Duggan and Martinelli, 2011; Willens, 2013; Fu and Li, 2014; and Dewan and Hortala-Vallve, 2014). In the benchmark interpretation of our model, the IP controls the flow of information about policy payoffs, given the exogenous information about valence available to voters. However, in Section 5.3, we reinterpret our model as the IP controlling the flow of information about valence, given the exogenous information about policy payoffs. Bernecker, Boyer and Gathmann (2015) consider a model in which politicians use their choice of policy experiment to signal competence and test it with data from the 1996 US Welfare Reform. While theirs is a “signaling” model of competence, their finding that governors with high reputation are less likely to experiment is consistent with our results in Proposition 1 and Corollary 1.

Some papers consider how information revelation about valence during an election influences the strategies of parties and politicians. For instance, Carrillo and Mariotti (2001) study how information about valence affects parties’ choice of whether to run a known or an untried candidate; Carrillo and Castanheira (2008) study how this information affects parties’ investments in improving candidate or policy quality; and Boleslavsky and Cotton (2015) study how this information affects parties’ incentives to run moderate candidates. Our paper complements these papers by studying how information about valence affects policy experimentation by the incumbent. In all four papers, through different, but related, mechanisms, an increase in information leads to a change in party strategy, which can lead to a decrease in voter welfare.

Another strand considers the effect of policy experimentation on voters’ learning about policies. For instance, Callander (2011) and Callander and Hummel (2014) study the incentives of politicians to engage in trial-and-error experimentation, while Callander and Harstad (2015) consider the effect of learning spillovers on the incentives of heterogeneous districts to experiment. Millner, Ollivier, and Simon (2014) show that a policy-motivated party — in or-

der to show the opposite party that its belief is “wrong” and to *reduce* belief disagreement—may over-experiment when politicians have heterogeneous prior beliefs. In contrast, in our model, the purely office-motivated IP strategically discloses information to *increase* belief disagreement and influence elections.

Bayesian Persuasion: Our paper relates to the recent literature on Bayesian persuasion that follows KG. In Alonso and Câmara (2016b), the goal of the “sender” is also to sway elections in favor of her preferred alternative. However, in their model, there is no uncertainty over voters’ preferences after the results of the experiment are realized. Therefore, the only information that is relevant for electoral outcomes is whether voters prefer one policy or the other — how strongly they do so is irrelevant. In our probabilistic voting model, however, this intensity is crucial — the IP wins re-election with a probability that increases in the degree of political disagreement. Therefore, the IP would like to convince voters from the majority group not only that its candidate supports a good policy, but also that the minority candidate supports a bad policy.

Kolotilin et al. (2015) study a Bayesian persuasion model with a single receiver that has private information about his type and a sender with a payoff that is a linear increasing function of the expected state. Although their setup and focus are quite different from ours, they find (Theorem 2) that if the receiver’s type has a log-concave (log-convex) p.d.f., then it is optimal to use an upper (lower) censorship signal. Their proof relies on a mechanism-design approach, while our proof of Proposition 3 is closer to the concave-closure approach of KG.

Polarization and Disagreement: A number of papers argue that access to information can increase polarization and disagreement (e.g., Dixit and Weibull, 2007; Van den Steen, 2011; and Alonso and Câmara, 2016c). In most papers, a higher disagreement is a somewhat unintended side effect of the actions of individuals generating information, such as the media catering to the demand of biased voters. In our main extension (Section 6), the IP generates information with the sole purpose of increasing belief disagreement to benefit its candidate.

2 Model

Overview: There are two parties and two groups of voters. Party \mathcal{A} represents voters in group A and party \mathcal{B} represents voters in group B , where group A is larger than B . In our

benchmark model, party \mathcal{A} holds office at the beginning of the game (Section 7 presents the opposite case). The incumbent party (IP) strategically designs a policy experiment to influence the next election. Voters observe the experiment's results and update their beliefs about policy payoffs. Voters then observe a (possibly noisy) signal about the valence of untried candidate B — voters already know the valence of incumbent A . The election takes place; the elected candidate implements a policy; payoffs are realized; and the game ends.

Voters' Preferences: Voters care about the policy choice and the valence (i.e., competence) of the elected official. If elected, the candidate has to choose one policy x from the compact, convex set $X \subset \mathbb{R}^d$, with a finite $d \geq 1$. For example, X can represent the set of feasible governmental budget allocations across d projects, the government's policy on a left-right Downsian model, or a proportional income tax rate. Each citizen's payoff from policy x depends on an unknown state $\theta \in \Theta \equiv \{\theta_1, \dots, \theta_N\}$, with a finite $N \geq 2$. To simplify presentation, let $\Theta \subset \mathbb{R}$ and $\theta_1 < \dots < \theta_N$. Players share a common prior belief p in the interior of the simplex $\Delta(\Theta)$. Citizens within each group are homogeneous, but groups differ in their policy preferences. Formally, each citizen in group $i \in \{A, B\}$ has preferences over policies characterized by the von Neumann-Morgenstern utility function $u^i(x, \theta)$, where u^i is a differentiable function of x . Each candidate is also endowed with a valence $v \in \mathbb{R}$, which we discuss momentarily. For a voter in group i , the total payoff from electing a politician with valence v who implements policy x when state θ is realized is

$$U^i(v, x, \theta) = v + u^i(x, \theta).$$

Political Parties: We model each party as a primarily office-motivated institution (or, similarly, party leaders and bureaucrats as purely office-motivated individuals), with ties to the policy interests of a particular group of voters. Formally, each party receives payoff one if its candidate is elected and zero otherwise. If elected, party \mathcal{A} implements the policy that maximizes the expected payoff of voters in group A , while party \mathcal{B} implements the best policy for voters in group B . Our assumption is equivalent to assuming that parties are both office- and policy-motivated, but myopic (or having lexicographic preferences): they first maximize the probability of winning the election, and, once elected, they implement their preferred policy. Consequently, the preferences of each party and those of the voters it represents are

only partially aligned. Party \mathcal{A} always strictly prefers to elect its own candidate, independently of policies and valences. However, given parties’ policies, voters in group A prefer to elect the candidate from party \mathcal{B} if she is sufficiently more competent than the candidate from party \mathcal{A} . See Section 7 for further discussion on policy-motivated parties.

Strategic Policy Experimentation: The IP controls the government and has the monopoly over a *policy experiment* (a public signal that is correlated with the state). By strategically designing this experiment, the party can influence voters’ beliefs and electoral outcomes. Formally, prior to the election, the IP chooses a policy experiment π , consisting of a finite realization space S and a family of distributions 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, so that $q(s|\pi, p)$ is voters’ updated posterior belief after observing realization $s \in S$ of π . To simplify notation, we use q or $q(s)$ as shorthand for $q(s|\pi, p)$.

Our learning technology follows important assumptions from KG: the IP has the monopoly over the experiment; it has no private information; it can choose any experiment that is correlated with the state; and experiments are costless to the IP. As in our model, Callander (2011) and Callander and Hummel (2014) consider a learning technology in which the incumbent party has the monopoly over the policy experiment and has no private information. However, they consider a different learning technology — one related to a Brownian process. In order to learn, the incumbent must implement a new policy, and all players (including the IP) incur the resulting policy payoff of this experiment. Thus, we interpret these as models of “full-scale” policy experimentation. In our benchmark model, we view the experiment as a small-scale policy trial, that does not directly affect the payoff of the IP. The IP controls the informativeness of the trial by strategically designing its protocol (designing treatment and control groups, evaluation tools, etc.). In Section 5, we consider alternative interpretations of this public signal.⁷

Candidate’s Policy: We refer to the candidates from parties \mathcal{A} and \mathcal{B} as candidates A and B , respectively. Parties and their candidates have the same preferences: they want to max-

⁷In Section B.7 of online Appendix B, we consider costs that increase in the experiment’s informational content.

imize the probability of winning the election, and, if elected, they will implement their preferred policy. There are no exogenous commitment devices available to politicians. However, since the candidates' party affiliations and the experiment's results are common knowledge, in equilibrium, voters can correctly anticipate the policy that each candidate would choose. If elected, candidate $i \in \{A, B\}$ will implement policy $x^{i*}(q) \equiv \arg \max_{x \in X} \sum_{\theta \in \Theta} q_{\theta} u^i(x, \theta)$. We refer to $x^{i*}(q)$ as the “preferred policy” of candidate i .

The only distinction that we make is that candidates are endowed with valence, which we define next, while party bureaucrats control the flow of information in the government. We could also assume that the incumbent politician directly chooses the policy experiment.

Candidate's Valence: Besides the policy dimension, candidates also differ in a valence dimension. All players already know the valence v^A of incumbent A since they observe her performance in office. After the IP chooses its experiment, but before the election, voters observe valence v^B of untried candidate B . Our timing assumption is rooted in the fact that it takes time to set up and implement policy experiments, while the identity (and, hence, the actual valence) of the challenger is only defined much closer to the election. Hence, at the time that the IP chooses an experiment, there is significant uncertainty over the valence of the next challenger.⁸

We assume that challenger's valence v^B is a random variable distributed according to the cumulative distribution function F , with probability density function f . In this paper, we focus on two cases. We first assume that:

(A1) F is twice differentiable and has full support on the real numbers;⁹ and f is log-concave.

Condition **(A1)** holds, for example, for the normal, logistic, and extreme value distributions. See Bagnoli and Bergstrom (2005) for a discussion of the properties of log-concave density functions. In Section 4.1, we consider the case in which f is log-convex and show that the main equilibrium features are reversed. We show how this sharp contrast between the two

⁸Valence is a preference shock that smooths out the probability of victory, as is standard in other probabilistic voting models. We can reinterpret valence as other types of preference shock — see Section 5.3.

⁹The full-support assumption simplifies presentation, as it avoids corner solutions in which expected victory probabilities are either zero or one. When this support is bounded, but sufficiently large, our qualitative results continue to hold if we restrict attention to preference parameters such that solutions are interior.

cases helps us better understand the IP's equilibrium incentives to design the experiment.

The model is easily extended to the case in which the incumbent politician is not running for re-election. The incumbent party \mathcal{A} then runs with an untried candidate, and voters simultaneously observe valences v^A and v^B of the untried candidates. Although we say that voters observe candidates' "true" valences, the model can easily be reinterpreted as voters observing a noisy, exogenous signal about the valence of each candidate (e.g., information from media coverage during the campaign). In this case, variables v^A and v^B are interpreted as the new expected valence of each candidate, after voters observe the implicit realization of the signals about valence.¹⁰

Election: At the time of the election, voters can predict candidates' policies $x^{A*}(q)$ and $x^{B*}(q)$. Voters also observe the realized valences v^A and v^B . Thus, for a citizen in group i , the total expected payoff of electing candidate j is

$$\mathcal{U}^{ij}(q, v^A, v^B) = v^j + \sum_{\theta \in \Theta} q_{\theta} u^i(x^{j*}(q), \theta). \quad (1)$$

To rule out uninteresting equilibria, we eliminate weakly dominated voting strategies. This implies that each voter votes for the candidate who provides him with the highest expected utility¹¹. The candidate who wins the majority of the votes is elected and then implements her preferred policy. Voters in group A are decisive since the group encompasses a majority of voters. That is, a candidate wins if and only if she receives the support of the majority group.

2.1 Political Disagreement

The previous discussion implies that a voter from group i votes for the candidate from group A if and only if¹²

$$\begin{aligned} \mathcal{U}^{iA}(q, v^A, v^B) &\geq \mathcal{U}^{iB}(q, v^A, v^B) \\ \iff \sum_{\theta \in \Theta} q_{\theta} [u^i(x^{A*}(q), \theta) - u^i(x^{B*}(q), \theta)] &\geq -(v^A - v^B). \end{aligned} \quad (2)$$

¹⁰Defining the new random variable $\xi \equiv v^B - v^A$, our assumption **(A1)** refers to the distribution of ξ . In this case, our results on changes in v^A would then refer to location shifts of the distribution of ξ .

¹¹Voters 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.

¹²We abstract from abstentions. One could extend our model so that a citizen is less likely to abstain if his expected payoff difference between the candidates is higher, similar to Matsusaka (1995).

The RHS of (2) captures the realized valence differential. The LHS of (2) captures the *degree of political disagreement* between the two groups. That is, from the point of view of a voter in group i , it captures the expected policy-payoff difference from electing the different candidates. Define the political disagreement from the point of view of group A voters as

$$D(q) \equiv \sum_{\theta \in \Theta} q_{\theta} [u^A(x^{A^*}(q), \theta) - u^A(x^{B^*}(q), \theta)]. \quad (3)$$

Majority group A is decisive: after an experiment outcome that induces belief q , candidate A wins the election if and only if she receives the support of voters in group A , $D(q) \geq -v^A + v^B$. If the realized v^B is sufficiently high, then even voters from group A vote for candidate B , and vice-versa.¹³ Since $v^B \sim F$, given v^A , the majority candidate wins with probability

$$W(q; v^A) \equiv F(D(q) + v^A). \quad (4)$$

Therefore, candidate A wins the election with a probability that increases in the degree of political disagreement — candidate A has a “policy advantage” because a majority of voters believe that she has the “correct” preference, and, hence, she will implement the “correct” policy.

In order to guarantee the existence of an optimal experiment and simplify notation, throughout the paper, we maintain the following assumption:

(A2) *Political disagreement D is upper semicontinuous in $\Delta(\Theta)$ and differentiable at the prior belief.*

Condition **(A2)** holds for a large class of models, including the applications that we study throughout this paper. Differentiability of F and **(A2)** imply that W is upper semicontinuous in $\Delta(\Theta)$ and differentiable at the prior belief.¹⁴

¹³Voters for which the final vote goes in consonance with valence preferences, rather than with policy preferences, are dubbed “Stokes voters” by Groseclose (2001).

¹⁴Assumption **(A2)** implicitly establishes the following. Given q , if there are multiple optimal policies $x^{i^*}(q)$, then we select an optimal policy such that D is upper semicontinuous. Moreover, it implicitly implies that we restrict attention to language-invariant equilibria — see Alonso and Càmara (2016a) for a discussion of language-invariant equilibria.

2.2 Notational Conventions

For vectors $q, w \in \mathbb{R}^J$, we denote by $\langle q, w \rangle$ the standard inner product in \mathbb{R}^J — i.e., $\langle q, w \rangle = \sum_{j=1}^J q_j w_j$ — and we denote by qw the component-wise product of vectors q and w — i.e., $(qw)_j = q_j w_j$.

For an arbitrary real-valued function g , define \tilde{g} as the concave closure of g ,

$$\tilde{g}(q) = \sup \{y \mid (q, y) \in \text{co}(g)\},$$

where $\text{co}(g)$ is the convex hull of the graph of g .

We use $\pi \succ \pi'$ to denote that experiment π is Blackwell more informative than experiment π' . Finally, $\text{card}(S)$ denotes the cardinality of the set S .

2.3 Party's Expected Payoff

The incumbent party's problem is to choose an experiment π that maximizes the expected probability of victory $E_\pi[W(q; v^A)]$. Upper semicontinuity of W ensures the existence of an optimal experiment, and choosing an optimal experiment is equivalent to choosing a probability distribution σ over q that maximizes $E_\sigma[W(q; v^A)]$, subject to the constraint $E_\sigma[q] = p$ (see KG). That is, the supremum of the expected victory probability is

$$W^* = \sup_{\sigma} E_\sigma[W(q; v^A)], \quad \text{s.t. } E_\sigma[q] = p.$$

The following remarks follow immediately from KG:

(R1) An optimal experiment exists.

(R2) There exists an optimal experiment with $\text{card}(S) \leq N$.¹⁵

(R3) The IP's maximum expected payoff is $W^* = \tilde{W}(p; v^A)$.

(R4) The value of persuasion is $W^* - W(p; v^A) = \tilde{W}(p; v^A) - W(p; v^A)$.

¹⁵Note that, in the original setup of KG, there exists an optimal straightforward signal that directly recommends an action to the receiver. In our setup, the pivotal majority voter has a binary action space: vote for candidate A or B . However, when $N > 2$ in our model, an optimal experiment might require more than two realizations. This is so because, from the point of view of the IP, before the valence shock is realized, the voting behavior is probabilistic rather than binary. That is, voting behavior can be interpreted ex ante as a continuous “action” (probability of electing A) in the interval $[0, 1]$ rather than a binary choice.

2.4 Application: Spatial Policy Model

Although we prove our main results using the general setup described above, for concreteness throughout the paper, we illustrate our results using the following application.

Consider a spatial policy model in which the state $\theta \in \Theta \subset \mathbb{R}$ captures voters' uncertainty over the optimal policy in a left-right dimension. Let $X = [-\bar{x}, +\bar{x}]$, with \bar{x} sufficiently large. Voters in group A have a quadratic policy payoff $u^A(x, \theta) = -(x - \theta)^2$. From the point of view of majority voter A , with belief q , the optimal policy is linear on the expected value of the state, $x^{A*}(q) = E[\theta|q]$. Let $x^{B*}(q)$ be the optimal policy from the point of view of minority voter B . Political disagreement (3) is

$$\begin{aligned} D(q) &= \sum_{\theta \in \Theta} q_\theta \left[u^A(x^{A*}(q), \theta) - u^A(x^{B*}(q), \theta) \right] \\ &= \sum_{\theta \in \Theta} q_\theta \left[- (E[\theta|q] - \theta)^2 + (x^{B*}(q) - \theta)^2 \right] \\ &= (E[\theta|q] - x^{B*}(q))^2. \end{aligned} \tag{5}$$

From (5), political disagreement translates naturally into the degree of disagreement over optimal policies, $D(q) = (x^{A*}(q) - x^{B*}(q))^2$.

The shape of the disagreement function D depends fundamentally of the nature of preference misalignment between the two groups. We next present three examples, using different payoff functions for group B . In Example 1, disagreement endogenously becomes a strictly convex function of beliefs; therefore, any experiment π increases the expected political disagreement, $E_\pi[D(q)] \geq D(p)$. The opposite is true in Example 2: since disagreement is strictly concave, information, on average, decreases disagreement. In Example 3, disagreement is neither concave nor convex. In these examples, we consider a binary state space $\Theta = \{0, 1\}$, and let q_2 be the probability that the state is $\theta = +1$. Formally,

Example 1 — Suppose that $u^B(x, \theta) = -(x - \frac{1}{2}\theta)^2$. Then, $x^{B*}(q) = \frac{1}{2}E[\theta|q]$, and disagreement (5) becomes $D(q) = \frac{1}{4}E[\theta|q]^2$.

Example 2 — Suppose that $u^B(x, \theta) = -(\frac{x^2}{2} - \theta)^2$. Then, $x^{B*}(q) = \sqrt{2E[\theta|q]}$, and disagreement (5) becomes $D(q) = \left(E[\theta|q] - \sqrt{2E[\theta|q]} \right)^2$.

Example 3 — Suppose that $u^B(x, \theta) = -(x - \theta)^3$. Then, $x^{B*}(q) = \frac{q_2 - \sqrt{q_2(1-q_2)}}{2q_2 - 1}$, and

disagreement (5) becomes $D(q) = \left(q_2 - \frac{q_2 - \sqrt{q_2(1-q_2)}}{2q_2 - 1} \right)^2$.

Figure 1 illustrates these examples. The three figures on the top contrast the optimal policy $x^{A^*}(q) = q_2$ (dashed lines) and the different optimal policies $x^{B^*}(q)$ (solid lines). The three figures on the bottom depict the corresponding political disagreement.

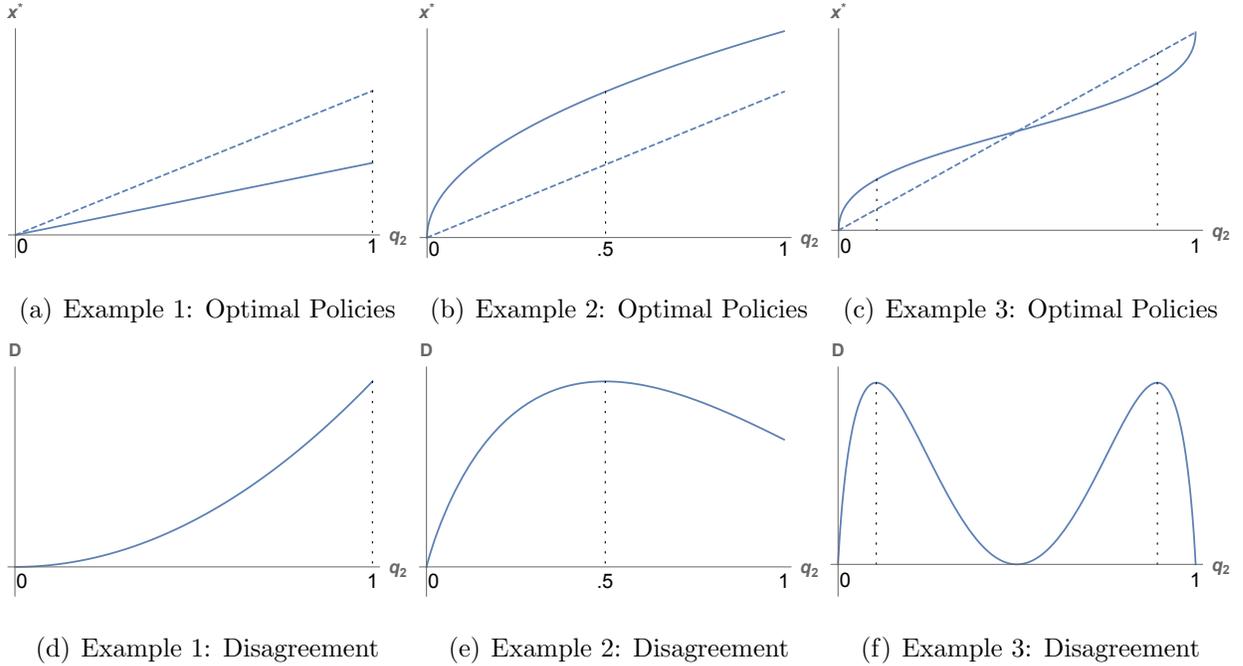


Figure 1: Top: Optimal policies x^{A^*} (solid line) and x^{B^*} (dashed line); Bottom: Political disagreement D , with $\Theta = \{0, 1\}$, $q_2 = Pr(\theta = 1)$.

3 Valence and Information

In this section, we show that the incumbent party's gain from *any given experiment* π has a single-crossing property with respect to the incumbent's valence. This property leads to a monotone behavior of the informativeness of optimal experiments: as we increase the incumbent's competence v^A , her party does not benefit from providing a more-informative experiment.

3.1 Single-Crossing

In our model, the incumbent party seeks to maximize its candidate’s chances of re-election. Following (4), the likelihood that candidate A wins the election increases in the degree of political disagreement — a larger D implies that, in the eyes of group A voters, the minority candidate B is expected to implement a much “worse policy” than A . As the outcome of the experiment can change the policy championed by each candidate, as well as voters’ expected payoff from these policies, it follows that policy experimentation can change the degree of political disagreement. As a result, the IP’s choice of an experiment is driven by its desire to uncover information that increases political disagreement.

As the underlying state θ is independent of both candidates’ valences, the IP’s choice of experiment cannot affect the distribution of the challenger’s valence. Nevertheless, if the IP has access to an experiment that, on average, increases disagreement, as in the example in Figure 1(d), then it is not clear why the IP would not gain from this experiment independently of v^A . The next lemma shows that, for any experiment π , this gain actually satisfies a single-crossing condition: If the IP prefers not to experiment rather than provide experiment π when its candidate’s valence is v^A , then the IP continues to find no experimentation better than experiment π for any higher valence $v^{A'} > v^A$.

Lemma 1 *Suppose that (A1) and (A2) hold. Consider any experiment π and incumbent’s valence v^A . If, for the IP, no experimentation is better than experiment π when the incumbent has valence v^A , then no experimentation continues to be better for all higher valences. That is, if $E_\pi[W(q; v^A)] \leq W(p; v^A)$, then $E_\pi[W(q; v^{A'})] \leq W(p; v^{A'})$ for all $v^{A'} > v^A$.*

To understand Lemma 1, note that the effect of changing disagreement by an amount Δ is that it changes the probability of victory by $F(z + \Delta) - F(z)$, with $z = D(p) + v^A$. If $\Delta > 0$, then the benefit in increasing victory probability, relative to the likelihood that the challenger’s valence induces z , is given by

$$\frac{F(z + \Delta) - F(z)}{f(z)} = \int_0^\Delta \frac{f(z + s)}{f(z)} ds. \quad (6)$$

If $\Delta < 0$, then the cost of decreasing victory probability relative to $f(x)$ is

$$\frac{F(z) - F(z + \Delta)}{f(z)} = \int_\Delta^0 \frac{f(z + s)}{f(z)} ds. \quad (7)$$

Lemma 1 then follows from the fact that, for log-concave probability density functions, the ratio $f(z + \Delta)/f(z)$ decreases in z if $\Delta > 0$, but increases in z if $\Delta < 0$. That is, the relative benefit (6) of increasing victory probability decreases in z — hence, in the incumbent’s competence v^A — while the relative cost (7) increases in z . Integrating over all possible realizations of Δ generated by experiment π , we then have that the relative gain from an experiment π weakly decreases in the incumbent’s competence. In other words, if the IP does not gain from experiment π when the incumbent’s valence is v^A , this is still true for an incumbent candidate of higher valence. Notice that this property is satisfied irrespective of whether, in the absence of the IP’s experiment, the incumbent is expected to win the election ($F(z) > 1/2$) or the minority candidate is the frontrunner ($F(z) < 1/2$).

The next proposition builds upon Lemma 1 to show that, if we increase the competence of the majority candidate, then the IP does not benefit from providing a more-informative experiment.

Proposition 1 *Suppose that (A1) and (A2) hold. Suppose, also, that π^* is an optimal experiment given incumbent’s valence v^A . Then, for any higher valence, experiment π^* is weakly better than any Blackwell more informative experiment. That is, for every $v^{A'} > v^A$ and every $\pi' \succ \pi^*$, we have*

$$E_{\pi^*}[W(q; v^{A'})] \geq E_{\pi'}[W(q; v^{A'})]. \quad (8)$$

In the proof of the proposition, we first rewrite the Blackwell more informative experiment π' as a payoff equivalent grand experiment. In this grand experiment, voters first observe realization s of π^* , and then they observe an additional experiment π_s conditional on s . When the incumbent’s valence is v^A , optimality of π^* implies that the IP does not benefit from disclosing any additional information π_s after each realization s of π^* . We then apply Lemma 1 to each posterior belief q^* in the support of π^* : if the IP does not benefit from disclosing information in addition to π^* when the incumbent’s valence is v^A , then the IP does not benefit from disclosing any information in addition to π^* when the incumbent’s valence is higher.¹⁶

¹⁶Although Lemma 1 holds for any experiment, the result in Proposition 1 is deeply rooted in the endogenous properties of optimal experiments. In general, two Blackwell-ordered experiments do not enjoy this single-crossing property. If $\pi' \succ \pi$ for some non-optimal pair of experiments, then it might be the case that

Next, we apply Proposition 1 to characterize the relationship between the IP's optimal level of transparency and the incumbent's valence.

Corollary 1 *Suppose that (A1) and (A2) hold. There are cutoffs v_1^A and v_2^A in the extended real line, with $v_1^A \leq v_2^A$, such that:*

- (i) *a fully informative experiment is optimal if $v^A < v_1^A$;*
- (ii) *a partially informative experiment is optimal if $v_1^A < v^A < v_2^A$; and*
- (iii) *an uninformative experiment is optimal if $v_2^A < v^A$.*

Corollary 1 defines partitions on the expected competence of the majority candidate. When the incumbent party's candidate is sufficiently incompetent, it prefers to be completely transparent about policies, and engages in fully informative experimentation; the IP is partially transparent for intermediate levels of competence and is completely opaque (forgoes experimentation) when its candidate is sufficiently competent.

Corollary 1 does not guarantee that cutoffs v_1^A and v_2^A are finite.¹⁷ Proposition B.1 in online Appendix B provides sufficient conditions so that v_1^A and v_2^A are finite.

3.2 Examples

We next provide some examples to illustrate the effects of the incumbent's valence v^A on the IP's payoff function W and on the optimal experiment.

Recall that $W(q; v^A) = F(D(q) + v^A)$. Figure 2 illustrates how a higher v^A increases W for each q and changes the overall curvature of W . It assumes that F follows a Normal Distribution and uses the political disagreement D from the spatial policy model in Figure 1(d).

Recall that we can derive the optimal experiment from the concave closure of W (see KG for details). In particular, whether W is concave or convex is important to define whether or not the IP benefits from implementing an informative experiment. Although in Figure 1(d) disagreement D is strictly convex, the resulting payoff W might be locally concave or

the IP prefers the less informative π when valence is low and prefers the more informative π' when valence is high: $E_\pi[W(q; v^A)] > E_{\pi'}[W(q; v^A)]$ and $E_\pi[W(q; v^{A'})] < E_{\pi'}[W(q; v^{A'})]$ for some $v^{A'} > v^A$. See Section B.2 in online Appendix B for details.

¹⁷E.g., in Example 2 from Section 2.4, if the prior belief already maximizes disagreement, $p_2=0.5$, then no information disclosure is optimal for all values of v^A , so that $v_1^A = v_2^A = -\infty$.

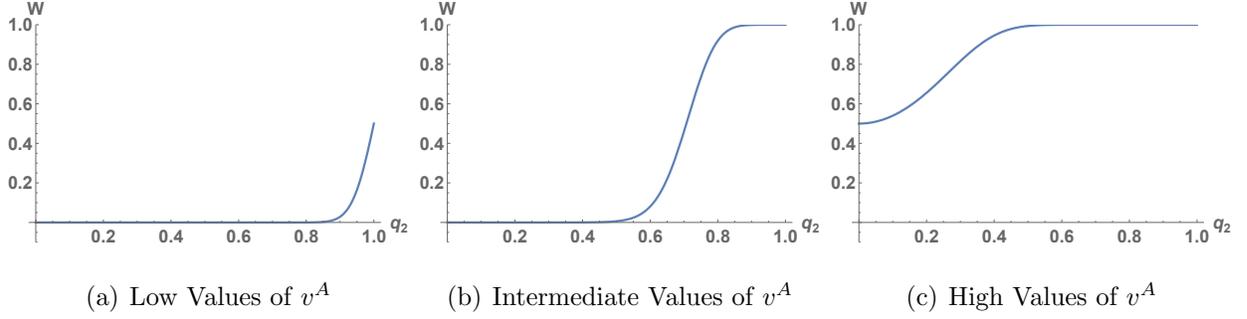


Figure 2: Effects of v^A on victory probability W , using disagreement D from Figure 1(d).

locally convex, depending on belief q_2 and on valence v^A . Log-concavity of f implies that $F(D(q) + v^A)$ is locally concave for sufficiently high values of $D(q) + v^A$ and locally convex for sufficiently low values. The red solid lines in Figure 3 depict the concave closure of W . We next use Figure 3 to derive an optimal experiment.

First, suppose that v^A is sufficiently low, as in Figure 3(a). The IP's payoff W is everywhere strictly convex; hence, any optimal experiment must be fully informative, independently of the prior belief.

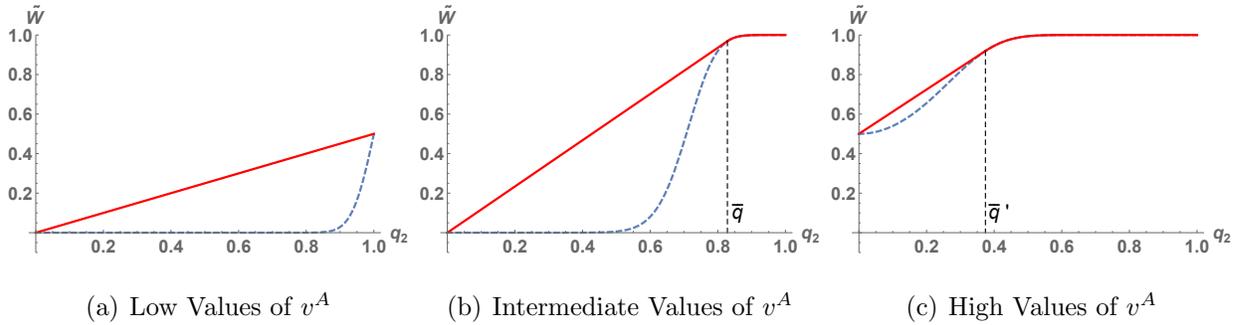


Figure 3: Concave closure of W from Figure 2.

Now suppose that v^A is intermediate, as in Figure 3(b). The concave closure \widetilde{W} is given by a straight line in the set of beliefs $q_2 \leq \bar{q}$, and by W itself for $q_2 \geq \bar{q}$. Consequently, no experimentation is optimal for all priors $p_2 \geq \bar{q}$. When $p_2 \geq \bar{q}$, although any informative experiment increases average disagreement (D is strictly convex), any informative experiment is strictly worse for the IP than no information disclosure. Signal realizations that increase *political disagreement* increase *victory probability* by only a small amount, while signal realizations that decrease *political disagreement* decrease *victory probability* by a relatively large amount. Now suppose that $p_2 \leq \bar{q}$. Since in this set $\widetilde{W}(q; v^A) > W(q; v^A)$, policy

experimentation is valuable. Every optimal experiment is partially informative and induces exactly two posterior beliefs, $q_2 = 0$ and $q_2 = \bar{q}$. Finally, for each prior belief $p_2 \in (0, 1)$, optimal experiments are less informative in Figure 3(b) than in Figure 3(a).

As we further increase v^A , the cutoff \bar{q} decreases to \bar{q}' — see Figure 3(c). Therefore, no experimentation is optimal for a larger set of prior beliefs. Moreover, for the prior beliefs in the set $p_2 \leq \bar{q}'$, every optimal experiment is supported only on the posterior beliefs $q_2 = 0$ and $q_2 = \bar{q}'$. Consequently, the partially informative experiment in Figure 3(c) is less informative than the partially informative experiment in Figure 3(b).

What if political disagreement is everywhere strictly concave, as in Figure 4? Figures 5 and 6 use a normally distributed v^B to illustrate the corresponding victory probability W , which might be locally concave or convex, depending on the incumbent’s valence. See online Appendix B (Section B.4.1) for a detailed discussion of this example.

We conclude by highlighting that the IP might find it optimal to experiment even when its candidate is the frontrunner and might find it optimal not to experiment even when its candidate is the underdog. For example, in Figure 3(b), if the prior belief is $p_2 = 0.8$, then, without experimentation, the majority candidate wins with a very high probability, above 90%. Nevertheless, it is optimal for the incumbent party to provide a partially informative experiment, because it increases its candidate’s expected victory probability even further. In Figure 6(a), if the prior belief is $p_2 = 0.45$, then, without information disclosure, the majority candidate wins with a very low probability, around 26%. Nevertheless, *any* informative experiment decreases the candidate’s expected victory probability even further .

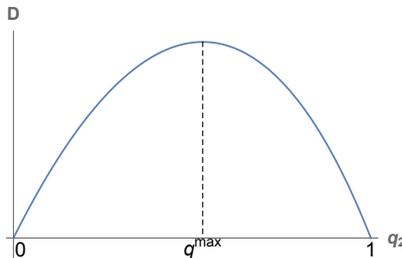


Figure 4: Strictly Concave Political Disagreement.

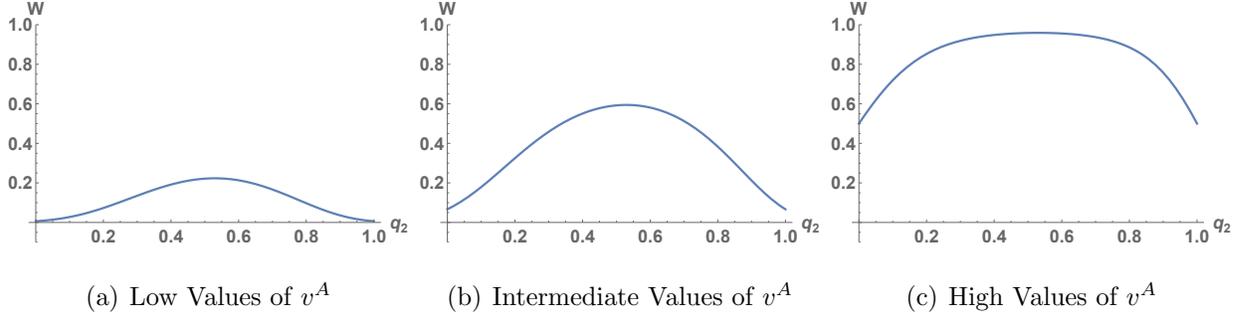


Figure 5: Effects of v^A on victory probability W , using disagreement D from Figure 4.

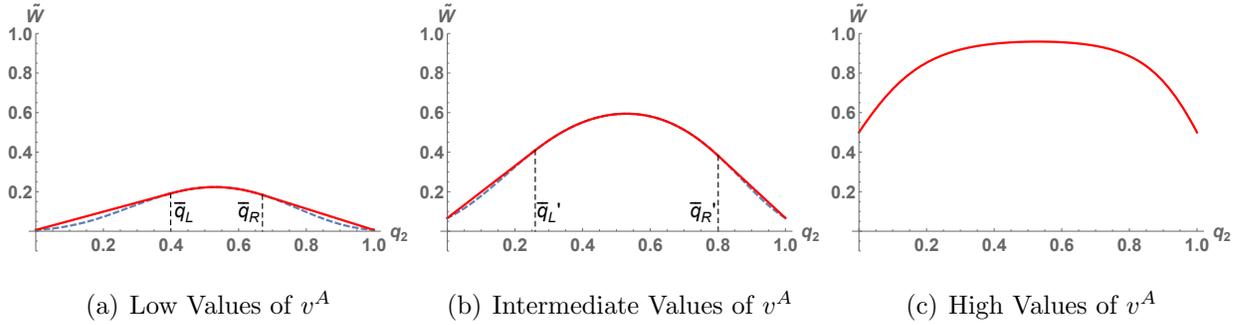


Figure 6: Concave closure of W from Figure 5.

3.3 Information and Voter Welfare

We next apply our results to highlight the possible negative effects of the incumbents' policy experimentation on voters' welfare.

In many voting models, interested parties strategically provide information to voters. In some cases, this information can adversely affect voters' equilibrium welfare — voters' payoff would be higher if they made uninformed choices. For instance, in Alonso and Câmara (2016b), the information provided by the IP always weakly decreases the expected payoff of a majority of voters under a simple majority voting rule. This is so because the optimal experiment has signal realizations targeting different winning coalitions of voters.

In our model, the IP cannot target different winning coalitions because voters in group A are representative. Nevertheless, the next proposition shows that the IP's optimal experiment may hurt all voters in majority group A .

Proposition 2 *Suppose that (A1) holds. Consider the spatial policy model from Section 2.4, with payoffs $u^i(x, \theta) = -(x - \beta^i \theta)^2$. If either $\beta^A \beta^B < 0$ or $|\beta^B| > 2|\beta^A|$, then there exists a finite cutoff \bar{v}^A such that, for any $v^A < \bar{v}^A$, the IP's optimal experiment strictly*

decreases the expected payoff of all voters in majority group A .

The result follows from the different interests of the IP and voters A , and from the fact that the IP benefits from promoting disagreement. The IP’s goal is to elect candidate A , not to ensure that the elected candidate implements a good policy for voters A . Candidate A is more likely to win when information leads candidate B to adopt a new policy that voters A consider worse. This worse policy benefits the IP, but it hurts voters A when candidate B is elected.

Under the conditions of the proposition, the IP implements a fully informative experiment, and the information (on average) leads candidate B to implement a worse policy for voters A . This is the case here, as preference misalignment is sufficiently severe: If $\beta^A\beta^B < 0$ (candidates want policies with opposing signals) or $|\beta^B| > 2|\beta^A| \geq 0$ (from A ’s point of view, candidate B “overreacts” to information), then voter A strictly prefers an elected candidate B not to have access to any informative experiment. When v^A is sufficiently low, candidate A is unlikely to win (the benefit of providing information to candidate A is small), while candidate B is likely to win (the loss of providing information to candidate B is large). Hence, voters in group A are strictly worse off because of the IP’s experiment. They would prefer no experiment over the IP’s equilibrium experiment.

4 Disagreement as a Function of the Expected State

To derive a sharper characterization of optimal experiments, in this section, we focus on models in which political disagreement is a strictly increasing function of the expected value of some unknown state. Formally, we assume:

(A2’) *Political disagreement takes the form $D(q) = H(E[\theta|q])$, where H is twice differentiable and strictly increasing. Moreover, the ratio $\frac{H''}{(H')^2}$ is non-increasing.*

Assumption **(A2’)** holds in many important cases. For example, it holds if disagreement is a power function of expectation $D(q) = \gamma E[\theta|q]^\rho$, with $\gamma > 0$, $\theta_1 \geq 0$ and $\rho \geq 1$. It also holds in the spatial policy model of Section 2.4, when voters have quadratic payoffs.¹⁸ Later in this section, we study another relevant application in which **(A2’)** holds (choice of an income tax).

¹⁸If $u^A(x, \theta) = -(x - \beta^A\theta)^2$ and $u^B(x, \theta) = -(x - \beta^B\theta)^2$, where β^A and β^B are known preference

Given **(A2')**, political disagreement increases if voters learn that the realized state is “high,” which benefits the incumbent, and disagreement decreases if they learn that the state is “low.” One could then conjecture that the incumbent party would prefer to hide information about low-disagreement states, and to fully disclose information about high states. However, Proposition 3 shows that the opposite is true. Borrowing from the statistics literature, we define an upper-censoring experiment (or right-censoring experiment) as one that fully reveals low-disagreement states and pools high-disagreement states. Formally:

Definition: Experiment π is *upper-censoring* at cutoff state θ_k if it has a realization space $S = \{s_1, \dots, s_k, s_{\text{pooling}}\}$ and the following holds. For each $n < k$, state θ_n induces signal realization s_n with probability one. For each $n > k$, state θ_n induces signal realization s_{pooling} with probability one. Cutoff state θ_k induces realization s_{pooling} with some probability $\alpha_k \in [0, 1]$ and induces realization s_k with probability $1 - \alpha_k$.

Proposition 3 *Suppose that **(A1)** and **(A2')** hold. Then, there exists an optimal experiment π^* that is upper-censoring at some cutoff state θ_k . Moreover, cutoff state θ_k weakly decreases with the incumbent’s valence v^A .*

In the proof of Proposition 3, we show that for each optimal experiment π^* , there exists a payoff-equivalent upper-censoring experiment. The intuition behind the result is as follows. Under **(A1)** and **(A2')**, given v^A , the IP’s payoff $W(q; v^A) = F(H(E[\theta|q]) + v^A)$ is concave if $E[\theta|q]$ is high and strictly convex if $E[\theta|q]$ is low — see the example in Figure 7(a). Strict convexity implies that the IP always strictly benefits from providing additional information if the initial experiment yields a non-degenerate belief corresponding to a low expected state. Therefore, outcomes under optimal experiments that indicate the state to be low must be fully revealing. Conversely, concavity of the incumbent’s payoffs implies that the IP cannot be made worse off by an experiment that pools all outcomes corresponding to high expected states into a single realization. That is, the incumbent then (weakly) gains from bundling all states in the concave (high-disagreement) region: they all induce signal s_{pooling} with probability one, resulting in a single posterior belief q^+ and a high expectation $E[\theta|q^+]$.

parameters and $\theta_1 \geq 0$, then disagreement is proportional to the square of the expectation of the state, $D(q) = (\beta^A - \beta^B)^2 E[\theta|q]^2$.

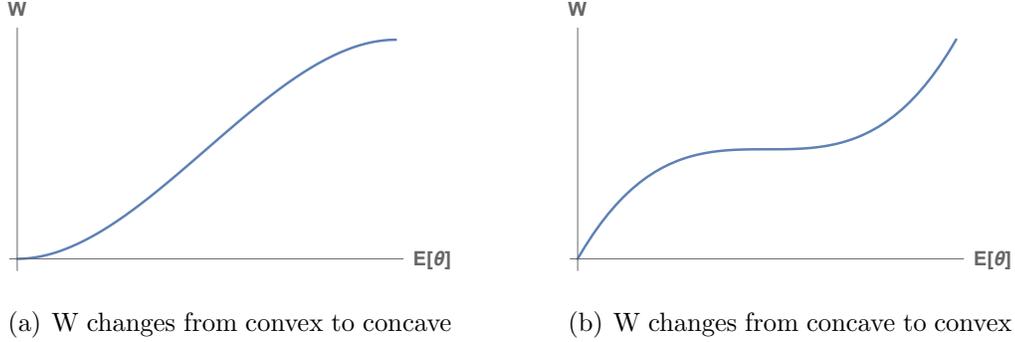


Figure 7: Re-election Probability as a Function of the Expected State.

While the IP does not gain from designing an experiment that pools together only states in the convex region, it may gain from “hiding” some low-disagreement states, such that these states induce signal s_{pooling} with positive probability. Of course, pooling low-disagreement states would make s_{pooling} more likely but would reduce expected disagreement if s_{pooling} occurs. Still, the incumbent must decide which disagreement states should be pulled in s_{pooling} . Suppose that θ_l and θ_h are in the convex region, with $\theta_l < \theta_h$. Should θ_l or θ_h be the incumbent’s first choice to be mixed with the high-disagreement signal s_{pooling} ? The IP now faces an important tradeoff. On the one hand, pooling θ_h leads to a lower reduction in posterior disagreement resulting from s_{pooling} . On the other hand, disclosing θ_l to voters is worse than disclosing θ_h ; thus, “hiding” θ_l by pooling it with s_{pooling} is more important than hiding θ_h . The proof of Proposition 3 shows that, given **(A1)** and **(A2’)**, the first effect always dominates: the IP’s optimal decision must be a cutoff on θ , independent of prior beliefs, the incumbent’s valence, and the other parameters of the model — these values are relevant only for defining the actual cutoff state.

Finally, the cutoff state defined by Proposition 3 monotonically decreases with the incumbent’s valence v^A . This implies that the set of optimal upper-censoring experiments that we construct is Blackwell-ordered: experiments become less Blackwell-informative as the majority candidate become more competent.

It is important to note that the logic behind the proof of Proposition 3 applies to a broad class of models. Consider a Bayesian persuasion game between a sender and a receiver, as in KG. Suppose that the sender’s payoff can be written as a twice differentiable, strictly increasing function of the expected state. If the derivative of the sender’s payoff function

is single-peaked, then there exists an optimal experiment that is upper-censoring. In our model, conditions **(A1)** and **(A2')** simply imply that this derivative is log-concave, hence single-peaked. See Proposition A.1 in Appendix A for details.

4.1 Log-convex Valence Distribution

Our previous results depend fundamentally on the assumption that the p.d.f. of the challenger's valence distribution is log-concave. The results are reversed if we change **(A1)** so that f is log-convex.¹⁹ In the log-convex case, the single-crossing property goes in the opposite direction: lower values of the incumbent's competence v^A induce less experimentation, while higher competence induces more experimentation.

Moreover, suppose that political disagreement is a strictly increasing function of the expected state, as in Section 4. With a log-convex f , if we change assumption **(A2')** so that the ratio $\frac{H''}{(H')^2}$ is non-decreasing,²⁰ then the optimal experiment is lower-censoring at some cutoff state θ_k . Furthermore, this cutoff state *decreases* with the incumbent's valence v^A ; hence, the experiment becomes more informative.

The reason for the sharp change in results is rooted in the change in the curvature of the incumbent's victory probability as a function of political disagreement and valence v^A . Loosely speaking, in the log-concave case, it is as if the IP features increasing absolute risk aversion (IARA). When disagreement and the incumbent's valence are low, the IP benefits from gambling on disagreement. That is, the IP benefits from implementing a risky experiment that might increase or decrease disagreement. When disagreement and valence are high, the IP prefers to avoid these gambles. In the log-convex case, it is as if the IP features DARA, and the reverse results hold.

Figure 7 illustrates our point with an example in which disagreement equals the expected state. In Figure 7(a), with a single-peaked p.d.f., the re-election probability changes from convex to concave. Hence, the incumbent wants to disclose information for low states and

¹⁹With a log-convex p.d.f., we can no longer assume full support on the real numbers. To simplify presentation, we want to avoid corner solutions — that is, cases such that the victory probability $F(D(q)+v^A)$ equals one or zero. To this end, we assume that the support of the challenger's valence v^B is large, and we restrict attention to valence v^A and disagreement D such that $F(D(q)+v^A) \in (0, 1)$ for all $q \in \Delta(\Theta)$.

²⁰For example, if disagreement is a power function $D(q) = \gamma E[\theta|q]^\rho$, with $\gamma > 0$, $\theta_1 \geq 0$ and $\rho \in (0, 1]$.

hide information for high states. In Figure 7(b), with a single-dipped p.d.f., the re-election probability changes from concave to convex; thus, we find the opposite incentives.

We next present an application of our model, in which political disagreement endogenously becomes a function of the expected state, and Proposition 3 applies.

4.2 Application: Optimal Tax

Consider the following model, in which the elected politician must choose a proportional income tax $x \in [0, 1]$. Voters care about the consumption of a private good and a public good. Each voter in group $i \in \{A, B\}$ is endowed with income $\beta^i > 0$, where $\beta^A \neq \beta^B$. Given the implemented tax rate x , voter i consumes $(1-x)\beta^i$ units of the private good. The government uses all tax revenues to produce the public good. The production technology is such that the government produces x^ψ units of the public good, where $\psi \in (0, 1)$ is a known technology parameter.²¹ Voters' policy payoff is $u^i(x, \theta) = (1-x)\beta^i + \theta x^\psi$, where state θ represents the unknown marginal value of the public good, with $0 \leq \theta_1 < \dots < \theta_N < \max\{\beta^A, \beta^B\}^{\frac{1}{\psi}}$. Given belief q , the optimal tax rate of voter i is $x^{i*}(q) = \left(\frac{\psi E[\theta|q]}{\beta^i}\right)^{\frac{1}{1-\psi}}$. Both groups want higher taxes if the marginal value of the public good is higher. However, voters agree on the optimal tax if and only if the expected marginal value of the public good is zero. Political disagreement increases with $E[\theta|q]$:

$$\begin{aligned} D(q) &= (1 - x^{A*}(q))\beta^A + E[\theta|q](x^{A*}(q))^\psi - (1 - x^{B*}(q))\beta^B - E[\theta|q](x^{B*}(q))^\psi \\ &= \gamma E[\theta|q]^\rho, \end{aligned}$$

where $\gamma \equiv \psi^{\frac{\psi}{1-\psi}} \left\{ (1-\psi)(\beta^A)^{\frac{-\psi}{1-\psi}} + \beta^A \psi (\beta^B)^{\frac{-1}{1-\psi}} - (\beta^B)^{\frac{-\psi}{1-\psi}} \right\} > 0$ and $\rho \equiv \frac{1}{1-\psi} > 1$.

Independently of whether the majority group is richer or poorer than the minority group (β^A is higher or lower than β^B), disagreement is a power function of the expected state and satisfies the conditions of Proposition 3. To maximize the majority candidate's victory probability, the IP's optimal experiment either partially reveals that the public good is sufficiently important, or fully reveals that the public good has a low marginal value.

²¹Without loss of generality, let $m^A \beta^A + m^B \beta^B = 1$, where m^i is the number of voters in group i . Total tax revenue is then x .

5 Alternative Interpretations

In this section, we emphasize that our model can take on many other interpretations, and, hence, our results apply to a wide set of economic models. In particular, we can change the interpretation and structure of the policy space, the set of voters, the public signal, and the random shock on voters' preferences. To illustrate the flexibility of our model, we next present three alternative interpretations. In these applications, we study multiple policy dimensions; we replace the assumption of a majority and a minority group by one decisive median voter; we interpret the policy experiment as disclosure of information about a major policy reform; and we consider the case in which the IP strategically controls information about its competence.

5.1 The Relative Importance of Policy Dimensions

There is a single policy issue in the spatial policy model of Section 2.4 and in the optimal tax model of Section 4.2. Moreover, information about θ induces politicians to re-evaluate their beliefs and choose a new policy. In other important cases, the policy issue is multidimensional, and voters and politicians are convinced about what the optimal policy is, but they are uncertain about the relative importance of different policy dimensions.

To study these cases, consider the following alternative model. There are $d \geq 2$ policy dimensions (e.g., public education, public health, national defense, etc.). A policy is a d -dimensional vector $x = (x_1, \dots, x_d)$. The preferences of voter $i \in \{A, B\}$ are captured by the preference vector $\beta^i = (\beta_1^i, \dots, \beta_d^i)$ and by the loss function l , with $l(0) = 0$ and $l' > 0$. Voters' policy payoff is

$$u^i(x, \theta) = \sum_{j=1}^d -\lambda_j(\theta)l(|x_j - \beta_j^i|),$$

where each function $\lambda_j(\theta)$ captures the relative importance of policy dimension j given state θ , with $\sum_{j=1}^d \lambda_j(\theta) = 1$ and $\lambda_j(\theta) > 0$. Note that voters' preferred policies are independent of beliefs about θ , $x_j^{i*}(q) = \beta_j^i$. Although voters know their preferred policies for education and national defense, they are uncertain about which policy issue will be more important during the next term.

The degree of political disagreement, from the point of view of voters in group A , is

simply the expected (weighted) loss from the policy of candidate B ,

$$\begin{aligned} D(q) &= E \left[\sum_{j=1}^d -\lambda_j(\theta)l(|\beta_j^A - \beta_j^A|) - \sum_{j=1}^d -\lambda_j(\theta)l(|\beta_j^B - \beta_j^A|) \right] \\ &= E \left[\sum_{j=1}^d \lambda_j(\theta)l(|\beta_j^B - \beta_j^A|) \right]. \end{aligned}$$

Given the valence of the incumbent and the valence distribution of the challenger, victory probability is $F(D(q) + v^A)$, as before. To apply Proposition 3, rewrite the unknown state as follows. For each $\theta \in \Theta$, compute $\theta' \equiv \sum_{j=1}^d \lambda_j(\theta)l(|\beta_j^B - \beta_j^A|)$. Define a new state space Θ' as the collection of θ' . We can then rewrite disagreement simply as the expected value of θ' , so that victory probability becomes $F(E[\theta'] + v^A)$ and we can apply Proposition 3.

In summary, voters have a fundamental disagreement over the optimal policy but are uncertain about how important each policy dimension will be. For instance, suppose that there are only two issues, and voters disagree relatively more on national defense and less on education (that is, $|\beta_j^B - \beta_j^A|$ is larger for the national defense dimension). The incumbent's optimal experiment pools together states that attach more weight to national defense and fully reveals states that attach more weight to education. That is, the optimal experiment either reveals that the controversial national defense issue will be "sufficiently important" in the upcoming years, or it fully reveals that the more agreed-upon education issue will be more important.

We can extend this model to study wedge issues.²² For example, suppose that there are three groups of voters, $i \in \{A, B, C\}$. No single group forms a majority, but any pair of groups forms a majority. The incumbent party \mathcal{A} is committed to the preferred policy of group A , while party \mathcal{B} is committed to the preferred policy of group B . Voters in group C have preferences β_j^C that are the same as voters' in group A in some dimensions, but that are the same as voters' in group B in the remaining dimensions. Therefore, in equilibrium, voter C servers as a decisive median voter: for any realized valences v^A and v^B , and for any belief q , if voters in group C prefer candidate A over candidate B , then voters in group A also prefer candidate A ; if group C prefers candidate B over candidate A , then group B also prefers candidate B . Consequently, the incumbent party would like to convince the decisive voter C that the issues on which they share the same preferences are important, while the

²²We thank Ben Golub for suggesting this question.

issues on which they disagree are not important. The optimal experiment of the incumbent party is designed to endogenously select the wedge issues that it wants to emphasize.

5.2 The Rollout of a Major Reform

Suppose that the incumbent has implemented a major policy reform (e.g., a major change in the healthcare system). However, it takes time for voters to observe the true long-run payoff consequences of this complex new law. During the rollout, the government can choose which information to publicly collect about the initial effects of the reform. Let the state space Θ represent how much information the government can potentially collect, in the short-run, about the long-run consequences of the new policy. Let $E[\theta|q]$ be the expected payoff for the majority voter from keeping the reform, and normalize to zero the payoff from reversing the reform. To simplify exposition, suppose that, for all short-run information in Θ , the incumbent party wants to keep the reform, and the opposing party wants to reverse the reform. Given the valences, the incumbent's re-election probability is then $F(E[\theta|q] + v^A)$ and we can apply Proposition 3. An incumbent with higher valence is less transparent about the early effects of her policy reform. Similarly, government transparency decreases in players' optimism about the reform — if players' prior beliefs regarding θ are high (in the sense of a location shift in the prior distribution), then the government becomes less transparent.

5.3 Strategic Information about Competence

Our main results continue to hold if we invert the interpretation of the model: the incumbent party uses the experiment to strategically reveal information about its competence, while voters observe an exogenous signal about the relative payoff of parties' policies. Moreover, instead of having majority and minority groups of voters, we can assume a continuum of voters with a representative median voter.

To illustrate this equivalence, consider the following alternative model. The incumbent party is committed to a policy x^A , while the opposing party is committed to a different policy x^B . Without loss of generality, let the incumbent party be the right-wing party $x^A = +1$ and the challenger be the left-wing party $x^B = -1$. There is a continuum of voters indexed by their ideology $y \in [\underline{y}, \bar{y}]$. Let y_m be the voter with the median ideology. For voter y , the

payoff from electing candidate $j \in \{A, B\}$ is $v^j - (x^j - (y - \theta))^2$, where v^j is the valence of candidate j and θ is the current state of the economy.

The incumbent was recently elected and players do not know her realized valence $v^A \in V^A \equiv \{v_1^A, \dots, v_N^A\}$, with a finite $N \geq 2$ and $v_j^A < v_{j+1}^A$. Players share a common prior belief p in the interior of the simplex $\Delta(V^A)$. The incumbent can then design a public signal π that will reveal information about her competence v^A . After voters observe the signal realization s of π , they reach the common posterior belief q and update their belief about the incumbent's competence $E[v^A|q]$. Voters then observe the realized state of the economy θ , distributed according to the c.d.f. F , which satisfies **(A1)**. The expected valence of the challenger is \bar{v}^B , and players do not observe further information about it.

Given any posterior belief q , it is straightforward to verify that the median voter is decisive, and the incumbents' victory probability is

$$W(q; y_m) = F\left(\frac{E[v^A|q] - \bar{v}^B}{4} + y_m\right).$$

Hence, we can apply Proposition 3: the incumbent's optimal choice is an upper-censoring experiment on competence v^A . Moreover, as we move the median voter's ideology y_m to the right, it increases the advantage of the right-wing incumbent. Hence, the right-wing incumbent chooses to be Blackwell less informative about its competence when facing a more right-wing voting district, and more informative when facing a more left-wing district.

6 The Role of Belief Disagreement

Heterogeneous prior beliefs play an important role in politics — see Millner, Ollivier, and Simon (2014) for a recent review of the literature on heterogeneous priors in politics. We now extend our analysis to the case in which voters in the same group share a common prior belief, but voters in opposite groups openly disagree over the likelihood of state θ .

Formally, voters in group i have a common prior belief p^i in the interior of the simplex $\Delta(\Theta)$, but priors differ across groups, $p^A \neq p^B$. Each party shares its affiliates' beliefs. Preferences and prior beliefs are common knowledge — voters “agree to disagree.” If we interpret θ as describing the mapping between policy x and outcomes, then different priors represent differences in voters' views about the outcomes that different government policies produce.

Let q^A and q^B be the posterior beliefs of voters in groups A and B after observing the experiment's results. Voters can correctly predict the policies $x^{A*}(q^A)$ and $x^{B*}(q^B)$ that each candidate would implement if elected. From the point of view of voters in group A , political disagreement (3) becomes

$$\mathcal{D}(q^A, q^B) \equiv \sum_{\theta \in \Theta} q_{\theta}^A [u^A(x^{A*}(q^A), \theta) - u^A(x^{B*}(q^B), \theta)]. \quad (9)$$

As in the common priors case, voters in majority group A are decisive, and candidate A wins with probability $F(\mathcal{D}(q^A, q^B) + v^A)$. Again, candidate A wins the election with a probability that increases in the degree of political disagreement — candidate A has a “policy advantage” because a majority of voters believe that she has not only the “correct” preference, but also the “correct” belief, and, thus, she will implement the “correct” policy.

Let $r_{\theta} \equiv \frac{p_{\theta}^B}{p_{\theta}^A}$ and $r \equiv (r_{\theta})_{\theta \in \Theta}$ capture the likelihood ratio of prior beliefs. We can then use the results from Alonso and Câmara (2016a) to express disagreement $\mathcal{D}(q^A, q^B)$ as a function $D(q^A)$, which depends only on the beliefs of voters in group A :

$$D(q^A) \equiv \mathcal{D}\left(q^A, q^A \frac{r}{\langle q^A, r \rangle}\right).$$

Victory probability then becomes $F(D(q^A) + v^A)$, and our results continue to hold for this new function — see online Appendix B for a more detailed analysis.

6.1 Increasing Belief Disagreement

To shed some light on the role of belief disagreement, we now focus on cases in which all voters share the same preferences, so that political disagreement is zero when voters hold a common belief. As Callander (2011, pg. 657) notes, “[o]n some policy issues it is conceivable that we all share common outcome preferences (or at least similar preferences), yet we disagree as to how best to go about achieving the desired outcome. [...] Viewed this way, much political disagreement is over beliefs rather than outcomes.”

For example, consider the spatial policy model from Section 2.4. Suppose that voters share the same payoff function $u^A(x, \theta) = u^B(x, \theta) = -(x - \theta)^2$. Recall that the optimal policy is $x^{i*}(q^i) = E[\theta|q^i]$. Political disagreement (9) translates naturally into the degree of

belief disagreement over expectations,

$$\begin{aligned}
\mathcal{D}(q^A, q^B) &= \sum_{\theta \in \Theta} q_{\theta}^A [u^A(x^{*A}(q^A), \theta) - u^A(x^{*B}(q^B), \theta)] \\
&= \sum_{\theta \in \Theta} q_{\theta}^A [-(E[\theta|q^A] - \theta)^2 + (E[\theta|q^B] - \theta)^2] \\
&= (E[\theta|q^A] - E[\theta|q^B])^2.
\end{aligned} \tag{10}$$

Similarly, consider the tax model from Section 4.2. Suppose that voters have the same income normalized to one, $\beta^A = \beta^B = 1$, and the production technology of the public good is $\psi = \frac{1}{2}$. The optimal tax rate becomes $x^{i*}(q^i) = \frac{E[\theta|q^i]^2}{4}$, and political disagreement takes the simple form $\mathcal{D}(q^A, q^B) = \frac{1}{4}(E[\theta|q^A] - E[\theta|q^B])^2$. In this case, disagreement over the optimal tax derives solely from the belief disagreement over the marginal value of the public good.

When voters have the same payoff function but different beliefs, can the IP increase political disagreement? As in these two applications, suppose that political disagreement is a strictly increasing function of the difference between voters' expectation over the state. Although voters share a common payoff function, we show in the next proposition that if the state space is rich enough, then the IP can generically design an experiment that increases political disagreement with probability one.

Proposition 4 *Suppose that political disagreement strictly increases in the degree of belief disagreement over expectations, $\mathcal{D}(q^A, q^B) = R(|E[\theta|q^A] - E[\theta|q^B]|)$, where $R \geq 0$ and $R' > 0$. If $N \geq 4$, then the IP can generically²³ design an experiment that increases political disagreement with probability one. Consequently, if F has full support on the real numbers, the value of persuasion is positive for each finite incumbent's valence v^A .*

The following example illustrates how the IP can guarantee a higher disagreement.²⁴

Example 4 — Increasing Belief Disagreement: Let $\Theta = \{1, 2, 3, 4\}$. Consider priors $p^A = (.05, .45, .45, .05)$ and $p^B = (.45, .05, .05, .45)$, so that $E[\theta|p^A] = E[\theta|p^B] = 2.5$. Although prior beliefs are different, initial political disagreement is zero. The following binary experiment $S = \{s_1, s_2\}$ is optimal for the IP. States 1 and 3 induce signal s_1 with

²³Genericity is interpreted over the space of pairs of prior beliefs.

²⁴See Kartik, Lee and Suen (2015) for conditions such that a Blackwell more informative experiment, on average, brings posterior beliefs closer to each other.

probability one, while states 2 and 4 induce signal s_2 with probability one. After observing signal s_1 beliefs become $E[\theta|q^A] = 2.8$ and $E[\theta|q^B] = 1.2$, while s_2 induces $E[\theta|q^A] = 2.2$ and $E[\theta|q^B] = 3.8$. Therefore, every realization induces a strictly higher belief disagreement. From the point of view of voters A , candidate B not only “overreacts” to the information (updates her policy “too much”), but also moves the policy in the wrong direction. \square

7 Extensions

In this section, we consider other extensions of our model. We discuss the case in which the IP supports the minority candidate, the case in which parties are both office- and policy-motivated, and the impact of competition in information provision — when the opposing party can also generate some information about the state. We discuss additional extensions in online Appendix B (costly experiments, post-election information, and valence shocks that are independent across voters).

IP Supports the Minority: So far, we have assumed that the IP supports the majority candidate — that is, candidate A is the incumbent. Now suppose that the minority party \mathcal{B} is in power (hence controls the experiment) and supports the incumbent candidate B . Since the political advantage of the majority candidate is due solely to political disagreement, the IP now benefits from *decreasing* political disagreement. The results from Section 3 now apply to the valence v^B of the incumbent: the IP uses less-informative experiments when the minority incumbent is more competent (v^B is high) and more-informative experiments when she is less competent. Interestingly, the optimal experiment in Proposition 3 becomes *lower-censoring*: the minority party pools low disagreement states and fully reveals high disagreement states.

Moreover, consider the models of Section 6.1, in which citizens share the same payoff function but hold different prior beliefs. In these cases, regardless of priors, full information disclosure is *always* optimal for the minority candidate. Complete transparency eliminates political disagreement and the policy advantage of the majority candidate, thus increasing the chances of the minority candidate. Therefore, for policy issues in which political disagreement derives solely from belief disagreement, we should empirically observe that policy experiments

by minority incumbents are more informative than those of majority incumbents.

Policy-Motivated Parties: In this paper, we focus on a purely office-motivated incumbent party, whose primary concern is to be re-elected. Consider, now, the opposite case: a purely policy-motivated party. Suppose that each party has the same payoff function as the voter it represents, and parties do not receive any direct benefit from holding office. That is, when party $i \in \{A, B\}$ is the incumbent, it chooses the experiment that maximizes the expected payoff of voter i .

Intuitively, if the payoff functions of the two groups are sufficiently aligned, then a fully informative experiment is optimal, independently of the incumbent's valence. This is in spite of the fact that a fully informative experiment might reduce the probability of re-electing the incumbent. That is, the policy-motivated IP is willing to sacrifice its re-election probability in order to generate more information to both candidates and guarantee a better policy.

However, the incumbent faces a more intricate problem when there is a large conflict of preferences (for example, see the preferences in Proposition 2). In this case, the IP prefers a more-informative experiment when the incumbent's valence is large and she is almost sure to win. In this case, it is valuable to provide information to the likely winner. In contrast, when the incumbent is very incompetent and likely to lose, the IP prefers to implement a less-informative experiment. Again, the policy-motivated IP is willing to forgo the possibility of implementing a more-informative experiment — and, hence, increase its re-election probability — simply to avoid detrimental policies from the opposing party.

Interestingly, it seems that the purely policy-motivated case resembles the case of a purely office-motivated party, with a log-convex valence distribution. If, empirically, one finds that more competent incumbents tend to implement weakly more-informative experiments, then it is harder to say whether politicians are policy-motivated or office-motivated with a log-convex distribution of valence. However, if, empirically, one finds that more competent incumbents choose less-informative experiments, then it is more likely that parties are purely office-motivated with a log-concave valence distribution.

It would be interesting to consider an alternative model in which parties are both policy- and office-motivated. We leave this promising agenda to future work.

Competition-in-persuasion: In this paper, we have focused on the case in which the

incumbent has the monopoly over the information that reaches voters. What happens if the challenger can launch her own public investigation? We next describe some results on this “competition-in-persuasion” game (see the online Appendix for details).

The timing of this extended game is as follows. The incumbent party implements an experiment π , and its outcome becomes public. The opposing party then chooses an experiment, and its result becomes public.²⁵ The valence of the challenging candidate is realized and becomes public information. The election takes place. We consider two cases. In the first case, the challenger is unconstrained — she has access to *every* experiment that is correlated with the state. In the second case, the challenger is constrained on her access to experiments.

When the challenger is unconstrained, we show that there is always a subgame perfect equilibrium in which the incumbent selects a fully informative experiment. The intuition behind the result follows from the fact that parties have opposite preferences over the disclosure of information. Loosely speaking, if the incumbent benefits from “garbling” certain information, then the challenger benefits from disclosing it. Therefore, the incumbent can do no better than fully disclosing the state.

In practice, however, the incumbent typically has access to a richer set of experiments than the challenger does since the incumbent directly controls the government. How do constraints in the challenger’s access to experiments alter, in equilibrium, the information that reaches voters? To provide some insights into this question, in online Appendix B we consider an information technology in which the challenger’s ability to launch a fully informative investigation is captured by an exogenous technology parameter $\alpha \in (0, 1)$. We show that, in equilibrium, voters have access to more information if the challenger has easier access to the government’s information (a higher α).

8 Conclusion

In this paper, we study the strategic control of information by an incumbent who wants to be re-elected. The incumbent, through her direct control of the government, is in a privileged position to control the information that reaches voters. For example, she can run a small-scale

²⁵We believe that the most natural assumption for our model is to have the incumbent playing first. See Gentzkow and Kamenica (2016) for a model in which players choose experiments simultaneously.

pilot test of a novel policy, design an experiment to evaluate unobserved effects of existing policies, decide which information is released during the early stages of a complex policy reform, and establish disclosure rules for government agencies. In all these cases, the public information generated by the government can affect the incumbent’s re-election probability.

In our benchmark model, information changes the degree of political disagreement and sways future elections — experimental outcomes that increase disagreement increase the victory probability of the candidate whose preferences and beliefs are similar to those of a majority of voters. Therefore, an incumbent supported by the majority benefits from policy experiments that create more dissent between the majority and the minority. We derive conditions such that more-competent politicians are less informative than incompetent politicians, and conditions for an upper-censoring experiment to be optimal — it fully reveals low disagreement states and pools high disagreement states. Finally, we consider cases in which all voters share the same payoff function, so that political disagreement is due solely to belief disagreement. We show that, even in these cases, policy experiments can be used to increase disagreement.

In this paper, we focus on an incumbent who has the monopoly over the information generated by the government. In this case, she often chooses to hide some information, in order to maximize her re-election probability. We then extend our model to consider competition in information provision between the incumbent and the challenger. We show that if the challenger has full access to the same government information as the incumbent, then there always exists an equilibrium in which competition forces the incumbent to be fully transparent. However, we believe that, in most cases, the challenger has only limited access to government information. In this case, we show that the incumbent retains the incentives to hide some information. That is, if the challenger has less-than-perfect access to information, then voters might not become fully informed. We hope that our model can be further extended to study the roles of different political institutions in providing the correct incentives to incumbent politicians. For instance, in a related model, Bernecker, Boyer and Gathmann (2015) provide empirical evidence that term limits affect incumbents’ incentives to experiment. It would be important to also understand how the decentralization of information generation, from the federal government to the states, affects the behavior of governors from different parties, who are concerned with re-election.

A Appendix

Before we present the proof of Lemma 1, we provide the following lemma.

Lemma A.1 *Fix any $a, b, c \in \mathbb{R}$. Define*

$$G(a, b, v^A) \equiv \frac{F(b + v^A) - F(a + v^A)}{f(a + v^A)}. \quad (11)$$

If F satisfies (A1), then $G(a, b, v^A)$ is non-increasing in v^A .

Proof of Lemma A.1: We first rewrite the function G as

$$G(a, b, v^A) = \int_0^{b-a} \frac{f(a + v^A + z)}{f(a + v^A)} dz.$$

Since f is log-concave, it exhibits decreasing ratios in the sense that for every $z > 0$ and $v^A \geq v^{A'}$ we have

$$\frac{f(a + v^{A'} + z)}{f(a + v^{A'})} \geq \frac{f(a + v^A + z)}{f(a + v^A)}. \quad (12)$$

Suppose first that $b > a$. Then integrating both sides of (12) between 0 and $b - a$ shows that $G(a, b, v^{A'}) \geq G(a, b, v^A)$. Now suppose that $a > b$. Then for any $z \in [0, a - b]$ we can rewrite (12) as

$$\frac{f(a + v^A - z)}{f(a + v^A)} \geq \frac{f(a + v^{A'} - z)}{f(a + v^{A'})}.$$

Integrating between 0 and $b - a$ we conclude that

$$-G(a, b, v^A) = \int_0^{a-b} \frac{f(a + v^A - z)}{f(a + v^A)} dz \geq \int_0^{a-b} \frac{f(a + v^{A'} - z)}{f(a + v^{A'})} dz = -G(a, b, v^{A'}),$$

or, in other words, $G(a, b, v^{A'}) \geq G(a, b, v^A)$. ■

Proof of Lemma 1: Consider an experiment π that generates a distribution $\sigma \in \Delta(\Delta(\Theta))$ over posterior beliefs. Note that this distribution is independent of valences. For any q in the support of σ , the change in the victory probability of the majority candidate is

$$W(q; v^A) - W(p; v^A) = F(D(q_A) + v^A) - F(D(p) + v^A) = f(D(p) + v^A)G(D(p), D(q), v^A),$$

where G is defined by (11). Therefore, the expected change in victory probability from experiment π can be written as

$$E_\pi[W(q; v^A) - W(p; v^A)] = f(D(p) + v^A) \int_{q \in \text{supp}(\sigma)} G(D(p), D(q), v^A) d\sigma.$$

Because $f > 0$ rewrite

$$\frac{E_\pi[W(q; v^A) - W(p; v^A)]}{f(D(p) + v^A)} = \int_{q \in \text{supp}(\sigma)} G(D(p), D(q), v^A) d\sigma. \quad (13)$$

From Lemma A.1, we know that G is non-increasing in v^A , hence the LHS of (13) is non-increasing in v^A . This implies that if $E_\pi[W(q; v^A) - W(p; v^A)] \leq 0$ then $E_\pi[W(q; v^{A'}) - W(p; v^{A'})] \leq 0$ for any $v^{A'} > v^A$, concluding the proof. ■

Proof of Proposition 1: Suppose π^* is an optimal experiment given valence v^A . Take any $v^{A'} > v^A$ and any π' that is Blackwell more informative than π^* . The proof has two steps. In the first step, we construct a sequential experiment $\{\pi^*, \{\pi_s\}_{s \in S}\}$ that is payoff equivalent to π' . In the second step, we show that since π^* is weakly better than $\{\pi^*, \{\pi_s\}_{s \in S}\}$ when valence is v^A , then π^* is weakly better than $\{\pi^*, \{\pi_s\}_{s \in S}\}$ when the valence is higher. Consequently, π^* is weakly better than π' for any $v^{A'} > v^A$.

Step 1: If π' , with realizations $z_{\pi'} \in Z_{\pi'}$, is Blackwell more informative than π^* , with realizations $s \in S$, then there exist an stochastic transformation $\gamma(s|z_{\pi'})$ such that

$$\Pr[\theta, s] = \sum_{z_{\pi'} \in Z_{\pi'}} \gamma(s|z_{\pi'}) \Pr[\theta, z_{\pi'}]. \quad (14)$$

Let $\tau(z_{\pi'}; s)$ be

$$\tau(z_{\pi'}; s) = \frac{\gamma(s|z_{\pi'}) \Pr[z_{\pi'}]}{\sum_{\hat{z}_{\pi'} \in Z_{\pi'}} \gamma(s|\hat{z}_{\pi'}) \Pr[\hat{z}_{\pi'}]} = \gamma(s|z_{\pi'}) \frac{\Pr[z_{\pi'}]}{\Pr[s]} \quad (15)$$

and define the experiment π_s as the experiment that, when voters have belief q_s , it leads to a posterior $q_{z_{\pi'}}$ with probability $\tau(z_{\pi'}; s)$. In other words, experiment π_s is described by the conditional probabilities $\pi_s(z_{\pi'}|\theta) = q_{z_{\pi'}}^\theta \tau(z_{\pi'}; s) / q_s^\theta$ (see KG). We now show that this experiment is well defined. First, it is immediate that $\tau(z_{\pi'}; s) \geq 0$ with $\sum_{z_{\pi'} \in Z_{\pi'}} \tau(z_{\pi'}; s) = 1$. Next, using (15) we have

$$q_s^\theta = \frac{\Pr[\theta, s]}{\Pr[s]} = \frac{\sum_{z_{\pi'} \in Z_{\pi'}} \gamma(s|z_{\pi'}) \Pr[\theta, z_{\pi'}]}{\sum_{\hat{z}_{\pi'} \in Z_{\pi'}} \gamma(s|\hat{z}_{\pi'}) \Pr[\hat{z}_{\pi'}]} = \sum_{z_{\pi'} \in Z_{\pi'}} \tau(z_{\pi'}; s) \frac{\Pr[\theta, z_{\pi'}]}{\Pr[z_{\pi'}]} = \sum_{z_{\pi'} \in Z_{\pi'}} \tau(z_{\pi'}; s) q_{z_{\pi'}}^\theta,$$

so that experiment π_s is Bayes feasible.

Finally, we show that the observation by voters of the outcomes of the sequential experiment $\{\pi^*, \{\pi_s\}_{s \in S}\}$ leads to the same joint distribution over posteriors and the state as experiment π' . Let $\cup_{s \in S} \{s, z_{\pi'}\}$ be the event in which voters have posterior $q_{z_{\pi'}}$. We have

$$\begin{aligned} \Pr[\theta, \cup_{s \in S} \{s, z_{\pi'}\}] &= \sum_{s \in S} \pi_s(z_{\pi'} | \theta) \Pr[\theta, s] = \sum_{s \in S} \frac{q_{z_{\pi'}}^\theta \tau(z_{\pi'}; s)}{q_s^\theta} \Pr[\theta, s] \\ &= \sum_{s \in S} q_{z_{\pi'}}^\theta \tau(z_{\pi'}; s) \Pr[s] = q_{z_{\pi'}}^\theta \sum_{s \in S} \gamma(s | z_{\pi'}) \Pr[z_{\pi'}] \\ &= q_{z_{\pi'}}^\theta \Pr[z_{\pi'}] = \Pr[\theta, z_{\pi'}]. \end{aligned}$$

Step 2: When the incumbent's valence is v^A , optimality of π^* implies that IP does not benefit from further disclosing information after each signal realization s of π^* . That is, for every posterior belief q_s in the support of π^* and every experiment π_s , we have

$$E_{\pi_s}[W(q; v^A) | q_s] \leq W(q_s; v^A). \quad (16)$$

Apply Lemma 1-(i) to (16): for each posterior belief q_s in the support of π^* , for every $v^{A'} > v^A$, and every experiment π_s , we have $E_{\pi_s}[W(q; v^{A'}) | q_s] \leq W(q_s; v^{A'})$. Taking expectations over the realizations of π^* yields $E_{\{\pi^*, \{\pi_s\}_{s \in S}\}}[W(q; v^{A'})] \leq E_{\pi^*}[W(q; v^{A'})]$. ■

Proof of Corollary 1: Suppose that for some v_2^A a completely uninformative experiment is optimal, and note that every possible experiment is Blackwell more informative than no information. Then Proposition 1 immediately implies that a completely uninformative experiment is weakly better than every other experiment for any $v^A > v_2^A$.

Suppose that for some v_1^A the fully informative experiment π^{FD} is optimal. Alonso and Câmara (2016a, Corollary 2) show that a fully informative experiment is optimal if and only if $E_{\pi^{FD}}[W(q'; v_1^A) | q] \geq W(q; v_1^A)$ for all $q \in \Delta(\Theta)$. Lemma 1 implies that²⁶ for every $v^A < v_1^A$, we have $E_{\pi^{FD}}[W(q'; v^A) | q] \geq W(q; v^A)$ for all $q \in \Delta(\Theta)$. Hence, π^{FD} is optimal for all $v^A < v_1^A$. ■

Proof of Proposition 2: Suppose $\beta^A \beta^B < 0$ or $|\beta^B| > 2|\beta^A|$. This implies that $\beta^A \neq \beta^B$ and $\beta^B(2\beta^A - \beta^B) < 0$.

²⁶Lemma 1 implies that if $E_\pi[W(q; v^A) - W(p; v^A)] \geq 0$, then $E_\pi[W(q; v^{A'}) - W(p; v^{A'})] \geq 0$ for any $v^{A'} < v^A$.

We first show that, from the point of view of voter A , the expected policy payoff if candidate B is elected is strictly lower if the candidate observes a fully informative experiment, compared to no information. That is, full information makes candidate B choose a strictly worse policy on average.

Without further information, candidate B chooses policy $\beta^B E[\theta|p]$, which yields expected policy payoff $E[-(\beta^B E[\theta|p] - \beta^A \theta)^2|p]$ to voter A . With a fully informative signal, candidate B chooses policy $\beta^B \theta$ after learning that the state is θ . This yields an expected policy payoff $E[-(\beta^B \theta - \beta^A \theta)^2|p]$ to voter A . No information yields a strictly higher payoff than full information if and only if

$$\begin{aligned} E[-(\beta^B E[\theta|p] - \beta^A \theta)^2|p] &> E[-(\beta^B \theta - \beta^A \theta)^2|p] \\ -(\beta^B)^2 E[\theta|p]^2 + 2\beta^A \beta^B E[\theta|p]^2 - (\beta^A)^2 E[\theta^2|p] &> -(\beta^B)^2 E[\theta^2|p] + 2\beta^A \beta^B E[\theta^2|p] - (\beta^A)^2 E[\theta^2|p] \\ (2\beta^A \beta^B - (\beta^B)^2) E[\theta|p]^2 &> (2\beta^A \beta^B - (\beta^B)^2) E[\theta^2|p] \\ 0 &> \beta^B (2\beta^A - \beta^B) (E[\theta^2|p] - E[\theta|p]^2). \end{aligned}$$

Since the variance $(E[\theta^2|p] - E[\theta|p]^2)$ is strictly positive given any interior prior belief, the inequality holds if and only if $0 > \beta^B (2\beta^A - \beta^B)$, concluding this step of the proof.

Disagreement is a convex function of the posterior belief, $D(q) = (\beta^B - \beta^A)^2 (E[\theta|q])^2$. Consequently, if v^A is sufficiently low, then the IP's optimal experiment is fully informative. From the point of view of voter A , compared to no information, full information leads candidate B to choose a worse policy on average, while it leads candidate A to choose a better policy when elected. Moreover, if v^A is sufficiently low, then candidate B is sufficiently likely to win the election, and the strictly negative effect of a worse policy from candidate B dominates the positive effect from the better policy from candidate A . ■

In Proposition A.1 below, we show that upper-censoring is an optimal experiment for a large class of Bayesian persuasion games. Then, in the proof of Proposition 3, we show that our model satisfies the conditions of Proposition A.1.

Proposition A.1 *Consider a Bayesian persuasion game between a sender and a receiver, as in KG. Suppose that the sender's payoff $u_S(a, \theta)$ and receiver's optimal action $a^*(q)$ satisfy*

$$\sum_{\theta \in \Theta} q_\theta u_S(a^*(q), \theta) = K(E[\theta|q]), \quad (17)$$

where $K(\cdot)$ is twice differentiable and strictly increasing, and $E[\theta|q]$ denotes the expected state given posterior belief q . If K' is single-peaked (single-dipped), then there exists an optimal signal that is upper-censoring (lower-censoring).

Proof of Proposition A.1:

Note that $U_S(q) = \sum_{\theta \in \Theta} q_\theta u_S(a^*(q), \theta)$ is the sender's expected payoff, as a function of posterior belief q . Therefore, (17) implies that the sender's expected utility depends on posterior beliefs only through the posterior expectation of the state. Furthermore, function K is assumed to be twice differentiable, with a strictly positive derivative $K'(E) \equiv \left. \frac{dK(e)}{de} \right|_{e=E} > 0$.

First, suppose that K' is single-peaked — that is, there exists an \bar{E} in the extended real line such that $K''(E) \geq 0$ for all $E < \bar{E}$, and $K''(E) \leq 0$ for all $E > \bar{E}$. Consequently, K is locally convex in the range $E < \bar{E}$, and locally concave in the range $E > \bar{E}$. Since K' might be “flat” at its peak, we define \bar{E} as the lowest expectation at the peak. That is, \bar{E} is defined such that $K'(E) < K'(\bar{E})$ for all $E < \bar{E}$, and $K'(\bar{E}) \geq K'(E)$ for all $E \geq \bar{E}$.

Since $\theta_1 < \dots < \theta_N$, players' posterior expectation of the state must be in $[\theta_1, \theta_N]$. If $\bar{E} \geq \theta_N$, then the sender's payoff is everywhere convex and a fully informative experiment is optimal; if $\bar{E} \leq \theta_1$, then the sender's payoff is everywhere concave and a completely uninformative experiment is optimal (see KG for details). Note that full disclosure and no disclosure are the extreme cases of upper-censoring, with cutoff states θ_N and θ_1 , respectively.

Now consider the remaining case: $\theta_1 < \bar{E} < \theta_N$. We next construct an optimal experiment that is upper-censoring. The proof has two steps.

Step 1) We first show that, among the class of optimal experiments, there is always one that induces at most one non-degenerate posterior belief. To see this, take any optimal experiment π^* and let σ^* be the distribution of posterior beliefs induced by this experiment. All beliefs q^- in the support of σ^* such that $E[\theta|q^-] < \bar{E}$ are in the locally convex region of K . Hence, the sender weakly benefits from further disclosing some information. All beliefs q^+ in the support of σ^* such that $E[\theta|q^+] \geq \bar{E}$ are in the locally concave region of K . Hence, the sender weakly benefits from combining all these beliefs into a single belief. Repeated use of this argument implies the following. There exists an experiment π' that (i) is weakly better than π^* , hence π' is also optimal, and (ii) letting σ' be the distribution of posterior beliefs induced by π' , there is at most one non degenerate belief in the support of σ' . This non degen-

erate belief is in the concave region of K , while every belief in the convex region is degenerate. **Step 2)** We now solve for the optimal experiment in the class of experiments that induce at most one non degenerate belief. Given Step 1, this experiment is then optimal for the sender when she is unconstrained in her choice of experiment.

Consider any experiment that induces at most one non degenerate belief. Without loss of generality, define the signal space as $S \equiv \{s_{\theta_1}, \dots, s_{\theta_N}, s_{\text{pooling}}\}$. Each state $\theta \in \Theta$ induces the pooling signal s_{pooling} with probability $\alpha_\theta \in [0, 1]$, and induces the fully revealing signal s_θ with probability $1 - \alpha_\theta$. Given $\alpha = (\alpha_{\theta_1}, \dots, \alpha_{\theta_N})$, let $q^+(\alpha) \equiv \left(\frac{\alpha_\theta p_\theta}{\sum_{\theta' \in \Theta} \alpha_{\theta'} p_{\theta'}} \right)_{\theta \in \Theta}$ be the updated posterior belief after observing s_{pooling} , and $E^+(\alpha) \equiv \langle q^+(\alpha), \theta \rangle = \frac{\sum_{\theta \in \Theta} \alpha_\theta p_\theta \theta}{\sum_{\theta \in \Theta} \alpha_\theta p_\theta}$ be the updated expectation of θ . The sender's problem then simplifies to choosing α that maximizes her expected payoff:

$$\max_{\alpha_\theta \in [0,1], \theta \in \Theta} \Pi(\alpha) \equiv \left(\sum_{\theta \in \Theta} \alpha_\theta p_\theta \right) K(E^+(\alpha)) + \sum_{\theta \in \Theta} (1 - \alpha_\theta) p_\theta K(\theta).$$

We now solve for an optimal α^* and show that the optimal experiment is upper-censoring, that is, there exists a cutoff state θ_k such that $\alpha_\theta^* = 0$ if $\theta < \theta_k$ and $\alpha_\theta^* = 1$ if $\theta > \theta_k$.

From Step 1, the pooling belief $q^+(\alpha^*)$ is in the concave region, $E^+[\alpha^*] \geq \bar{E}$. Moreover, since posterior beliefs in the convex region are degenerate, we have $\alpha_\theta^* = 1$ for all $\theta \geq \bar{E}$. Now consider the convex region $\theta < \bar{E}$. Taking the derivative of the objective function with respect to $\alpha_{\theta'}$ for each state $\theta' < \bar{E}$, and noting that $\frac{\partial E^+(\alpha)}{\partial \alpha_{\theta'}} = \frac{p_{\theta'}}{\sum_{\theta \in \Theta} \alpha_\theta p_\theta} [\theta' - E^+(\alpha)]$, we have

$$\begin{aligned} \frac{\partial \Pi(\alpha)}{\partial \alpha_{\theta'}} &= p_{\theta'} K(E^+(\alpha)) - p_{\theta'} K(\theta') + \left(\sum_{\theta \in \Theta} \alpha_\theta p_\theta \right) K'(E^+(\alpha)) \frac{\partial E^+(\alpha)}{\partial \alpha_{\theta'}} \\ &= p_{\theta'} \left\{ K(E^+(\alpha)) - K(\theta') + K'(E^+(\alpha)) [\theta' - E^+(\alpha)] \right\} \\ &= p_{\theta'} \int_{\theta'}^{E^+(\alpha)} [K'(E) - K'(E^+(\alpha))] dE. \end{aligned}$$

Since $p_{\theta'} > 0$, the derivative $\frac{\partial \Pi(\alpha)}{\partial \alpha_{\theta'}}$ has the same sign as

$$A(\theta') \equiv \int_{\theta'}^{E^+(\alpha)} [K'(E) - K'(E^+(\alpha))] dE.$$

Suppose that $\alpha_{\theta'}^* < 1$ for some $\theta' < \bar{E}$, which implies that $A(\theta') \leq 0$. Single-peakedness of K' implies that K' is increasing for $\theta < \bar{E}$. Therefore, $K'(\theta) \leq K'(\theta')$ for all $\theta < \theta'$, which

leads to $A(\theta) < A(\theta') \leq 0$ for all $\theta < \theta'$. This establishes that, for all $\theta < \theta'$, we must have $\alpha_{\theta}^* = 0$.

The same steps of the proof show the opposite (lower-censoring) result for the case of a single-dipped K' . ■

Proof of Proposition 3: Suppose that (A.1) and (A.2') hold. Then we can write the IP's payoff as a strictly increasing, twice differentiable function of the expected state, $K(E) = F(H(E) + v^A)$. Moreover, the derivative $K'(E) = f(H(E) + v^A)H'(E)$ is log-concave, therefore it is single-peaked. Consequently, the conditions of Proposition A.1 hold and there is an optimal experiment that is upper-censoring. The fact that the optimal censoring cutoff weakly decreases in v^A follows immediately from Proposition 1: strictly increasing the censoring cutoff increases the informativeness of the experiment, but the IP does not benefit from a more informative experiment if v^A increases. ■

Proof of Proposition 4: The proof has two steps.

Step 1) Define the vector $v = r(\theta - E[\theta|q^B])$, the linear subspaces $W_1 = \{x \in \mathbb{R}^{card(\Theta)} : \langle x, 1 \rangle = 0\}$ and $W_{\theta-v} = \{x \in \mathbb{R}^{card(\Theta)} : \langle x, \theta - v \rangle = 0\}$. In this first step, we prove that if the projections of θ and r are not negatively collinear with respect to $W_1 \cap W_{\theta-v}$, then there exists an experiment π where all signal realizations increase political disagreement.

Since $q_{\theta}^B = \frac{q_{\theta}^A r_{\theta}}{\langle q^A, r \rangle}$ (see Alonso and Câmara (2016a)), we can rewrite

$$\mathcal{D}(q^A, q^B) = R(|E[\theta|q^A] - E[\theta|q^B]|) = R\left(\left|\langle q^A, \theta \rangle - \frac{\langle q^A r, \theta \rangle}{\langle q^A, r \rangle}\right|\right) \equiv D(q^A).$$

Define $q^A = \varepsilon\lambda + p^A$, with $\lambda \in W_1 = \{x : \langle x, 1 \rangle = 0\}$ and $\varepsilon \in \mathbb{R}$, and let

$$L(\varepsilon; \lambda) = \langle q^A, \theta \rangle - \frac{\langle q^A r, \theta \rangle}{\langle q^A, r \rangle} = \varepsilon \langle \lambda, \theta \rangle + E[\theta|q^A] - \frac{\varepsilon \langle \lambda, r \theta \rangle + E[\theta|q^B]}{\varepsilon \langle \lambda, r \rangle + 1}.$$

Disagreement is a strictly increasing function of the absolute value of $L(\varepsilon; \lambda)$. First suppose that $L(\varepsilon; \lambda) \geq 0$. We will show that under the conditions of the proposition one can always find a vector of “marginal beliefs” λ' such that L achieves a local minimum with respect to ε at $\varepsilon = 0$. This means that along the line λ' and in a neighborhood of 0, any belief $q^A = \varepsilon\lambda' + p^A$ with $\varepsilon > 0$ increases L , and thus $D(q^A) > D(p^A)$, while any belief $q^A = \varepsilon\lambda' + p^A$ with $\varepsilon < 0$ also increases L , yielding $D(q^A) > D(p^A)$. That is, we have found collinear beliefs that can average to the prior and that increase D .

First, we have

$$\begin{aligned}\frac{dL}{d\varepsilon} &= \langle \lambda, \theta \rangle - \frac{\langle \lambda, r\theta \rangle - \langle \lambda, r \rangle E[\theta|q^B]}{(\varepsilon \langle \lambda, r \rangle + 1)^2}, \\ \frac{d^2L}{d\varepsilon^2} &= \frac{2 \langle \lambda, r \rangle [\langle \lambda, r\theta \rangle - \langle \lambda, r \rangle E[\theta|q^B]]}{(\varepsilon \langle \lambda, r \rangle + 1)^3}.\end{aligned}$$

For $L(\varepsilon; \lambda)$ to achieve a local minimum at $\varepsilon = 0$, it is sufficient to exist $\lambda \in W$ such that

$$\left. \frac{dL}{d\varepsilon} \right|_{\varepsilon=0} = 0 \Rightarrow \langle \lambda, \theta \rangle = \langle \lambda, r(\theta - E[\theta|q^B]) \rangle, \quad (18)$$

$$\left. \frac{d^2L}{d\varepsilon^2} \right|_{\varepsilon=0} > 0 \Rightarrow \langle \lambda, r \rangle \langle \lambda, r(\theta - E[\theta|q^B]) \rangle > 0. \quad (19)$$

Since θ and r are not negatively collinear with respect to $W_1 \cap W_{\theta-v}$, then there exists $\lambda' \in W_1 \cap W_{\theta-v}$ with $\langle \lambda', \theta \rangle \langle \lambda', r \rangle > 0$ — see Alonso and Câmara (2016a). Since $\lambda' \in W_{\theta-v}$ then λ' satisfies (18). Then, given (18), the fact that $\langle \lambda', \theta \rangle \langle \lambda', r \rangle > 0$ implies that λ' also satisfies (19). Therefore, $L(\varepsilon; \lambda')$ achieves a local minimum at $\varepsilon = 0$.

Now consider the remaining case, $L(\varepsilon; \lambda) < 0$. Since disagreement strictly increases in the absolute value of L , we now can increase disagreement by decreasing L . The same steps of the proof above can be used to show that under the conditions of the proposition one can always find a vector of “marginal beliefs” λ'' such that L achieves a local maximum with respect to ε at $\varepsilon = 0$. This follows as the fact that θ and r are not negatively collinear with respect to $W_1 \cap W_{\theta-v}$ implies the existence of $\lambda'' \in W_1 \cap W_{\theta-v}$ with $\langle \lambda'', \theta \rangle \langle \lambda'', r \rangle < 0$ (see Alonso and Câmara 2016a), so that $L(\varepsilon; \lambda'')$ is locally concave at $\varepsilon = 0$. This concludes the first step of the proof.

Step 2) The previous step shows that if the projection of θ and r are not negatively collinear with respect to $W_1 \cap W_{\theta-v}$ then persuasion is valuable. We now show that negative collinearity of θ and r with respect to $W_1 \cap W_{\theta-v}$ is a non-generic property if $N \geq 4$. First note that $W_1 \cap W_{\theta-v}$ has at least dimension $N - 2$, and thus the projections of θ and r also have dimension $N - 2 \geq 2$. As collinearity is a non-generic property with vectors of dimension at least 2, this concludes the proof. ■

References

- [1] ALONSO, R. AND O. CÂMARA (2016a): “Bayesian Persuasion with Heterogeneous Priors,” *Journal of Economic Theory*, 165: 672-706.

- [2] ALONSO, R. AND O. CÂMARA (2016b): “Persuading Voters,” *American Economic Review*, forthcoming.
- [3] ALONSO, R. AND O. CÂMARA (2016c): “Common Information and Belief Disagreement,” *mimeo*.
- [4] BAGNOLI, MARK, AND TED BERGSTROM (2005): “Log-concave probability and its applications,” *Economic Theory*, 26(2): 445-469.
- [5] BIGLAISER, G., AND C. MEZZETTI (1997): “Politicians’ decision making with re-election concerns,” *Journal of Public Economics*, 66(3): 425-447.
- [6] BOLES LAVSKY, R. AND C. COTTON (2015): “Information and extremism in elections,” *AJEA: Microeconomics*, 7(1): 165-207.
- [7] BERNECKER, A., P. C. BOYER AND C. GATHMANN (2015): “Trial and Error? Re-election Concerns and Policy Experimentation during the US Welfare Reform,” *mimeo*.
- [8] CALLANDER, S. (2011): “Searching for good policies,” *American Political Science Review*, 105(04): 643-662.
- [9] CALLANDER, S. AND P. HUMMEL (2014): “Preemptive policy experimentation,” *Econometrica*, 82(4): 1509-1528.
- [10] CALLANDER, S. AND B. HARSTAD (2015): “Experimentation in Federal Systems,” *Quarterly Journal of Economics*, 130(2): 951-1002.
- [11] CARRILLO, J. D. AND M. CASTANHEIRA (2008) “Information and Strategic Political Polarisation,” *Economic Journal*, 118 (530): 845-874.
- [12] CARRILLO, J. D. AND T. MARIOTTI (2001) “Electoral competition and politician turnover,” *European Economic Review*, 45(1): 1-25.
- [13] DEWAN, T. AND R. HORTALA-VALLVE (2014): “Electoral Competition, Control and Learning,” LSE, *mimeo*.
- [14] DIXIT, A. K. AND J. W. WEIBULL (2007): “Political polarization,” *Proceedings of the National Academy of Sciences*, 104(18): 7351-7356.
- [15] DOWNS, A. (1957): *An Economic Theory of Democracy*, Harper and Row, New York.
- [16] DOWNS, G. W. AND D. M. ROCKE (1994): “Conflict, agency, and gambling for resurrection: The principal-agent problem goes to war,” *American Journal of Political Science*, 38(2): 362-380.

- [17] DUGGAN, J., AND C. MARTINELLI (2011): “A Spatial Theory of Media Slant and Voter Choice,” *Review of Economic Studies*, 78(2): 640-666.
- [18] FU, Q., AND M. LI (2014): “Reputation-concerned policy makers and institutional status quo bias,” *Journal of Public Economics*, 110: 15-25.
- [19] GENTZKOW, M., AND E. KAMENICA (2016): “Competition in Persuasion,” *Review of Economic Studies*, forthcoming.
- [20] GREENBERG, D. AND M. SCHRODER (2004): *The Digest of Social Experiments*, 3rd Edition. Washington D.C. Urban Institute Press.
- [21] GROSECLOSE, T. (2001): “A model of candidate location when one candidate has a valence advantage,” *American Journal of Political Science*, 45(4), pp. 862-886.
- [22] IYENGAR, S. AND A. F. SIMON (2000): “New perspectives and evidence on political communication and campaign effects.” *Annual review of psychology*, 51(1), pp. 149-169.
- [23] KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, pp. 2590-2615.
- [24] KARTIK, N., F. X. LEE, AND W. SUEN (2015): “A Theorem on Bayesian Updating and Applications to Signaling Games,” *mimeo*.
- [25] KOLOTILIN, A., M. LI, T. MYLOVANOV, AND A. ZAPECHELNYUK (2015): “Persuasion of a Privately Informed Receiver,” *mimeo*.
- [26] MAJUMDAR, S. AND S.W. MUKAND (2004): “Policy gambles,” *American Economic Review*, 94(4), pp. 1207-1222.
- [27] MATSUSAKA, J. G. (1995): “Explaining voter turnout patterns: An information theory,” *Public Choice*, 84(1-2), pp. 91-117.
- [28] MILLNER, A., H. OLLIVIER, AND L. SIMON (2014): “Policy experimentation, political competition, and heterogeneous beliefs,” *Journal of Public Economics*, 120, pp. 84-96.
- [29] STOKES, D. E. (1963) “Spatial models of party competition,” *American Political Science Review*, 57(02), pp. 368-377.
- [30] VAN DEN STEEN, E. (2011) “Overconfidence by Bayesian-Rational Agents,” *Management Science*, 57(5), pp. 884-896.
- [31] WILLENS, T. (2013) “Political Accountability and Policy Experimentation: Why to Elect Left-Handed Politicians,” Oxford U. Discussion Paper Series.