Retrospective Voting and Party Polarization*

Ignacio Esponda (WUSTL)

Demian Pouzo (UC Berkeley)

June 14, 2016

Abstract

We provide a new and favorable perspective on voter naiveté and party polarization. We study a model where two parties compete by committing to policies and voters subsequently vote for their preferred party. We contrast sophisticated with naive voting. The former is embodied by Nash equilibrium while the latter is formalized using the notion of a retrospective voting equilibrium (Esponda and Pouzo, forthcoming). Retrospective voters do not understand the mapping between states and outcomes induced by a policy; instead, they simply vote for the party that has delivered the best performance in the past. We show that parties have an incentive to polarize under retrospective, compared to Nash, voting. Moreover, this polarization often results in higher welfare due to a better match between policies and fundamentals.

*We thank Heski Bar-Isaac, Ernesto Dal Bó, Erik Eyster, Alessandro Lizzeri, Santiago Oliveros, Jesse Rothstein, Andy Skrzypacz, and several seminar participants for their comments. Esponda: Olin Business School, Campus Box 1133, Washington University, 1 Brookings Drive, Saint Louis, MO 63130, iesponda@wustl.edu; Pouzo: Department of Economics, UC Berkeley, 530 Evans Hall #3880, Berkeley, CA 94704, dpouzo@econ.berkeley.edu.

Contents

1 Introduction

There are two stylized facts in political environments that are often perceived to result in adverse policies. The first is that voters are not very sophisticated. The empirical evidence suggests that voters are often poorly informed and have little understanding of ideology and policy (e.g., Delli Carpini and Keeter (1997) and Converse (2000)). Consistent with the evidence, political scientists often view voters as boundedly rational individuals who vote "retrospectively" and reward or punish politicians and their parties based on their past performance. The second fact is that parties are often (and sometimes increasingly) polarized, both in the U.S. (McCarty et al., 2006) and around the world (Benoit et al., 2006).

Voter naiveté is often viewed as a detriment to a healthy democracy. If voters do not understand policy and are not well informed, the argument goes, how can elections aggregate information and yield desirable outcomes? Similarly, polarization is also viewed as leading to dysfunctional politics and bad policies.

We do not question that there are many negative aspects of voter naiveté and party polarization. The objective in this paper, however, is to point out two key elements that are missing from the analysis. The first element is that the policies available to voters are endogenously selected by the parties or candidates. So, while it is true that naive voters make worse choices than sophisticated voters for a *fixed* set of policies, less is known about the parties' incentives to offer different policies to naive and sophisticated voters. The second element is that optimal policies are often state contingent, and so the question is not about polarization per se but rather if there is an adequate match between policies and states. For example, a liberal economic policy is often preferred during recessions and a conservative policy during booms. The fact that these two polarized policies are available may make us better off than if we just had a neutral policy at all times.

To study these missing elements, we consider a simple model of two-party competition. In the first stage, the parties, $Left$ and *Right*, commit to certain policies. In the simplest version of the model, the $Left$ party can choose a Left or a Neutral policy and the *Right* party can choose a Right or a Neutral policy. The Left and Right policies represent polarized policies, while the Neutral policy represents a convergent platform at the center of the political spectrum. The restriction on the policy space captures an environment with two well-defined and ideologically opposed parties. Our objective is not to understand the origins of this political environment, but rather to understand if parties that are ideologically constrained have incentives to choose more or less polarized policies.

The best policies depend on the state of the world, with the Left policy being best in left states, the Neutral policy best in center states, and the Right policy in right states. This assumption captures the point we made above that the desirability of a policy often depends on the state of the world. To illustrate, suppose that the state of the economy ranges from recession to boom. The Left policy represents expansionary fiscal policy while the Right policy is contractionary. The Left (i.e., expansionary) policy does best in a recession but hurts in a boom, while the opposite is true for the Right (i.e., contractionary) policy. There is also a Neutral, hands-off policy that neither helps nor hurts the economy.

In the second stage, after parties commit to policies, a large number of voters observe these commitments, receive private information about the state of the world, and subsequently cast their vote. The party with a majority of votes wins the election, and we assume that parties want to maximize their chance of winning the election and have no preference over the polarized vs. neutral policies.

We contrast two different assumptions regarding voting behavior: sophisticated and naive voting. Sophisticated voting is embodied by the notion of Nash equilibrium (NE). In our setting, this type of sophistication requires that voters vote for the best party conditional on their private information and conditional on the (negligible) event that they are pivotal, given the strategies followed by the entire electorate. We view this benchmark case not necessarily as realistic but rather as embodying the notion of perfect sophistication.

Naive voting, on the other hand, may be modeled in different ways. Our objective is to capture the behavior of voters who do not understand policy (i.e., do not understand the mapping between states and outcomes implied by different policies) but who nevertheless want to select the best party or policy and use past, observed performance to predict future performance. In an earlier paper (Esponda and Pouzo, forthcoming), we formalized this type of behavior via an equilibrium notion that we called retrospective voting equilibrium (RVE). In this paper, we apply RVE to model the behavior of naive voters.[1]

[1] An RVE was inspired by Downs' (1957) view of retrospective voting as a way to predict how parties will perform in the future rather than as a way to simply punish or reward the party for past performance; see also Key (1966) and Fiorina (1981). This view is different from the principal-agent perspective started by Barro (1973) and Ferejohn (1986), which studies elections as incentive

To put the model in perspective, note that, in the special case where voters do not observe private information, our model collapses to a Downsian model of two-party competition where the payoff from policies is random. Even an infinitesimal amount of private information is enough to deliver new results. The reason is that the presence of private information implies a specific match between policies and states, an issue that is not present in the standard Downsian model. For example, suppose that one signal is indicative of recession and another of boom, and that the $Left$ party chooses a Left, expansionary policy and the *Right* party chooses a Right, contractionary one. If voters are more likely to vote for the $Left$ party after observing a recessionary signal (which will be true in equilibrium, under both NE and RVE), then the *Left* party will be more likely to be elected in states where expansionary policies are more likely to be beneficial. Thus, signals about the state of the world, even if their precision is almost negligible, will affect the matching of policies with states.

We find that, when voters play NE, the policy choices of the parties converge to the Neutral policy. The logic behind this result is similar to the standard convergence result (Downs, 1957). The idea is that polarization hurts the chances of a party not only in states that are in the opposite extreme of the policy but also in intermediate states. Thus, the parties end up converging to a common, Neutral policy.

When voters play RVE, however, they evaluate parties based on observed, not counterfactual, performance, and the standard logic no longer applies. A party has an incentive to choose relatively extreme policies that work well in those states in which it is elected into office, since those are the states that retrospective voters use to evaluate its performance. In the previous example, under RVE voting, the $Left$ party chooses an expansionary policy and tends to be elected during recessions and the *Right* party chooses a contractionary policy and tends to be elected during booms. The intuition is that deviations to the middle are not profitable because they would lead to lower observed performance. The incentive to polarize that we identify under RVE often leads to a better match between states and policies, and, consequently, to higher welfare, compared to NE, under which implemented policies do not respond to fundamentals.

These results provide both a novel and favorable perspective on voter naiveté and polarization. The main insight, while new in this political context, is analogous to the standard idea of specialization or division of labor. Parties specialize in certain policies

mechanisms to hold politicians accountable.

and they are elected into office when these policies tend to be best. The findings also highlight that the presence of a boundedly rational electorate with little understanding of the counterfactual effect of unobserved policies can be more of a blessing than a curse.

One general implication of these results is that, in order to evaluate the functioning of an electoral system, it is misleading to focus exclusively on whether voters are sophisticated or well informed and to ignore the incentives of the parties. When policies are exogenous, it is not surprising that NE voting is more efficient compared to boundedly-rational RVE voting. But, when the parties' incentives to choose policies are taken into account, a simple retrospective voting heuristic may lead to higher welfare than sophisticated voting.

Our paper is part of a large literature that, motivated by the empirical evidence on polarization, relaxes the assumptions of the Downsian framework to explain non-convergent policies. One of the earliest and most common auxiliary assumptions to explain polarization is that candidates are policy motivated and are uncertain about median voter preferences (Wittman (1977), Calvert (1985)). Other auxiliary assumptions include: the threat of entry by a third party (Palfrey, 1984), the effect of executive-legislative compromise (Alesina and Rosenthal, 2000), lack of policy commitment (Osborne and Slivinski (1996), Besley and Coate (1997)), candidates with "valence" attributes, (Aragones and Palfrey (2002), Gul and Pesendorfer (2009), Kartik and McAfee (2007)), differentiation in the presence of multiple constituencies (Eyster and Kittsteiner, 2007), and convex voter preferences (Kamada and Kojima, 2014).

There are two main features that distinguish our paper from this literature: (1) we relax the assumption of rationality; and (2) we allow the performance of policies to depend on the state of the world. Both relaxing rationality and allowing for state-contingent payoffs are important features of real-life elections that are currently neglected in the previous literature.[2] Of course, the mechanism highlighted in the paper does not preclude any of these other motives for polarization, and we view our work as complementary to this other literature.

Our paper is also related to a growing political economy literature that relaxes the assumption of voter rationality. Spiegler (2013) studies a dynamic model of reforms in

[2] McMurray (2015) also allows for state contingent policies. He studies a pure common value setting and shows that Nash voting leads to convergent policies when parties can commit and are office-motivated, which is inefficient because policies do not match the state.

which an infinite sequence of policy makers care about the public evaluation of their interventions. The public follows a simple attribution rule and (mistakenly) attributes changes in outcomes to the most recent intervention. Levy and Razin (2015) study a setting where voters have two correlated pieces of private information but naive voters fail to account for their correlation. They find that correlation neglect might lead to higher welfare under fixed policies and to less polarized policies when parties choose policies. Lizzeri and Yariv (2015) and Bisin et al. (2015) highlight potentially harmful effects of paternalistic policies. They study models with time inconsistent voters and show how political forces driven by these behavioral voters induce outcomes that are different to those preferred by a benevolent social planner.

In the next section, we use a simple example to illustrate the intuition behind our main results. In Section 3, we present the formal model. We characterize voting behavior (both under NE and RVE assumptions) for fixed policies in Section 4, and we solve for the equilibrium policies chosen by the parties in Section 5. In Section 6, we discuss alternative assumptions and extensions. We conclude in Section 7.

2 Illustrative example

There are three policies, Left, Neutral, and Right, and three equally likely corresponding regions of states, the left, center, and right regions. Voters are all identical and their payoffs are described in Figure 1. In particular, Left is best for states in the left region, Neutral is best for states in the center, and Right is best for states in the right region.

There are two political parties, $Left$ and $Right$, and each party simultaneously commits to a policy that is consistent with their platform but varies in the degree of polarization. The $Left$ party can choose Left or Neutral, and the $Right$ party can choose Right or Neutral. The parties make their choices without knowing the state of the world. If the parties make the same, Neutral choice, then they are elected with equal probability.

After the parties make their choices, a large number of voters receive private signals about the state of the world and must cast their votes in a majority election. For simplicity, we assume in this discussion that the informativeness of these signals is negligible. We contrast two types of behavioral assumptions on voting behavior: Nash equilibrium (NE) and retrospective voting equilibrium (RVE). In particular, we

	left states	center states	right states
Left	1	-1/2	-1
Neutral	0	0	0
Right	-1	-1/2	1

Figure 1: Illustrative example

compare the Nash equilibrium strategies of the parties under each of these two voting assumptions.

Consider first NE voting. Suppose that the $Left$ party chooses Left and the *Right* party chooses Right. Feddersen and Pesendorfer (1997) show that, as the number of voters goes to infinity, NE voting is characterized by full information equivalence. This means that the NE outcome of the election is asymptotically identical to the outcome that arises in a model where all voters observe the realization of the state. Thus, the $Left$ party will win in the left states and the *Right* party will win in the right states. Each party, however, would do better with a unilateral deviation to the Neutral policy. For example, the $Left$ party could switch to the Neutral policy and win in both the left and center states. We will show that this incentive to move to the middle of the political spectrum implies, under natural assumptions, policy convergence under NE.

Consider next RVE voting, a solution concept developed by Esponda and Pouzo (forthcoming) and motivated by the evidence that voters make choices based on observed performance, without trying to assess counterfactual performance or without having a sophisticated understanding of policy.[3] If both parties choose the Neutral policy, then they tie in all states. But consider what happens, say, if the $Left$ party deviates to the Left policy. Even a negligible amount of private information implies that the proportion of people voting for the *Right* party is lowest in the left states and highest in the right states. Therefore, there is a cutoff state such that, for states below the cutoff, the proportion of votes for *Right* is below 50% and the $Left$ party wins and, for states above the cutoff, the proportion is over 50% and the *Right* party

[3]The idea that voters do not try to correct for unobserved counterfactuals is consistent with the empirical findings of Achen and Bartels (2004), Leigh (2009), and Wolfers (2007), who find that voters punish politicians for events that are outside of their control. Healy and Malhotra (2010) find that punishment is related to the politician's response to these events. Our model allows voters to be fairly sophisticated and to condition their learning on private signals, such as campaign platforms, media reports, and economic indicators. In a lab experiment, Esponda and Vespa (2015) show that subjects fail to account for selection. See Esponda and Pouzo (forthcoming) for further discussion of the evidence.

wins.

It is not too difficult to see that the cutoff state must be in the right region if the *Left* party deviates to the Left policy and the *Right* party chooses the Neutral policy; thus, polarization helps the *Left* party win in both left and center states with probability 1. To see this point, it is essential to recall that a retrospective voter is someone who votes for the alternative with highest observed performance. In an equilibrium where both alternatives are elected with positive probability, it must then be the case that the observed performance of both parties is equalized in expectation; otherwise, voters would always prefer one party over the other. In the context of the current example, equalized performances means that it is not possible for the *Left* party (with its Left policy) to be elected in the left states and for the *Right* party (with its Neutral policy) to be elected in the center and right states. If such were the election outcome, then retrospective voters would observe a payoff of 1 every time the *Left* party is elected and a payoff of 0 every time the *Right* party is elected, and they would all want to vote for *Left*. Similarly, it cannot be the case that the *Left* party is elected only in the left and center states. If that were the case, voters would observe an expected payoff of $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1/2) = 1/4$, which is still greater than the payoff from the *Right* party. So it must be that the cutoff state lies in the right region.

Finally, the example also illustrates that incentives to polarize are not necessarily a bad thing in this environment. Voters are in expectation better off when the *Left* party chooses the Left policy and wins in the left states and the *Right* party chooses the Right policy and wins in the right states (the payoffs are 1, -1/2, and 1 in the left, center, and right states of the world, respectively, for an expected payoff of 1/2) compared the case where both parties follow the Neutral policy (and the payoff is zero in all states).

This simple example shows that parties have an incentive to move to the middle under NE voting and an incentive to polarize under RVE voting. Note also that the restrictions on the policy space, which prevented parties from crossing platforms, played no role in the intuition above. The role of these restrictions will be to rule out other RVE equilibria, e.g., equilibria where both parties choose the Left policy. As discussed in Section 6, these other equilibria can also be ruled out by assuming that there is a reputation cost of crossing platforms. The goal of the next sections is to formalize these ideas in a general setting.

3 The setup

There are two parties, $Left$ and $Right$, and three policies, $x \in \{L, N, R\}$, where L stands for the Left policy, R for the Right policy, and N is a Neutral policy. We refer to policies L and R as *polarized policies*. Each party can choose either the Neutral policy or the polarized policy that accords with its ideology (L for the $Left$ party and R for the $Right$ party). Given this restriction, the objective of each party is to maximize its probability of winning the election.

The timing of the game is as follows:

1. The $Left$ and $Right$ parties simultaneously propose a policy $x_{Left} \in \{L, N\}$ and $x_{Right} \in \{N, R\}$, respectively. We order the policies by assuming that $L < N < R$.

2. The state of the world $\omega \in \Omega = [-1, 1]$ is drawn according to a probability distribution with cdf G.

3. Voters observe the proposed policies, but not the state of the world. Conditional on a state $W = \omega$, each voter independently observes a private signal $S = s \in \{s_L, s_R\}$ with probability $q(s \mid \omega)$. Each voter then simultaneously casts her vote for party $Left$ or $Right$.

4. The party with a majority of votes wins the election and implements its proposed policy.

5. Each voter receives payoff $u(x, \omega)$, where $x_j \in \{R, N, L\}$ is the policy of the winning party $j \in \{Left, Right\}$ and $\omega \in \Omega$ is the realized state.

We focus on voting equilibria in large elections, i.e., in the limit as the number of voters goes to infinity. Voters vote according to one of two solution concepts, Nash equilibrium voting (NE voting) or retrospective equilibrium voting (RVE voting), and we contrast the implications of these different solution concepts.

Under NE voting, voters vote for the party that is best conditional on their own private information and any information that can be inferred from the event that they are pivotal; these inferences are required to be correct and depend on the strategies of all other voters. As shown by Feddersen and Pesendorfer (1997), NE voting in an election with a large number of voters is characterized by full information equivalence, meaning that, as the number of voters goes to infinity, the NE voting outcome corresponds to the outcome of an election where voters *know* the state of the world.[4]

RVE, on the other hand, is a solution concept introduced in an earlier paper (Esponda and Pouzo, forthcoming), where it was shown to characterize the steady-state behavior of a large number of voters who update their beliefs about the performance of the parties based on their past *observed* performance.

For the special case in which both parties choose the same, Neutral policy, the monotone assumptions on the primitives used to characterize NE and RVE do not hold. We assume that, if both parties choose the Neutral policy, then the *Right* party wins in states $\omega > 0$ and the $Left$ party wins in states $\omega < 0$.

We make the following assumptions on the primitives.

A1. (Technical) (i) $u(L,\cdot)$, $u(R,\cdot)$, and $u(N,\cdot)$ are all bounded and continuously differentiable in Ω; (ii) G has a density function g, where $\inf_{w\in\Omega} g(\omega) > 0$; (iii) there exists $d > 0$ such that $q(s \mid \omega) > d$ for all $s \in \{s_R, s_L\}$ and $\omega \in \Omega$; (iv) $q(s \mid \cdot)$ is continuous for all $s \in \{s_R, s_L\}$.

A2. (MLRP) For all $\omega' > \omega$:

$$\frac{q(s_R|\omega')}{q(s_R|\omega)} > \frac{q(s_L|\omega')}{q(s_L|\omega)}.$$

A3. (Monotone policies) (i) $u(L,\cdot)$ is decreasing, $u(R,\cdot)$ is increasing, and $u(N,\cdot)$ is constant. We normalize payoffs by letting $u(N,\omega) = 0$ for all $\omega \in \Omega$; (ii) $u(L,0) < 0$ and $u(R,0) < 0$.

A4. (Potential value of polarization)

$$E[u(L,W) \mid W < 0, S = s_R] > 0$$
$$E[u(R,W) \mid W > 0, S = s_L] > 0$$

Assumption A1 collects some technical assumptions, A2 is the standard monotone likelihood ratio assumption, and A3 is a monotonicity assumption on preferences. These assumptions are used to characterize NE and RVE voting behavior in large

[4]This result was proven for the case where voters are assumed to play a *symmetric* Nash equilibrium, a restriction that we also make in this paper.

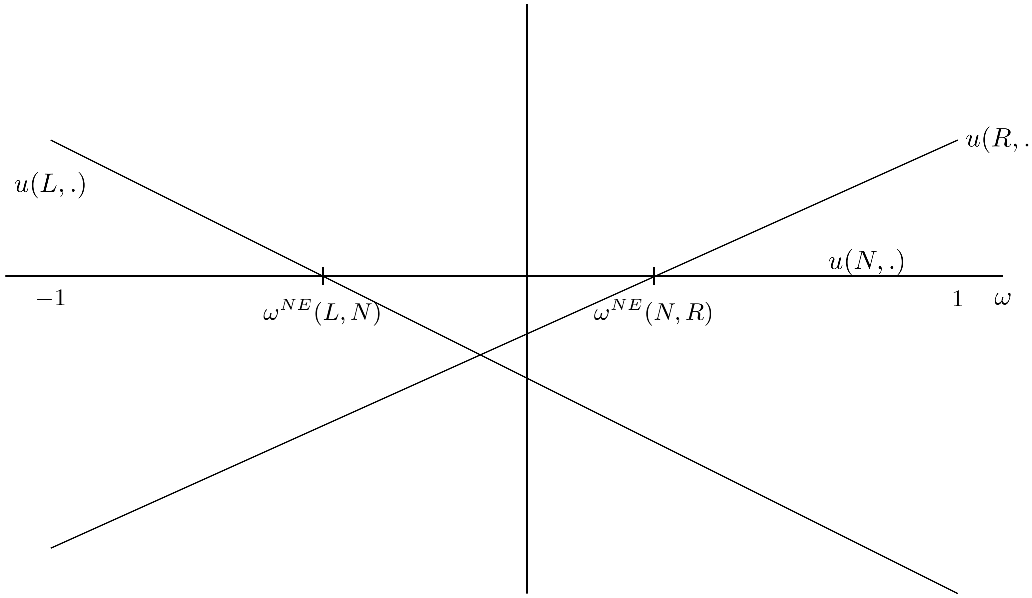

Figure 2: Payoffs from Left, Neutral, and Right policies.
The figure shows that policies divide the state space into three regions: policy Left is best for $\omega < \omega^{NE}(L,N) < 0$, Neutral is best for $\omega^{NE}(L,N) < \omega < \omega^{NE}(N,R)$, and Right is best for $\omega > \omega^{NE}(N,R) > 0$.

elections when the parties choose different policies. Assumption A3 and the assumption that the utility functions are continuous imply that the Neutral policy is best in those states of the world that are close to the center, i.e., close to state $\omega = 0$. Assumption A4 plays two roles. First, it implies that the state can be divided into three regions, as illustrated by Figure 2. In the Left region ($\omega < \omega^{NE}(L,N)$ in the figure), policy L is best; in the Center region ($\omega^{NE}(L,N) < \omega < \omega^{NE}(N,R)$ in the figure), policy N is best; and in the Right region ($\omega > \omega^{NE}(N,R)$ in the Figure), policy R is best. Assumption A4 also says that polarized policies are better than the Neutral policy when evaluated in the states of the world that are favorable to the corresponding party (e.g., these are states $\omega < 0$ for the $Left$ party, due to the assumption that, if both parties were to choose the same, Neutral policy, then the $Left$ party would win in these states) and conditional on the signal that makes those states less likely to happen. This assumption guarantees that polarized policies are sufficiently good provided they are matched with the appropriate states of the world.

We refer to a Nash equilibrium policy profile (x_{Left}, x_{Right}) as a **policy equilib-**

rium.[5] Our objective is to characterize and compare policy equilibria under both the NE and RVE voting assumptions.[6]

The following definition characterizes the information precision of the environment.

Definition 1. An environment has $z \in (0,1)$ information precision if, for all $\omega, \omega' \in \Omega$, $\left|q(s_R \mid \omega) - q(s_R \mid \omega')\right| \leq z$.

We will obtain stronger results for environments where signals are not very informative in the sense that they have z information precision and z is sufficiently small. In the limit, as $z \to 0$, the signals become completely uninformative about the state of the world and $q(s_R \mid \cdot)$ is constant.

We conclude by discussing several examples to illustrate the environment.

Unemployment policies. For example, the primitives may represent an election between two candidates who offer solutions to fight unemployment. The state of the world captures the true underlying cause for unemployment. In high states of the world, unemployment is mostly due to weak demand; in low states of the world, it is due to workers lacking the right skills. Voters observe information correlated with the true cause for unemployment, such as reasons for job loss, types of job listings, or the current premium for skilled labor.

The Left policy corresponds to spending resources in education and training, which is a good policy to lower unemployment in low states of the world. The Right policy corresponds to lowering corporate taxes to incentivize employment, which is a good policy in high states. There is also a Neutral policy that mitigates the costs of unemployment by a magnitude that does not depend on whether unemployment is due to weak demand or poor skills. A typical example is welfare policy intended to bring people out of poverty. In particular, the Neutral policy does not depend on the state of the world, and we normalize its payoff to zero.

Crime policies. Alternatively, the model may capture an election between two district attorneys in a county plagued by drug-related crime. There is a choice between

[5] We restrict parties to choose pure strategies because the interpretation of mixed strategies is not very appealing in this setting. It turns out, however, that equilibrium is in dominant strategies and the main results would go through with mixed strategies. In addition, in Section 6 we show that the main results go through when parties can choose from a continuum of policies.

[6] Our focus on Nash equilibrium of the policy game does *not* implicitly require nor preclude that parties are sophisticated and understand whether voters behave according to NE or RVE. In particular, it is possible to adhere to a learning interpretation of Nash equilibrium (e.g., Fudenberg and Levine, 1998) and to view the parties as naively responding to past experience.

a "Left" intervention that targets the supply for drugs and a "Right" intervention that targets demand. There is also a Neutral policy under which prosecution efforts are not increased for drug-related crimes. Drug-related crime is constant across states, but crime is mostly driven by the demand side in high states and by the supply side in low states.

In the above examples, the payoff of the Left and Right policies is monotone in the state. The next two examples illustrate that the model is also applicable to non-monotone policies provided that voters evaluate parties with respect to a benchmark. The benchmark can be the payoff in a control group or the payoff before the party in control enacts its policy; for empirical evidence of the use of such comparisons, see Healy and Malhotra (2010).

Union election. Consider two candidates competing in a local union election. Workers/voters have the quadratic utility $\Pi(x, \omega) = -(x - \omega)^2$. The $Left$ candidate can make a relatively tough demand $x = L < 0$ and the *Right* candidate can make a relatively soft demand $x = R > 0$. Both candidates can also make a Neutral demand of $x = 0$. The interpretation is that, the higher the state of the world, the higher the firm's bargaining power and, therefore, the softer is the optimal demand by the union. Workers also observe, as a benchmark, the payoffs $\Pi(0, \omega)$ of a non-unionized sector that is equivalent to implementing a Neutral policy $x = 0$. Workers evaluate their union representative against this benchmark,

$$u(x, \omega) = \Pi(x, \omega) - \Pi(0, \omega) = -x^2 + 2x\omega.$$

In particular, higher states make it more desirable to adopt softer demands, so that the benchmark comparison naturally leads to two policies, L and R, with payoffs that are monotone in the state, and to a Neutral policy with payoff that is independent of the state of the world.

Expansionary vs. contractionary policies.[7] The natural rate of unemployment is given by a function $\bar{U}(\omega)$ that is decreasing in the state. The actual unemployment rate is

$$U(x, \omega) = \bar{U}(\omega) + x > 0,$$

where x is the policy. A policy $x = L < 0$ is a fiscal stimulus and decreases unemploy-

[7] This example is based on a model by Persson and Tabellini (2000, p. 426).

ment; a policy $x = R > 0$ is contractionary (e.g., expenditure reduction) and increases unemployment. The Neutral policy $x = 0$ results in the natural unemployment rate. Voters dislike both unemployment and increases in government expenditure, and their utility is given by

$$\Pi(x, \omega) = -\left(U(x, \omega)\right)^2 + x.$$

At the beginning of a period, the Neutral policy of $x = 0$ is in place and voters observe the effects of this benchmark policy. Then, the party in power implements its chosen policy and voters observe the effects of this policy. Voters then assess the extent to which the policy implemented by the party was beneficial; thus

$$u(x, \omega) = \Pi(x, \omega) - \Pi(0, \omega) = -x^2 - 2x\bar{U}(\omega) + x,$$

so that $u(L, \cdot)$ is decreasing and $u(R, \cdot)$ is increasing. In particular, higher states represent better economic fundamentals and make $x = R > 0$ policies more desirable; similarly, $x = L < 0$ policies are more desirable when fundamentals are bad.

4 Voting equilibria with exogenous policies

We begin by characterizing voter behavior of the non-partisan voters for all cases where the parties follow different policies $l < r$, where $l, r \in \{L, N, R\}$; i.e., $(x_{Left}, x_{Right}) \in \{(L, N), (L, R), (N, R)\}$.

4.1 NE voting with exogenous policies

The fact that (symmetric) NE voting is characterized by full information equivalence implies the following result.

Proposition 1. *NE voting outcome for fixed policies (l, r), $l < r$: There is a unique cutoff state $\omega^{NE}(l, r)$ such that the $Left$ party wins the election in all states $\omega < \omega^{NE}(l, r)$ and the Right party wins the election in all states $\omega > \omega^{NE}(l, r)$. Moreover,*

$$\omega^{NE}(L, N) < 0 < \omega^{NE}(N, R)$$

and

$$\omega^{NE}(L,N) < \omega^{NE}(L,R) < \omega^{NE}(N,R).$$

Proof. See the Appendix. □

Figure 2 in page 10 illustrates this result. The intersection of $u(L,\cdot)$ with 0, the payoff of the Neutral policy, is at the cutoff state $\omega^{NE}(L,N) < 0$. Similarly, the intersection of $u(R,\cdot)$ with 0 is at the cutoff state $\omega^{NE}(N,R) > 0$. Finally, the intersection of $u(L,\cdot)$ and $u(R,\cdot)$ must be strictly in between the previous two cutoffs. Note that the party choosing the low policy l is preferred for states below the cutoff and the party choosing the high policy r is preferred for states above the cutoff. So the behavior of perfectly informed voters is completely characterized by these cutoffs. The fact that full information equivalence characterizes Nash equilibrium in large elections yields the desired result.

4.2 RVE voting with exogenous policies

We begin by providing the definition of RVE for fixed policies (l,r), where $l < r$ and $l, r \in \{L, N, R\}$. For simplicity, we sometimes drop (l,r) from the notation. Let $\sigma : \{s_L, s_R\} \to [0,1]$ denote a strategy of a representative voter, where $\sigma(s)$ is the probability of voting for the higher policy r. For a fixed strategy σ and for any state $\omega \in \Omega$, define

$$\kappa(\omega;\sigma) \equiv \sum_{s\in\{s_L,s_R\}} q(s \mid \omega)\sigma(s).$$

to be the proportion of votes in favor of r.

In addition, for any signal $s \in \{s_L, s_R\}$ and state $\omega \in \Omega$, define

$$v(s;\omega) \equiv E\left[u(r,W) \mid W \geq \omega, S=s\right] - E\left[u(l,W) \mid W \leq \omega, S=s\right]. \tag{1}$$

Definition 2. A state ω^{RVE} is a retrospective voting equilibrium (RVE) cutoff for fixed policies (l,r), $l < r$, if there exists a strategy σ satisfying the following conditions:

(i) Cutoff outcome: $\kappa(\tilde{\omega};\sigma) \geq 1/2$ for all $\tilde{\omega} > \omega^{RVE}$ and $\kappa(\tilde{\omega};\sigma) \leq 1/2$ for all $\tilde{\omega} < \omega^{RVE}$.

(ii) Optimality: $v(s, \omega^{RVE}) > 0$ implies $\sigma(s) = 1$ and $v(s, \omega^{RVE}) < 0$ implies $\sigma(s) = 0$.

The first condition in Definition 2 says that the election outcome is characterized by a cutoff state ω^{RVE} with the property that less than 50% of the electorate vote for *Right* if $\omega < \omega^{RVE}$ and more than 50% vote for *Right* if $\omega > \omega^{RVE}$. In particular, the party choosing policy l wins in states $\omega < \omega^{RVE}$ and the party choosing r wins in states $\omega > \omega^{RVE}$. The cutoff state ω^{RVE}, in turn, is determined by the strategy followed by the voters, σ. The second condition says that the strategy followed by the voters, σ, must be optimal given their perceptions, as captured by v in equation (1). Their perceptions are derived from the parties' actual observed equilibrium performance, which in turn depends on the cutoff state ω^{RVE}.

Esponda and Pouzo (forthcoming) provide an explicit learning foundation for Definition 2. In particular, the perceptions captured by v in equation 1 are a *characterization* of the steady state beliefs of voters in an environment with repeated elections. In other words, voters do not explicitly compute the conditional expectations in (1). Instead, voters simply keep track of the past performance of the parties, conditional on their private information, and vote in each period for the party that has exhibited the best performance. To see the intuition, consider a retrospective voter who observes signal s in an election. The way she tries to predict the performance of policy r is by looking at past elections where r was implemented and she also had observed signal s. As she accumulates more past data, her beliefs converge to the first term in the RHS of equation (1). This term is simply the expected payoff of policy r, conditional on signal s and also on the event that policy r is implemented and, therefore, observed. A similar intuition explains the second term in the RHS of equation (1).

Esponda and Pouzo (forthcoming) also provide the following characterization result for an RVE cutoff. For each signal $s \in \{s_L, s_R\}$, define the *personal cutoff*

$$c(s) \equiv \arg\min_{\omega \in \Omega} |v(s; \omega)|. \tag{2}$$

Since Ω is compact and $v(s; \cdot)$ is continuous and increasing (by A2-A3), the personal cutoffs are unique and satisfy $c(s_R) \le c(s_L)$. For a given RVE cutoff ω^{RVE}, optimal behavior is characterized by the relationship between ω^{RVE} and the personal cutoffs: If $c(s) < \omega^{RVE}$, this means that $v\left(s; \omega^{RVE}\right) > 0$ and, therefore, it is optimal to vote for the *Right* party (i.e., policy r) after observing signal s. Similarly, if $c(s) > \omega^{RVE}$

it is optimal to vote for the $Left$ party (i.e., policy l). In particular, for any election cutoff $\omega \in \Omega$,

$$\overline{\kappa}(\omega) \equiv \sum_{\{s \in \{s_L, s_R\} : c(s) < \omega\}} q(s \mid \omega) \tag{3}$$

may be interpreted as the proportion of players that vote for the *Right* party in state ω when the RVE cutoff is also given by ω.[8]

Esponda and Pouzo (forthcoming) show that there is a unique RVE cutoff and that it is essentially given by the state where the proportion of votes for the *Right* party, as captured by the function $\overline{\kappa}$, equals 1/2.

Theorem 1. *(Esponda and Pouzo forthcoming) For fixed policies (l, h), $l < h$, there is a unique RVE voting cutoff, given by $\omega^{RVE} = \inf\{\omega \in \Omega : \overline{\kappa}(\omega) \geq 1/2\}$.*

By Theorem 1, an RVE cutoff can be found in four simple steps. Figures 3 and 4 depict these steps for fixed policies (L, R). First, we compute the belief functions $v(s; \cdot)$, $s \in \{s_L, s_R\}$. Second, we use this function to find the personal cutoffs $c(s)$, which solve $v\left(s; c(s)\right) = 0$, as illustrated by Figure 3. Third, we compute the vote share for R, $\bar{\kappa}$. This function satisfies

$$\overline{\kappa}(\omega) = \begin{cases} 0 & \text{if}\, \omega \leq c(s^R) \\ q(s_R \mid \omega) & \text{if}\, c(s^R) < \omega \leq c(s^L) \\ 1 & \text{if}\, \omega > c(s^L) \end{cases} \tag{4}$$

Finally, we intersect $\bar{\kappa}(\cdot)$ with 1/2 to find the RVE cutoff ω^{RVE}, as illustrated by Figure 4.

To see intuitively why ω^{RVE} is the equilibrium cutoff with retrospective voters, consider any other potential cutoff, say $\omega' < \omega^{RVE}$ depicted in Figure 4. If the outcome were characterized by ω', this means that L would be chosen for $\omega < \omega'$ and R would be chosen for $\omega > \omega'$. A voter's perception of the difference in expected performance of R vs. L, satisfies $v(s_R; \omega') > 0$ when the signal is s_R and $v(s_L; \omega') < 0$ when the signal is s_L. In particular, a voter would find it optimal to vote for R after observing s_R and to vote for L after observing s_L. If voters vote in this manner, the proportion of votes for the *Right* party would be $q(s_R \mid \omega') < 1/2$ at state ω'. By

[8]The interpretation is exact except when ω is one of the personal cutoffs.

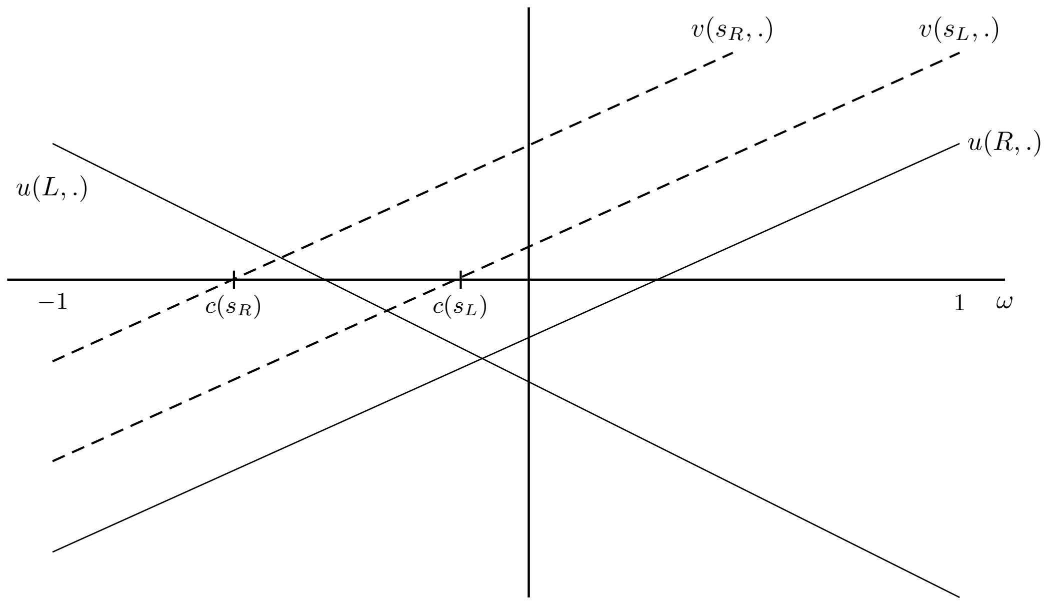

Figure 3: Payoffs, beliefs, and personal cutoffs for fixed policies (L, R).
The figure shows the payoffs $u(L, \cdot)$ and $u(R, \cdot)$, the belief functions $v(s_L; \cdot)$ and $v(s_R; \cdot)$, and the corresponding personal cutoffs $c(s_L)$ and $c(s_R)$ for fixed policies (L, R).

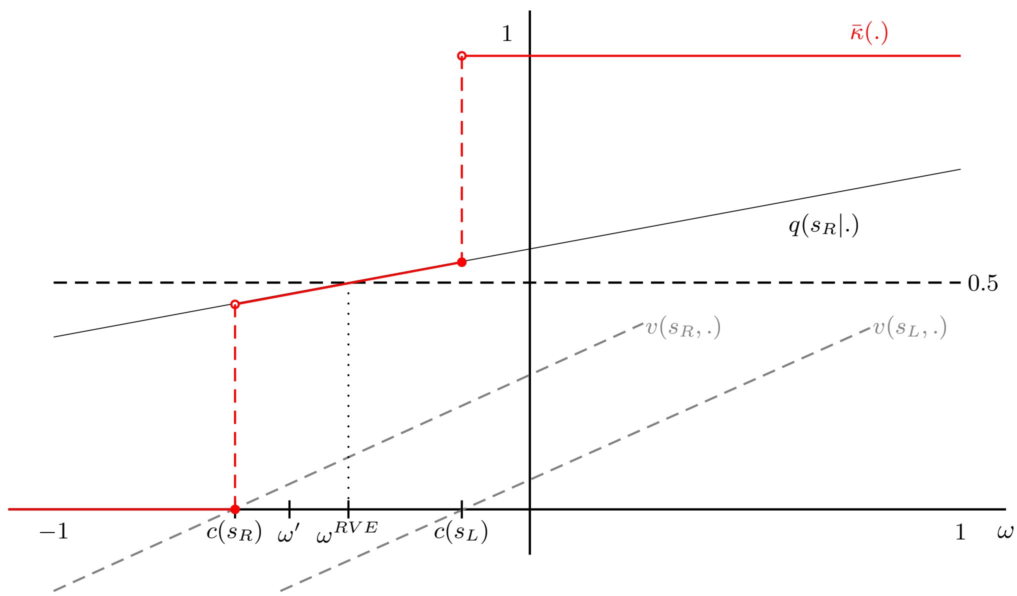

Figure 4: Retrospective voting equilibrium (RVE) for fixed policies (L, R).
The figure shows how to use the personal cutoffs to construct the vote share function $\bar{\kappa}$, and how to then find the equilibrium cutoff ω^{RVE} by intersecting the vote share function with $1/2$, which corresponds to majority rule.

continuity of $q(s_R \mid \cdot)$, the proportion of votes for *Right* would also be below 1/2 for states slightly above ω'. But this means that *Left* will be chosen with probability 1 for states slightly above ω', therefore contradicting that ω' can be an RVE cutoff to begin with.

EXAMPLE 1. The state is uniformly distributed in $[-1, 1]$ and the utility functions are $u(L, \omega) = -\omega - 1/3$ and $u(R, \omega) = \omega - 1/4$. The *Left* party chooses policy L and the *Right* party chooses R. In particular, $c^{FB} = -1/24$ is the first-best election cutoff, i.e., R is preferred in states $\omega > c^{FB}$ and L in states $\omega < c^{FB}$. In addition, each voter privately observes signal s_R with probability $q(s_R \mid \omega) = .6 + z\omega/2$, where $z \in (0, .8)$ is a parameter representing information precision.

Simple algebra yields

$$\begin{aligned} v\big(s_R; \omega\big) &= E\big[W \mid W \geq \omega, S = s_R\big] - E\big[-W \mid W < \omega, S = s_R\big] + 1/12 \\ &= \frac{.3(1-\omega^2) + \frac{z}{6}(1-\omega^3)}{.6(1-\omega) + \frac{z}{4}(1-\omega^2)} + \frac{.3(\omega^2-1) + \frac{z}{6}(\omega^3+1)}{.6(\omega+1) + \frac{z}{4}(\omega^2-1)} + 1/12, \end{aligned}$$

and, similarly,

$$v\left(s_L; \omega\right) = \frac{.2(1-\omega^2) - \frac{z}{6}(1-\omega^3)}{.4(1-\omega) - \frac{z}{4}(1-\omega^2)} + \frac{.2(\omega^2-1) - \frac{z}{6}(\omega^3+1)}{.4(\omega+1) - \frac{z}{4}(\omega^2-1)} + 1/12.$$

Solving for $v\left(s; \cdot\right) = 0$ for $s \in \{s_L, s_R\}$, we obtain the personal cutoffs $c(s_R; z)$ and $c(s_L; z)$, which we index by the the information parameter z. Finally, since $\bar{\kappa}(\cdot)$ is characterized by (4), we know that the equilibrium cutoff will either be $c(s_R; z)$, $c(s_L; z)$, or the solution to $q(s_R \mid \omega) = 1/2$, which in this case is given $\max\{-.2/z, -1\}$. In this example, $-.2/z < c(s_R; z)$ for all $z \in (0, .8)$, and so the intersection of $\bar{\kappa}(\cdot)$ and 1/2 occurs at $c(s_R; z)$. Therefore, the RVE cutoff is $\omega^{RVE}(L, R; z) = c(s_R; z)$ for all $z \in (0, .8)$. Moreover, it is easy to check that $c(s_R; z)$ is decreasing in z and that $\lim_{z\to 0} c(s_R; z) < c^{FB}$. In particular, the *Right* party is elected too often compared to the first-best outcome (or, equivalently, the NE outcome), and this bias becomes more pronounced as the precision of information increases.[9] $\square$

We now apply Theorem 1 to obtain a characterization of RVE voting outcomes for different policy profiles.

[9] Esponda and Pouzo (forthcoming) had already pointed out that welfare may decrease with better information under RVE; see that paper for more on comparative statics.

Proposition 2. *RVE voting outcome for fixed policies* (l, r), $l < r$*: There is a unique cutoff state* $\omega^{RVE}(l, r)$ *such that the* $Left$ *party wins the election in all states* $\omega < \omega^{RVE}(l, r)$ *and the Right party wins the election in all states* $\omega > \omega^{RVE}(l, r)$*. Moreover,*

$$\omega^{RVE}(N, R) < 0 < \omega^{RVE}(L, N)$$

and

$$\omega^{RVE}(N, R) \leq \omega^{RVE}(L, R) \leq \omega^{RVE}(L, N).$$

Finally, there exists $\bar{z} > 0$ *such that, in any environment with information precision* $z \leq \bar{z}$ *and satisfying* $E\big[u(R, W) \mid S = s_R\big] < u(L, -1)$ *and* $E\big[u(L, W) \mid S = s_L\big] < u(R, 1)$*, the previous two inequalities are strict.*

Proof. See the Appendix. □

We now discuss the intuition behind Proposition 2. Consider first the statement that $\omega^{RVE}(N, R) < 0$. Consider what would happen if, contrary to the statement, the cutoff were given by $\omega^* \geq 0$. We argue that, if that were the case, then voters would prefer to vote for the *Right* party for every signal, and so the *Right* party would be elected with probability 1 in all states, i.e., $\omega^* = -1$, which is a contradiction. The reason voters would want to vote for *Right* for every signal is that, even for signal s_L, which is the least favorable for *Right*,

$$E\big[u(R, W) \mid W \geq \omega^*, S = s_L\big] \geq E\big[u(R, W) \mid W \geq 0, S = s_L\big] > 0,$$

where the first inequality follows from the facts that $u(R, \cdot)$ is increasing and $\omega^* \geq 0$ and the second by A4. Another way of stating this result is that it must be the case that $c_{(N,R)}(s_L) < 0$ because, by definition, the personal cutoff $c_{(N,R)}(s_L)$ is the state at which the expected utility of R conditional on the state being higher and signal s_L is exactly zero. Since $c_{(N,R)}(s_R) < c_{(N,R)}(s_L) < 0$, it follows by Theorem 1 that the RVE cutoff, which is always between $c_{(N,R)}(s_R)$ and $c_{(N,R)}(s_L)$, must be strictly negative. A similar arguments establishes that $\omega^{RVE}(L, N) > 0$.

Next, fix a voter with signal s and suppose that we start from policy positions (N, R), where we argued above that the personal cutoff satisfies $c_{(N,R)}(s) < 0$, and consider a change to policies (L, R). The effect of this change is illustrated by Figure 5. The voters' evaluation of the *Right* party if the election cutoff were given by $c_{(N,R)}(s)$

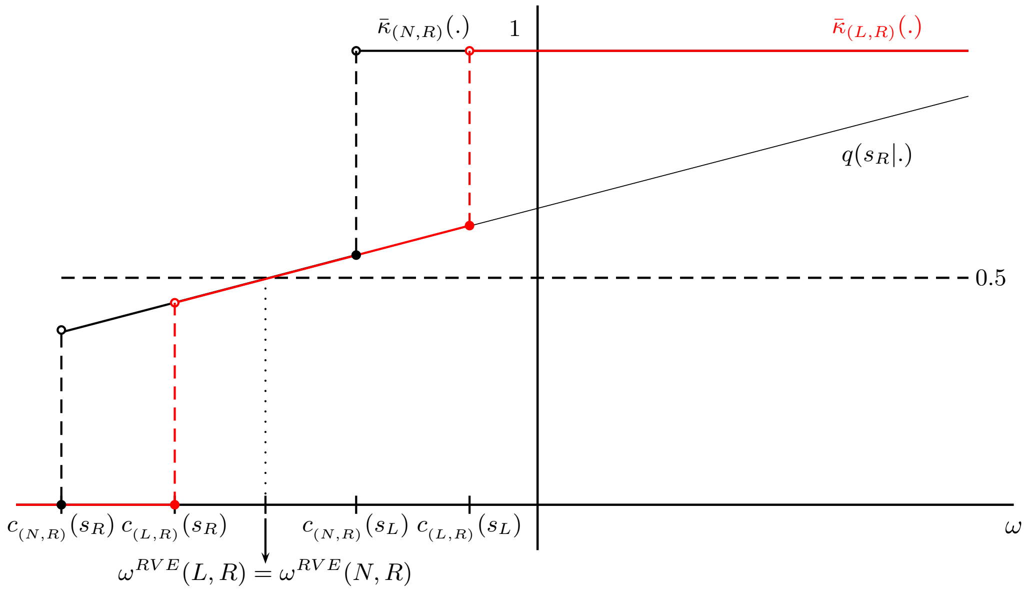

Figure 5: Comparing RVE cutoffs for (L, R) vs. (N, R).
Going from policies (N, R) to (L, R) shifts the function $\bar{\kappa}(\cdot)$ to the right and, consequently, the RVE cutoff, thus implying that $\omega^{RVE}(N, R) \leq \omega^{RVE}(L, R)$. The figure illustrates that it is possible that $\omega^{RVE}(N, R) = \omega^{RVE}(L, R)$.

would continue to be zero, because the *Right* party continues to choose policy R. But now the evaluation of the *Left* party would go from 0 to a strictly positive value, because A4 says that the expected payoff of Left conditional on being chosen in states $\omega < 0$ is strictly positive; since $u(L, \cdot)$ is decreasing and $c_{(N,R)}(s) < 0$, the expected payoff must also be strictly positive conditional on Left being chosen in states below $c_{(N,R)}(s)$. Therefore, fixing the cutoff $c_{(N,R)}(s)$, the *Left* party is perceived to be better by switching from policy N to L. Consequently, the personal cutoffs under (L, R) will be to the right of the original cutoffs. A move to the right increases the observed performance of the *Right* party and decreases the observed performance of the *Left* party, bringing them once again to be perceived to be equal at $c_{(L,R)}(s) > c_{(N,R)}(s)$. The above argument is true for all s. Therefore, the change in personal cutoffs shifts the $\bar{\kappa}(\cdot)$ function to the right and, by Theorem 1, it also shifts the equilibrium cutoff to the right, and we obtain $\omega^{RVE}(N, R) \leq \omega^{RVE}(L, R)$. As illustrated by Figure 5, however, the inequality may not be strict. A similar argument establishes that $\omega^{RVE}(L, N) \geq \omega^{RVE}(L, R)$.

Finally, we argue that the previous two inequalities comparing the RVE cutoffs for (N, R), (L, N), and (L, R) are *both* strict under additional assumptions—that at least one of them is strict follows trivially from the first claim in the proposition. The first assumption is that the signals are sufficiently uninformative. To see the intuition, consider the limiting case where there are no signals (outside the limit, the argument goes through due to continuity). Then, because the personal cutoffs are strictly ranked (e.g., $c_{(L,R)} > c_{(N,R)}$), it follows that the equilibrium cutoff, which coincides with the personal cutoff in this case, will also be strictly ranked, except possibly if one of these cutoffs is already -1 or 1. The second set of assumptions, $E\big[u(R, W) \mid S = s_R\big] < u(L, -1)$ and $E\big[u(L, W) \mid S = s_L\big] < u(R, 1)$, require that neither L or R are sufficiently superior to the other policy, thus ruling out extreme cutoffs -1 and 1.

5 Endogenous policies and polarization

The following proposition characterizes equilibrium under both NE and RVE voting when the parties can choose policies.

Proposition 3. *Policy equilibrium:*

(i) Under NE voting, $x_{Left} = x_{Right} = N$ is the unique policy equilibrium.

(ii) Under RVE voting, $x_{Left} = x_{Right} = N$ is not an equilibrium and $x_{Left} = L$, $x_{Right} = R$ is the unique policy equilibrium in weakly undominated strategies. Moreover, there exists $\bar{z} > 0$ such that, in any environment with information precision $z \leq \bar{z}$ and satisfying $E\big[u(R, W) \mid S = s_R\big] < u(L, -1)$ and $E\big[u(L, W) \mid S = s_L\big] < u(R, 1)$, $x_{Left} = L$, $x_{Right} = R$ is the unique policy equilibrium.

Proof. See the Appendix. □

Proposition 3 is the main result of the paper. It says that, under the assumption of NE voting, the two parties choose convergent policies but, under the assumption of RVE voting, the two parties choose polarized policies. This result follows immediately from the characterization for fixed policies derived in Propositions 1 and 2 in Section 4, but it is still convenient to reiterate the main intuition.

Consider first the case of NE voting, where the result says that both parties choose the Neutral policy in equilibrium. This result extends the standard Downsian logic of the median voter theorem to a setting where there is uncertainty about the best alternative. The intuition is as follows. Under NE voting, the election outcome coincides with the full information outcome (Feddersen and Pesendorfer, 1997). Thus, it is as if voters were able to compare, for each state, not only the payoff of the elected party but also the correct payoff that would result from electing the other party. If both parties are choosing the Neutral policy, then, by assumption, the *Right* party wins in states $\omega > 0$ and the $Left$ party wins in states $\omega < 0$. Suppose that, say, the *Right* party deviates to the polarized policy R. Then the election cutoff changes from 0 to $\omega^{NE}(N, R) > 0$, and so the *Right* party now loses in all states between 0 and $\omega^{NE}(N, R)$; see Figure 2 in page 10. In other words, by polarizing, the *Right* party continues to win in states closest to 1 and to lose in states closest to -1, but it now also loses in center states where the Neutral policy is superior to the Right policy. Thus, parties do not have incentives to polarize under NE and end converging to a common, middle platform.

The previous logic does not apply under RVE voting. Suppose that both parties were choosing the Neutral policy, corresponding to an equilibrium cutoff of 0. A deviation by, say, the *Right* party to the polarized policy R moves the equilibrium cutoff to the left, therefore increasing the set of states where *Right* wins the election. Intuitively, by deviating to R, the *Right* party manages to increase its observed performance. The reason is that the policy R is much better suited for the high states of the world in which the *Right* party is elected. The better observed performance of the deviating party prompts voters to increase their support towards it. Thus, the equilibrium cutoff moves to the left of the original cutoff.

The key feature driving these different results is that, under NE, voters are sophisticated and understand counterfactual payoffs. Thus, voters favor Neutral policies because they realize that polarized policies, while being attractive in some extreme states of the world, would produce bad outcomes both in intermediate states and states on the opposite extreme. Under RVE, in contrast, voters judge parties exclusively on observed performance, and polarization is attractive because it increases observed performance.

We conclude by comparing welfare under both NE and RVE voting. It is not surprising that, given fixed policies, NE voting yields higher welfare than RVE voting. But, once parties choose their policies endogenously, this prediction can be reversed. In particular, we show that, in environments with sufficiently low information precision, voter welfare is higher under RVE compared to NE voting.

Proposition 4. *There exists $\hat{z} > 0$ such that, in any environment with information precision $z \leq \hat{z}$, any policy equilibrium under RVE yields strictly higher expected welfare than the unique policy equilibrium under NE.*

Proof. See the Appendix. □

Broadly speaking, the reason why welfare is higher under RVE voting is that parties choose polarized policies, resulting in a better match between policies and the state of the world. In particular, when the $Left$ party chooses policy L and the *Right* party chooses R, the $Left$ party is elected in the "left" states of the world, where policy L tends to be best, and the *Right* party is elected in the "right" states of the world, where policy R tends to be best. In contrast, while Nash voting is efficient in aggregating information for fixed policies, it does poorly if policies are endogenous when compared to retrospective voting.

The intuition behind Proposition 4 is best seen by focusing attention on the RVE equilibrium in weakly undominated strategies, (L, R), and by considering the limiting case where voters have no private information (outside the limit, the argument goes through due to continuity). Under RVE voting and essentially no private information, there is essentially a unique personal cutoff $c(s_R) \approx c(s_L)$ and this cutoff is, by Theorem 1, the equilibrium cutoff, ω^{RVE}. In particular, it must be the case that, if both parties are elected with positive probability, the observed performance of both parties is equalized.[10] Otherwise, voters would always want to vote for one of the parties, and this party would get elected with probability 1 in all states. In particular, if the policies are (L, R), then it must be the case that

$$E\big[u(R, W) \mid W \geq \omega^{RVE}\big] = E\big[u(L, W) \mid W \leq \omega^{RVE}\big]. \tag{5}$$

If $\omega^{RVE} \geq 0$, then the first of these terms must be strictly positive, because A2, A3(i) and A4 imply that $E\big[u(R, W) \mid W \geq 0\big] > 0$ and A3(i) says that $u(R, \cdot)$ is increasing. Thus, by the equality in (5), the second term must also be strictly positive. A similar argument shows that both terms are also strictly positive if $\omega^{RVE} < 0$.

[10]The result also holds if one of the parties is elected with probability 1 in all states, but we ignore this case when providing the intuition. See the proof for the complete argument.

Expected voter welfare in an RVE equilibrium with policies (L, R) is given by

$$G(\omega^{RVE})E\big[u(L,W) \mid W \le \omega^{RVE}\big] + (1 - G(\omega^{RVE}))E\big[u(R,W) \mid W \ge \omega^{RVE}\big],$$

and so it follows that expected voter welfare in this RVE policy equilibrium is strictly positive because it is the weighted average of two strictly positive terms. In contrast, the unique policy equilibrium under NE is (N, N), and voter welfare is equal to 0 in all states under the NE policies.

EXAMPLE 1, continued from page 18. For a hypothetical equilibrium cutoff $\omega \in (-1, 1)$, voter welfare in the example is given by

$$\begin{aligned} Welfare(\omega) &\equiv \Pr(W \ge w)\big(E\big[W \mid W \ge w\big] - \frac{1}{4}\big) + \Pr(W < w)\big(E\big[W \mid W < w\big] - \frac{1}{3}\big) \\ &= -\omega^2/2 - \omega/24 + 5/24. \end{aligned}$$

The function $Welfare(\cdot)$ is strictly concave, maximized at the first-best cutoff $\omega = c^{FB} = -1/24$, and strictly higher than zero in the interval $[-.68, .60]$. Previously, we established that the RVE cutoffs for $z \in (0, .8)$ are strictly lower than $c^{FB} < 0$ and higher than $\lim_{z \to .8} c(s_R; z)$. Since $\lim_{z \to .8} c(s_R; z) = -.21 > -.68$, it follows that welfare is strictly greater than 0 under RVE for all z.

In example 1, welfare under RVE is strictly greater than the welfare under NE (which is zero) for *all* values of the information precision parameter, ι. In the following example, we illustrate the importance of the condition in Proposition 4 by showing that the opposite result may be true if signals are sufficiently informative.

EXAMPLE 2. The state is uniformly distributed in $[-1, 1]$ and the utility functions are $u(L, \omega) = -0.0001\omega - b_L$ and

$$u(R, \omega) = \begin{cases} -\frac{b_R}{(1+\omega)^{1.1}} & \text{if } \omega \le 0 \\ 3\omega - b_R & \text{if } \omega > 0 \end{cases}.$$

where $b_R = 0.7941$ and $b_L = 0.0004$.[11] In addition, each voter privately observes signal

[11] Our choice of $u(R, .)$ formally violates A1(i) because it is not differentiable and is not bounded at $\omega = -1$. By continuity arguments, the results in this example continue to hold if we approximate $u(R, .)$ by a truncated and smooth function. To keep the exposition as simple as possible, we do not

s_R with probability

$$q(s_R \mid \omega) = \begin{cases} 0.52\left(1 - z(-\omega)^8\right) & \text{if } \omega \le 0 \\ 0.52 + z0.48\omega^{1/8} & \text{if } \omega > 0 \end{cases}.$$

where $z \in (0, 1]$ is a parameter representing information precision. Simple but tedious algebra yields, for any $\omega \le 0$,

$$\begin{aligned} v\left(s_R;\omega\right) =& \frac{-b_R \times 0.52 \times \int_\omega^0 \frac{1-z(-x)^8}{(1+x)^{1.1}}dx + 0.7059 \times 0.52 + 0.3388 \times z}{0.52(1-\omega) + z\frac{0.52}{9}\omega^9 + z\frac{0.48}{1.125}} \\ &+ 0.0009\frac{(\omega^2-1) - z\left(\omega^{10}-1\right)}{9(\omega+1) - z(\omega^9+1)} + b_L. \end{aligned}$$

Also, voter welfare is given by

$$Welfare(\omega) = 0.5\left(10b_R(1-(1+\omega)^{-0.1}) + 0.7056 - b_L\omega - \frac{0.0001}{2}\omega^2\right).$$

For all $\omega \in [-1, 0]$, $Welfare(\cdot)$ is strictly concave and increasing. Numerical computations indicate that welfare is positive for values of z less than $\hat{z} \approx 0.99$ and negative for higher values. For example, consider the case where $z = 0.999$. Solving for $v\left(s_R;\cdot\right) = 0$, we obtain the personal cutoff $c(s_R; z) \approx -0.6434$. The solution to $q(s_R \mid \omega^o) = 1/2$, is given by $\omega^o \approx -0.6654$. Since, $c(s_R; z) > \omega^o$, the RVE cutoff is given by $c(s_R; z)$. At this value, welfare is $Welfare(c(s_R; z)) \approx -0.0848$. □

6 Discussion and extensions

The model above was deliberately kept simple to investigate the importance of allowing policies to be both endogenous and state dependent when assessing the consequences of having a naive electorate. The objective of this section is to argue that the main insights continue to be relevant in richer environments.

HETEROGENEOUS VOTER PREFENCES. For simplicity, we assumed that all voters have the same preferences. Alternatively, we can assume an ordering over preferences such that the equilibrium cutoff under NE and RVE is determined exclusively by the

follow this approach. Also, the parameters (b_R, b_L) were chosen to ensure that A4 holds given the signal structure in this example.

preferences of the median voter. We can then interpret the analysis in this paper as being done with respect to the median voter and derive the same polarization and welfare results, where welfare corresponds to the welfare of the median voter.

TIMING: COMMITMENT AND OBSERVABILITY OF THE STATE. An important assumption is that parties must commit to a non-state contingent policy during the campaign stage. This assumption is both common and natural in models of political economy. Consider first the assumption of non-state contingent policies. For candidates to be able to make contingent promises, voters would need to be able to understand and agree on the definition of a state and also observe the state ex post to corroborate if parties indeed followed their promises. Not surprisingly, precise *contingent* promises are not common in the real world.

Consider next the assumption that candidates can commit in the campaign stage. Without commitment, a key issue is whether or not elected politicians can observe the state of the world before choosing policy. With perfect information about the state of the world and no commitment power, politicians would of course be able and willing to choose the best policy after being elected. But the polarization result is still robust to intermediate cases where there is partial commitment or observability. For example, suppose that there is partial commitment in the sense that politicians incur a cost of reneging on their campaign promises. Once elected, if they happen to observe additional information about the state of the world, they might learn that there is a better policy. But they would keep the original promise provided that the cost of reneging is higher than the benefit from choosing a better policy. Note also that the benefit from choosing a better policy is rather limited under RVE, because it is already the case that parties are being chosen in states of the world in which their policies tend to be best.

To summarize, polarization will be observed in settings where elected politicians remain uncertain about the state of the world when choosing policy or in settings where politicians have little uncertainty but where they are required to propose policies during the campaign and face significant costs of breaking these promises. There are several policy issues for which little additional information will be revealed between the election and the actual exercise of power. For example, in the recent financial crisis, there was (and still is) a lot of uncertainty about the optimal policy. The evidence also suggests that candidate repositioning is costly and, thus, the assumption of commitment captures a realistic aspect of the world.[12]

CONSTRAINED PLATFORMS. In the model, we did not allow parties to cross platforms. For example, the *Right* party was not allowed to choose the Left policy. We view this assumption as capturing a realistic aspect of electoral systems with two ideologically opposed parties. In practice, candidates first choose to belong to a party platform (such as the Republican or Democratic parties in the U.S.) and then run for office. The party labels then constrain candidates to choose policies that are coherent with the platform they have chosen, and the empirical evidence appears consistent with this view.[13]

While candidates are often constrained by their parties' platforms, they can still decide to choose policies that are closer to the center or policies that are more polarized and may appeal more to the base. It is important, however, to take this constraint into account when interpreting our results. Our goal is limited to explaining whether candidates faced with these constraints will decide to exacerbate their positions or not. The objective of this paper is not to explain the emergence of a two party system to begin with, and we view this emergence as being orthogonal to whether or not voters are sophisticated, as embodied by NE, or naive, as embodied by RVE.

It is nevertheless interesting to consider what would happen if parties were able to cross platforms. To do this, we need to make an assumption about the outcome of the election if the parties chose the same policies, (R, R) or (L, L). Suppose, to be consistent with our earlier assumption, that the $Left$ party would get elected in states $\omega < 0$ and the *Right* party in states $\omega > 0$.[14] Then it is still true that (N, N) is a policy equilibrium under NE, since a deviation to L or R by *either* party would make that party worse off. In fact, by arguments similar to those in Section 5, it follows that (N, N) continues to be an equilibrium under NE even if parties can cross platforms.

[12]Tomz and Van Houweling (forthcoming) provide evidence that it is costly for candidates to change course and reposition. They also look at several electoral debates to study how politicians have used rhetoric to criticize, deny, or excuse changes in position. They find that candidates typically cannot use rhetoric to erase the costs of repositioning and that voters dislike repositioning and punish candidates for breaking pledges.

[13]Ansolabehere et al. (2001) show that local politicians in the U.S. are constrained by the national parties' ideologies and Gerber and Morton (1998) show that the fact that many candidates have to go through primaries and appeal to more ideological voters also constrains the types of positions they can take. In addition, a literature in political science argues that political parties play the role of *brand labels*. These brand names are valuable to voters and candidates, because they help voters make decisions and they help candidates win elections (e.g., Snyder and Ting, 2002).

[14]Alternatively, we could consider a symmetric model where ω follows a uniform distribution and assume that each party is equally likely to win in each state.

For the case of RVE voting, it is still the case that (N, N) is not an equilibrium, since either party can do better by deviating to a polarized policy, following the arguments in the paper. But it is possible that now (L, L) or (R, R) could constitute an equilibrium, depending on the assumption that we make regarding what happens when both parties chooses the same polarized policies. Thus, it is possible to have both parties choosing the *same* polarized policy. Of course, in a richer model, one may expect entry by a third party advocating a policy on the opposite side of the ideological spectrum. Or, alternatively, one could imagine a world where there are two parties targeting opposite extremes of the ideological spectrum and a non-negligible fraction of *partisan* voters. These partisan voters automatically vote for their ideologically preferred party provided that such party does not cross platforms, while the non-partisan voters behave as assumed in the previous sections. For example, "right partisans" vote for the *Right* party if the *Right* party chooses policies N or R, but abstain if the *Right* party chooses policy L. We can show that if the fraction of partisan voters is high enough, then (L, L) and (R, R) are not equilibria and that the only equilibrium under RVE is (L, R).

CONTINUUM OF POLICIES AND TIE-BREAKING RULE. For tractability, we assumed that there were only three policies, L, N, and R, and that, if both parties choose the same policy N, then the equilibrium cutoff is 0. It is possible to extend our results to a more general model where there is a continuum of policies available to each party and where the tie-breaking rule is not arbitrary but rather arises endogenously from the choice of policy space. For example, suppose that policies are represented by the $[-1, 1]$ interval and that the $Left$ party can choose any policy $L \leq -\varepsilon$ and the $Right$ party can choose any policy $R \geq -\varepsilon$, where $\varepsilon > 0$ captures the distance to the Neutral policy. The utility of voters is $u(x, \omega)$, where $x \in [-1, 1]$ is the policy and ω is the state of the world. Suppose that $u(L, \cdot)$ is decreasing for all $L < 0$, that $u(R, \cdot)$ is decreasing for all $R > 0$, and that $u(0, \omega) = 0$ for all ω. In particular, policies $L < 0$ are Left policies, policies $R > 0$ are Right policies, and the policy of 0 is what we called the Neutral policy. By restricting parties to choose different policies, we do not have to specify a tie-breaking rule. But we can naturally define a policy equilibrium to be the limit of policy equilibria as ε goes to zero. We show in Appendix B that, together with some assumptions that are the analog of assumptions A3-A4 in the paper, this model continues to deliver policy convergence under NE and policy divergence under RVE.

7 Conclusion

We provide a new perspective on two issues that are often viewed negatively: voter naiveté and party polarization. We study a model where two parties compete by committing to policies and voters subsequently vote for their preferred party. We contrast sophisticated with naive voting. The former is embodied by Nash equilibrium while the latter is formalized via the notion of a retrospective voting equilibrium. Retrospective voters do not understand the mapping between states and outcomes induced by a policy; instead, they simply vote for the party that has delivered the best performance in the past. We find that parties have an incentive to polarize under retrospective, compared to Nash, voting. Moreover, this polarization often results in higher welfare due to a better match between policies and fundamentals. These results provide a new and favorable perspective on voter naiveté and party polarization. The results also imply that it is important to take the responses of the parties into account before concluding if voter naiveté can be detrimental to democracy.

References

Achen, C.H. and L.M. Bartels, "Blind retrospection: Electoral responses to drought, flu, and shark attacks," *working paper*, 2004.

Alesina, A. and H. Rosenthal, "Polarized platforms and moderate policies with checks and balances," *Journal of Public Economics*, 2000, *75* (1), 1–20.

Ansolabehere, Stephen, James M Snyder, and Charles Stewart III, "Candidate positioning in US House elections," *American Journal of Political Science*, 2001, pp. 136–159.

Aragones, E. and T. Palfrey, "An Economic Model of Representative Democracy," *Journal of Economic Theory*, 2002, *103*, 131–161.

Barro, R., "The Control of Politicians: An Economic Model," *Public Choice*, 1973, *14*, 19–42.

Benoit, Kenneth, Michael Laver et al., *Party policy in modern democracies*, Routledge, 2006.

Besley, T. and S. Coate, "An Economic Model of Representative Democracy*," *Quarterly Journal of Economics*, 1997, *112* (1), 85–114.

Bisin, Alberto, Alessandro Lizzeri, and Leeat Yariv, "Government policy with time inconsistent voters," *The American Economic Review*, 2015, *105* (6), 1711–1737.

Calvert, R. L., "Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence," *American Journal of Political Science*, 1985, pp. 69–95.

Converse, P.E., "Assessing the capacity of mass electorates," *Annual Review of Political Science*, 2000, *3* (1), 331–353.

Delli Carpini, M.X. and S. Keeter, *What Americans know about politics and why it matters*, Yale Univ Press, 1997.

Downs, A., *An economic theory of democracy*, New York: Harper and Row, 1957.

Esponda, I. and D. Pouzo, "Conditional Retrospective Voting in Large Elections," *forthcoming in American Economic Journal: Microeconomics*, forthcoming.

_ **and E.I. Vespa**, "Endogenous Sample Selection: A Laboratory Study," *working paper*, 2015.

Eyster, E. and T. Kittsteiner, "Party platforms in electoral competition with heterogeneous constituencies," *Theoretical Economics*, 2007, (2), 41–70.

Feddersen, T. and W. Pesendorfer, "Voting behavior and information aggregation in elections with private information," *Econometrica*, 1997, pp. 1029–1058.

Ferejohn, J., "Incumbent performance and electoral control," *Public choice*, 1986, *50* (1), 5–25.

Fiorina, M.P., *Retrospective voting in American national elections*, Yale University Press, 1981.

Fudenberg, D. and D.K. Levine, *The theory of learning in games*, Vol. 2, The MIT press, 1998.

Gerber, Elisabeth R and Rebecca B Morton, "Primary election systems and representation," *Journal of Law Economics and Organization*, 1998, *14* (2), 304–324.

Gul, F. and W. Pesendorfer, "Partisan politics and election failure with ignorant voters," *Journal of Economic Theory*, 2009, *144* (1), 146–174.

Healy, A. and N. Malhotra, "Random Events, Economic Losses, and Retrospective Voting: Implications for Democratic Competence," *Quarterly Journal of Political Science*, 2010, *5* (2), 193–208.

Kamada, Yuichiro and Fuhito Kojima, "Voter preferences, polarization, and electoral policies," *American Economic Journal: Microeconomics*, 2014, *6* (4), 203–236.

Kartik, N. and R. P. McAfee, "Signaling character in electoral competition," *The American economic review*, 2007, pp. 852–870.

Key, V.O., *The responsible electorate*, Vintage Books New York, 1966.

Leigh, A., "Does the World Economy Swing National Elections?," *Oxford Bulletin of Economics and Statistics*, 2009, *71* (2), 163–181.

Levy, G. and R. Razin, "Does Polarisation of Opinions Lead to Polarisation of Platforms? The Case of Correlation Neglect," *Quarterly Journal of Political Science*, 2015, *10*, 321–355.

Lizzeri, Alessandro and Leeat Yariv, "Collective Self Control," *working paper*, 2015.

McCarty, N.M., K.T. Poole, and H. Rosenthal, *Polarized America: The dance of ideology and unequal riches*, MIT Press Cambridge, MA, 2006.

McMurray, J. C., "Ideology as Opinion: Polarization in a Spatial Model of Common-value Elections," *working paper*, 2015.

Osborne, M. J. and A. Slivinski, "A model of political competition with citizen-candidates," *The Quarterly Journal of Economics*, 1996, *111* (1), 65–96.

Palfrey, T. R., "Spatial equilibrium with entry," *The Review of Economic Studies*, 1984, *51* (1), 139–156.

Snyder, James M and Michael M Ting, "An informational rationale for political parties," *American Journal of Political Science*, 2002, pp. 90–110.

Spiegler, R., "Placebo reforms," *The American Economic Review*, 2013, *103* (4), 1490–1506.

T., Persson and G. Tabellini, *Political economics - Explaining economic policy*, MIT Press: Cambridge, 2000.

Tomz, Michael and Robert Van Houweling, *Political Repositioning* forthcoming.

Wittman, D., "Candidates with policy preferences: A dynamic model," *Journal of Economic Theory*, 1977, *14* (1), 180–189.

Wolfers, J., "Are Voters Rational? Evidence from Gubernatorial Elections," *working paper*, 2007.

Appendix A: Proof of results

Proof of Proposition 1. By A3, A4, and continuity of the utility functions (A1), there is a unique cutoff state $\omega^{NE}(l,r)$ that solves $u(l,\omega) = u(r,\omega)$. In addition, by A3, $u(l,\omega) < u(r,\omega)$ for all $\omega < \omega^{NE}(l,r)$ and $u(l,\omega) > u(r,\omega)$ for all $\omega > \omega^{NE}(l,r)$. By full information equivalence (Feddersen and Pesendorfer, 1997), the NE voting outcome in large elections is asymptotically equal to the outcome arising under perfect information about the state. Thus, the party choosing l wins in states $\omega < \omega^{NE}(l,r)$ and the party choosing r wins in states $\omega > \omega^{NE}(l,r)$.

We now compare these cutoff states for different policies. For policies (L,N), the cutoff solves $u(L,\omega) = 0$ and, by A3, $\omega^{NE}(L,N) < 0$. A similar argument establishes that $\omega^{NE}(N,R) > 0$. Finally, for policies (L,R), the cutoff solves $u(L,\omega) = u(R,\omega)$ and, by A3, $\omega^{NE}(L,R) \in (\omega^{NE}(L,N), \omega^{NE}(N,R))$. □

Proof of Proposition 2. Consider first fixed policies (N,R), so that $v(s;\omega) = E\big[u(R,W) \mid W \geq w, S = s\big]$. By A4, $v(s_L;0) > 0$. Because $v(s_L;\cdot)$ is increasing, then the personal cutoff $c(s_L) < 0$. By Theorem 1, $\omega^{RVE}(N,R) \leq c(s_L)$. Therefore, $\omega^{RVE}(N,R) < 0$. A similar argument yields $\omega^{RVE}(L,N) > 0$.

Next, we compare the personal cutoffs for fixed policies (N,R) and (L,R). Let $v_{(x_{Left},x_{Right})} \equiv v$ and $c_{(x_{Left},x_{Right})}(s)$ denote the corresponding belief function and the personal cutoff given s for policies (x_{Left}, x_{Right}). Note that $v_{(N,R)}(s;\omega) = E\big[u(R,W) \mid W \geq w, S = s\big]$ and $v_{(L,R)}(s;\omega) = E\big[u(R,W) \mid W \geq w, S = s\big] - E\big[u(L,W) \mid W \leq w, S = s\big]$. By A2, A3(i), and A4, for all s, $v_{(N,R)}(s;0) > 0$ and $v_{(N,R)}(s;\omega) - v_{(L,R)}(s;\omega) = E\big[u(L,W) \mid W \leq w, S = s\big] > 0$ for all $\omega \leq 0$. Thus, $c_{(N,R)}(s) \leq c_{(L,R)}(s)$ for all s. By Theorem 1, it follows that $\omega^{RVE}(N,R) \leq \omega^{RVE}(L,R)$. A similar argument yields $\omega^{RVE}(L,N) \geq \omega^{RVE}(L,R)$.

Finally, consider an environment with information precision z and compare policies (N,R) and (L,R). By the definition of information precision and A1(iii), $\big|q(s|\omega)/q(s|\omega')-$

$1\big| \leq z/d$ for all $\omega, \omega' \in \Omega$ and $s \in \{s_L, s_R\}$. Thus, for any $z < d$,

$$\begin{aligned}
&E\big[u(R,W) \mid W \geq \omega, S = s_R\big] - E\big[u(R,W) \mid W \geq \omega, S = s_L\big] \\
&= \int_\omega^1 u(R,\tilde{\omega})\Big(\frac{q(s_R|\tilde{\omega})}{\int_\omega^1 q(s_R \mid x)g(x)dx} - \frac{q(s_L|\tilde{\omega})}{\int_\omega^1 q(s_L \mid x)g(x)dx}\Big)G(d\tilde{\omega}) \\
&\leq \int_\omega^1 \big|u(R,\tilde{\omega})\big|\Big|\frac{1}{\int_\omega^1 \frac{q(s_R|x)}{q(s_R|\tilde{\omega})}g(x)dx} - \frac{1}{\int_\omega^1 \frac{q(s_L|x)}{q(s_L|\tilde{\omega})}g(x)dx}\Big|G(d\tilde{\omega}) \\
&\leq \big(\frac{1}{1-z/d} - \frac{1}{1+z/d}\big)\int_\omega^1 \big|u(R,\tilde{\omega})\big|(1-G(\omega))^{-1}G(d\tilde{\omega}) \\
&= \frac{2z/d}{(1-z/d)\,(1+z/d)}E\big[\big|u(R,W)\big| \mid W \geq \omega\big].
\end{aligned}$$

Since the utility function is bounded, the above expression goes to zero as z goes to zero. This fact and the fact, implied by A3(i)-A4, that $E\big[u(L,W) \mid W \leq \omega, S = s_R\big] > 0$ for all $\omega \leq 0$, imply that there exists $\bar{z}_1 > 0$ such that, for all $z < \bar{z}_1$, $v_{(N,R)}(s_L;\omega) > v_{(L,R)}(s_R;\omega)$ for all $\omega \leq 0$. Thus, in such environments, $c_{(N,R)}(s_L) \leq c_{(L,R)}(s_R)$, with equality holding if and only if both personal cutoffs are equal to -1. The assumption that $E\big[u(R,W) \mid S = s_R\big] < u(L,-1)$, however, implies that $c_{(L,R)}(s_R) \neq -1$. Thus, $c_{(N,R)}(s_L) < c_{(L,R)}(s_R)$ and, by Theorem 1, $\omega^{RVE}(N,R) < \omega^{RVE}(L,R)$ for environments with $z \leq \bar{z}_1$. A similar argument shows that there exists $\bar{z}_2 > 0$ such that, for all $z \leq \bar{z}_2$, $\omega^{RVE}(L,N) > \omega^{RVE}(L,R)$ if $E\big[u(L,W) \mid S = s_L\big] < u(R,1)$ holds. The claim follows by setting $\bar{z} \equiv min\{\bar{z}_1, \bar{z}_2\}$. □

Proof of Proposition 3. Note that, by the full support assumption over Ω (A1), the payoff of the $Left$ [$Right$] party increases as the equilibrium cutoff increases [decreases] under both NE and RVE voting.

NE voting: We show that for each party, policy N is a dominant strategy. If the $Left$ party chooses L, then the $Right$ party prefers N to R because, by Proposition 1, $\omega^{NE}(L,N) < \omega^{NE}(L,R)$. And, if the $Left$ party chooses N, then the $Right$ party prefers N to R because, by Proposition 1, $\omega^{NE}(N,N) = 0 < \omega^{NE}(N,R)$. Thus, N is a dominant strategy for the $Right$ party. Similarly, the facts that $\omega^{NE}(N,R) > \omega^{NE}(L,R)$ and $\omega^{NE}(N,N) = 0 > \omega^{NE}(L,N)$ (see Proposition 1) imply that N is a dominant strategy for the $Left$ party.

RVE voting: By Proposition 2, $\omega^{RVE}(L,R) \leq \omega^{RVE}(L,N)$ and $\omega^{RVE}(N,R) < 0 = \omega^{RVE}(N,N)$, implying that R weakly dominates N for the $Right$ party. Similarly, by Proposition 2, $\omega^{RVE}(L,N) > 0 = \omega^{RVE}(N,N)$ and $\omega^{RVE}(L,R) \geq \omega^{RVE}(N,R)$, implying that L weakly dominates N for the $Left$ party. Moreover, by Proposition 2, there exists $\bar{z} > 0$ such that, in any environment with information precision $z \leq \bar{z}$ and satisfying $E\big[u(R,W) \mid S = s_R\big] < u(L,-1)$ and $E\big[u(L,W) \mid S = s_L\big] < u(R,1)$, the previous inequalities regarding the RVE cutoffs are strict. Therefore, L and R are strictly dominant strategies for the $Left$ and $Right$ parties, respectively, and so (L,R) is the unique policy equilibrium. □

Proof of Proposition 4. By Proposition 3, the unique policy equilibrium under NE is (N,N) and, therefore, welfare is equal to zero at every state. For comparison, consider (L,R), which, by Proposition 3, is a policy equilibrium under RVE in weakly undominated strategies, (L,R). Suppose that $\omega^{RVE}(L,R) \leq 0$; the argument for the other case is analogous and, therefore, omitted. By Theorem 1, $v(s_R; \omega^{RVE}(L,R)) \geq 0$. The previous inequality, A3(i), A4, and $\omega^{RVE}(L,R) \leq 0$ imply that $E\big[u(L,W) \mid W \leq \omega^{RVE}(L,R), s_R\big] > 0$ and $E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R), s_R\big] > 0$. The first of these inequalities and A2 imply that

$$E\big[u(L,W) \mid W \leq \omega^{RVE}(L,R)\big] > 0.$$

Moreover, by algebra similar to that in the proof of Proposition 2, it follows that, for an environment with information precision $z < d$,

$$\begin{aligned} E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R)\big] \geq {} & E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R), s_R\big] \\ & - \frac{z/d}{1 - z/d} E\big[\lvert u(R,W)\rvert \mid W \geq \omega^{RVE}(L,R)\big]. \end{aligned}$$

The above inequality, the previous result that $E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R), s_R\big] > 0$, and the assumption that $u(R,\cdot)$ is bounded imply that there exists $\hat{z}$ such that, for an environment with information precision $z \leq \hat{z}$,

$$E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R)\big] > 0.$$

Since expected welfare under (L,R) is given by the weighted average of $E\big[u(L,W) \mid W \leq \omega^{RVE}(L,R)\big]$ and $E\big[u(R,W) \mid W \geq \omega^{RVE}(L,R)\big]$, the fact that each of these two terms is strictly positive implies that expected welfare is strictly positive, and so strictly higher than welfare under (N,N).

By Proposition 3, it remains to consider (N, R) and (L, N) as potential equilibria under RVE. Suppose that (N, R) is an equilibrium policy under RVE. By Theorem 1, the lowest possible equilibrium cutoff is $c_{(N,R)}(s_R)$ and satisfies

$$v(s_R; c_{(N,R)}(s_R)) = E\big[u(R, W) \mid W \geq c_{(N,R)}(s_R), s_R\big] \geq 0.$$

This inequality and A3(i)-A4 imply that $E\big[u(R, W) \mid W \geq \omega^{RVE}(N, R)\big] > 0$ for any equilibrium cutoff $\omega^{RVE}(N, R) \geq c_{(N,R)}(s_R)$. Thus, expected welfare is given by a weighted average of the previous term, which is strictly positive, and zero, from policy N. Moreover, the weight on the strictly positive term is strictly positive because, by Proposition 2, $\omega^{RVE}(N, R) < 0$. Therefore, the expected welfare under the RVE policy (N, R) is strictly positive. A similar argument establishes that expected welfare under an RVE policy (L, N) is also strictly positive. $\square$

Appendix B: Continuum of policies

In this Appendix, we present the model with a continuum of policies spelled out in Section 6. For simplicity, we restrict attention to the case of no private information (which should be interpreted as the limiting case of information precision going to zero; see Esponda and Pouzo (forthcoming)).[15] We replace assumptions A3 and A4 with the following assumptions:

A3'. (Monotone policies) (i) $u(L,\cdot)$ is decreasing for all $L<0$, $u(R,\cdot)$ is increasing for all $R>0$, and $u(0,\omega)=0$ for all $\omega\in\Omega$; (ii) $u(x,0)<0$ for all policies $x\neq 0$; (iii) For all $\omega\in\Omega$: If $u(R,\omega)<0$, then $u(R',\omega)<u(R,\omega)$ for all $R'>R$; similarly, if $u(L,\omega)<0$, then $u(L',\omega)<u(L,\omega)$ for all $L'<L$.

A4'. (Potential value of polarization) There exist policies $\bar{L}<0$ and $\bar{R}>0$ such that $E\big[u(\bar{L},W)\mid W\leq 0\big]>0$ and $E\big[u(\bar{R},W)\mid W\geq 0\big]>0$.

These assumptions are the analogs of A3 and A4. The added assumption is A3'(iii), which provides an ordering for polarized policies. It roughly says that, the more polarized a policy, then the worse it is in "neutral" states. For example, if a polarized policy R is worse than the neutral policy, then a more polarized policy $R'>R$ must be even worse.

As mentioned in the text, party $Left$ is restricted to $L\leq-\varepsilon$ and party $Right$ to $R\geq\varepsilon$, where $\varepsilon>0$. A policy equilibrium (under NE or RVE) is defined to be the limit of a sequence of policy equilibria as ε goes to zero.

Proposition 5. *The Neutral policy profile* $(0,0)$ *is the unique policy equilibrium when voters play NE, but it is not an equilibrium when voters play RVE. Moreover, any policy equilibrium under RVE yields strictly higher welfare than the unique policy equilibrium under NE.*

Proof. Suppose voters play NE. Consider an ε-game, where $L\leq-\varepsilon$ and $R\geq\varepsilon$, and the policies $\bar{L}$ and $\bar{R}$ defined in A4' are feasible, i.e., $\varepsilon\leq\max\left\{-\bar{L},\bar{R}\right\}$. First, we show that $L=-\varepsilon$ is the unique best response to any R such that the NE cutoff $\omega^{NE}(-\varepsilon,R)$ is interior. For any deviation $L<-\varepsilon$,

$$0>u\left(R,\omega^{NE}(-\varepsilon,R)\right)=u\left(-\varepsilon,\omega^{NE}(-\varepsilon,R)\right)>u\left(L,\omega^{NE}(-\varepsilon,R)\right), \qquad (6)$$

[15] By continuity, the results extend to settings with sufficiently low information precision.

where the first inequality follows because, by monotonicity of u (A3'(i)) and by A3'(ii), $u(R,\cdot)$ and $u(L,\cdot)$ must be strictly negative at their point of intersection $\omega^{NE}(-\varepsilon, R)$; the equality follows by full information equivalence and the assumption that $\omega^{NE}(-\varepsilon, R)$ is interior; and the last inequality follows by A3'(iii) and the result that $u\left(-\varepsilon, \omega^{NE}(-\varepsilon, R)\right) < 0$. Equation (6) and (B1) imply that $\omega^{NE}(L, R) < \omega^{NE}(-\varepsilon, R)$, so that the *Left* party is strictly worse off from deviating to any policy that is more polarized than $L = -\varepsilon$.

Second, we show that $L = -\varepsilon$ is the (not necessarily unique) best response to any R such that the NE cutoff $\omega^{NE}(-\varepsilon, R)$ is *not* interior. Note that $\omega^{NE}(-\varepsilon, R) = -1$ is not possible, since that would imply that $u(-\varepsilon, -1) < 0$, which then implies, by A3'(iii), that $u(L, -1) < 0$ for all $L \leq -\varepsilon$, thus contradicting that polarization can potentially be beneficial (A4'). So it must be the case that, if the cutoff is not interior, $\omega^{NE}(-\varepsilon, R) = 1$. But then the *Left* party cannot do better and $-\varepsilon$ is a best response to R.

Third, note that a similar argument establishes that $R = \varepsilon$ is a best response to any L, and it is a unique best response provided that the cutoff $\omega^{NE}(L, \varepsilon)$ is interior. In particular, $(-\varepsilon, \varepsilon)$ is a policy equilibrium of the ε-game.

Fourth, we establish that $(-\varepsilon, \varepsilon)$ is the unique equilibrium. Suppose that there were another equilibrium, $(L, R) \neq (-\varepsilon, \varepsilon)$. Suppose, without loss of generality, that $L \neq -\varepsilon$. Then, it must be the case that $\omega^{NE}(L, R) = 1$ (otherwise, by the previous steps, $-\varepsilon$ would be the unique best response). But then the *Right* party has a profitable deviation to $\bar{R}$ (defined in A4'). To see this, note that $u(\bar{R}, 1) > 0$ and that, by A3'(i) and A3'(ii), $u(L, 1) < 0$, which then implies, by A3'(i), that $\omega^{NE}(L, \bar{R}) < 1$. Thus, (L, R) is not an equilibrium, and there is a unique equilibrium $(-\varepsilon, \varepsilon)$. The result that $(0, 0)$ is the unique policy equilibrium under NE follows by taking $\varepsilon \to 0$.

Now suppose that voters play RVE. Consider a sequence of policy equilibria $(L_\varepsilon, R_\varepsilon)$. Suppose, in order to obtain a contradiction, that there exists a subsequence (which we still denote as ε) along which $\lim_{\varepsilon\to 0}(L_\varepsilon, R_\varepsilon) = (0, 0)$. By A4', there exist $\bar{R}$ and $\bar{L}$ such that $\bar{u}_R \equiv E\big[u(\bar{R}, W)|W \geq 0\big] > 0$ and $\bar{u}_L \equiv E\big[u(\bar{L}, W)|W \leq 0\big] > 0$. Let $\bar{u} \equiv \min\{\bar{u}_R, \bar{u}_L\}$. Continuity of payoffs in policies and the assumption that $\lim_{\varepsilon\to 0}(L_\varepsilon, R_\varepsilon) = (0, 0)$ imply that

$$\lim_{\varepsilon\to 0} E\big[u(R_\varepsilon, W)|W \geq 0\big] = \lim_{\varepsilon\to 0} E\big[u(L_\varepsilon, W)|W \leq 0\big] = 0.$$

In particular, there exists $\bar{\varepsilon}$ such that, for all $\varepsilon \leq \bar{\varepsilon}$, $E\big[u(R_\varepsilon, W)|W \geq 0\big] \leq \bar{u}/2$ and $E\big[u(L_\varepsilon, W)|W \leq 0\big] \leq \bar{u}/2$. Suppose that $\varepsilon \leq \bar{\varepsilon}$ and that the equilibrium cutoff $\omega^{RVE}(L_\varepsilon, R_\varepsilon) \geq 0$. If party *Right* deviates to $\bar{R}$, then

$$E\big[u(\bar{R}, W)|W \geq 0\big] - E\big[u(L_\varepsilon, W)|W \leq 0\big] \geq \bar{u}/2 > 0.$$

It follows from Theorem 1 that $\omega^{RVE}(L_\varepsilon, \bar{R}) < 0$, which shows that the *Right* party has a profitable deviation and, therefore, contradicts the assumption that $(L_\varepsilon, R_\varepsilon)$ is an equilibrium. A similar logic (with the *Left* party deviating to $\bar{L}$) applies to the case $\omega^{RVE}(L_\varepsilon, R_\varepsilon) < 0$.

Finally, to compare welfare, consider any policy equilibrium under RVE, (L, R), with election cutoff $\omega^* \equiv \omega^{RVE}(L, R) \leq 0$. Theorem 1 implies that the equilibrium welfare of voters is $Welfare(\omega^*) = E\big[u(R, W)|W \geq \omega^*\big]$. Suppose, in order to obtain a contradiction, that $Welfare(\omega^*) \leq 0$, where 0 is the payoff in the policy equilibrium under NE. Then, by A4', there exists $\bar{L}$ such that

$$\begin{aligned} E\big[u(R, W)|W \geq \omega^*\big] &\leq 0 < E\big[u(\bar{L}, W)|W \leq 0\big] \\ &\leq E\big[u(\bar{L}, W)|W \leq \omega^*\big], \end{aligned}$$

where the last inequality follows by A3'(i). Then $\bar{L}$ is a profitable deviation for the *Left* party because, by A3'(i), $\omega^{RVE}(\bar{L}, R) > \omega^{RVE}(L, R)$, thus contradicting that (L, R) is an equilibrium under RVE. The case where $\omega^{RVE}(L, R) > 0$ is similar and, therefore, omitted. □