Conditional Retrospective Voting in Large Elections

By Ignacio Esponda and Demian Pouzo*

We introduce a solution concept in the context of large elections with private information by embedding a model of boundedly rational voters into an otherwise standard equilibrium setting. A retrospective voting equilibrium *(RVE) formalizes the idea that voters evaluate alternatives based on past performance. Since counterfactual outcomes are not observed, the sample from which voters learn is potentially biased, leading to systematically biased beliefs in equilibrium. We provide an explicit learning foundation for RVE and contrast it to standard solution concepts in the literature.*
JEL: C72, D72, D83.
Keywords: bounded rationality, retrospective voting, sample selection, biased learning

In economics, voters are usually portrayed as sophisticated individuals who have well-defined preferences, can solve complicated signal-extraction problems, and have correct expectations about the distribution of (counterfactual) payoffs. The empirical evidence, however, often finds that voters are poorly informed and have little understanding of ideology and policy.[1] 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 main contribution of this paper is to embed a model of boundedly rational voters who learn from the previous performance of policies or parties into an otherwise standard equilibrium setting. By doing so, we are able to capture an important feature of elections that is often overlooked in the literature. To illustrate this feature, consider an election between a Republican and a Democratic candidate in the United States. Voters are likely to use information about the past performance of the parties to predict their future performance and determine which party to vote for. For example, voters who are currently unemployed may favor a Democratic candidate if, while being unemployed in the past, they experienced better results from previously elected Democratic, compared to Republican, administrations. This tendency to learn from the past is not limited to political elections. When shareholders vote on takeover proposals, they ben-

* 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. We thank Heski Bar-Isaac, Ernesto Dal Bo, Erik Eyster, Frederico Finan, Gregory Huber, Navin Kartik, Alessandro Lizzeri, Salvatore Nunnari, Santiago Oliveros, Marco Ottaviani, Wolfgang Pesendorfer, Jesse Rothstein, Andy Skrzypacz, and several seminar participants, anonymous referees, and the editor for their comments.

[1] See, e.g., Delli Carpini and Keeter (1997) and Converse (2000).

efit from learning the outcome of previous takeovers in the same or comparable firms. A similar phenomenon occurs with legislators choosing whether to vote along party lines, union members voting to accept or reject negotiated contracts, and residents voting whether to approve additional funding for school districts.

A key feature in these examples is that, when using past information to evaluate alternatives, voters only observe the performance of the *elected* alternatives, so that counterfactual outcomes are not observable. For example, we will never find out how Romney would have performed had he been elected President of the U.S. in 2012 instead of Obama. Similarly, shareholders will not learn the benefits of a takeover that is not approved. Consequently, the sample from which voters learn is potentially biased. The reason is that the selection of alternatives is not randomized: To the extent that voters have some private information, they will elect alternatives that are likely to perform better. Thus, voters who fail to account for other voters' private information will end up with systematically biased beliefs.

There is a large experimental literature in common value auctions and elections that shows that many subjects do not account for the informational content of other people's actions (e.g., Thaler (1988), Kagel and Levin (2002), Charness and Levin (2009), Esponda and Vespa (2014)). Recently, Esponda and Vespa (2015) find that essentially no subject correctly accounts for the sample selection problem that arises in a setting with private information and lack of counterfactual information.

The idea that voters do not account for unobserved counterfactuals is also consistent with the empirical findings of Achen and Bartels (2004), Leigh (2009), and Wolfers (2009), 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 our setup, there is a continuum of voters and two alternatives. One of the alternatives wins the election if it receives a high enough proportion of votes; otherwise the other alternative wins. Voters have payoffs that are increasing in the state of the world for one alternative and decreasing for the other. In addition, voters have some information about the state of the world. This environment is essentially the one considered by the literature on voting and information aggregation (e.g., Feddersen and Pesendorfer (1997)). For example, in an election between two political parties, the state can represent the fundamentals of the economy. One of the parties might be best at governing during recessions and the other during booms (perhaps because of their different positions on monetary and fiscal policy).

We propose a new solution concept, *retrospective voting equilibrium* (RVE), to formalize the idea that voters learn from a biased sample and have systematically biased beliefs. An RVE consists of a strategy profile and an election cutoff that satisfy two conditions: (i) there is a tie at the cutoff, with one alternative being elected above and the other below the cutoff; (ii) the strategy profile must be *optimal* given the election cutoff. Optimality is defined in terms of retrospective voting: Voters' perceptions of the benefits of each alternative derive from the observed performance of each alternative, which depends on the states in which each alternative is elected, and, therefore, on the election cutoff. This parsimonious characterization of retrospective voting in large elections is a major advantage of the framework.

We then contrast RVE to the two most prominent solution concepts for voting games, Nash equilibrium (NE) and sincere voting (SV). By capturing the sample selection problem, RVE provides different insights about the information aggregation properties of elections while exhibiting the more attractive features of these other solution concepts: Behavior is endogenous (as in NE, but unlike SV), but outcomes depend on both the electoral rule and the precision of information, and individual voting behavior depends on private information for a significant fraction of the electorate (as in SV, but unlike NE).

We also provide a foundation for RVE by studying a dynamic voting environment with a finite number of players. We first characterize steady state behavior when players follow a simple retrospective learning rule that uses *observed* performance to update beliefs about the alternatives and fails to account for sample selection. We do so by adapting the results of Fudenberg and Kreps (1993), who provide a learning foundation for Nash equilibrium in strategic-form games of complete information. In our setting, however, there is private information and players do not observe the (information-contingent) strategies of other players. Instead, players learn the average performance of the two alternatives from the biased sample of elected alternatives. Finally, we show that, in the limit, as the number of players goes to infinity, the steady state of the dynamic environment is characterized by our notion of RVE.

This paper follows the literature on learning in games (see Fudenberg and Levine (1998) and Fudenberg and Levine (2009) for surveys) as well as the theoretical literature on bounded rationality that studies mistakes in learning, coarse thinking, and selection (e.g., Rubinstein (1993), Osborne and Rubinstein (1998), Barberis et al. (1998), Rabin and Schrag (1999), Rabin (2000), Jehiel (2005), Eyster and Rabin (2005), Jehiel and Samet (2007), Jehiel and Koessler (2008), Mullainathan et al. (2008), Gabaix (2014), Schwartzstein (2014), Spiegler (2014), Esponda and Pouzo (2015)).[2] Esponda (2008) was the first to propose a general solution concept where players fail to account for endogenous sample selection.

Spiegler (2013) studies a dynamic model of reforms in 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

[2]In voting contexts, the most common solution concept, other than Nash equilibrium, is sincere voting (generalized by Eyster and Rabin (2005)). See also Osborne and Rubinstein (2003), who apply their notion of sampling equilibrium to a voting context.

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. Bendor et al. (2010, 2011) postulate a dynamic model of retrospective voting where voters follow a satisficing rule and vote for the incumbent if it has performed well given their endogenous aspiration level. These papers focus on other interesting aspects of bounded rationality and not on the type of sample selection problem that motivates our paper.[3]

Our work is conceptually different to the formal literature in retrospective voting, beginning with Barro (1973) and Ferejohn (1986), which studies elections as incentive mechanisms to hold politicians accountable. Instead, our model follows 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)).[4]

In Section I, we introduce the framework and the solution concept, and compare it to Nash equilibrium and sincere voting. In Section II, we provide the equilibrium foundation.

I. Voting framework

A. *Setup*

A continuum of voters participate in an election between two alternatives, R (Right) and L (Left). A state $\omega \in \Omega = [-1, 1]$ is first drawn according to a probability distribution G and, conditional on the state, each player observes an independently-drawn private signal. Players then simultaneously submit a vote for either R or L. Votes are aggregated according to an electoral rule $\rho \in (0, 1)$: Alternative R is elected if the proportion of votes in favor of R is greater or equal than ρ; otherwise, L is elected.

For expositional clarity, we assume that all voters have the same preferences and information; the extension to heterogenous voters is straightforward and relegated to Online Appendix A. Conditional on a state $W = \omega$, players independently draw a signal $S = s$ from a finite, nonempty set $\mathbb{S} \subset \mathbb{R}$ with probability $q(s \mid \omega)$; let s^L and s^R denote the lowest and highest signals in $\mathbb{S}$. The payoff of each voter is $u(o, \omega)$, where $o \in \{L, R\}$ is the winner of the election.

Let $\sigma : \mathbb{S} \to [0, 1]$ denote the strategy of an (average) voter, where $\sigma(s)$ is the probability of voting for alternative R after observing signal s.[5] A strategy σ is *nondecreasing* if $\sigma(s') \geq \sigma(s)$ for all $s' > s$.

We maintain the following assumptions throughout the paper:

[3] Callander (2011) studies a model of dynamic policy-making where "rational" voters learn the mapping between policies and outcomes.

[4] In the words of Fiorina (1981, p. 5), voters "need *not* know the precise economic or foreign policies of the incumbent administration in order to see or feel the *results* of those policies."

[5] As shown in Online Appendix B, σ represents the *average* strategy in the population and so we are not restricting voters to play symmetric strategies.

A1. (i) $u(R,\cdot):\Omega\to\mathbb{R}$ is nondecreasing, $u(L,\cdot):\Omega\to\mathbb{R}$ is nonincreasing, and one of them is strictly monotone; (ii) $u(R,\cdot)$ and $u(L,\cdot)$ are continuously differentiable, except possibly in a finite number of points, and $\sup_{(o,\omega)\in\{L,R\}\times\Omega}|u(o,\omega)|\leq K<\infty$.

A2. MLRP: For all $\omega'>\omega$, and $s'>s$:

$$\frac{q(s'|\omega')}{q(s'|\omega)}-\frac{q(s|\omega')}{q(s|\omega)}>0.$$

A3. (i) G has a density function g, where $\inf_{\omega\in\Omega}g(\omega)>0$; (ii) there exists $d>0$ such that $q(s|\omega)>d$ for all $s\in\mathbb{S}$ and $\omega\in\Omega$; (iii) $q(s\mid\cdot)$ is continuous for all $s\in\mathbb{S}$.

Assumptions A1-A2 provide an ordering between states, information, and players' preferences.[6] Note that A2 is trivially satisfied if there is no private information (i.e., $\mathbb{S}$ is a singleton). The case without private information should be viewed as the limiting case of a class of MLRP environments with private information in which information precision vanishes (see Online Appendix B.B2, Remark 2, for details).[7] Assumption A3 rules out "strong signals" in the sense of Milgrom (1979).

Example 1. The state is uniformly distributed in $[-1,1]$ and there is a continuum of identical voters with payoffs $u(R,\omega)=\omega-1/3$, $u(L,\omega)=-\omega-1/3$, so that the payoff from the Right [Left] policy is increasing [decreasing] in the state. In particular, $c^{FB}=0$ is the first-best election cutoff, i.e., everyone prefers R in states $\omega>c^{FB}$ and L in states $\omega<c^{FB}$. In addition, each voter privately observes a binary signal from $\mathbb{S}=\{s^L,s^R\}$ with probability $q(s^R\mid\omega)=.5+\iota\omega$, where $\iota\in(0,.5]$ is the precision of information. $\square$

B. *Retrospective voting equilibrium*

Let

$$\kappa(\omega;\sigma)\equiv\sum_{s\in\mathbb{S}}q(s\mid\omega)\sigma(s)$$

denote the proportion of votes in favor of alternative R under state ω when voters follow strategy σ. Assumption A2 implies that $\kappa(\cdot;\sigma)$ is nondecreasing if σ is nondecreasing. In the case where the strategy depends on private information, so that $\sigma(\cdot)$ is not a constant function, then $\kappa(\cdot;\sigma)$ is increasing and the outcome of the election can be characterized by a cutoff: R is elected if and only if $\kappa(\omega;\sigma)\geq\rho$,

[6]One difference with the standard setup (e.g., Feddersen and Pesendorfer (1997)) is that, to characterize Nash equilibrium, it suffices to require that $u(R,\cdot)-u(L,\cdot)$ is monotone, as opposed to each individual term being monotone.

[7]For the game-theoretic foundation of RVE in Section II, we rely on assumption A5, which strengthens A2 by requiring a bound on the rate at which the likelihood ratio changes.

or, equivalently, for all sufficiently high states. This observation motivates the following definition.[8]

DEFINITION 1: *A state $\omega \in \Omega$ is an* election cutoff *given a strategy σ if $\kappa(\tilde{\omega};\sigma) \geq \rho$ for all $\tilde{\omega} > \omega$ and $\kappa(\tilde{\omega};\sigma) \leq \rho$ for all $\tilde{\omega} < \omega$.*

When making her decision, each voter takes the cutoff as given. A cutoff determines the set of states for which each alternative is chosen, and, consequently, each voter's evaluation of the benefits of electing each alternative. For a given cutoff $\omega \in \Omega$, the difference in benefits from electing R over L that is perceived by a voter who observes signal s is

$$v(s;\omega) \equiv E\left(u(R,W) \mid W > \omega, S = s\right) - E\left(u(L,W) \mid W < \omega, S = s\right). \tag{1}$$

To interpret the above expression, note that, for election cutoff ω, alternative R is elected whenever $W > \omega$. So a voter's retrospective evaluation of R is given by expected observed performance of R, conditional on her own signal, which is the first term in the right hand side of (1). A similar interpretation holds for the second term. Assumptions A1-A2 guarantee that $v(\cdot;\omega)$ is increasing.

The following definition captures the idea that each voter votes for the alternative that she sincerely believes to have the highest perceived benefit.

DEFINITION 2: *A strategy σ is* optimal *given an election cutoff ω if, for all $s \in \mathbb{S}$, $v(s;\omega) > 0$ implies $\sigma(s) = 1$ and $v(s;\omega) < 0$ implies $\sigma(s) = 0$.*

Note that the fact that $v(\cdot;\omega)$ is increasing implies that optimal strategies must be nondecreasing.

DEFINITION 3: *A* retrospective voting equilibrium *(RVE) is a strategy σ^* and an election cutoff ω^* such that: (i) σ^* is optimal given ω^*, and (ii) ω^* is an election cutoff given σ^*.*

A retrospective voting equilibrium requires players to choose a strategy σ^* that is optimal given their beliefs about the expected benefits of the alternatives. Beliefs, however, are not necessarily correct but rather determined by the average *observed* performance of the alternatives, $v(s,\omega^*)$, which depends on the election cutoff ω^*. Thus, voters do not try to account for the fact that they do not observe the performance of an alternative that is not elected. In addition, the election cutoff ω^* is endogenously determined by the equilibrium strategy σ^*. In particular, unlike the standard notion of sincere voting, voting behavior depends endogenously on the aggregate behavior of all voters.

Notice that voters form beliefs *conditional on their signals*, and so their choices may depend on the observed signals. For instance, in the context of a political

[8] When σ, and, therefore, $\kappa(\cdot;\sigma)$ are constant functions, this definition is motivated by the limiting case where signals satisfy MLRP but become uninformative; see Online Appendix B.B2.

election, voters observe signals (i.e., news) about the state of the economy, foreign affairs, etc. A voter, for example, may learn that the performance of an economic policy is correlated with her chances of becoming unemployed.

In Section II, we argue that retrospective voting equilibrium (RVE) corresponds to a steady state of a dynamic environment in which a *large* number of voters are assumed to follow a particular retrospective voting rule. In particular, voters do *not* compute conditional expectations to reach expression (1); in contrast, voters follow a particular retrospective voting rule and their beliefs happen to be characterized in the steady state by expression (1).

C. *Characterization of RVE*

We now characterize retrospective voting equilibrium. For each signal s, define the *personal cutoff*

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

which depends only on the primitives of the environment. Since Ω is compact and $v(s;\cdot)$ is continuous and increasing (by A1-A3), there exists a unique solution $c(s)$; moreover, $c(\cdot)$ is nonincreasing because $v(\cdot;\omega)$ is nondecreasing for all $\omega \in \Omega$. Let $\underline{c} \equiv c(s^R)$ and $\overline{c} \equiv c(s^L)$ denote the lowest and highest personal cutoffs.

If we knew the equilibrium election cutoff ω^*, then it would be straightforward to characterize the equilibrium strategy: a voter with signal s such that $c(s) < \omega^*$ must satisfy $v(s;\omega^*) > 0$ and, therefore, she will optimally vote for R; similarly, if $c(s) > \omega^*$, then she will optimally vote for L. For example, consider a voter with two signals and personal cutoffs $c(s^R) < c(s^L)$, as depicted in Figure 1. If the equilibrium election cutoff were lower than $c(s^R)$, this voter would always vote for L. Similarly, if the election cutoff were higher than $c(s^L)$, she would always vote for R. In the case where the election cutoff were between her personal cutoffs $c(s^R)$ and $c(s^L)$, this voter would vote her signal.[9]

We now characterize the set of equilibrium cutoffs. For any election cutoff $\omega \in \Omega$,

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

may be interpreted as the proportion of players that vote for R in state ω when the cutoff is also given by ω.[10]

The next result says that there is a unique equilibrium cutoff and that it is essentially given by the state where the proportion of votes for R, as captured by the function $\overline{\kappa}$, coincides with the electoral rule ρ. In particular, the proportion

[9] The analysis is easily extended to the case where there is a continuum of signals.
[10] The interpretation is exact except when ω is one of the personal cutoffs.

of votes for R is higher than ρ for states above this intersection and lower than ρ for states below this intersection.

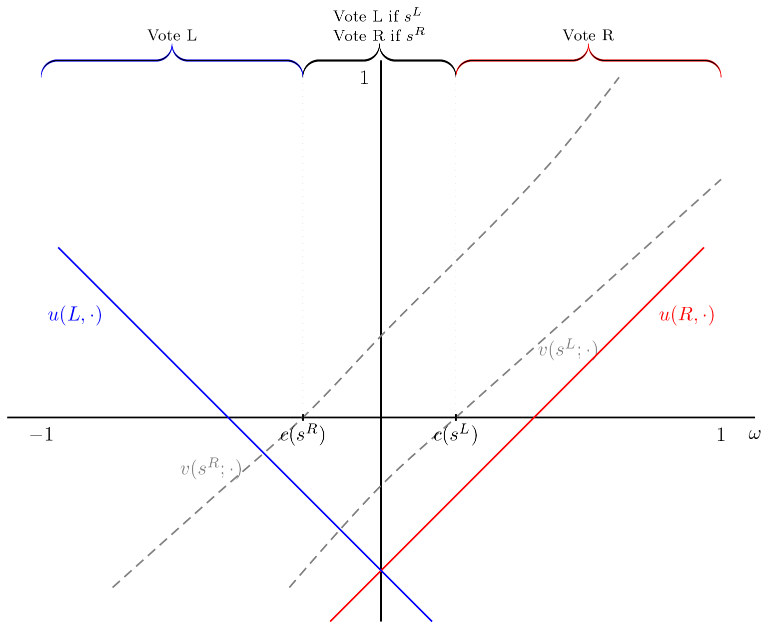

FIGURE 1. BELIEFS, PERSONAL CUTOFFS, AND VOTING BEHAVIOR.

Note: The figure shows the personal cutoffs $c(s^R)$ and $c(s^L)$ that result from equating the perceived benefits from electing R over L, represented by $v(s^R;\cdot)$ and $v(s^L;\cdot)$, to zero. The personal cutoffs determine voting behavior as a function of the equilibrium cutoff.

THEOREM 1: *For any electoral rule $\rho \in (0,1)$, there exists a unique retrospective voting equilibrium cutoff ω^* and it is given by $\omega^* = \bar{\kappa}^{-1}(\rho) \in [\underline{c}, \bar{c}]$.*[11]

PROOF:

See the Appendix.

The following examples illustrate how to find a retrospective voting equilibrium.

Example 1 (continued from pg. 5). An RVE can be found in four simple steps, depicted in Figure 2 for the case $\rho > 1/2$ and $\iota = 1/2$. First, we obtain the "belief function"

$$\begin{aligned} v\left(s;\omega\right) &= E\left(W \mid W > \omega, s\right) - E\left(-W \mid W < \omega, s\right) \\ &= \frac{\frac{1}{4}(1-\omega^2) + I_s \frac{\iota}{3}(1-\omega^3)}{\frac{1}{2}(1-\omega) + I_s \frac{\iota}{2}(1-\omega^2)} + \frac{\frac{1}{4}(\omega^2-1) + I_s \frac{\iota}{3}(\omega^3+1)}{\frac{1}{2}(\omega+1) + I_s \frac{\iota}{2}(\omega^2-1)}, \end{aligned}$$

[11]The inverse function $\bar{\kappa}^{-1} : (0,1) \to [\underline{c}, \bar{c}]$ is defined as $\bar{\kappa}^{-1}(\rho) = \inf\{\omega \in \Omega : \bar{\kappa}(\omega) \geq \rho\}$.

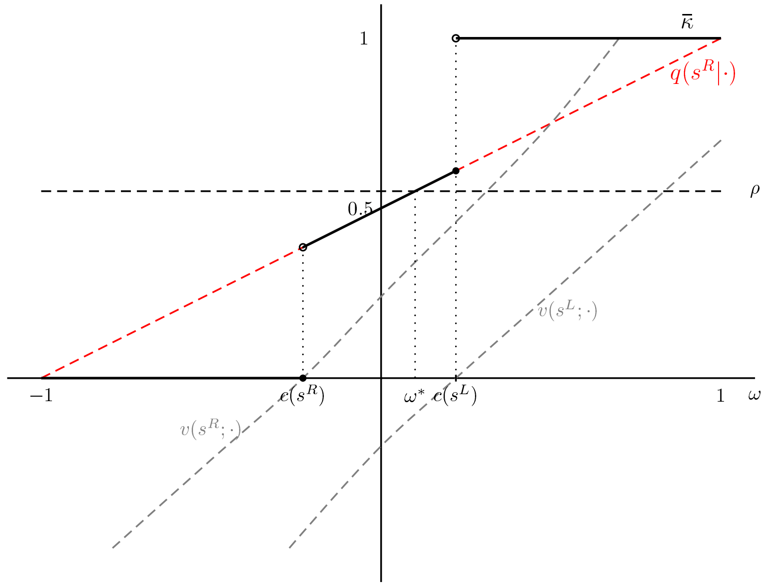

FIGURE 2. EXAMPLE 1. FINDING A RETROSPECTIVE VOTING EQUILIBRIUM.

Note: The figure shows how to use the personal cutoffs to construct the vote share function $\overline{\kappa}$, and how to then find the equilibrium cutoff ω^* by intersecting the vote share function with the threshold rule ρ.

where $I_{s^R} = -I_{s^L} = 1$. Second, we compute the personal cutoffs $c(s)$, which solve $v(s; c(s)) = 0$. Since $v(s^R; 0) > 0 > v(s^L; 0)$, then $c(s^R) < c^{FB} < c(s^L)$.

Third, we compute the vote share for R,

$$\overline{\kappa}(\omega) = \begin{cases} 0 & \text{if } \omega \leq c(s^R) \\ .5 + \iota\omega & \text{if } c(s^R) < \omega \leq c(s^L) \\ 1 & \text{if } \omega > c(s^L) \end{cases}$$

Finally, we intersect the vote share for R with the electoral rule. The equilibrium cutoff as a function of the electoral rule ρ is then

$$\omega^* = \begin{cases} c(s^R) & \text{if } \rho \leq .5 - (-\iota c(s^R)) \\ \frac{1}{\iota}(\rho - .5) & \text{if } .5 - (-\iota c(s^R)) < \rho < .5 + \iota c(s^L) \\ c(s^L) & \text{if } \rho \geq .5 + \iota c(s^L) \end{cases}$$

Thus, the first-best outcome can be obtained with our boundedly rational voters if and only if the electoral rule is $\rho = 1/2$. In contrast, a rule that requires a supermajority to elect one of the alternatives will inefficiently elect the other alternative too often in equilibrium (as shown in Figure 2 for $\rho > 1/2$, alternative R is elected in states higher than ω^* and L is elected in lower states, which implies that L is being inefficiently elected in the interval $(0, \omega^*)$). $\square$

Example 2. *(State-dependent vs. state-independent payoffs)* A representative voter with uncertain gross income $y(\omega)$ that is increasing in the state chooses between two policies. Under a *full stabilization policy* (Left), taxes $t(\omega) = y(\omega) - \bar{y}$ are set to smooth recessions and booms and to obtain a constant disposable income $\bar{y}$; hence, the policy leads to state-independent payoffs. Under a *budget balance policy* (Right), taxes $t(\omega) = e$ are set to balance a fixed amount of expenditure e; hence, the policy leads to state-dependent payoffs. Figure 3 depicts the disposable income from each policy, where $\bar{y}$ is normalized to zero and higher states are associated with higher gross income (the figure plots the case of three signals for concreteness). The first-best election cutoff is $c^{FB} = 0$. But, since $E(u(R, W) \mid W > 0, s) > E(u(L, W) \mid W < 0, s) = 0$ for any signal s, it follows that all personal cutoffs are negative and, therefore, the budget balance policy (Right) will be excessively elected in equilibrium relative to the first-best outcome. The intuition is that, since voters tend to elect the state-dependent policy (Right) in those states in which it is best, they will overestimate its value and will be biased towards voting for it. As shown by Figure 3, this bias can be partially mitigated by choosing a high enough threshold (above ρ^*) for electing the policy with state-dependent payoffs, thus providing a new normative rationale for requiring supermajorities to adopt alternatives with uncertain payoffs.[12] $\square$

D. Comparison to Nash and sincere voting

Beliefs in an RVE (i.e., expression (1) in Section I.B) can be compared to corresponding expressions for the cases of Nash equilibrium and sincere voting. In the case of Nash equilibrium (NE), the conditioning events in (1) are *not* the events that an alternative is elected, i.e., $\{W > \omega\}$ or $\{W < \omega\}$, but rather the event that a voter is pivotal, i.e.,

$$(4) \qquad E_n\left(u(R, W) \mid pivotal, S = s\right) - E_n\left(u(L, W) \mid pivotal, S = s\right),$$

where n is the number of voters (note that expression (1) already represents the limit as $n \to \infty$; see Section II for details). Feddersen and Pesendorfer (1997) show that, as the number of voters goes to infinity, the NE outcome in state ω is essentially the outcome that would arise if voters had perfect information about ω and, therefore, voted according to $u(R, \omega) - u(L, \omega)$. This result, which is far from being obvious, is known as *full information equivalence* and is true irrespective of the level of information precision.[13]

[12] See, e.g., Buchanan and Tullock (1967), Caplin and Nalebuff (1988), Dal Bo (2006), and Holden (2009) for alternative justifications of conservatism

[13] We restrict attention to the symmetric Nash equilibrium that is characterized by Feddersen and Pesendorfer (1997).

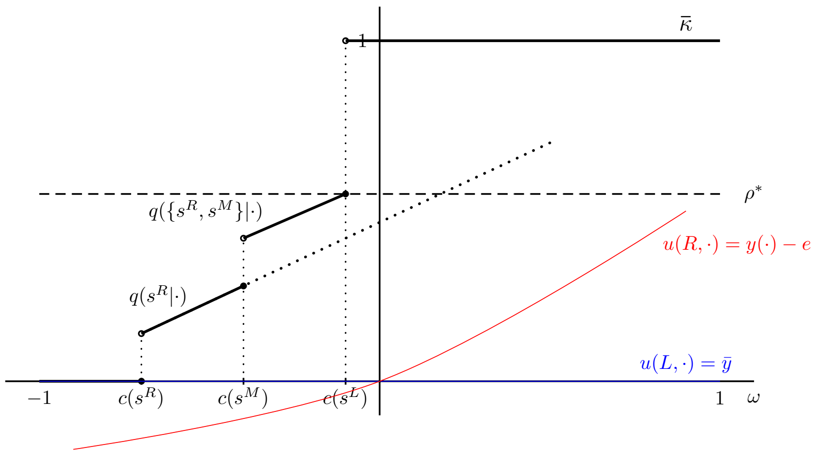

Figure 3. Example 2. RVE with state-dependent and state-independent payoffs.

Note: All personal cutoffs are negative and, therefore, the uncertain policy (Right) is excessively elected in equilibrium relative to the first best cutoff $c^{FB} = 0$. This welfare loss is mitigated by choosing any majority rule $\rho \geq \rho^*$, leading to an equilibrium cutoff $\omega^* = c(s^L)$.

In contrast, in the case of sincere voting (SV), voters choose the alternative with the highest expected payoff conditional on their private signal. Thus, a voter's belief is simply given by

$$E\left(u(R,W) \mid S = s\right) - E\left(u(L,W) \mid S = s\right), \tag{5}$$

irrespective of the number of voters. Hence, in large elections with sincere voting, the proportion of people voting for R in state ω is approximately given by the conditional probability of getting a signal s such that expression (5) is positive.

Example 1 illustrates how RVE differs from NE and SV, while still exhibiting some of the more attractive features of these solution concepts: Behavior is endogenous (as in NE, but unlike SV), but outcomes depend on both the electoral rule and the precision of information (as in SV, but unlike NE).[14]

Example 1 (continued from pgs. 5 and 8). Under SV, voters vote for R if they observe s^R and for L if they observe s^L. The election cutoff is given by the intersection of the electoral rule with the proportion of voters choosing R in each state, which is given by $q(s^R \mid \cdot)$. One can see from Figure 4 that SV is efficient (i.e., aggregates information) if and only if majority rule is used, $\rho = 1/2$. This result is an illustration of Condorcet's famous "jury theorem". Under NE,

[14]In Online Appendix A, we allow voters to be heterogeneous and show that individual voting behavior under RVE depends on private information for a significant fraction of the electorate (as in SV, but unlike NE).

by full information equivalence, the NE outcome is efficient and R is elected for $\omega > 0$ and L is elected for $\omega < 0$. The striking aspect of NE is that this is true for *any* (non-unanimous) election rule and for any precision of information $\iota > 0$, no matter how small. In contrast, changes in the election rule or information precision affect outcomes both under SV and RVE.

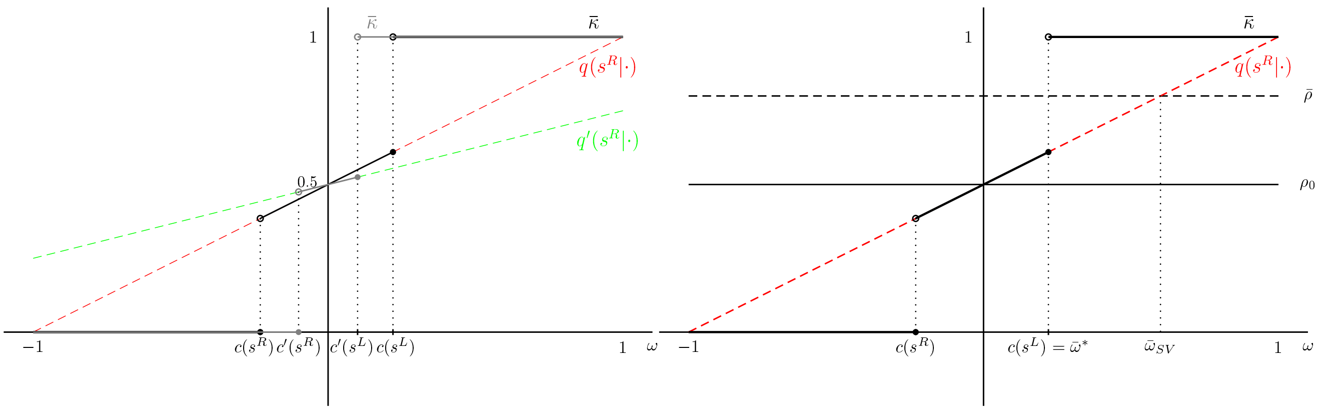

FIGURE 4. EXAMPLE 1. COMPARATIVE STATICS AND COMPARISON TO SV AND NE.

Note: The right panel shows that an increase in the electoral rule from ρ_0 to $\bar{\rho}$ leads to an increase in the election cutoff to $\bar{\omega}^*$ and $\bar{\omega}_{SV}$ under RVE and SV, respectively. The left panel shows that a decrease in information precision (i.e., from $q(s^R|\cdot)$ to $q'(s^R|\cdot)$) has two opposing effects under RVE: it leads to more extreme cutoffs along the flatter schedule $q'(s^R|\cdot)$ but the personal cutoffs also get closer to zero. The final effect on welfare is ambiguous. Under SV, only the first effect is present and lower information precision results in lower welfare.

Changes in information precision. Under SV, a decrease in information precision flattens $q(s^R \mid \cdot)$, thus leading to more extreme election cutoffs and lower welfare. The situation is more subtle under RVE. On the one hand, a flatter $q(s^R \mid \cdot)$ leads to a flatter $\bar{\kappa}(\cdot)$ over some range. This effect leads to more extreme equilibrium cutoffs. On the other hand, the personal cutoffs get closer to zero as information decreases (see $c'(s^R)$ and $c'(s^L)$ in Figure 4), therefore bringing the equilibrium cutoff closer to the first-best cutoff of zero. Thus, as can be seen from the left panel of Figure 4, information has an ambiguous welfare effect. This result makes sense because voters learn from a biased sample and have systematically biased beliefs, so there is no reason why better information should mitigate this bias.

Changes in election rules. Suppose that the electoral rule increases from ρ_0 to $\bar{\rho}$, as shown in the right panel of Figure 4. Under SV, behavior is exogenous and voters continue to vote in the same way, so that the election cutoff increases from zero to ω_{SV}. In contrast, the RVE cutoff increases from zero to $\bar{\omega}^* = c(s^L) < \bar{\omega}_{SV}$, thus mitigating welfare losses from naive voting. Intuitively, the out of equilibrium dynamics implied by the dynamic retrospective rule described in Section II is as follows. As the rule increases to $\bar{\rho}$, voters initially do not react to this change but the election outcome of course changes: R is now elected for states higher than $\bar{\omega}_{SV}$ and L is elected for all lower states. As a result, R's observed performance improves and L's observed performance worsens, so that voters start voting for R even if they get s^L signals. But this change in voting behavior implies that the election cutoff will decrease and that R will begin to be chosen in states lower than $\bar{\omega}_{SV}$. This change in cutoff in turn makes R less desirable, and this process stops at the new RVE cutoff $\bar{\omega}^* = c(s^L)$, where voters who receive signal s^R vote R and voters who receive signal s^L are indifferent and randomize. In particular, under RVE, voter behavior is disciplined by the performance of the parties, and dismal performances (e.g., an extreme election cutoff) produces changes in behavior that in turn affects the cutoff and mitigates welfare losses from changes in the primitives. □

Example 2 illustrates how biased behavior in the direction suggested by the selection problem arises endogenously under RVE while not arising under NE or arising only exogenously, for certain primitives, under SV.

Example 2 (continued from pg. 10). Under RVE, the logic of sample selection implies that the alternative with state-dependent payoffs, R, is chosen too often relative to the first-best solution. Under NE, in contrast, full information equivalence implies that the first-best cutoff of zero is attained: R is elected if $\omega > 0$ and L is elected if $\omega < 0$. Under SV, both excessive risk-taking and excessive conservatism are possible depending on the primitives. To illustrate this point, suppose that there are three signals, $s^L < s^M < s^R$, as illustrated in Figure 3. If expression (5) is positive for $s \in \{s^R, s^M\}$ and negative for s^L, then R is chosen too often under SV (i.e., the cutoff is at the intersection of $q(\{s^R, s^M\} \mid \omega)$ and ρ^*); if expression (5) is positive for s^R and negative for $s \in \{s^M, s^L\}$, then L is chosen too often under SV (i.e., the cutoff is at the intersection of $q(\{s^R\} \mid \omega)$ and ρ^*). □

An alternative solution concept is to postulate that voters have beliefs that are a convex combination of beliefs under NE and SV, i.e., equations (4) and (5). This alternative corresponds to what Eyster and Rabin (2005) call a partially cursed equilibrium. This approach, however, does not capture the key feature of our equilibrium concept, which is that voters naively fail to account for sample selection in elections in which *counterfactuals are not observed.* In Online Appendix C, we characterize partially cursed equilibrium as the number of voters goes to infinity. As expected, the insights about sample selection that emerge from RVE and not from either NE or SV, also don't emerge from a convex combination of NE and SV.

II. Equilibrium foundation

We first define equilibrium for a finite number of players, rather than for a continuum. Then, in Section II.B, we provide a learning foundation for equilibrium by showing that it corresponds to the steady state of a dynamic environment where players follow a particular retrospective voting rule. Finally, in Section II.C, we take the number of players to infinity and show that the resulting solution concept is characterized by the definition of RVE in Section I.B.

A. Voting game with finite number of players

Consider the game described in Section I but with a finite number of identical players indexed by $i = 1, ..., n$ and a threshold of k votes required to elect alternative R. Player i's payoff when the election outcome is $o \in \{R, L\}$ is now $u(o, \omega) + 1\{o = L\}\nu_i$, where $\nu_i \in \mathbb{V} \subseteq \mathbb{R}$ is a privately-observed payoff perturbation drawn independently for each player from a probability distribution F. Recall that K is the uniform bound on payoffs postulated in assumption A1. In addition to A1-A3, we assume:

A4. F is absolutely continuous and satisfies $F(-2K) > 0$ and $F(2K) < 1$; its density f satisfies $\inf_{x\in[-2K,2K]} f(x) > 0$.

A5. $\mathbb{S}$ has at least two elements and there exists $z > 0$ such that for all $\omega' > \omega$ and $s' > s$,

$$\frac{q(s'|\omega')}{q(s'|\omega)} - \frac{q(s|\omega')}{q(s|\omega)} \geq z(\omega' - \omega).$$

Assumption A4 guarantees that each alternative is voted with positive probability. It implies that the probability that players are pivotal (i.e., that their vote decides the election) becomes negligible as $n \to \infty$.[15] Assumption A5 is a strengthening of MLRP that establishes a bound on the rate at which the likelihood ratio changes.

An action plan for player i is a function that describes player i's signal-contingent vote, L or R, as a function of her realized payoff perturbation.[16] We restrict attention to weakly undominated strategies, so that, irrespective of her signal, voter i votes for R if her perturbation satisfies $\nu_i < -2K$ and for L if $v_i > 2K$. Following Harsanyi (1973), for each action plan we can integrate over the perturbations to obtain a (mixed) strategy, $\sigma : \mathbb{S} \to [0, 1]$, where $\sigma(s)$ is the probability of voting for R after observing signal s. We restrict attention to symmetric equilibria

[15] A4 also yields a refinement, which is standard in the literature, that rules out equilibria where everyone votes for the same alternative because a unilateral deviation cannot change the outcome. As shown in the proof of Theorem 2, the perturbations are also important for providing a learning foundation for equilibrium.

[16] The restriction to *pure* action plans is justified because F is absolutely continuous (Harsanyi, 1973).

and show in Theorem 2 that this restriction is without loss of generality—i.e., the steady states are symmetric—because the environment is symmetric.[17] In addition, a symmetric strategy profile $\boldsymbol{\sigma} = (\sigma, ..., \sigma)$, together with the primitives of the game, induces a distribution $P^n(\boldsymbol{\sigma})$ over the profile of votes, signals, and the state, $\mathbb{Z} = \{L, R\}^n \times \mathbb{S}^n \times \Omega$.[18]

DEFINITION 4: *A strategy profile $\boldsymbol{\sigma} = (\sigma, ..., \sigma)$ is a (symmetric)* naive equilibrium *of the voting game if for every $s \in \mathbb{S}$,*

$$\sigma(s) = F\left(\Delta(P^n(\boldsymbol{\sigma}), s)\right),$$

where $\Delta(P^n(\boldsymbol{\sigma}), s) \equiv E_{P^n(\boldsymbol{\sigma})}\left(u(R, W) \mid o = R, S = s\right) - E_{P^n(\boldsymbol{\sigma})}\left(u(L, W) \mid o = L, S = s\right)$.[19] *We refer to $P^n(\boldsymbol{\sigma}) \in \Delta(\mathbb{Z})$ as a* naive equilibrium distribution.

In equilibrium, each player best responds to a belief that depends endogenously on everyone's strategy and that is consistent with observed equilibrium outcomes. Players' beliefs are given by the observed equilibrium performance of each alternative conditional on their signal. Players, however, do not account for the correlation between others' votes and the state of the world (conditional on their own private information). Existence of equilibrium follows from a standard application of Brouwer's fixed point theorem; the proof is omitted.[20]

B. Retrospective learning foundation

We provide a learning foundation for Definition 4. A group of n players play the stage game described above for each discrete time period $t = 1, 2, ...$ At time t, the state is denoted by $\omega_t \in \Omega$, the signals by $\mathbf{s}_t = (s_{1t}, ..., s_{nt}) \in \mathbb{S}^n$, and the votes by $\mathbf{x}_t = (x_{1t}, ..., x_{nt}) \in \{L, R\}^n$. The outcome of the election at time t is denoted by $o_t \in \{L, R\}$. Player i's utility is $u(o_t, \omega_t) + 1\{o_t = L\}\, v_{it}$, where v_{it} is the payoff perturbation drawn independently (across players and time) from F.

Let $h^t = (\mathbf{z}_1, ..., \mathbf{z}_{t-1})$ denote the history of the game up to time $t-1$, where $\mathbf{z}_t = (\mathbf{x}_t, \mathbf{s}_t, \omega_t) \in \mathbb{Z}$ is the time-t outcome. Let $\mathcal{H}^t$ denote the set of all time-t histories and let $\mathcal{H}$ be the set of infinite histories.[21] We define a *retrospective voting rule* for player i, $\phi_i = (\phi_{i1}, ..., \phi_{it}, ...)$, where $\phi_{it} : \mathcal{H} \times \mathbb{V} \to \{L, R\}^{\mathbb{S}}$ for all

[17] The extension to asymmetric players is straightforward and provided in Online Appendix B.
[18] Definition 4 extends naive behavioral equilibrium (Esponda, 2008) to mixed strategies for the specific voting game that we consider.
[19] Whenever an expectation E_P has a subscript P, this means that the probabilities are taken with respect to the distribution P.
[20] The definition of naive equilibrium does not rely on the monotonicity assumptions on payoff functions and the information structure. Bhattacharya (2013) relaxes monotonicity restrictions for the case of Nash equilibrium.
[21] The payoff perturbations are not part of the history, implicitly assuming that players understand that the perturbations are independent payoff shocks that are unrelated to the learning problem.

t is defined as follows:

$$\phi_{it}(h, v_{it})(s) = \begin{cases} R & \text{if } \mu_{Rit}(h)(s) \geq \mu_{Lit}(h)(s) + v_{it} \\ L & \text{if } \mu_{Rit}(h)(s) < \mu_{Lit}(h)(s) + v_{it} \end{cases}$$

where, for each alternative $o \in \{L, R\}$, $\mu_{oit} : \mathcal{H} \to \mathbb{R}^{\mathbb{S}}$ represents player i's belief at time t about the expected utility of alternative o conditional on her signal s and is defined as follows: Let

$$\mathcal{Z}_{iso} = \left\{(\mathbf{x}', \mathbf{s}', \omega') \in \mathbb{Z} : o(\mathbf{x}') = o, s'_i = s\right\}$$

denote the event that player i observes $s \in \mathbb{S}$ and the elected outcome is $o \in \{L, R\}$. For any history $h \in \mathcal{H}$, let[22]

$$\mu_{oit}(h)(s) = \frac{\sum_{\tau=1}^{t-1} \mathbf{1}_{\mathcal{Z}_{iso}}(\mathbf{z}_\tau) u(o, \omega_\tau)}{\sum_{\tau=1}^{t-1} \mathbf{1}_{\mathcal{Z}_{iso}}(\mathbf{z}_\tau)} \tag{6}$$

for every $s \in \mathbb{S}$, and $t \geq 2$ whenever the denominator is greater than zero. If the denominator is zero, then $\mu_{oit}(h)(s) \in (-2K, 2K)$.[23] In words, players believe that the expected payoff of an alternative given a signal is given by the *observed* empirical average payoff. The idea is that players do not observe counterfactuals and take the information they see at face value without attempting to account for the informativeness of others' votes.

EXAMPLE 1: *To illustrate the retrospective voting rule, Table 1 shows data for eight elections from the point of view of one particular voter (perturbations are omitted for simplicity). For simplicity, we suppose that the voter knows that L always yields a payoff of zero, and so she only learns about R (i.e., $\mu_{Lit} = 0$). Suppose that there is an election in period 9 in which this voter observes signal r prior to voting. A* retrospective voter *behaves as follows. First, she uses past elections to form a belief about the performance of each alternative* conditional *on signal r. In this example, alternative L is known to always deliver a payoff of zero, while alternative R delivers an average payoff of $\mu_{Ri9}(h)(r) = (-1+1+1)/3 = 1/3$ when the signal is r (i.e., in periods 1, 2, and 6). Then, in period 9, if the voter observes signal r, she votes for R, i.e., $\phi_{i9}(h)(r) = R$.* $\square$

We now characterize the steady state outcomes under the assumption that players follow the above retrospective voting rule. Given a *retrospective voting rule*

[22] Throughout the paper, $\mathbf{1}$ stands for the indicator function, i.e., $\mathbf{1}_A(z) = 1$ if $z \in A$ and $\mathbf{1}_A(z) = 0$ if $z \notin A$.

[23] This assumption guarantees that posteriors always belong to $(-2K, 2K)$; hence, the perturbations guarantee that both alternatives are chosen with positive probability.

TABLE 1—ILLUSTRATION OF RETROSPECTIVE VOTING RULE

time	signal	vote	election outcome	observed payoff	
1	**r**	R	R	**-1**	
2	**r**	L	R	**1**	
3	**r**	R	L	0	→ counterfactual not observed
4	**l**	L	R	-1	
5	**l**	L	L	0	
6	**r**	R	R	**1**	
7	**r**	R	L	0	→ counterfactual not observed
8	**l**	L	R	1	

profile $\boldsymbol{\phi} = (\phi_1, ..., \phi_n)$, let $\mathbf{P}^{\phi}$ denote the unconditional probability distribution over histories, which we can construct by Kolmogorov's extension theorem.

For simplicity, we assume that the set of states Ω is finite, and so the set $\mathbb{Z}$ is also finite. For $t \geq 2$, define the sequence of random variables $\overline{P}_t : \mathcal{H} \to \Delta(\mathbb{Z})$, where

$$\overline{P}_t(h)(\mathbf{z}) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbf{1}_{\{\mathbf{z}\}}(\mathbf{z}_\tau)$$

is the frequency distribution over outcomes in the dynamic game.

We focus attention on frequency distributions that eventually stabilize around a steady-state distribution over outcomes. The following definition of stability accounts for the probabilistic nature and possible multiplicity of steady states.

DEFINITION 5: $P \in \Delta(\mathbb{Z})$ *is a* stable outcome distribution *of the dynamic game under a policy profile* $\boldsymbol{\phi}$ *if for all* $\varepsilon > 0$ *there exists* t_ε *such that*[24]

$$\mathbf{P}^{\phi}\left(\left\|\overline{P}_t(h) - P\right\| < \varepsilon \textit{ for all } t \geq t_\varepsilon\right) > 0.$$

This definition of stability captures the idea that after a *finite* number of periods, there is a positive probability that the frequency distribution over outcomes $\overline{P}_t$ remains forever close to P.

THEOREM 2: *If P is stable under a retrospective voting rule profile, then P is a naive equilibrium distribution of the stage game.*

[24]The norm $||\cdot||$ is defined as $||f|| = \max_{y \in Y} |f(y)|$.

PROOF:

See Online Appendix B.

Theorem 2 provides a justification for Definition 4: Any profile that is not a naive equilibrium generates an outcome distribution that is not stable. The proof of Theorem 2 adapts the arguments by Fudenberg and Kreps for games with complete information (1993, Proposition 7.5) to our asymmetric-information setting.

DISCUSSION. The mistake that our retrospective voters make is that they fail to account for the correlation between others' votes and the state of the world. Because the counterfactual payoff of the non-elected alternative is not observed, this correlation in turn implies that voters face a sample selection problem. In Table 1, the voter also observes signal r in periods 3 and 7, but, since L is elected, the voter does not observe the performance of R in those periods. If L and R were randomly chosen each period, the fact that the performance of R is not observed in periods 3 and 7 should not affect beliefs in the long run. The problem, however, is that the election outcome depends on private information that is correlated with performance. In particular, it is likely that the reason why R was not elected in periods 3 and 7 is that voters obtained signals that were relatively unfavorable to R. So, if our voter had been somehow able to observe the counterfactual performance of R in periods 3 and 7, she would have probably observed a relatively bad performance. Thus, the fact that counterfactual performances are not observed means that naive voters, who miss the correlation between others' votes and the state of the world, will end up overestimating the value of electing R in this example.

Two things are crucial here: retrospective voting and the fact that counterfactual outcomes are not observed. Suppose, instead, that counterfactual outcomes were observed (this is realistic in many contexts, though less so in the applications that we have in mind, as motivated in the introduction). Then steady-state behavior with retrospective voting would correspond to sincere voting (generalized to general game-theoretic settings by Eyster and Rabin's (2005) fully cursed equilibrium).

Finally, consider the case of sophisticated voters. Suppose that counterfactual outcomes are not observed (as in this paper) but voters understand the selection problem and follow a *sophisticated* voting rule that is just like the retrospective voting rule described above, except that voters only use data from elections in which they were pivotal. Then it is possible to show that steady-state behavior corresponds to Nash equilibrium. Of course, one may want to tell other stories for how voters reach Nash.

C. Large elections

We now consider the naive equilibrium in Definition 4 (which is the steady-state of the retrospective rule described above, by Theorem 2) and characterize its limit as the number of voters goes to infinity. The resulting characterization is what we called RVE in Definition 3.[25]

We explain intuitively why equation (1) characterizes limiting beliefs with a large number of players; see Online Appendix B.B2 for the formal results. Due to the retrospective ruled described above, player i's equilibrium belief is given by

$$E_{P^n(\boldsymbol{\sigma})}\left(u(R,W) \mid o=R, S=s\right) - E_{P^n(\boldsymbol{\sigma})}\left(u(L,W) \mid o=L, S=s\right).$$

As the number of voters goes to infinity, the events $o=R$ (i.e., R is elected) and $o=L$ (i.e., L is elected) are equivalent to $\{W>\omega\}$ and $\{W<\omega\}$, respectively, where ω is the equilibrium cutoff, thus explaining the origin of equation (1). The intuition is as follows. Suppose that voter behavior converges to strategy σ. Then the probability that a randomly chosen player votes for Right, conditional on ω, converges to $\kappa(\omega;\sigma)$. By standard asymptotic arguments, the proportion of votes for Right becomes concentrated around $\kappa(\omega;\sigma)$. So, for states where $\kappa(\omega;\sigma)>\rho$, the probability that the outcome is Right converges to 1. Similarly, for states where $\kappa(\omega;\sigma)<\rho$, the probability that the outcome is Right converges to 0. Finally, the key is to show that σ is increasing, which then implies that there is at most one (measure zero) state such that $\kappa(\omega;\sigma)=\rho$; thus, the outcome is characterized by a cutoff.

The proof that σ is increasing is standard for Nash equilibrium, where it relies on the fact that, by MLRP, higher signals convey more favorable information about Right. In our context, higher signals also have a second, indirect effect, because, to the extent that a player can be pivotal and affect the outcome of the election, her beliefs about the alternatives also depend on her own strategy. In fact, this indirect effect may go in the opposite direction of the standard effect.[26] However, we establish that the probability of being pivotal goes to zero as the number of players increases and, therefore, the indirect effect eventually vanishes and becomes dominated by the direct effect (provided a uniform version of MLRP holds). Thus, equilibrium strategies are increasing in the limit as the number of players goes to infinity.

[25] As a comparison, a similar exercise was carried out by Feddersen and Pesendorfer (1997) for the case of Nash equilibrium.

[26] To see this claim, fix a player and a signal and suppose that she votes for Right with probability close to 1. Then, most often, Right is the outcome of the election whenever at least $k-1$ or more of the other players have voted for Right. Compare this case to the case where she votes for Left with probability close to 1. Then, most often, Right is the outcome of the election whenever at least k or more of the other players have voted for Right. If strategies are increasing, then, by MLRP, the first event conveys less favorable information about Right. Therefore, a higher signal leads this player to vote more for Right, which then makes her less favorable about Right.

REFERENCES

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

Barberis, N., A. Shleifer, and R. Vishny, "A model of investor sentiment," *Journal of financial economics*, 1998, *49* (3), 307–343.

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

Bendor, J., D. Diermeier, Siegel D.A., and M.M. Ting, *A behavioral theory of elections*, Princeton Univ Press, 2011.

_ **, S. Kumar, and D.A. Siegel**, "Adaptively Rational Retrospective Voting," *Journal of Theoretical Politics*, 2010, *22* (1), 26.

Bhattacharya, S., "Preference monotonicity and information aggregation in elections," *Econometrica*, 2013, *81* (3), 1229–1247.

Buchanan, J.M. and G. Tullock, *The calculus of consent: Logical foundations of constitutional democracy*, Vol. 100, University of Michigan Press, 1967.

Callander, S., "Searching for good policies," *American Political Science Review*, 2011, *105* (4), 643–662.

Caplin, A. and B. Nalebuff, "On 64%-majority rule," *Econometrica*, 1988, *56* (4), 787–814.

Charness, G. and D. Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–236.

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

Dal Bo, E., "Committees with supermajority voting yield commitment with flexibility," *Journal of Public Economics*, 2006, *14*, 573–599.

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., "Behavioral equilibrium in economies with adverse selection," *The American Economic Review*, 2008, *98* (4), 1269–1291.

_ **and D. Pouzo**, "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Forthcoming Econometrica*, 2015.

_ **and E. Vespa**, "Hypothetical Thinking and Information Extraction in the Laboratory," *American Economic Journal: Microeconomics*, 2014, *6* (4), 180–202.

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

Eyster, E. and M. Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

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. Kreps, "Learning Mixed Equilibria," *Games and Economic Behavior*, 1993, *5*, 320–367.

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

_ **and** _ , "Learning and Equilibrium," *Annual Review of Economics*, 2009, *1*, 385–420.

Gabaix, X., "A Sparsity-Based Model of Bounded Rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

Harsanyi, J.C., "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International Journal of Game Theory*, 1973, *2* (1), 1–23.

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.

Holden, R., "Supermajority voting rules," 2009.

Jehiel, P., "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

_ **and D. Samet**, "Valuation equilibrium," *Theoretical Economics*, 2007, *2* (2), 163–185.

_ **and F. Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Kagel, J.H. and D. Levin, *Common value auctions and the winner's curse*, Princeton Univ Press, 2002.

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.

Milgrom, P.R., "A convergence theorem for competitive bidding with differential information," *Econometrica*, 1979, pp. 679–688.

_ , "Good news and bad news: Representation theorems and applications," *The Bell Journal of Economics*, 1981, pp. 380–391.

Mullainathan, S., J. Schwartzstein, and A. Shleifer, "Coarse thinking and persuasion," *Quarterly Journal of Economics*, 2008, *123* (2), 577–619.

Osborne, M.J. and A. Rubinstein, "Games with procedurally rational players," *American Economic Review*, 1998, *88*, 834–849.

_ **and** _ , "Sampling equilibrium, with an application to strategic voting," *Games and Economic Behavior*, 2003, *45*, 434–441.

Rabin, M., "Inference by believers in the law of small numbers," *Quarterly journal of Economics*, 2000, *117* (3), 775–816.

_ **and J. L. Schrag**, "First impressions matter: A model of confirmatory bias," *Quarterly journal of Economics*, 1999, *114* (1), 37–82.

Rubinstein, A., "On price recognition and computational complexity in a monopolistic model," *Journal of Political Economy*, 1993, pp. 473–484.

Schwartzstein, J., "Selective Attention and Learning," *Journal of the European Economic Association*, 2014, *12* (6), 1423–1452.

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

_ , "Bayesian networks and boundedly rational expectations," 2014.

Thaler, R.H., "Anomalies: The winner's curse," *The Journal of Economic Perspectives*, 1988, *2* (1), 191–202.

Wolfers, J., "Are Voters Rational? Evidence from Gubernatorial Elections," 2009.

Appendix

The following result is used in the proof of Theorem 1.

LEMMA 1: $\overline{\kappa} : \Omega \to [0,1]$ is left-continuous, increasing over the subdomain $(\underline{c}, \overline{c})$, and satisfies: $\overline{\kappa}(\omega) = 0$ if $\omega \leq \underline{c}$ and $\overline{\kappa}(\omega) = 1$ if $\omega > \overline{c}$.

PROOF:

$\bar{\kappa}(\cdot)$ is left-continuous: Since there are a finite number of personal cutoffs (defined by equation (2)), then for each $c \in (-1,1)$ there exists $\omega' < c$ such that all personal cutoffs are outside the interval $[\omega', c)$. Then, for all $\hat{\omega}$, $\sum_{\{s:c(s)<\omega\}} q\left(s \mid \hat{\omega}\right) = \sum_{\{s:c(s)<c\}} q\left(s \mid \hat{\omega}\right)$ for all $\omega \in [\omega', c]$. In addition, $q(s \mid \cdot)$ is continuous by A3(iii). Therefore, $\lim_{\omega \uparrow c} \sum_{\{s:c(s)<\omega\}} q\left(s \mid \omega\right) = \sum_{\{s:c(s)<c\}} q\left(s \mid c\right)$.
$\bar{\kappa}(\cdot)$ is increasing over $(\underline{c}, \overline{c})$: Let $\underline{c} < \omega < \omega' < \overline{c}$. Then

$$\begin{aligned} \sum_{\{s:c(s)<\omega'\}} q\left(s \mid \omega'\right) &\geq \sum_{\{s:c(s)<\omega\}} q\left(s \mid \omega'\right) \\ &> \sum_{\{s:c(s)<\omega\}} q\left(s \mid \omega\right), \end{aligned} \tag{A1}$$

where the last inequality follows because, since $c(\cdot)$ is nonincreasing, the event $\{c(s) < \omega\}$ is equivalent to $\{s \geq s(\omega)\}$ for some threshold $s(\omega)$, and there is strict MLRP (A2); see (Milgrom, 1981).

Finally: If $\omega \leq \underline{c}$, then $\{c(s) < \omega\} = \emptyset$, so that $\bar{\kappa}(\omega) = 0$. Similarly, if $\omega > \overline{c}$, then $\{c(s) < \omega\} = \mathbb{S}$, so that $\bar{\kappa}(\omega) = 1$.

Proof of Theorem 1. The proof relies on the following claim.

Claim 1.1 Suppose that σ is optimal given election cutoff ω^*. Then

$$\kappa(\omega;\sigma) = \sum_{\{s:c(s)<\omega^*\}} q(s \mid \omega) + \sum_{\{s:c(s)=\omega^*\}} q(s \mid \omega)\sigma(s) \tag{A2}$$

for all $\omega \in \Omega$. In addition, $\bar{\kappa}(\omega) \geq \kappa(\omega;\sigma)$ for $\omega > \omega^*$ and $\bar{\kappa}(\omega) \leq \kappa(\omega;\sigma)$ for $\omega < \omega^*$.

PROOF:

Since σ is optimal given ω^*, then

$$\sigma(s) = \begin{cases} 0 & \text{if } c(s) > \omega^* \\ 1 & \text{if } c(s) < \omega^* \end{cases} \tag{A3}$$

and equation (A2) follows. By equation (A2), for all $\omega > \omega^*$,

$$\kappa(\omega;\sigma) \leq \sum_{\{s:c(s)\leq\omega^*\}} q(s \mid \omega)$$
$$\leq \sum_{\{s:c(s)<\omega\}} q(s \mid \omega) = \bar{\kappa}(\omega).$$

Similarly, for all $\omega < \omega^*$, $\kappa(\omega;\sigma) \geq \bar{\kappa}(\omega)$.

We now prove Theorem 1. Fix $\rho \in (0,1)$ and let

$$\omega^* = \kappa^{-1}(\rho) = \inf\{\omega \in \Omega : \bar{\kappa}(\omega) \geq \rho\}. \tag{A4}$$

Note that, by Lemma 1, $\omega^* \in [\underline{c}, \overline{c}]$. We begin by showing that there exists σ^* such that (σ^*, ω^*) is a voting equilibrium. Let σ^* satisfy (A3). It remains to specify $\sigma^*(s)$ for s such that $c(s) = \omega^*$. First, suppose that $\omega^* \notin \{-1, 1\}$. If ω^* is the election cutoff, then s such that $c(s) = \omega^*$ is indifferent between R and L, and, therefore, $\sigma^*(s) = \alpha$ is optimal for any $\alpha \in [0,1]$. Let σ^*_α denote the strategy profile that satisfies (A3) and $\sigma^*(s) = \alpha$. We now pick α such that ω^* is an election cutoff given σ^*_α. By Claim 1.1,

$$\kappa(\omega^*;\sigma^*_\alpha) = \left(\sum_{\{s:c(s)<\omega^*\}} q(s \mid \omega^*) + \sum_{\{s:c(s)=\omega^*\}} q(s \mid \omega^*)\alpha \right),$$

which is continuous in α. First, we establish that $\kappa(\omega^*;\sigma^*_0) \leq \rho$. Suppose not, so that $\kappa(\omega^*;\sigma^*_0) = \bar{\kappa}(\omega^*) > \rho$. Since $\bar{\kappa}$ is left-continuous (Lemma 1), then there exists $\omega' < \omega^*$ such that $\bar{\kappa}(\omega') > \rho$. But then (A4) is contradicted. Second, we establish that $\kappa(\omega^*;\sigma^*_1) \geq \rho$. Suppose not, so that $\kappa(\omega^*;\sigma^*_1) = \lim_{\omega\downarrow\omega^*} \bar{\kappa}(\omega) < \rho$. Then, there exists $\omega'' > \omega^*$ such that $\bar{\kappa}(\omega'') < \rho$. But, since $\bar{\kappa}(\cdot)$ is nondecreasing (Lemma 1), then (A4) is contradicted. Since $\kappa(\omega^*;\sigma^*_0) \leq \rho$ and $\kappa(\omega^*;\sigma^*_1) \geq \rho$, by continuity of $\alpha \mapsto \kappa(\omega^*;\sigma^*_\alpha)$ there exists α^* such that $\kappa(\omega^*;\sigma^*_{\alpha^*}) = \rho$. Since $\kappa(\cdot;\sigma^*_{\alpha^*})$ is nondecreasing (because $\sigma^*_{\alpha^*}$ is nondecreasing), then ω^* is an election cutoff given $\sigma^*_{\alpha^*}$. Hence, $(\sigma^*_{\alpha^*}, \omega^*)$ is a voting equilibrium. Next, suppose that $\omega^* = -1$ (the case $\omega^* = 1$ is similar and, therefore, omitted). Now let $\alpha^* = 1$; in particular, σ^*_1 is optimal given ω^* (note it would not necessarily be optimal for $\alpha^* \neq 1$). Note that the argument provided above to show $\kappa(\omega^*;\sigma^*_1) \geq \rho$ for all $\omega^* \notin \{-1, 1\}$ also holds for $\omega^* = -1$. Thus, $\kappa(-1;\sigma^*_1) \geq \rho$. Since $\kappa(\cdot;\sigma^*_1)$ is nondecreasing, it follows that $\kappa(\omega;\sigma^*_1) \geq \rho$ for all $\omega \in \Omega$, implying that $\omega^* = -1$ is a cutoff given σ^*_1.

Finally, we show that, for all $\omega \neq \omega^*$, there exists no σ such that (σ, ω) is a voting equilibrium. Suppose, in order to obtain a contradiction, that (σ, ω) is a voting equilibrium, where $\omega < \omega^*$ (the case $\omega > \omega^*$ is similar and, therefore, omitted). Let $\omega' \in (\omega, \omega^*)$. Then $\bar{\kappa}(\omega') \geq \kappa(\omega'; \sigma) \geq \rho$, where the first inequality follows from Claim 1.1 and the fact that σ is optimal given ω, and the second from the fact that ω is an election cutoff given σ. But then (A4) is contradicted. □

Online Appendix

Online appendix "Conditional Retrospective Voting in Large Elections, " by Ignacio Esponda and Demian Pouzo.

RVE with heterogeneous voters

We show that it is straightforward to define an RVE with heterogenous voters. We model heterogeneity (in preferences and information) by assuming that each voter is of a particular type θ, where φ is the full-support probability distribution over the set of types $\Theta \subset \mathbb{R}$. Conditional on a state $W = \omega$, players of type θ independently draw a signal $S_\theta = s$ from a finite, nonempty set $\mathbb{S}_\theta \subset \mathbb{R}$ with probability $q_\theta(s \mid \omega)$; let s_θ^L and s_θ^R denote the lowest and highest signals in $\mathbb{S}_\theta$. The payoff of type θ is given by $u_\theta(o, \omega)$, where $o \in \{L, R\}$ is the winner of the election.

Let $\sigma_\theta : \mathbb{S}_\theta \to [0, 1]$ denote the strategy of type θ, where $\sigma_\theta(s)$ is the probability that type θ votes for alternative R after observing signal s. A strategy σ_θ is *nondecreasing* if $\sigma_\theta(s') \geq \sigma_\theta(s)$ for all $s' > s$. A strategy profile $\sigma = (\sigma_\theta)_{\theta\in\Theta}$ is nondecreasing if σ_θ is nondecreasing for each θ.

Assumptions A1, A2, and A3 are now required for all $\theta \in \Theta$. The following additional assumption guarantees uniqueness of the equilibrium cutoff and is made only for convenience.

A6. $\Theta \subset \mathbb{R}$ is a compact interval and $\mathbb{S}_\theta = \mathbb{S}$; the functions $(\theta, \omega) \mapsto u_\theta(R, \omega)$, $(\theta, \omega) \mapsto u_\theta(L, \omega)$, and $(\theta, \omega) \mapsto q_\theta(s \mid \omega)$ are jointly continuous in $\Theta \times \Omega$ for all $s \in \mathbb{S}$.

Let $\kappa(\omega; \sigma) \equiv \int_\Theta \sum_{s\in\mathbb{S}} q_\theta(s \mid \omega)\sigma_\theta(s)\varphi(d\theta)$ and

$$v_\theta(s; \omega) \equiv E(u_\theta(R, W) \mid W > \omega, S_\theta = s) - E(u_\theta(L, W) \mid W < \omega, S_\theta = s)$$

be the natural counterparts of the expressions defined in the text for the case of homogenous voters. Optimality is now required to hold for all $\theta \in \Theta$, and the definition of RVE is the same as in the text (Definition 3).

In order to characterize RVE, let $c_\theta(s) \equiv \arg\min_{\omega\in\Omega} |v_\theta(s; \omega)|$ denote the personal cutoff of type θ with signal s, and let $\underline{c} \equiv \min_\theta c_\theta(s^R)$ and $\bar{c} = \max_\theta c_\theta(s^L)$ be the lowest and highest personal cutoffs across all types (which exist by A1-A3 and A6). Let $\bar{\kappa}(\cdot)$ in equation (3) now be defined by

$$\bar{\kappa}(\omega) \equiv \int_\Theta \sum_{\{s:c_\theta(s)<\omega\}} q_\theta(s \mid \omega)\varphi(d\theta)$$

for all $\omega \in \Omega$. The characterization of RVE (i.e., Theorem 1) then holds exactly as stated. The proof is almost identical, with the obvious difference that we need to integrate over Θ in several places. The only claim that requires a new proof is the claim that the inequality in equation (A1) holds strictly. This is trivially true if there exists a positive φ-measure of types with personal cutoffs in $[\omega, \omega')$, so suppose that is not the case. Since, by A6, $c_\theta(s^L)$ is continuous in θ and Θ is a compact interval, the union of $c_\theta(s^L)$ over all $\theta \in \Theta$ is a compact interval. Given that there is no positive measure of types with personal cutoffs in $[\omega, \omega')$, then, the facts that φ has full support and $\omega' < \bar{c}$ implies that, for all $\theta \in \Theta$, $c_\theta(s^L) \geq \omega' > \omega$ and, therefore, $\{c_\theta(s) < \omega\} \neq \mathbb{S}$. Then, because MLRP holds strictly (by A2), the second inequality in (A1) is strict.

We provide an example to illustrate how easy it is to incorporate heterogeneity into the model (and to once again compare RVE to NE and SV).

Example 3. (*Heterogeneous voters*) Consider an environment identical to Example 1, with the exception that preferences are heterogeneous: Payoffs of type θ are $u_\theta(R, \omega) = \omega - 1/2 + \theta$ and $u_\theta(L, \omega) =$

$-\omega - 1/2$ for all $\omega \in \Omega$, and types are distributed uniformly on the interval $[-1, 1]$. In particular, higher types get higher payoffs under R. For concreteness, suppose majority rule, $\rho = .5$, and information precision $\iota = .5$.

Figure A1 illustrates equilibrium with heterogeneous voters for all solution concepts. Under RVE, the vote share function is given by

$$\overline{\kappa}(\omega) = q(s^R \mid \omega) \Pr\left(c_\theta(s^R) < \omega\right) + q(s^L \mid \omega) \Pr\left(c_\theta(s^L) < \omega\right),$$

where $c_\theta(s)$ is the personal cutoff of type θ for signal s. The equilibrium cutoff is $\omega^* = 0$ and all types lower than $-.22$ always vote L, all types higher than $.22$ always vote R, but all types between $-.22$ and $.22$ vote their signal. In particular, there is a significant fraction of "independent types" that vote according to their signal.

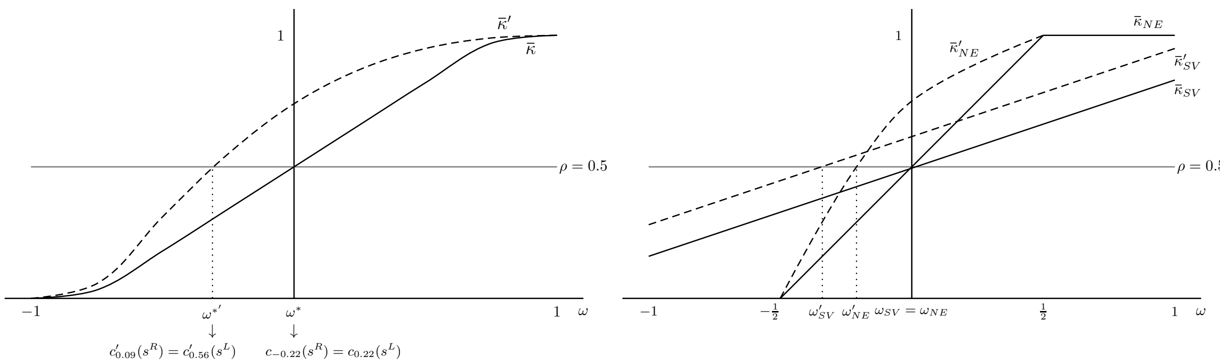

FIGURE A1. EXAMPLE 3: HETEROGENEOUS VOTERS.

Note: The vote share functions are smoothed out with heterogenous voters. The figure shows the effect of an increase (in first order stochastic sense) in the distribution of voters who prefer R (left panel shows RVE and right panel shows SV and NE).

Under SV, a voter votes R whenever she either observes s^R and has type $\theta > -2/3$ or she observes s^L and has type $\theta > 2/3$. Thus, the proportion of R-votes in state ω is

$$\bar{\kappa}^{SV}(\omega) = q(s^R \mid \omega) \Pr(\theta > -2/3) + q(s^L \mid \omega) \Pr(\theta > 2/3) = .5 + \omega/3$$

and the election cutoff is $\omega_{SV} = 0$. Types below $-2/3$ always vote L, types above $2/3$ always vote R, and types in between $-2/3$ and $2/3$ vote their signal.

Under NE, it suffices to characterize the outcome under full information. Since type θ prefers R whenever $u_\theta(R, \omega) > u_\theta(L, \omega)$, or, equivalently, $\omega > -.5\theta$, then the proportion of votes for R is $\bar{\kappa}_{NE}(\omega) = .5 + \omega$ for $\omega \in [-.5, .5]$. Thus, the NE election cutoff is also $\omega_{NE} = 0$. Moreover, Feddersen and Pesendorfer (1997) show that NE voting behavior is as follows. For a given $\varepsilon > 0$, there is a sufficiently large number of voters such that, for all larger elections, voters that have a type in an ε-neighborhood of $\theta = 0$ vote their signal, but everyone else votes always R or always L.

Shift in the distribution of preferences. Suppose that there is an increase (in the first order stochastic dominance sense) in the distribution of voters who prefer R. For concreteness, let the new probability density function of types be $\varphi'(\theta) = .5(1 + \theta)$ for $\theta \in [-1, 1]$. Figure A1 shows the new functions $\bar{\kappa}'_{SV}$, $\bar{\kappa}'_{NE}$, and $\bar{\kappa}'$, under SV, NE, and RVE. As expected, in all cases there is an upward shift in the proportion of votes for R, and the cutoff moves to the left, so that R is chosen in more states of the world. But the effect on voting behavior is different in each case.

Under SV, voting behavior is exogenous and so every type behaves exactly as before. The increase in the proportion of votes for R is driven by the fact that there are more high types and these types vote for R. In contrast, behavior is endogenously affected by the change in equilibrium cutoff under both NE and RVE. Under NE, the new "marginal type" is $\theta = .42$: lower types always vote L and higher types always vote R (while a vanishing fraction of types around .42 vote their signal). Under RVE, types lower than .09 always vote L, types higher than .56 always vote R, and types in between .09 and .56 vote their signal. In particular, more popular alternatives are supported by voters with more extreme preferences under NE and RVE, which helps mitigate the preference shift in favor of the more popular alternative.[27]
□

Foundation for RVE with heterogeneous players

In Section B.B1, we provide the proof of the learning foundation for a finite number of players (Theorem 2 in the text), but allow players to be heterogenous. In Section B.B2, we consider the case where the number of players goes to infinity, therefore providing a game-theoretic foundation for RVE.

We now allow players to be asymmetric and assume for simplicity that the set of types Θ (see Online Appendix A) and the set of states Ω are both finite. Recall that player i is of type $\theta_i \in \Theta$. The general definition of naive equilibrium with a finite number of players is the following.

DEFINITION 6: *A strategy profile* $\boldsymbol{\sigma} = (\sigma_1, ..., \sigma_n)$ *is a* naive equilibrium *of the voting game if for every player* $i = 1, ..., n$ *and for every* $s \in \mathbb{S}_{\theta_i}$,

$$\sigma_i(s) = F_{\theta_i}\left(\Delta_i(P^n(\boldsymbol{\sigma}), s)\right),$$

where $\Delta_i(P^n(\boldsymbol{\sigma}), s) \equiv E_{P^n(\boldsymbol{\sigma})}\left(u_{\theta_i}(R, W) \mid o = R, S_i = s\right) - E_{P^n(\boldsymbol{\sigma})}\left(u_{\theta_i}(L, W) \mid o = L, S_i = s\right)$. *We refer to* $P^n(\boldsymbol{\sigma})$ *as a* naive equilibrium distribution.

B1. Proof of Theorem 2

Throughout the proof, we fix a stable outcome distribution P and retrospective voting rule profile $\boldsymbol{\phi}$. The proof compares "strategies" in the dynamic game with strategies in the stage game. To define the former, let

$$\mathcal{A}_i = \left\{\alpha_i \in \mathbb{R}^{\mathbb{S}} : F_{\theta_i}(-2K) \leq \alpha_i(s) \leq F_{\theta_i}(2K)\ \forall s \in \mathbb{S}_{\theta_i}\right\}$$

denote a player's strategy space and define the vector-valued random variable $\boldsymbol{\alpha}_t = (\alpha_{1t}, ..., \alpha_{nt}) : \mathcal{H} \to \prod_{i=1}^{n} \mathcal{A}_i$ denote a time-t strategy profile, where

$$\alpha_{it}(h)(s) = \int \mathbf{1}_{\{v_i : \phi_{it}(h, v_i)(s) = R\}}(\nu_i) dF_{\theta_i} \tag{B1}$$

is the probability that player i votes for R when observing signal s, conditional on history h^t.

[27]In the case of NE, the reason for more extreme supporters of R is that the pivotal voter believes that the state is given by the cutoff state, since this is where the proportions voting for R and L are equal. When preferences shift and the cutoff decreases, then the pivotal voter believes the state is lower. In order to be indifferent between voting for L or R, then its type must be higher. In the case of RVE, the reason for more extreme supporters of R is that, if R is elected more often, then it must be elected in worse states and its observed performance must be lower. Therefore, types that were marginally willing to vote for R will no longer desire to vote for R.

Finally, let $\boldsymbol{\sigma} = (\sigma_1, ..., \sigma_n) \in \Pi_{i=1}^n \mathcal{A}_i$ be such that

$$\sigma_i(s) = F_{\theta_i}(\Delta_i(P, s)) \tag{B2}$$

is the probability that player i votes for R if she optimally responds to beliefs $\Delta_i(P, s)$.

The proof of Theorem 2 follows from the following claims, which are proven at the end of this section.

Claim 2.1 For all $\varepsilon > 0$, there exists H_ε with $\mathbf{P}^\phi(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$, there exists $t_{\varepsilon,h}$ such for all $t \geq t_{\varepsilon,h}$, $\|\boldsymbol{\alpha}_t(h) - \boldsymbol{\sigma}\| < \varepsilon$ and $\|\overline{P}_t(h) - P\| < \varepsilon$.

Claim 2.2 $||P - P^n(\boldsymbol{\sigma})|| = 0$

Claim 2.1 establishes that stability of P implies that beliefs for each player i, conditional on $s \in \mathbb{S}_{\theta_i}$, eventually remain close to $\Delta_i(P, s)$, thus implying that the time-t strategy profile $\boldsymbol{\alpha}_t$ eventually remains close to $\boldsymbol{\sigma}$. The key of the proof is that players' payoff perturbations are independently drawn from an atom-less distribution, implying that if beliefs settle down, then strategies must also settle down, not just in an average sense, but actually in a per-period sense. In particular, Claim 2.1 implies that any correlation in players' strategies induced by a common history eventually vanishes. In Claim 2.2, we show that the fact that time-t strategies remain close to $\boldsymbol{\sigma}$ implies that $P = P^n(\boldsymbol{\sigma})$. Recall that $P^n(\boldsymbol{\sigma})$ is the distribution over outcomes $\mathbb{Z} \equiv \{L, R\}^n \times \prod_{i=1}^n \mathbb{S}_{\theta_i} \times \Omega$, i.e., $P^n(\boldsymbol{\sigma})(\mathbf{x}, \mathbf{s}, \omega) = g(\omega) \prod_{i=1}^n \sigma_i(s_i)^{\mathbf{1}_{\{R\}}(x_i)} (1 - \sigma_i(s_i))^{\mathbf{1}_{\{L\}}(x_i)} q_{\theta_i}(s_i \mid \omega)$. Both claims rely on a straightforward generalization of a technical result by Fudenberg and Kreps (1993, Lemma 6.2); this result allows us to apply the law of large numbers in a context where a sequence of random variables is not independently distributed, but where the distributions conditional on past history are eventually very close to some common distribution.

Claim 2.2 and equation (B2) imply that, for all i and $s \in \mathbb{S}_{\theta_i}$,

$$\sigma_i(s) = F_{\theta_i}(\Delta_i(P^n(\boldsymbol{\sigma}), s)),$$

so that $\boldsymbol{\sigma}$ is a naive equilibrium of the stage game. Therefore, $P = P^n(\boldsymbol{\sigma})$ is a naive equilibrium distribution. $\square$

The proof of Claims 2.1 and 2.2 used above rely on the following two lemmas. The statement and proof of Lemma 2 are straightforward adaptations of a result by Fudenberg and Kreps (1993).

LEMMA 2: *(cf. Fudenberg and Kreps, 1993, Lemma 6.2) Let $(z_t)_t$ be a sequence of random variables with range on a finite set Z. Fix a set-function $\pi : 2^Z \to [0, 1]$ (not necessarily a probability measure) and fix $\varepsilon \in \mathbb{R}$. Let H_ε be a subset of infinite histories such that for all $h \in H_\varepsilon$ there exists $t_{\varepsilon,h}$ such that for all $t \geq t_{\varepsilon,h}$, the distribution of each z_t conditional on $h^t = (z_1, ..., z_{t-1})$, denoted $\pi_t(\cdot \mid h^t)$, satisfies*

$$\max_{Z' \in \mathcal{Z}} \pi_t(Z') - \pi(Z') > -\varepsilon, \tag{B3}$$

where $\mathcal{Z} \subseteq 2^Z$ is a set of subsets of Z.[28]

Then

$$\liminf_{t \to \infty} \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}_{Z'}(z_\tau) \geq \pi(Z') - \varepsilon \tag{B4}$$

for all $Z' \in \mathcal{Z}$, almost surely on H_ε. Moreover, if (B3) is replaced by $\max_{Z' \in \mathcal{Z}} \pi_t(Z') - \pi(Z') < \varepsilon$, then the conclusion in (B4) is replaced by $\limsup_{t \to \infty} \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}_{Z'}(z_\tau) \leq \pi(Z') + \varepsilon$.

[28]If H_ε has zero probability, the lemma is taken to be vacuous.

PROOF:

First note that $\#Z < \infty$ and thus any subset of $\mathcal{Z} \subset 2^Z$ has also finitely many elements. Therefore, it suffices to show the result for any (arbitrary) subset $Z' \in \mathcal{Z}$ since there are only finitely many of them (roughly speaking, Z' is what a is in FK93, i.e., Fudenberg and Kreps, 1993). Since Z is finite we can order the elements as $(z_1, ...z_{\#Z})$, and WLOG we set the first $\#Z'$ elements of Z to be the elements of Z'. Just as FK 93, let $(\omega_t)_t$ be an independent sequence of uniform random variables and let $y_t : \Omega \to Z$ be a new random variable.

As in FK 93, we construct $(y_t(\omega_t))_t$ as follows. For $t = 1$, $y_1(\omega_1) = z_m$ iff $\sum_{n=1}^{m-1} \pi_1(z_n) \leq \omega_1 < \sum_{n=1}^{m} \pi_1(z_n)$. For $t = \tau$, let $y_\tau(\omega_\tau) = z_m$ iff $\sum_{n=1}^{m-1} \pi_\tau(z_n|y_1, ..., y_{\tau-1}) \leq \omega_\tau < \sum_{n=1}^{m} \pi_\tau(z_n|y_1, ..., y_{\tau-1})$. Moreover, by construction the probability over h^t coincides with the probability over $(\omega_\tau)_{\tau \leq t}$; we thus can use both interchangeably. In particular, the set of ω for which $y_t(\omega_t) \in Z'$ is the set of $\{\omega : \omega_t \leq \sum_{n=1}^{\#Z'} \pi_t(z_n|y_1, ..., y_{t-1}) = \pi_t(Z'|y_1, ..., y_{t-1})\}$ (recall that Z' consists of the first $\#Z'$ elements in Z).

Under equation (B3) the latter set includes the set $\{\omega : \omega_t \leq \pi(Z') - \varepsilon\}$; thus $\mathbf{1}_{\{\omega:\omega_t \leq \pi(Z')-\varepsilon\}} \leq \mathbf{1}_{\{\omega:\omega_t \leq \pi_t(Z'|y_1,...,y_{t-1})\}} = \mathbf{1}_{\{\omega:y_t(\omega_t)\in Z'\}} = \mathbf{1}_{\{z_\tau \in Z'\}}$.

Let $\nu_t(r, \omega)$ be the number of times $\omega_t \leq r$. Then $\nu_t(\pi(Z') - \varepsilon) \leq \sum_{\tau=1}^{t} \mathbf{1}_{Z'}(z_\tau)$. By the strong law of large numbers, $\lim_{t\to\infty} \nu_t(\pi(Z') - \varepsilon) = \pi(Z') - \varepsilon$ a.s. on H_ε (this is under the measure of $(\omega_t)_t$, which by construction is equal to the measure associated with $(z_t)_t$). Therefore it must follow that $\liminf_{t\to\infty} t^{-1} \sum_{\tau=1}^{t} \mathbf{1}_{Z'}(z_\tau) \geq \pi(Z') - \varepsilon$.

Similarly, under equation (B3), the set $\{\omega : \omega_t \leq \pi_t(Z'|y_1, ..., y_{t-1})\}$ is included in the set $\{\omega : \omega_t \leq \pi(Z') + \varepsilon\}$. By a similar argument as before, $\limsup_{t\to\infty} t^{-1} \sum_{\tau=1}^{t} \mathbf{1}_{Z'}(z_\tau) \leq \pi(Z') + \varepsilon$.

LEMMA 3: *There exists H' with $\mathbf{P}^\phi(H') = 1$, such that for all $\varrho > 0$ and for all $h \in H'$, there exists $t_{\varrho,h}$ such that for all $t \geq t_{\varrho,h}$ and all i, $s_i \in \mathbb{S}_{\theta_i}$, and $o \in \{L, R\}$, $\overline{P}_t(h)(\mathcal{Z}_{iso}) > K_p - \varrho$, where*

$$K_p \equiv \frac{1}{2} \min_{i,s_i} \left\{ \Psi_i \times \int_\Omega q_{\theta_i}(s_i \mid \omega) G(d\omega) \right\} > 0, \tag{B5}$$

and, for all i, $\Psi_i \equiv min \left\{ \prod_{j=1}^{n} F_{\theta_j}(-2K), \prod_{j=1}^{n} \left(1 - F_{\theta_j}(2K)\right) \right\}$.

PROOF:

By definition of assessment rules, $\mu_{Rit}(h)(s) - \mu_{Lit}(h)(s) \in [-K, K]$, and, therefore, $\alpha_{it}(h)(s) \in \left[F_{\theta_i}(-2K), F_{\theta_i}(2K)\right]$ for all $i, s \in \mathbb{S}_{\theta_i}$, for all h, and for all t. Hence, for all $i, s \in \mathbb{S}_{\theta_i}$, for all h, and for all t,

$$\mathbf{P}_t^\phi\left(\boldsymbol{z}_t \in \mathcal{Z}_{isR} \mid h^t\right) \geq \prod_{j=1}^{n} F_{\theta_j}(-2K) \int_\Omega q_{\theta_i}(s \mid \omega) G(d\omega) > K_p,$$

and, similarly, $\mathbf{P}_t^\phi\left(\boldsymbol{z}_t \in \mathcal{Z}_{isL} \mid h^t\right) > K_p$, where $\mathbf{P}_t^\phi(\cdot \mid h^t)$ denotes the probability distribution over histories, conditional on history up to time t, $h^t \in \mathcal{H}^t$.

An application of Lemma 2 (by setting $\pi(\mathcal{Z}_{isR}) = K_p$—the case of $\mathcal{Z}_{isL}$ is analogous and thus omitted—$\varepsilon = 0$, and $H_\varepsilon = H$) implies that $\liminf_{t\to\infty} \overline{P}_t(h)(\mathcal{Z}_{iso}) \geq K_p$ $\mathbf{P}^\phi$- a.s. on H. Therefore, there exists a $H' \subseteq H$ with $\mathbf{P}^\phi(H') = 1$ such that for all $\varrho > 0$ and all $h \in H'$, there exists a $t_{\varrho,h}$ such that for all $t \geq t_{\varrho,h}$, $\overline{P}_t(h)(\mathcal{Z}_{iso}) > K_p - \varrho$.

Proof of Claim 2.1. By continuity of F_{θ_i}, it suffices to show that for all $\varepsilon > 0$, there exist $\gamma(\varepsilon) > 0$ with $\lim_{\varepsilon\to 0} \gamma(\varepsilon) = 0$ and H_ε with $\mathbf{P}^\phi(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$, there exists $t_{\varepsilon,h}$ such for all $t \geq t_{\varepsilon,h}$, all i, and all $s \in \mathbb{S}_{\theta_i}$,

$$F_{\theta_i}(\Delta_i(P, s) - \gamma(\varepsilon)) \leq \alpha_{it}(h)(s) \leq F_{\theta_i}(\Delta_i(P, s) + \gamma(\varepsilon)) \tag{B6}$$

and

$$\left|\overline{P}_t(h)(\boldsymbol{z}) - P(\boldsymbol{z})\right| < \varepsilon \tag{B7}$$

for all $\boldsymbol{z} \in \mathbb{Z}$.

Note that equation (6) can be written as

$$\mu_{oit}(h)(s) = \frac{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{iso}} \overline{P}_t(h)(\boldsymbol{x},\boldsymbol{s},\omega)u_{\theta_i}(o,\omega)}{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{iso}} \overline{P}_t(h)(\boldsymbol{x},\boldsymbol{s},\omega)} \tag{B8}$$

provided that $\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{iso}} \overline{P}_t(h)(\boldsymbol{x},\boldsymbol{s},\omega) > 0$.

Because P is stable, for all $\varepsilon > 0$, there exists t_ε and H_ε^* with $\mathbf{P}^{\boldsymbol{\phi}}(H_\varepsilon^*) > 0$ such that for all $h \in H_\varepsilon^*$ and $t \geq t_\varepsilon^*$, equation (B7) holds for all $\boldsymbol{z} \in \mathbb{Z}$. In addition, (B7) implies that

$$\left|\overline{P}_t(h)(Z') - P(Z')\right| < \varepsilon \times \#\mathbb{Z} \tag{B9}$$

for all $Z' \subset \mathbb{Z}$. Next, let $b = \min\left\{\mathbf{P}^{\boldsymbol{\phi}}(H_\varepsilon^*), .5K_p\right\} > 0$, where $K_p > 0$ is defined by equation (B5). By Lemma 3, for all $h \in H \setminus H^o$ (where H^o has zero measure) there exists $t_{b,h}$ such that for all $t \geq t_{b,h}$ and $o \in \{L, R\}$

$$\overline{P}_t(h)(\mathcal{Z}_{iso}) > K_p - b \geq .5K_p > 0. \tag{B10}$$

Let $H_\varepsilon = H_\varepsilon^* \cap H \setminus H^o$, and note that by our choice of H_ε^*, $\mathbf{P}^{\boldsymbol{\phi}}(H_\varepsilon) > 0$. Therefore, for all $\varepsilon > 0$, there exists H_ε with $\mathbf{P}^{\boldsymbol{\phi}}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$ and $t \geq t_{\varepsilon,h} \equiv \max\{t_\varepsilon^*, t_{b,h}\}$, all $i, s \in \mathbb{S}_{\theta_i}$,

$$|\mu_{Rit}(h)(s) - \mu_{Lit}(h)(s) - \Delta_i(P,s)| \leq \gamma(\varepsilon) \equiv \frac{(\varepsilon \times \#\mathbb{Z}) \times (0.5K_p)^2}{2K\,(1 + \#\Omega) + 0.5\,(\varepsilon \times \#\mathbb{Z})\,K_p} \underset{\varepsilon\to 0}{\longrightarrow} 0 \tag{B11}$$

where the inequality follows from (B8), (B9), (B10), the facts that utility is bounded by K and $\#\Omega < \infty$, and simple algebra that uses the fact that

$$\Delta_i(P, s_i) = \frac{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{isR}} P(\boldsymbol{x},\boldsymbol{s},\omega)u_{\theta_i}(R,\omega)}{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{isR}} P(\boldsymbol{x},\boldsymbol{s},\omega)} - \frac{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{isL}} P(\boldsymbol{x},\boldsymbol{s},\omega)u_{\theta_i}(L,\omega)}{\sum_{(\boldsymbol{x},\boldsymbol{s},\omega)\in\mathcal{Z}_{isL}} P(\boldsymbol{x},\boldsymbol{s},\omega)}.$$

Then, equation (B11) and the definition of the policy rules imply that

$$\phi_{it}(h, v_i)(s) = \begin{cases} R & \text{if } v_i \leq \Delta_i(P,s) - \gamma(\varepsilon) \\ L & \text{if } v_i > \Delta_i(P,s) + \gamma(\varepsilon) \end{cases},$$

so that (B6) holds by (B1). □

Proof of Claim 2.2. Note that for each $\boldsymbol{z} \in \mathbb{Z}$,

$$\mathbf{P}^{\boldsymbol{\phi}}(\boldsymbol{z}_t = \boldsymbol{z} \mid h^t) = P^n(\boldsymbol{\alpha}_t(h))(\boldsymbol{z}).$$

Then, Claim 2.1 and the fact that $P^n(\cdot)$ is continuous imply that for all $\varepsilon > 0$, there exists H_ε with $\mathbf{P}^{\boldsymbol{\phi}}(H_\varepsilon) > 0$ such that for all $h \in H_\varepsilon$, there exists $\hat{t}_{\varepsilon,h}$ such that for all $t \geq \hat{t}_{\varepsilon,h}$,

$$\left|\mathbf{P}^{\boldsymbol{\phi}}(\boldsymbol{z}_t = \boldsymbol{z} \mid h^t) - P^n(\boldsymbol{\sigma})(\boldsymbol{z})\right| < \varepsilon \tag{B12}$$

and

$$\left|\overline{P}_t(h)(\boldsymbol{z}) - P(\boldsymbol{z})\right| < \varepsilon \tag{B13}$$

for all $\boldsymbol{z} \in \mathbb{Z}$.

Then, by equation (B12) and Lemma 2 applied to all the singleton sets of $\mathbb{Z}$,

$$\limsup_{t\to\infty} \overline{P}_t(h)(\boldsymbol{z}) \le P^n(\boldsymbol{\sigma})(\boldsymbol{z}) + \varepsilon \quad \text{and} \quad \liminf_{t\to\infty} \overline{P}_t(h)(\boldsymbol{z}) \ge P^n(\boldsymbol{\sigma})(\boldsymbol{z}) - \varepsilon \tag{B14}$$

for all $\boldsymbol{z} \in \mathbb{Z}$, almost surely on H_ε.

By the triangle inequality, for any t,

$$||P - P^n(\boldsymbol{\sigma})|| \le ||P - \overline{P}_t(h)|| + ||\overline{P}_t(h) - P^n(\boldsymbol{\sigma})|| \tag{B15}$$

for any $h \in H_\varepsilon$; we pick one $h \in H_\varepsilon$ (outside the measure zero set). By equation (B14), the second summand in the RHS of (B15) is less than ε for all t sufficiently large; by equation (B13), the first summand of the RHS is also less than ε for all t sufficiently large. Hence, $||P - P^n(\boldsymbol{\sigma})|| \le \varepsilon$; since this holds for all $\varepsilon > 0$, then we obtain the desired result by taking $\varepsilon \to 0$. □

B2. Large number of players

In this section, we characterize (naive) equilibria of the voting game as the number of voters goes to infinity. We do so by studying sequences of voting games. We build such sequences by independently drawing infinite sequences of types $\xi = (\theta_1, \theta_2, ..., \theta_n, ...) \in \Xi$ according to the probability distribution $\varphi \in \Delta(\Theta)$; we denote the distribution over Ξ by Φ and we let $\theta_i(\xi)$ denote the type of player i, i.e., the ith component of ξ. We interpret each sequence of types as describing an infinite number of n-player games by letting the first n elements of ξ represent the types of the n players.

Let $\boldsymbol{\varsigma}$ denote a *strategy mapping* from sequences of types Ξ to sequences of strategy profiles–i.e., for all $\xi \in \Xi$, let $\boldsymbol{\varsigma}(\xi) = \left(\boldsymbol{\sigma}^1(\xi), ..., \boldsymbol{\sigma}^n(\xi), ...\right)$, where

$$\boldsymbol{\sigma}^n(\xi) = (\sigma_1^n(\xi), ..., \sigma_n^n(\xi))$$

is the strategy profile that is played in the n-player game with types $\theta_1, ..., \theta_n$. Let $P^n(\boldsymbol{\varsigma}(\xi))$ be the probability distribution over $\{L, R\}^n \times \Pi_{i=1}^n \mathbb{S}_{\theta_i} \times \Omega$ induced by the strategy profile $\boldsymbol{\sigma}^n(\xi)$ in the n-player game. We define two properties of strategy mappings.

DEFINITION 7: *A strategy mapping $\boldsymbol{\varsigma}$ is an ε-*equilibrium mapping *if, for a.e. $\xi \in \Xi$, there exists $n_{\varepsilon,\xi}$ such that for all $n \ge n_{\varepsilon,\xi}$*

$$\left\|\sigma_i^n(\xi) - F_{\theta_i(\xi)}\left(\Delta_i(P^n(\boldsymbol{\varsigma}(\xi)), \cdot)\right)\right\| \le \varepsilon \tag{B16}$$

for all $i = 1, ..., n$.[29] *A strategy mapping $\boldsymbol{\varsigma}$ is* asymptotically interior *if, for a.e. $\xi \in \Xi$,*

$$\liminf_{n\to\infty} P^n(\boldsymbol{\varsigma}(\xi))\,(o = R) > 0 \quad \textit{and} \quad \limsup_{n\to\infty} P^n(\boldsymbol{\varsigma}(\xi))\,(o = R) < 1. \tag{B17}$$

The first property in Definition 7 requires that, for large enough n, players play strategies that constitute an ε equilibrium. Our notion of limit equilibrium will require this property to hold for all $\varepsilon > 0$; while being slightly weaker than requiring strategies to constitute an equilibrium, this condition yields a full characterization of limit equilibrium.[30] The second property requires that the probabilities of choosing R and L remain bounded away from zero as the number of players increases. The reason for this

[29]The a.e. in "for a.e. $\xi \in \Xi$" stands for "almost every" and means that there is a set Ξ' with $\Phi(\Xi') = 1$ such that a condition is true for all $\xi \in \Xi'$. The results continue to hold if we only require $\Phi(\Xi') > 0$.

[30]Our result that a limit equilibrium is a fixed point of a particular correspondence remains true under the stronger requirement that strategies constitute an equilibrium. But the converse result, that any fixed point is also a limit equilibrium, relies on the notion of ε equilibrium.

restriction is that we can always obtain extreme equilibria where everyone votes for the same alternative, no information is obtained about the other alternative, and, therefore, beliefs about the other alternative can be arbitrary. The restriction to asymptotically interior strategies allows us to focus on equilibria where beliefs are not arbitrary.

In addition to characterizing the equilibrium cutoff, we characterize the profile of equilibrium strategies. Given a strategy mapping ς and a sequence of types $\xi \in \Xi$, let $\bar{\sigma}^n(\xi;\varsigma) : \Theta \to [0,1]^{\mathbb{S}}$ represent the average strategy of each type in the n-player game. Formally, for all $\theta \in \Theta$ and $s \in \mathbb{S}$,

$$\bar{\sigma}^n_\theta(\xi;\varsigma)(s) = \frac{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}\, \sigma^n_i(\xi)(s)}{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}} \tag{B18}$$

whenever $\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\} > 0$, and arbitrary otherwise. We call any element $\sigma : \Theta \to [0,1]^{\mathbb{S}}$ an *average strategy function* and say that σ is *increasing* if $s' > s$ implies $\sigma_\theta(s') > \sigma_\theta(s)$ for every type $\theta \in \Theta$.

DEFINITION 8: *An average strategy function $\sigma^* : \Theta \to [0,1]^{\mathbb{S}}$ is a* limit ε-equilibrium *if there exists an asymptotically interior ε-equilibrium mapping ς such that $\lim_{n\to\infty} \|\bar{\sigma}^n(\xi;\varsigma) - \sigma^*\| = 0$ for a.e. $\xi \in \Xi$. An average strategy function σ^* is a* limit equilibrium *if it is a limit ε-equilibrium for all $\varepsilon > 0$.*

The following result characterizes limit equilibria.

THEOREM 3: *σ^* is a limit equilibrium if and only if there exists a cutoff $\omega^* \in (-1,1)$ such that $\kappa(\omega^*;\sigma^*) = \rho$ and $\sigma^*_\theta(s) = F(v_\theta(s;\omega^*))$ for all $\theta \in \Theta$ and $s \in \mathbb{S}$.*

The intuition of Theorem 3 is as follows. Suppose that there is a sequence of average strategy functions $\bar{\sigma}^n$ that converges to an increasing function σ^*.[31] Then the probability that a randomly chosen player votes for R in state of the world ω converges to $\kappa(\omega;\sigma^*)$. By standard asymptotic arguments, the proportion of votes for R in state ω becomes concentrated around $\kappa(\omega;\sigma^*)$. So, for states where $\kappa(\omega;\sigma^*) > \rho$, the probability that R is elected converges to 1. Similarly, for states where $\kappa(\omega;\sigma^*) < \rho$, the probability that R is elected converges to 0. Since σ is increasing, then there is at most one (measure zero) state ω^* such that $\kappa(\omega^*;\sigma^*) = \rho$, so that the election outcome is characterized by an election cutoff ω^*. Moreover, the fact that the election outcome is characterized by a cutoff means that the beliefs of player i, Δ_i, can be approximated by the belief function v_{θ_i} defined in equation (1). Thus, the optimal strategy of a player of type θ who observes signal s is $\sigma^*_\theta(s) = F(v_\theta(s;\omega^*))$.

VANISHING PERTURBATIONS. — We now consider sequences of equilibria where the perturbations vanish. We index games by a parameter η that indexes the cdf F^η_θ from which perturbations are drawn.

DEFINITION 9: *A family of perturbations $\{\mathbf{F}^\eta\}_{\eta\in\mathbb{N}}$, where $\mathbf{F}^\eta = \{F^\eta_\theta\}_{\theta\in\Theta}$, is* vanishing *if for all $\theta \in \Theta$ and η: assumption A4 is satisfied and*

$$\lim_{\eta\to 0} F^\eta_\theta(\nu) = \begin{cases} 0 & \text{if } \nu < 0 \\ 1 & \text{if } \nu > 0 \end{cases}$$

Under a vanishing family of perturbations, the payoff perturbations converge to zero and we recover the original, unperturbed game. The next two results provide a foundation for the notion of RVE introduced in Section I.

[31] We show in Section B.B2 that optimal strategies are increasing when the number of players is sufficiently large.

THEOREM 4: *(i) Suppose that there exists a vanishing family of perturbations $\{\mathbf{F}^\eta\}_\eta$ and a sequence $(\sigma^\eta, \omega^\eta)_\eta$ such that $\lim_{\eta\to 0}(\sigma^\eta, \omega^\eta) = (\sigma^*, \omega^*)$ and where σ^η is a limit equilibrium and ω^η its corresponding cutoff for all η. Then (σ^*, ω^*) is a retrospective voting equilibrium.*

(ii) Suppose that (σ^*, ω^*) is a retrospective voting equilibrium with $\omega^* \in (-1, 1)$. Then there exists a vanishing family of perturbations $\{\mathbf{F}^\eta\}_\eta$ and a sequence $(\sigma^\eta, \omega^\eta)_\eta$ such that $\lim_{\eta\to 0}(\sigma^\eta, \omega^\eta) = (\sigma^*, \omega^*)$ and where σ^η is a limit equilibrium and ω^η its corresponding cutoff for all η.

The first part of Theorem 4 follows by standard continuity arguments and the second part by construction.[32]

REMARK 1: *Situations where one alternative is never chosen are easily justified: if an alternative is never chosen, then beliefs about its performance can be arbitrary. Our solution concept in Section I considers, say, $\omega^* = 1$ (i.e., Right is never chosen) to be an equilibrium cutoff only if equilibrium beliefs are such that Right yields the payoff at state $\omega = 1$ and $Left$ yields the unconditional payoff. The formal justification is that, if players follow symmetric increasing strategies such that the probability of Right being elected converges to zero, then the probability of state $\omega = 1$ conditional on Right being elected converges to 1.*

REMARK 2: *Our game-theoretic foundation uses assumption A5, which is stronger than A2 in Section I. In particular, A2 allows for the case where voters have no private information. We can provide a foundation for such a case by considering a sequence of voting games indexed by $r \in \mathbb{N}$, where $z^r > 0$ denotes the constant defined in assumption A5, and where $\lim_{r\to\infty} z^r = 0$. Therefore, the case of no private information must be viewed as the limiting case of an information structure that satisfies A5 but where informativeness vanishes.*

REMARK 3: *Finally, in Section I we assumed that Θ was a compact interval, rather than a finite set, in order to obtain uniqueness of equilibrium and facilitate the application of the framework. However, we can view the case where Θ is a compact interval as the limiting case of a sequence of environments where the finite number of elements in Θ goes to infinity.*

The remainder of this section provides the proofs of Theorems 3 and 4.

PRELIMINARY LEMMA. — The proof of Theorem 3 relies on the following lemma.

LEMMA 4: *Let $\boldsymbol{\varsigma}$ be such that $\lim_{n\to\infty} \|\bar{\sigma}^n(\xi; \boldsymbol{\varsigma}) - \sigma^*\| = 0$ for a.e. Ξ, where σ^* is increasing. Then there exists $\omega^* \in \arg\min_{\omega\in\Omega} |\kappa(\omega; \sigma^*) - \rho|$ such that for all $\epsilon > 0$ and a.e. $\xi \in \Xi$,*

$$\lim_{n\to\infty} \inf_{\omega\in\Omega:\omega\geq\omega^*+\epsilon} P^n(\boldsymbol{\varsigma}(\xi))(o = R \mid \omega) = 1 \tag{B19}$$

and

$$\lim_{n\to\infty} \sup_{\omega\in\Omega:\omega\leq\omega^*-\epsilon} P^n(\boldsymbol{\varsigma}(\xi))(o = R \mid \omega) = 0. \tag{B20}$$

Moreover, if $\omega^* \in (-1, 1)$, then for a.e. $\xi \in \Xi$ and for all $\varepsilon > 0$ there exists $n_{\xi,\varepsilon}$ such that for all $n \geq n_{\xi,\varepsilon}$,

$$\left\|\Delta_i(P^n(\boldsymbol{\varsigma}(\xi)), \cdot) - v_{\theta_i(\xi)}(\cdot; \omega^*)\right\| \leq \varepsilon \tag{B21}$$

for all $i = 1, ..., n$.

[32] As shown in the proof, the argument holds for *any* family of perturbations if ω^* is the unique equilibrium cutoff and $\varphi\{(\theta : c_\theta(s) = \omega^*,\ s \in \mathbb{S}_\theta)\} = 0$.

PROOF:

We use the following notation. Let $x_i \in \{L, R\}$ denote the vote of player i, let $\kappa_i^n(\omega;\xi) \equiv P^n\left(x_i = R \mid \omega\right)$ be the probability that player $i = 1, ..., n$ votes for R conditional on the state being ω, and let $\kappa^n(\omega;\xi) \equiv \frac{1}{n}\sum_{i=1}^n \kappa_i^n(\omega;\xi)$ be the average over all players.

First, note that, for a.e. $\xi \in \Xi$, for all $\omega \in \Omega$,

$$\begin{aligned}
\lim_{n\to\infty} \kappa^n(\omega;\xi) &= \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^n \sum_{\theta\in\Theta}\sum_{s\in\mathbb{S}} q_\theta(s|\omega) 1\{\theta_i(\xi)=\theta\}\sigma_i^n(\xi)(s) \\
&= \lim_{n\to\infty} \sum_{\theta\in\Theta}\sum_{s\in\mathbb{S}} q_\theta(s|\omega)\left\{\frac{1}{n}\sum_{i=1}^n 1\{\theta_i(\xi)=\theta\}\sigma_i^n(\xi)(s)\right\} \\
&= \sum_{\theta\in\Theta}\sum_{s\in\mathbb{S}} q_\theta(s|\omega)\left\{\lim_{n\to\infty}\bar{\sigma}_\theta^n(\xi;\varsigma)(s) \times \left(\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^n 1\{\theta_i(\xi)=\theta\}\right)\right\} \\
&= \sum_{\theta\in\Theta}\sum_{s\in\mathbb{S}} q_\theta(s|\omega)\sigma_\theta^*(s)\varphi(\theta) = \kappa(\omega;\sigma^*), \qquad \text{(B22)}
\end{aligned}$$

where we have used the assumption that $\lim_{n\to\infty}\|\bar{\sigma}^n(\xi;\varsigma) - \sigma^*\| = 0$ a.s.-Ξ and the strong law of large numbers applied to $\frac{1}{n}\sum_{i=1}^n 1\{\theta_i(\xi)=\theta\}$. Note also that, for all $\omega, \omega' \in \Omega$,

$$\begin{aligned}
|\kappa^n(\omega;\xi) - \kappa^n(\omega';\xi)| &\le \sum_{\theta\in\Theta}\sum_{s\in\mathbb{S}} \left|q_\theta(s|\omega) - q_\theta(s|\omega')\right| \left\{\bar{\sigma}_\theta^n(\xi;\varsigma)(s) \times \left(\frac{1}{n}\sum_{i=1}^n 1\{\theta_i(\xi)=\theta\}\right)\right\} \\
&\le \max_{\theta\in\Theta}\max_{s\in\mathbb{S}} \left|q_\theta(s|\omega) - q_\theta(s|\omega')\right|
\end{aligned}$$

and since $|\Theta| < \infty$ and $|\mathbb{S}| < \infty$, this display and A3(iii) imply that the family $\{\kappa^n(\cdot;\xi) : \Omega \to [0,1] \colon n = 1, 2, ...\}$ is equicontinuous for all $\xi \in \Xi$. This result, the one in (B22) and the fact that Ω is compact, implies that

$$\lim_{n\to\infty}\sup_{\omega\in\Omega} |\kappa^n(\omega;\xi) - \kappa(\omega;\xi)| = 0 \qquad \text{(B23)}$$

a.s.-Ξ.

Second, let $Y^n(\omega;\xi) \equiv n^{-1/2}\sum_{i=1}^n \left(1\{x_i^n = R\} - \kappa_i^n(\omega;\xi)\right)$. It follows that for all $\delta > 0$ and for a.e. ξ, there exists $n'(\delta,\xi)$ such that, for all $n \ge n'(\delta,\xi)$,

$$\begin{aligned}
P^n(\varsigma(\xi))(o = R \mid \omega) &= P^n(\varsigma(\xi))\left(Y^n(\omega;\xi) \ge \sqrt{n}(\rho - \kappa^n(\omega;\xi)) \mid \omega\right) \\
&\le P^n(\varsigma(\xi))\left(Y^n(\omega;\xi) \ge \sqrt{n}\delta \mid \omega\right) \\
&\le (n^2\delta^2)^{-1}\sum_{i=1}^n E\left[(1\{x_i^n = R\} - \kappa_i^n(\omega;\xi))^2 \mid \omega\right] \\
&\le 4(n\delta^2)^{-1},
\end{aligned}$$

for all $\omega \in \{\omega \in \Omega : \kappa(\omega;\sigma) \le \rho - \delta\}$, where the second line follows from (B23) and the third from the Markov inequality.

Third, the facts that $\kappa(\cdot;\sigma^*)$ is increasing (because σ^* is increasing) and continuous (by A3(iii)) imply that there exists $\omega^* \in [-1,1]$ such that $\omega^* \in \arg\min_{\omega\in\Omega}|\kappa(\omega;\sigma^*) - \rho|$ and that, for any $\epsilon > 0$, there exists a $\delta > 0$ such that $\kappa(\omega;\sigma^*) \le \rho - \delta$ for all $\omega \le \omega^* - \epsilon$. Hence, $\{\omega \in \Omega : \omega \le \omega^* - \epsilon\} \subseteq \{\omega \in \Omega : \kappa(\omega;\sigma^*) \le \rho - \delta\}$, and the previous argument implies that

$$\lim_{n\to\infty}\sup_{\omega\in\Omega:\omega\le\omega^*-\epsilon} P^n(\varsigma(\xi))(o = R \mid \omega) = 0. \qquad \text{(B24)}$$

By employing a similar argument, it follows that $\lim_{n\to\infty}\inf_{\omega\in\Omega:\omega\ge\omega^*+\epsilon} P^n(\varsigma(\xi))(o = R \mid \omega) = 1$ a.e.

$\xi \in \Xi$.

We now establish the second part of the lemma. Suppose that $\omega^* \in (-1,1)$. First note that the previous part of the proof implies that, for any $\omega \in \Omega$

$$\lim_{n\to\infty} P^n(\varsigma(\xi))(o = R \mid \omega) = 1\{\omega > \omega^*\}. \tag{B25}$$

Second, note that, for all n and all $\omega \in \Omega$,

$$P^n(\varsigma(\xi))(o = R \mid \omega) = \sum_{s\in\mathbb{S}} P^n(\xi)(o = R \mid \omega, S_i = s) q_{\theta_i(\xi)}(s|\omega) \tag{B26}$$

for all $i \leq n$. By (B25), (B26), and A3(ii), for a.e. $\xi \in \Xi$ and all $s \in \mathbb{S}$,

$$\lim_{n\to\infty} P^n(\varsigma(\xi))(o = R \mid \omega, S_i = s) = 0 \ (= 1) \tag{B27}$$

for $\omega < \omega^*$ ($\omega > \omega^*$), where convergence is uniform in $i \leq n$.[33] Therefore, for a.e. $\xi \in \Xi$ and all $s \in \mathbb{S}$, $\lim_{n\to\infty} E_{P^n(\varsigma(\xi))}\left(u_{\theta_i(\xi)}(R, W) \mid o = R, S_i = s\right) =$

$$\begin{aligned}
&= \lim_{n\to\infty} \frac{\int_\Omega P^n(\varsigma(\xi))\,(o = R \mid W, S_i = s)\, q_{\theta_i(\xi)}(s \mid W) u_{\theta_i(\xi)}(R, W) G(dW)}{\int_\Omega P^n(\varsigma(\xi))\,(o = R \mid W, S_i = s)\, q_{\theta_i(\xi)}(s \mid W) G(dW)} \\
&= \frac{\int_\Omega \lim_{n\to\infty} P^n(\varsigma(\xi))\,(o = R \mid W, S_i = s)\, q_{\theta_i(\xi)}(s \mid W) u_{\theta_i(\xi)}(R, W) G(dW)}{\int_\Omega \lim_{n\to\infty} P^n(\varsigma(\xi))\,(o = R \mid W, S_i = s)\, q_{\theta_i(\xi)}(s \mid W) G(dW)} \\
&= \frac{\int_\Omega 1\{W > \omega^*\} q_{\theta_i(\xi)}(s|W) u_{\theta_i(\xi)}(R, W) G(dW)}{\int_\Omega 1\{W > \omega^*\} q_{\theta_i(\xi)}(s|W) G(dW)} \\
&= E\left(u_{\theta_i(\xi)}(R, W) \mid W > \omega^*, S_i = s\right),
\end{aligned} \tag{B28}$$

where convergence is uniform in $i \leq n$. The first and fourth lines in (B28) follow by definition, the second line follows from the dominated convergence theorem and the fact that u_{θ_i} is bounded (and the denominator being greater than zero, as established next), and the third line follows from (B27) and the fact that G is absolutely continuous, so we can ignore the case $\{W = \omega^*\}$ (also, note the importance of $\omega^* < 1$ for the denominator to be well-defined). A similar argument holds for $E_{P^n(\varsigma(\xi))}\left(u_{\theta_i(\xi)}(L, W) \mid o = L, S_i = s\right)$, thus establishing the lemma.

Proof of Theorem 3. — proof:

Only if: Let σ^* be a limit equilibrium, so that σ^* is a limit ε-equilibrium for all $\varepsilon > 0$. Lemma OA in Section B.B2, shows that σ^* must be increasing. Fix any $\varepsilon > 0$ and let ς be the corresponding ε-equilibrium mapping that is asymptotically interior. Because ς is asymptotically interior, then $\omega^* \in (-1,1)$ and, therefore, (B21) holds by Lemma 4. Then, for all $\theta \in \Theta$, there exists $\xi \in \Xi$ and n' such that for all $n \geq n'$,

$$\begin{aligned}
\|\sigma^*_\theta - F_\theta(v_\theta(\cdot;\omega^*))\| &\leq \|\sigma^*_\theta - \bar{\sigma}^n_\theta(\xi;\varsigma)\| \\
&+ \left\| \frac{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}\, \sigma^n_i(\xi)(s)}{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}} - \frac{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}\, F_\theta\left(\Delta_i(P^n(\varsigma(\xi))), s)\right)}{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}} \right\| \\
&+ \left\| \frac{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}\, F_\theta\left(\Delta_i(P^n(\varsigma(\xi))), s)\right)}{\sum_{i=1}^n 1\{\theta_i(\xi) = \theta\}} - F_\theta\left(v_\theta\left(\cdot;\omega^*\right)\right) \right\| \\
&\leq \varepsilon + \varepsilon + \varepsilon,
\end{aligned}$$

[33]Formally, suppose that $\omega < \omega^*$. Then for all $\varepsilon > 0$ there exists $n_{\xi,\omega,\varepsilon}$ such that, for all $n \geq n_{\xi,\omega,\varepsilon}$, $P^n(\varsigma(\xi))(o = R \mid \omega, S_i = s) q_{\theta_i(\xi)}(s|\omega) \leq \varepsilon$ for all $i \leq n$ and $s \in \mathbb{S}$.

where the last inequality follows because: (i) σ^* being a limit equilibrium implies that $\lim_{n\to\infty}\|\bar{\sigma}^n(\xi;\varsigma)-\sigma^*\|=0$ for a.e. $\xi\in\Xi$; (ii) ς is an ε-equilibrium mapping; and (iii) equation (B21) and continuity of F_θ (A4). Since the above relationship holds for every $\varepsilon>0$, then $\left\|\sigma^*_\theta - F_\theta(v_\theta(\cdot;\omega^*))\right\|=0$ for all θ.

If: Consider the strategy mapping ς defined by letting players of type θ always play σ^*_θ–i.e., for all ξ,s,n, and $i\leq n$, $\sigma^n_i(\xi)(s)=\sigma^*_{\theta_i(\xi)}(s)$. First, note that $\bar{\sigma}^n=\sigma^*$ converges trivially to σ^*, and σ^* is increasing because F_θ and $v_\theta(\cdot;\omega^*)$ are increasing (by A1-A3 and A4). Moreover, $\omega^*\in(-1,1)$ by assumption. Then, equations (B19) and (B20) in Lemma 4 and the dominated convergence theorem imply that ς is asymptotically interior. In addition, for a.e. $\xi\in\Xi$ and for every $\varepsilon>0$, there exists $n_{\xi,\varepsilon}$ such that for all $n\geq n_{\xi,\varepsilon}$,

$$\begin{aligned}\left\|\sigma^n_i(\xi)-F_{\theta_i(\xi)}\left(\Delta_i(P^n(\varsigma(\xi)),\cdot)\right)\right\| &= \left\|\sigma^*_{\theta_i(\xi)}-F_{\theta_i(\xi)}\left(\Delta_i(P^n(\varsigma(\xi)),\cdot)\right)\right\|\\ &= \left\|F_{\theta_i(\xi)}\left(v_{\theta_i(\xi)}\left(\cdot;\omega^*\right)\right)-F_{\theta_i(\xi)}\left(\Delta_i(P^n(\varsigma(\xi)),\cdot)\right)\right\|\leq\varepsilon\end{aligned}$$

for all $i=1,...,n$, where the first line follows by construction of the strategy and the second line follows by (B21) and continuity of F_θ (A4). Thus, (σ^*,ω^*) is a limit equilibrium.

Proof of Theorem 4. — proof:

Part (i): Theorem 3 implies that $\sigma^*_\theta(s)=\lim_{\eta\to 0}\sigma^\eta_\theta(s)=\lim_{\eta\to 0}F^\eta_\theta\left(v_\theta(s;\omega^*)\right)$ for all $\theta\in\Theta$ and $s\in\mathbb{S}$. Since $\mathbf{F}^\eta$ is vanishing, then $\sigma_\theta(s)=1$ if $v_\theta(s;\omega^*)>0$ and $\sigma_\theta(s)=0$ if $v_\theta(s;\omega^*)<0$. Therefore, σ^* is optimal given ω^*. Next, fix any $\omega'<\omega^*$. Since $\omega^\eta\to\omega^*$, there exists $\bar{\eta}$ such that, for all $\eta<\bar{\eta}$, $\omega'<\omega^\eta$, and, by Theorem 3, $\kappa(\omega';\sigma^\eta)\leq\rho$. Since $\sigma^\eta\to\sigma^*$, continuity of $\kappa(\omega';\cdot)$ implies that $\kappa(\omega';\sigma^*)\leq\rho$. Similarly, $\kappa(\omega'';\sigma^*)\geq\rho$ for all $\omega''>\omega^*$. Therefore, ω^* is an election cutoff given σ^*.

Part (ii): For any family of vanishing perturbations $\{\mathbf{F}^\eta\}_\eta$, define

$$\bar{\kappa}^\eta(\omega)\equiv\sum_{\theta\in\Theta}\varphi(\theta)\sum_{s\in\mathbb{S}}q_\theta(s\mid\omega)F^\eta_\theta\left(v_\theta(s;\omega)\right).$$

Let (σ^*,ω^*) be a voting equilibrium with $\omega^*\in(-1,1)$. Because $\omega^*\in(-1,1)$ is an election cutoff given σ^* and $\kappa(\cdot;\sigma^*)$ is continuous, then $\kappa(\omega^*;\sigma^*)=\rho$. We split the proof into two cases: Either it is the case that all players vote for the same alternative (which may be different for each player) irrespective of their private information–so that $\kappa(\cdot;\sigma^*)$ is a constant function–or not–so that $\kappa(\cdot;\sigma^*)$ is increasing.

Case 1 ($\kappa(\cdot;\sigma^*)$ is increasing): Rewrite $\bar{\kappa}^\eta$ as

$$\begin{aligned}\bar{\kappa}^\eta(\omega)=\sum_{\theta\in\Theta}\varphi(\theta)\Bigg\{&\sum_{s:c_\theta(s)<\omega^*}q_\theta(s\mid\omega)F^\eta_\theta\left(v_\theta(s;\omega)\right)+\sum_{s:c_\theta(s)=\omega^*}q_\theta(s\mid\omega)F^\eta_\theta\left(v_\theta(s;\omega)\right)\\ &+\sum_{s:c_\theta(s)>\omega^*}q_\theta(s\mid\omega)F^\eta_\theta\left(v_\theta(s;\omega)\right)\Bigg\}\equiv T^\eta_1(\omega)+T^\eta_2(\omega)+T^\eta_3(\omega).\end{aligned}$$

Since $v_\theta(s;\cdot)$ is increasing and $\omega^*\in(-1,1)$, then: for all (θ,s) such that $c_\theta(s)\geq\omega^*$, $v_\theta(s;\omega)<0$ for all $\omega<\omega^*$ and, for all (θ,s) such that $c_\theta(s)\leq\omega^*$, $v_\theta(s;\omega)>0$ for all $\omega>\omega^*$. Therefore, since $\{\mathbf{F}^\eta\}_\eta$ is vanishing, $\lim_{\eta\to 0}T^\eta_2(\omega)+T^\eta_3(\omega)=0$ for all $\omega<\omega^*$ and $\lim_{\eta\to 0}T^\eta_1(\omega)+T^\eta_2(\omega)=\sum_{\theta\in\Theta}\varphi(\theta)q_\theta(c_\theta(S_\theta)\leq\omega^*\mid\omega)\geq\kappa(\omega;\sigma^*)$ for all $\omega>\omega^*$. In addition, $T^\eta_1(\omega)\leq\kappa(\omega;\sigma^*)$ and $T^\eta_3(\omega)\geq 0$ for all ω. Therefore, $\lim_{\eta\to 0}\bar{\kappa}^\eta(\omega)\leq\kappa(\omega;\sigma^*)<\kappa(\omega^*;\sigma^*)=\rho$ for all $\omega<\omega^*$ and $\lim_{\eta\to 0}\bar{\kappa}^\eta(\omega)\geq\kappa(\omega;\sigma^*)>\kappa(\omega^*;\sigma^*)=\rho$ for all $\omega>\omega^*$. Consequently, by continuity of $\kappa^\eta(\cdot)$, there exists $(\omega^\eta)_\eta$ such that $\omega^\eta\to\omega^*\in(-1,1)$ and $\bar{\kappa}^\eta(\omega^\eta)=\rho$ for all sufficiently small η. By letting $\sigma^\eta_\theta(s)=F^\eta_\theta(v_\theta(s;\omega^\eta))$ for all θ,s, it follows that $\kappa(\omega^\eta;\sigma^\eta)=\bar{\kappa}^\eta(\omega^\eta)=\rho$ for all sufficiently small η and, by Theorem 3, that σ^η is a limit equilibrium and ω^η its corresponding cutoff for all sufficiently small η. Finally, it remains to establish that $\sigma^\eta\to\sigma^*$. Consider a type and signal such that $c_\theta(s)<\omega^*$, so that $v_\theta(s;\omega^*)>0$. By continuity of $v_\theta(s;\cdot)$ and the fact that $\omega^\eta \to \omega^*$, it follows that $v_\theta(s;\omega^\eta) > 0$ for all sufficiently small η and, therefore, because $\{\boldsymbol{F}^\eta\}_\eta$ is vanishing, it also follows that $\lim_{\eta\to 0}\sigma_\theta^\eta(s) = 1 = \sigma_\theta^*(s)$, where the last equality follows since σ^* is optimal given ω^*–see equation (A3). A similar argument establishes that $\lim_{\eta\to 0}\sigma_\theta^\eta(s) = 0 = \sigma_\theta^*(s)$ for types and signals such that $c_\theta(s) > \omega^*$. Therefore, if $\{s : c_\theta(s) = \omega^*\} = \emptyset$ for all θ, we have shown that, for *any* family of vanishing perturbations, there exists a sequence of limit equilibria that converge to a voting equilibria. In the case where $\{s : c_\theta(s) = \omega^*\} \neq \emptyset$ for some θ, we construct a specific family of perturbations $\left\{\hat{\mathbf{F}}^\eta\right\}_\eta$ with the property that $\lim_{\eta\to 0}\hat{F}_\theta^\eta(v_\theta(s;\omega^\eta)) = \sigma_\theta(s)$ for all (θ, s) such that $c_\theta(s) = \omega^*$. The details that show existence of such a family are tedious but straightforward and are as follows.First, observe that $\omega^\eta \to \omega^*$ and thus, by continuity of $v_{\theta'}(s_{\theta'};\cdot)$, it follows that $v_{\theta'}(s_{\theta'};\omega^\eta) \to 0$. Since there are a finite number of such $(\theta', s_{\theta'})$, there exists a sequence $(r^\eta)_\eta$ such that $r^\eta \to 0$ and $|v_{\theta'}(s_{\theta'};\omega^\eta)| \leq r^\eta$ uniformly over $(\theta', s_{\theta'})$. Second, for each $\theta' \in \Theta$, let $\hat{F}_{\theta'}^\eta(0) = \sigma_{\theta'}(s_{\theta'})$ and for all $t \in [-r^\eta, r^\eta]$, $\hat{F}_{\theta'}^\eta(t) = \sigma_{\theta'}(s_{\theta'}) + t$ if $\sigma_{\theta'}(s_{\theta'}) + t \in (0,1)$ and either 0 or 1 if $\sigma_{\theta'}(s_{\theta'}) + t \leq 0$ or $\sigma_{\theta'}(s_{\theta'}) + t \geq 1$, respectively. For any $t \notin [-r^\eta, r^\eta]$ $\hat{F}_{\theta'}^\eta(t)$ can be chosen arbitrarily, provided it conforms with the properties of a cdf and $\hat{F}_{\theta'}^\eta(t) \to 0$ and $\hat{F}_{\theta'}^\eta(t) \to 1$ as $\eta \to 0$ for $t < -r^n$ and $t > r^n$, respectively. Note that since $r^\eta \to 0$, for each fixed $t \neq 0$, $\lim_{\eta\to 0}\hat{F}_{\theta'}^\eta(t) = 1\{t > 0\}$, and thus $\left\{\hat{\mathbf{F}}^\eta\right\}_\eta$ is vanishing.

Also note that, if $\sigma_{\theta'}(s_{\theta'}) \in (0,1)$, $\hat{F}_{\theta'}^\eta(v_{\theta'}(s;\omega^\eta)) = \sigma_{\theta'}(s_{\theta'}) + v_{\theta'}(s;\omega^\eta)$ for sufficiently small η and thus converges to $\sigma_{\theta'}(s_{\theta'})$. If $\sigma_{\theta'}(s_{\theta'}) = 1$, then $1 \geq \hat{F}_{\theta'}^\eta(v_{\theta'}(s;\omega^\eta)) \geq 1 + v_{\theta'}(s;\omega^\eta)$ and it also converges to $\sigma_{\theta'}(s_{\theta'}) = 1$. If $\sigma_{\theta'}(s_{\theta'}) = 0$ a similar result applies.

Case 2 ($\kappa(\omega;\sigma^*) = \rho$ for all ω): Without loss of generality, suppose that $\mathbb{S}_\theta \subset (0,\infty)$ for all θ. Let $\mathcal{T}_L = \{(\theta,s) : v_\theta(s;\omega^*) < 0 \,\text{or}\, (v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) = 0)\}$, $\mathcal{T}_R = \{(\theta,s) : v_\theta(s;\omega^*) > 0 \,\text{or}\, (v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) = 1)\}$, and $\mathcal{T}_0 = \{(\theta,s) : v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) \in (0,1)\}$. Note that, since (σ^*,ω^*) is a voting equilibrium, then $\sigma_\theta^*(s) = 0$ if $(\theta,s) \in \mathcal{T}_L$ and $\sigma_\theta^*(s) = 1$ if $(\theta,s) \in \mathcal{T}_R$. Define $X_L \equiv \sum_{(\theta,s)\in\mathcal{T}_L}\varphi(\theta)q(s \mid \omega^*)s \geq 0$, $X_R \equiv \sum_{(\theta,s)\in\mathcal{T}_R}\varphi(\theta)q(s \mid \omega^*)\frac{1}{s} \geq 0$, and $X_0 \equiv \sum_{(\theta,s)\in\mathcal{T}_0}\varphi(\theta)q(s \mid \omega^*) \geq 0$. The proof constructs a specific family of perturbations. For all η and all $\theta \in \Theta$ and $s \in \mathbb{S}_\theta$ and for any $(\zeta_L, \zeta_0, \zeta_R)$ let

$$F_\theta^\eta\left(v_\theta\left(s;\omega^*\right)\right) = \begin{cases} \zeta_L s\eta & \text{if}\, v_\theta(s;\omega^*) < 0 \,\text{or}\, (v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) = 0) \\ \sigma_\theta^*(s) + \zeta_0\eta & \text{if}\, \{v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) \in (0,1)\} \\ 1 - \frac{\zeta_R}{s}\eta & \text{if}\, v_\theta(s;\omega^*) > 0 \,\text{or}\, (v_\theta(s;\omega^*) = 0 \,\&\, \sigma_\theta^*(s) = 1) \end{cases}$$

By construction, for all $\zeta_j \in (0,\infty)$,$j = R, L$ and $\zeta_0 \in [0,\infty)$, and for all η sufficiently low, there exists a vanishing family $\{\mathbf{F}^\eta\}_\eta$ that satisfies the above restrictions; note that, by MLRP, for each θ there is at most one signal that satisfies $v_\theta(s;\omega^*) = 0$. Then, since $\omega^* \in (-1,1)$,

$$\begin{aligned} \bar{\kappa}^\eta(\omega^*) - \rho = \bar{\kappa}^\eta(\omega^*) - \kappa(\omega^*;\sigma^*) &= \sum_{(\theta,s)}\varphi(\theta)q(s \mid \omega^*)\left(F_\theta^\eta\left(v_\theta\left(s;\omega^*\right)\right) - \sigma_\theta^*(s)\right) \\ &= \eta\left(-\zeta_R X_R + \zeta_L X_L + \zeta_0 X_0\right). \end{aligned}$$

It is straightforward to check that we can always pick $\zeta_R, \zeta_L, \zeta_0$ such that $-\zeta_R X_R + \zeta_L X_L + \zeta_0 X_0 = 0$ and, therefore, $\bar{\kappa}^\eta(\omega^*) = \rho$ for all η sufficiently small. As in Case 1, by letting $\sigma_\theta^\eta(s) = F_\theta^\eta(v_\theta(s;\omega^\eta))$ for all (θ,s), it follows that σ^η is a limit equilibrium and ω^* its corresponding cutoff for all sufficiently small η. The proof is completed by noting that, by construction, $\lim_{\eta\to 0}\sigma^\eta = \sigma^*$.

Supplementary lemma: increasing strategies . —

Lemma OA. There exists $\bar{\varepsilon}$ such that for all $\varepsilon < \bar{\varepsilon}$: If σ is a limit ε-equilibrium, then it is increasing.

Proof: We use the following notation. Let $x_i \in \{L, R\}$ denote the vote of player i, let $\kappa_i^n(\omega;\xi) \equiv P^n(x_i = R \mid \omega)$ be the probability that player $i = 1, ..., n$ votes for R conditional on the state being ω, and let $\kappa^n(\omega;\xi) \equiv \frac{1}{n}\sum_{i=1}^n \kappa_i^n(\omega;\xi)$ be the average over all players.

Throughout the proof let Ξ' be the set in Definition 8 and fix $\xi \in \Xi'$ and a strategy mapping $\tilde{\boldsymbol{\varsigma}} = \left(\tilde{\boldsymbol{\sigma}}^1(\xi), ..., \tilde{\boldsymbol{\sigma}}^n(\xi), ...\right)$ such that 1.-3. in Definition 8 are satisfied. We drop ξ and $\tilde{\boldsymbol{\varsigma}}$ from the notation, let

$P^n \equiv P^n(\tilde{\boldsymbol{\varsigma}}(\xi))$ and, for each strategy σ_i^n, let $P_{\sigma_i}^n \equiv P^n\left(\sigma_i^n, \tilde{\boldsymbol{\sigma}}_{-i}^n(\xi)\right)$. The proof relies on the following claims; the proofs of the first three claims appear at the end of this section.

Claim OA.1: For all $\delta > 0$ and $\omega \in \Omega$, there exits $n_{\delta,\omega}$ such that for all $n \geq n_{\delta,\omega}$,
$\left|P_{\sigma_i}^n(o = R \mid \omega, s_i) - P_{\sigma_i'}^n\left(o = R \mid \omega, s_i'\right)\right| < \delta$ uniformly over $i, s_i, s_i', \sigma_i^n, {\sigma'}_i^n$.

Claim OA.2: For all $\delta > 0$ there exist n_δ such that for all $n \geq n_\delta$, $|\Delta_i(P^n, s_i) - \Delta_i(P_{\sigma_i}^n, s_i)| < \delta$ uniformly over i, s_i, σ_i^n.

Claim OA.3: There exists $c > 0$ and n_c such that for all $n \geq n_c$, $\Delta_i\left(P_{\sigma_i}^n, s_i'\right) - \Delta_i\left(P_{\sigma_i}^n, s_i\right) \geq c$ for all i and $s_i' > s_i$ such that $\sigma_i^n(s_i') = \sigma_i^n(s_i)$.

Claim OA: There exists $c' > 0$ and $n_{c'}$ such that for all $n \geq n_{c'}$, $\Delta_i\left(P^n, s_i'\right) - \Delta_i\left(P^n, s_i\right) \geq c'$ for all i and $s_i' > s_i$.

PROOF OF CLAIM OA:

Fix any σ_i^n such that $\sigma_i^n(s_i') = \sigma_i^n(s_i)$. By Claims OA.2 and OA.3, for all $n \geq \max\{n_c, n_\delta\}$

$$\begin{aligned}\Delta_i\left(P^n, s_i'\right) - \Delta_i\left(P^n, s_i\right) &\geq \left(\Delta_i\left(P_{\sigma_i}^n, s_i'\right) - \delta\right) - \left(\Delta_i\left(P_{\sigma_i}^n, s_i\right) + \delta\right)\\ &\geq c - 2\delta.\end{aligned}$$

The claim follows by setting $\delta = c/4$ and $c' = c/2 > 0$.

PROOF OF LEMMA OA:

The definition of ε-equilibrium implies that for all i, $s_i' > s_i$, $n \geq n_\varepsilon$,

$$\begin{aligned}\tilde{\sigma}_i^n(s_i') - \tilde{\sigma}_i^n(s_i) \geq\ & F_{\theta_i}\left(\Delta_i\left(P^n, s_i'\right)\right) - F_{\theta_i}\left(\Delta_i\left(P^n, s_i\right)\right) - 2\varepsilon.\\ & + F_{\theta_i}\left(\Delta_i\left(P^n, s_i\right) + c'\right) - F_{\theta_i}\left(\Delta_i\left(P^n, s_i\right) + c'\right),\end{aligned} \tag{B29}$$

where we have added and subtracted the same term to the RHS. Let $c' > 0$ be as defined in Claim OA. Since F_{θ_i} is absolutely continuous, then

$$F_{\theta_i}\left(\Delta_i\left(P^n, s_i\right) + c'\right) - F_{\theta_i}\left(\Delta_i\left(P^n, s_i\right)\right) = \int_{\Delta_i(P^n, s_i)}^{\Delta_i(P^n, s_i) + c'} f_{\theta_i}(t)\, dt \geq c'' > 0,$$

where the inequality follows from A4 and the fact that $c' > 0$. Hence, the sum of the second and fourth terms in the RHS of (B29) is at least $c'' > 0$. By Claim OA, the sum of the first and last terms in the RHS of (B29) is positive. Therefore, for all $i, s_i' > s_i$, $n \geq n_\varepsilon$,

$$\tilde{\sigma}_i^n(s_i') - \tilde{\sigma}_i^n(s_i) \geq c'' - 2\varepsilon > 0.$$

Since $\bar{\sigma}_\theta^n(\xi, \tilde{\boldsymbol{\varsigma}})$ are averages of the strategies, then for all $\theta, s' > s$, and $n \geq n_\varepsilon$, it follows that $\tilde{\sigma}_\theta^n(s') - \tilde{\sigma}_\theta^n(s) \geq c'' - 2\varepsilon$. Since $\lim_{n\to\infty} \|\bar{\sigma}^n - \sigma\| = 0$, then it follows that $\sigma_\theta(s') - \sigma_\theta(s) \geq c'' - 2\varepsilon > 0$, thus establishing that limit ε-equilibrium are increasing as long as $0 < \varepsilon < \bar{\varepsilon} \equiv c''/2 > 0$.

PROOF OF CLAIM OA.1:

The proof is divided into 3 steps.

Step 1. We first show that the probability of being pivotal goes to zero; i.e., for all $\omega \in \Omega$, for all i, $\lim_{n\to\infty} Piv_{\omega,i}^n = 0$, where

$$Piv_{\omega,i}^n \equiv P_1^n(o = R \mid \omega) - P_0^n(o = R \mid \omega),$$

where the "1" and "0" are understood as vectors of the same dimension as α_i. The sub-index "i" indicates that agent i is the one being pivotal.

By simple algebra,

$$Piv^n_{\omega,i} = P^n\left(\sqrt{n}K^n_\omega + \frac{\kappa^n_{i\omega} - 1}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \leq \frac{\sum_{j=1}^n Z^n_{j\omega}}{\sqrt{n}} < \sqrt{n}K^n_\omega + \frac{\kappa^n_{i\omega}}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \mid \omega\right),$$

where $Z^n_{j\omega} \equiv \frac{\{1\{x^n_j=R\}-\kappa^n_{j\omega}\}}{V^n_\omega}$, $V^n_\omega \equiv \sqrt{\frac{1}{n}\sum_{j=1}^n \kappa^n_{j,\omega}\left(1-\kappa^n_{j,\omega}\right)}$, and $K^n_\omega \equiv \frac{\rho-\kappa^n_\omega}{V^n_\omega}$. Note that, for a given n, $\{Z^n_{j\omega}\}_j$ are independent, they have zero mean and unit variance. Moreover, by Step 3 below, $\liminf_{n\to\infty} V^n_\omega > 0$, so that

$$\sum_{j=1}^n E\left[\left|\frac{Z^n_{j\omega}}{\sqrt{n}}\right|^3\right] \leq \frac{2}{\sqrt{n}\,(V^n_\omega)^3} \to 0 \; as \; n \to \infty,$$

Hence by Lindeberg-Feller CLT, it follows that, given ω, $\sum_{j=1}^n \frac{Z^n_{j\omega}}{\sqrt{n}} \Rightarrow N(0,1)$ as $n \to \infty$.

Note also that, $\frac{Z^n_{i\omega}}{\sqrt{n}} \to 0$ a.s. as $n \to \infty$ and this limit is uniform on i.

We divide the remainder of the proof in 3 cases: (a) $\sqrt{n}K^n_\omega \to -\infty$, (b) $\sqrt{n}K^n_\omega \to K \in (-\infty,\infty)$ or (c) $\sqrt{n}K^n_\omega \to \infty$ (if necessary, we take a subsequence that converges, which exists since $(V^n_\omega(\xi))_n$ and $(\kappa^n_\omega(\xi))_n$ are uniformly bounded).

We first explore case (a) (case (c) is symmetrical). Note that, since $\liminf_{n\to\infty} V^n_\omega > 0$, then $\frac{\kappa^n_{i\omega}}{V^n_\omega\sqrt{n}} \to 0$. Therefore, $\sqrt{n}K^n_\omega + \frac{\kappa^n_{i\omega}}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \to -\infty$, (and this limit holds uniformly for $i = 1,...,n$) so that we can take $n \geq n_{M,\epsilon}$ such that $\sqrt{n}K^n_\omega + \frac{\kappa^n_{i\omega}}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \leq -M$, where $\mathcal{L}_N(-M) < 0.5\epsilon$ (where $\mathcal{L}_N$ is the standard Gaussian cdf) for any ϵ. Therefore, for all $\epsilon > 0$ there exists $n_{\epsilon,\omega}$ such that for all $n \geq \max\{n_{\epsilon,\omega}, n_{M,\epsilon}\}$:

$$Piv^n_{\omega,i} \leq P^n\left(\frac{\sum_{j=1}^n Z^n_{j,\omega}}{\sqrt{n}} < -M \mid \omega\right) \leq 0.5\epsilon + \mathcal{L}_N(-M) < \epsilon$$

uniformly over $i = 1,...,n$, where the first inequality follows from the fact that $n \geq n_{M,\epsilon}$ and the second follows from CLT and our choice of M.

For case (b) (i.e., K finite). Let $\delta > 0$ be such that $\mathcal{L}_N(K+\delta) - \mathcal{L}_N(K-\delta) < 0.5\epsilon$. Note that since $\lim_{n\to\infty}(V^n_\omega\sqrt{n})^{-1} = 0$, there exists a $n_{\delta,\omega}$ such that $(V^n_\omega\sqrt{n})^{-1} < 0.5\delta$ for all $n \geq n_{\delta,\omega}$; also since $\frac{Z^n_{i\omega}}{\sqrt{n}} \to 0$ a.s. as $n \to \infty$, we can take $n_{\delta,\omega}$ such that $\left|\frac{Z^n_{i\omega}}{\sqrt{n}}\right| < 0.5\delta$ (note that $n_{\delta,\omega}$ does not depend on i since convergence is uniform on i). Then, it follows for all $\epsilon > 0$, there exists $n_{\epsilon,\omega}$ such that for all $n \geq \max\{n_{\epsilon,\omega}, n_{\delta,\epsilon}\}$:

$$\begin{aligned}
Piv^n_{\omega,i} \leq & P^n\left(\sqrt{n}K^n_\omega - \frac{1}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \leq \frac{\sum_{j=1}^n Z^n_{j\omega}}{\sqrt{n}} < \sqrt{n}K^n_\omega + \frac{1}{V^n_\omega\sqrt{n}} + \frac{Z^n_{i\omega}}{\sqrt{n}} \mid \omega\right) \\
\leq & P^n\left(K - \delta < \frac{\sum_{j=1}^n Z^n_{j\omega}}{\sqrt{n}} \leq K + \delta \mid \omega\right) \\
\leq & 0.5\epsilon + \mathcal{L}_N(K+\delta) - \mathcal{L}_N(K-\delta) < \epsilon,
\end{aligned}$$

where the third inequality follows from the CLT. We showed that for any convergent subsequence $(K^n_\omega)_n$, the associated subsequences of probabilities converge to zero, thus this result must hold for the whole sequence.

Step 2. Note that:

$$\begin{aligned}P^n_{\sigma_i}\left(o=R\mid\omega,s_i\right)=&\sigma_i^n(s_i)P_1^n\left(o=R\mid\omega\right)+(1-\sigma_i^n(s_i))P_0^n\left(o=R\mid\omega\right)\\=&P_0^n\left(o=R\mid\omega\right)\\&+\sigma_i^n(s_i)\left(P_1^n\left(o=R\mid\omega\right)-P_0^n\left(o=R\mid\omega\right)\right)\\\equiv&P^n\left(o=R\mid\omega\right)+\sigma_i^n(s_i)Piv^n_{\omega,i}\end{aligned}$$

Therefore

$$|P^n_{\sigma_i}\left(o=R\mid\omega,s_i\right)-P^n_{\sigma_i'}\left(o=R\mid\omega,s_i'\right)|\leq|\sigma_i^n(s_i)-{\sigma'}_i^n(s_i)|\cdot|Piv^n_{\omega,i}|.$$

By step 1, it follows that for all $n\geq n_{\delta,\omega}$: $|Piv^n_{\omega,i}|\leq\delta$. Since $|\sigma_i^n(s_i)-{\sigma'}_i^n(s_i)|\leq 1$ the desired result follows.

Step 3. We now show that for all $\omega\in\Omega$,

$$\liminf_{n\to\infty}\frac{1}{n}\sum_{j=1}^{n}\kappa^n_{j\omega}\left(1-\kappa^n_{j\omega}\right)>0. \tag{B30}$$

Fix any n and $j\leq n$. By assumption, $\sigma_j^n\left(s_j\right)\in\left[F_j\left(-2K\right),F_j\left(2K\right)\right]\subset(0,1)$ for all s_j. Therefore, $0<\kappa^n_{j\omega}<1$ for all ω, thus implying equation (B30).

PROOF OF CLAIM OA.2:

We prove that

$$\lim_{n\to\infty}\Big(E_{P^n}\left(u_{\theta_i}(R,W)\mid o=R,S=s_i\right)-E_{P^n_{\sigma_i}}\left(u_{\theta_i}(R,W)\mid o=R,S=s_i\right)\Big)=0;$$

the proof for $o=L$ is similar and therefore omitted. We first show that, for all i,s_i,σ_i,

$$E_{P^n_{\sigma_i}}\left(u_{\theta_i}(R,W)\mid o=R,S=s_i\right)=\frac{\int_\Omega P^n_{\sigma_i}\left(o=R\mid W,s_i\right)q_{\theta_i}(s_i\mid W)u_{\theta_i}(R,W)G(dW)}{\int_\Omega P^n_{\sigma_i}\left(o=R\mid W,s_i\right)q_{\theta_i}(s_i\mid W)G(dW)}$$

is well-defined for sufficiently large n. Fix any i. A3(ii) and the fact that $\boldsymbol{\varsigma}$ is asymptotically interior imply that there exists $\overline{n}$ such that for all $n\geq\overline{n}$, there exists s_i^* such that

$$P^n(o=R,s_i^*)=\int_\Omega P^n(o=R\mid W,s_i^*)q_{\theta_i}\left(s_i^*\mid W\right)G(dW)\geq c>0,$$

which implies that $\int_\Omega P^n(o=R\mid W,s_i^*)G(dW)\geq c>0$. By Claim OA.1, for each s_i,σ_i^n, $P^n\left(o=R\mid\omega,s_i^*\right)-P^n_{\sigma_i}\left(o=R\mid\omega,s_i\right)$ converges to zero as $n\to\infty$. Since both probabilities are bounded by one, then the dominated convergence theorem implies that $\int_\Omega\left(P^n\left(o=R\mid W,s_i^*\right)-P^n_{\sigma_i}\left(o=R\mid W,s_i\right)\right)G(dW)\to 0$ as $n\to\infty$, uniformly over σ_i. Therefore, there exists $n_{.5c}$ such that $\sup_{\sigma_i}|\int_\Omega[P^n\left(o=R\mid W,s_i^*\right)-P^n_{\sigma_i}\left(o=R\mid W,s_i\right)]G(dW)|<.5c$ for all $n\geq n_{.5c}$. So for all $n\geq\max\bar{n},n_{.5c}\equiv\overline{n}_c$,

$$\int_\Omega P^n_{\sigma_i}\left(o=R\mid W,s_i\right)q_{\theta_i}\left(s_i\mid W\right)G(dW)\geq d\int_\Omega P^n_{\sigma_i}\left(o=R\mid W,s_i\right)G(dW)>.5dc>0.$$

Hence, $E_{P^n_{\sigma_i}}\left(u_{\theta_i}(R,W)\mid o=R,S=s_i\right)$ is well defined.

By simple algebra, and letting $\Delta P^n_{\sigma_i}(R,\omega,s_i)\equiv P^n\left(o=R\mid\omega,s_i\right)-P^n_{\sigma_i}\left(o=R\mid\omega,s_i\right)$,

$$
\begin{aligned}
&\left| E_{P^n}\left(u_{\theta_i}(R,W) \mid o=R, S=s_i\right) - E_{P^n_{\sigma_i}}\left(u_{\theta_i}(R,W) \mid o=R, S=s_i\right)\right| \\
&\leq \frac{\left|\int_\Omega \Delta P^n_{\sigma_i}(R,W,s_i) q_{\theta_i}(s_i \mid W) u_{\theta_i}(R,W)\, G(dW)\right|}{\int_\Omega P^n_{\sigma_i}(o=R \mid W)\, q_{\theta_i}(s_i \mid W) G(dW)} \\
&+ \frac{\left|\int_\Omega \Delta P^n_{\sigma_i}(R,W,s_i) q_{\theta_i}(s_i \mid W) G(dW)\right| \int_\Omega P^n(o=R \mid W)\, q_{\theta_i}(s_i \mid W) u_{\theta_i}(R,W)\, G(dW)}{\int_\Omega P^n(o=R \mid W)\, q_{\theta_i}(s_i \mid W) G(dW) \int_\Omega P^n_{\sigma_i}(o=R \mid W)\, q_{\theta_i}(s_i \mid W) G(dW)}
\end{aligned}
$$

To establish the desired result, it is sufficient to show that each of the two absolute value terms in the numerator of the second and third line converge to zero as $n \to \infty$. However, this result follows by the dominated convergence theorem since $|u_{\theta_i}(R,\omega)| < K$, $q_{\theta_i}(s|\omega) \leq 1$, and pointwise convergence (for each ω) is obtained by Claim OA.1.

PROOF OF CLAIM OA.3:

For each $O \in \{R, L\}$: Let $G^n_{\sigma_i}(\omega \mid O, s_i) \equiv P^n_{\sigma_i}(\{W \leq \omega\} \mid o=O, s_i)$ denote the cdf of ω conditional on $o = O$ and s_i, and let $g^n_{\sigma_i}(\omega \mid O, s_i) \equiv P^n_{\sigma_i}(d\omega \mid o = O, s_i)$ denote the density. Let $\Delta g^n_{\sigma_i}(\omega \mid O, s_i', s_i) \equiv g^n_{\sigma_i}(\omega \mid O, s_i') - g^n_{\sigma_i}(\omega \mid O, s_i)$ and $\Delta G^n_{\sigma_i}(\omega \mid O, s_i', s_i) \equiv G^n_{\sigma_i}(\omega \mid O, s_i') - G^n_{\sigma_i}(\omega \mid O, s_i)$.

Then

$$
\begin{aligned}
\Delta_i\left(P^n_{\sigma_i}, s_i'\right) - \Delta_i\left(P^n_{\sigma_i}, s_i\right) &= \int_\Omega \left(u_{\theta_i}(R,W)\Delta g^n_{\sigma_i}(W \mid R, s_i', s_i) - u_{\theta_i}(L,W)\Delta g^n_{\sigma_i}(W \mid L, s'_i, s_i)\right) dW \\
&= \int_\Omega \left(\frac{du_{\theta_i}}{d\omega}(R,W)\Delta G^n_{\sigma_i}(W \mid R, s_i, s_i') - \frac{du_{\theta_i}}{d\omega}(L,W)\Delta G^n_{\sigma_i}(W \mid L, s_i, s_i')\right) dW \\
&\geq \int_{\Omega^n \subset \Omega} \frac{du_{\theta_i}}{d\omega}(R,W)\Delta G^n_{\sigma_i}(W \mid R, s_i, s_i') dW \\
&\geq c_M \int_{\Omega^n \subset \Omega} \frac{du_{\theta_i}}{d\omega}(R,W) dW \\
&\geq c_m \cdot c_M \inf_{W \in \bar{\Omega}} \frac{du_{\theta_i}}{d\omega}(R,W) \\
&\equiv c > 0
\end{aligned}
$$

for all $n \geq n'$ (where Ω^n, $c_m \cdot c_M > 0$, and n' are all defined in Claim OA.3.1 below), where the first line follows by definition, the second by integration by parts (note how the signals are inverted), the third by Claim OA.3.1(i) (see below) and the facts that that $\frac{du_{\theta_i}}{d\omega}(R,\omega) > 0$ and $\frac{du_{\theta_i}}{d\omega}(L,\omega) < 0$ for all ω, the fourth by Claim OA.3.1(ii). Finally, for the fifth line, let $\bar{\Omega} = \Omega \setminus \cup_{i=1}^N (\omega_i - \epsilon, \omega_i + \epsilon)$ where $(\omega_1, ..., \omega_N)$ are the discontinuity points of $\frac{du_{\theta_i}}{d\omega}(R, \cdot)$; by assumption there are finitely many, so $N < \infty$ and $\epsilon > 0$ is chosen such that $\epsilon < \min_{i \neq j} |\omega_i - \omega_j|$. It is easy to see that $\bar{\Omega}$ is compact and over it, $\frac{du_{\theta_i}}{d\omega}(R, \cdot)$ is well-defined and continuous. Since $c_m \cdot c_M > 0$ and $\inf_{\omega \in \bar{\Omega}} \frac{du_{\theta_i}}{d\omega}(R,\omega) = \min_{\omega \in \bar{\Omega}} \frac{du_{\theta_i}}{d\omega}(R,\omega) > 0$ where (because u_{θ_i} is continuously differentiable in $\bar{\Omega}$ and $\frac{du_{\theta_i}}{d\omega}(R,\omega) > 0$ for all ω).

Claim OA.3.1: For all i and $s_i' > s_i$ such that $\sigma_i^n(s_i) = \sigma_i^n(s_i')$: (i) For all n, $\Delta G^n_{\sigma_i}(\omega \mid O, s_i, s_i') \geq 0$ for all ω and $O \in \{R, L\}$; (ii) There exists n', $c_M > 0$, and $(\Omega^n)_n$ with $\Omega^n = [l_n, u_n] \subseteq \Omega$ and $\liminf_{n \to \infty} u_n - l_n = \beta_2 > 0$ such that, for all $n \geq n'$ and all $\tilde{\omega} \in \Omega^n \setminus \{-1, 1\}$,

$$\Delta G^n_{\sigma_i}(\tilde{\omega} \mid R, s_i, s_i') \geq c_M.$$

PROOF OF CLAIM OA.3.1:

There exists $z > 0$ such that for all n and all $\omega' > \omega$,

$$g^n_{\sigma_i}(\omega' \mid O, s'_i)g^n_{\sigma_i}(\omega \mid O, s_i) - g^n_{\sigma_i}(\omega' \mid O, s_i)g^n_{\sigma_i}(\omega \mid O, s'_i)$$

$$= \frac{P^n_{\sigma_i}(O \mid \omega', s_i)P^n_{\sigma_i}(O \mid \omega, s_i)g(\omega')g(\omega)}{P^n_{\sigma_i}(O, s'_i)P^n_{\sigma_i}(O, s_i)} \left[q_{\theta_i}\left(s'_i \mid \omega'\right) q_{\theta_i}\left(s_i \mid \omega\right) - q_{\theta_i}\left(s_i \mid \omega'\right) q_{\theta_i}\left(s'_i \mid \omega\right)\right]$$

$$\geq z \frac{P^n_{\sigma_i}(O \mid \omega', s_i)P^n_{\sigma_i}(O \mid \omega, s_i)g(\omega')g(\omega)q_{\theta_i}\left(s'_i \mid \omega\right) q_{\theta_i}\left(s_i \mid \omega\right)(\omega' - \omega)}{P^n_{\sigma_i}(O, s'_i)P^n_{\sigma_i}(O, s_i)}$$

(B31) $$\geq 0$$

where the first line uses the fact that $P^n_{\sigma_i}(O \mid \hat{\omega}, s_i) = P^n_{\sigma_i}(O \mid \hat{\omega}, s'_i)$ for all $\hat{\omega}$ (because of conditional independence and the fact that $\sigma^n_i(s_i) = \sigma^n_i(s'_i)$), the second line follows from A5, and the third line follows because $z > 0$ and $\omega' > \omega$. Therefore, it follows from Milgrom (1981, Proposition 1) that, for all n, $\Delta G^n_{\sigma_i}(\omega \mid O, s_i, s'_i) \geq 0$ for all ω.

(ii) From the proof of Claim OA.2, there exists n' and $c' > 0$ such that, for all $n \geq n'$,

$$\int_{\Omega} P^n_{\sigma_i}(o = R \mid W, s_i)G(dW) \geq c'$$

for all i, σ_i, s_i. For $a \in (0, 1)$, let

$$\omega^n_a = min\left\{\omega' : \int_{W \leq \omega'} P^n_{\sigma_i}(o = R \mid W, s_i)G(dW) \geq a \cdot c'\right\} \in \Omega.$$

Fix any $n \geq n'$. Then

$$c'/4 = \int_{\omega^n_{0.25} \leq W \leq \omega^n_{0.50}} P^n_{\sigma_i}(o = R \mid W, s_i)G(dW) \leq G\left(\omega^n_{0.50}\right) - G\left(\omega^n_{0.25}\right).$$

Therefore, the fact that G has no mass points implies that there exists $c_L > 0$ such that $\omega^n_{0.50} - \omega^n_{0.25} \geq c_L$. A similar argument establishes that here exists $c_R > 0$ such that $\omega^n_{0.75} - \omega^n_{0.50} \geq c_R$.

Let $\Omega^n = \left[\omega^n_{0.50} - c_m/2, \omega^n_{0.50} + c_m/2\right]$, where $c_m \equiv min\{c_L, c_R\} > 0$. Then, $u_n - l_n = c_m > 0$. In addition, fix any $\tilde{\omega} \in \Omega^n$. Then, by construction,

(B32) $$\int_{\omega < \tilde{\omega} - c_m/2} P^n_{\sigma_i}(o = R \mid W, s_i)G(dW) \geq c'/4$$

and

(B33) $$\int_{\omega > \tilde{\omega} + c_m/2} P^n_{\sigma_i}(o = R \mid W, s_i)G(dW) \geq c'/4.$$

By integrating each side of (B31) twice, first with respect to $G(d\omega)$ over $\omega \leq \tilde{\omega}$ and second with respect to $G(d\omega')$ over $\omega' > \tilde{\omega}$, we obtain

$$\Delta G^n_{\sigma_i}(\tilde{\omega} \mid R, s_i, s_i') =$$
$$= \frac{z}{P^n_{\sigma_i}(R, s_i')P^n_{\sigma_i}(R, s_i)}$$
$$\times \int_{W'>\tilde{\omega}} \int_{W<\tilde{\omega}} P^n_{\sigma_i}(R \mid \omega', s_i)P^n_{\sigma_i}(R \mid W, s_i)g(W')g(W)q_{\theta_i}\left(s_i' \mid W\right) q_{\theta_i}\left(s_i \mid W\right)(W' - W)dG(W)dG(W')$$
$$\geq z \int_{W'>\tilde{\omega}+\frac{c_m}{2}} \int_{W<\tilde{\omega}-\frac{c_m}{2}} P^n_{\sigma_i}(R \mid W', s_i)P^n_{\sigma_i}(R \mid W, s_i)g(W')g(W)q_{\theta_i}\left(s_i' \mid W\right) q_{\theta_i}\left(s_i \mid W\right)(W' - W)dG(W)dG(W')$$
$$\geq z \cdot c_m \cdot d^2 \int_{W'>\tilde{\omega}+\frac{c_m}{2}} P^n_{\sigma_i}(R \mid W', s_i)G(dW') \int_{W<\tilde{\omega}-\frac{c_m}{2}} P^n_{\sigma_i}(R \mid W, s_i)G(dW)$$
$$\geq z \cdot c_m \cdot d^2 \cdot \left(\frac{c'}{4}\right)^2 \equiv c_M > 0,$$

where the first inequality follows from $P^n_{\sigma_i}(R, s_i')P^n_{\sigma_i}(R, s_i) \leq 1$, the second from A3, and the third from (B32) and (B33).

Partially cursed equilibrium

In this section we characterize Eyster and Rabin's partially cursed equilibrium in the voting game (see Eyster and Rabin (2005) for the definition) as the number of voters goes to infinity. For simplicity, we restrict attention to the environment with homogeneous voters introduced in Section I.

For any $\chi \in (0, 1)$, define

$$v^{PC}(s; \omega) \equiv (1 - \chi)(u(R, \omega) - u(L, \omega)) + \chi\left(E\left(u(R, W) \mid s\right) - E\left(u(L, W) \mid s\right)\right).$$

This expression will be shown to represent the beliefs of a partially cursed voter who observes signal s in a large election in which the equilibrium cutoff is ω. Note that $\chi = 0$ corresponds to beliefs under Nash equilibrium (NE) and $\chi = 1$ corresponds to sincere voting (SV).

In a manner analogous to what we did for RVE, define the *partially cursed personal cutoff*

$$c^{PC}(s) \equiv \arg\min_{\omega \in \Omega} \left| v^{PC}(s; \omega) \right|$$

as well as

$$\bar{\kappa}_{PC}(\omega) \equiv \sum_{\{s: c^{PC}(s) < \omega\}} q\left(s \mid \omega\right),$$

which may be interpreted as the proportion of players that vote for R in state ω when the cutoff is also given by ω.[34]

Finally, define $\omega_{PC} \equiv \bar{\kappa}^{-1}_{PC}(\rho)$. The next result shows that ω_{PC} is essentially the election cutoff if there are a large number of partially cursed voters.

THEOREM 5: *Suppose that voters play a partially cursed equilibrium with $\chi \in (0, 1)$. For all $\varepsilon > 0$, there is an n_ε such that for all $n' > n_\varepsilon$, the following holds: if $\omega < \omega_{PC} - \varepsilon$, then L is elected with probability greater than $1 - \varepsilon$; if $\omega > \omega_{PC} + \varepsilon$, then R is elected with probability greater than $1 - \varepsilon$.*

PROOF:

We provide only a sketch of the proof and focus on equilibria in increasing strategies (even in the limit). Not surprisingly, the arguments are very similar to those used by Feddersen and Pesendorfer

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

(1997) to characterize NE in large elections.[35] First, it is easy to see that the definition of partially cursed equilibrium (see Eyster and Rabin 2005) implies that, for a given strategy profile σ^{-i} of other players, voter i's belief after observing signal s is a convex combination of NE and SV beliefs, i.e.,

(C1)
$$(1-\chi)\left[E_{n,\sigma^{-i}}\left(u(R,W)\mid pivotal,s\right)-E_{n,\sigma^{-i}}\left(u(L,W)\mid pivotal,s\right)\right]+\chi\left[E\left(u(R,W)\mid s\right)-E\left(u(L,W)\mid s\right)\right],$$

where the expectation in the first term depends on the number of voters and the strategies of these other voters. Second, because strategies are increasing, then the election cutoff in a large election will be uniquely given by some state ω at which the proportion of people voting for R equals the electoral rule ρ. Thus, conditional on being pivotal, a voter in a large election puts probability 1 on a neighborhood around the election cutoff ω, and so (C1) converges to $v^{PC}(s;\omega)$. By optimality, a voter who observes s votes R or L depending on the sign of $v^{PC}(s;\omega)$, and so $\overline{\kappa}_{PC}(\omega)$ is the proportion of votes for R at the election cutoff ω. But, as mentioned earlier, the cutoff must have the property that $\overline{\kappa}_{PC}(\omega)=\rho$; thus, the limiting cutoff is the unique state that solves this equation.

As mentioned in the text, partially cursed equilibrium also provides a middle ground between NE and SV. It does not, however, capture the failure to account for selection that motivates this paper. Not surprisingly, the insights about sample selection that emerge from RVE and not from either NE or SV, also don't emerge from a convex combination of NE and SV. We illustrate this point using Example 2 from the text.

Example 2 (continued from pg. 13). As mentioned in the text, the cutoff under NE is the efficient cutoff of zero and the cutoff under sincere voting can be either positive or negative depending on the two cases that depend on the primitives:

Case (1) If expression (5) is positive for $s\in\{s^R,s^M\}$ and negative for s^L, then the election cutoff is lower than zero and given by the intersection of $q(\{s^R,s^M\}\mid\cdot)$ and ρ^*.

Case (2) If expression (5) is positive for s^R and negative for $s\in\{s^M,s^L\}$, then the election cutoff is greater than zero and given by the intersection of $q(\{s^R\}\mid\cdot)$ and ρ^*.

It is straightforward to check that the partially cursed equilibrium cutoff is between the Nash and Sincere Voting cutoffs. In Case (1), the partially cursed equilibrium cutoff is greater than zero and converges to zero as χ converges to zero. In Case (2), the partially cursed equilibrium cutoff is less than zero and converges to zero as χ converges to zero. Note the difference with RVE, where the equilibrium cutoff is always less than zero and bounded away from zero due to the selection problem described in the main text. □

[35] The proof is actually much easier than their proof. The reason is that equilibrium strategies are always strictly monotone, even in the limit. In the case of NE (i.e., $\chi=1$), the main challenge is to show that, despite the fact that strategies are becoming flat and voters are almost ignoring their private information, this is happening at a sufficiently slow rate such that information is still being aggregated.