A Theory of Momentum in Sequential Voting*

S. Nageeb M. Ali†
Stanford University

Navin Kartik‡
UC San Diego

September 2006

Abstract

Elections with sequential voting, such as presidential primaries, are widely-thought to possess momentum effects, where the choices of early voters influence the behavior of later voters. Momentum can be subtle: it may take time to build, and depends on how candidates perform in each stage relative to expectations. This paper develops a rational theory of momentum in sequential elections that accounts for these phenomena. We analyze an election with two candidates in which some voters are uncertain about which candidate is more desirable. Voters obtain private signals and vote in a sequence, observing the history of votes at each point. We show that, regardless of the voting rule, voters can herd on a candidate with positive probability, and such a "bandwagon" can occur with probability approaching one in large electorates. Our theory is distinct from the standard information cascades literature because voting is a collective decision problem, and consequently voters have forward-looking incentives to consider the actions of those after them.

Keywords: sequential voting, sincere voting, information aggregation, bandwagons, momentum, information cascades, herd behavior

J.E.L. Classification: C7, D72

*We thank Llumi Arias, Susan Athey, Doug Bernheim, Vince Crawford, Nir Jaimovich, Fahad Khalil, Jon Levin, Mark Machina, Marc Meredith, Marco Ottaviani, Andres Santos, Joel Sobel, and numerous seminar and conference audiences for helpful comments and discussions. Ali is grateful to the Institute for Humane Studies, the John M. Olin Foundation, and Stanford's Economics Department for their financial support.

†Email: nageeb@gmail.com; Web: http://www.stanford.edu/∼nageeb; Address: 579 Serra Mall, Stanford, CA 94305.

‡Email: nkartik@ucsd.edu; Web: http://econ.ucsd.edu/∼nkartik; Address: Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093-0508.

"...when New Yorkers go to vote next Tuesday, they cannot help but be influenced by Kerry's victories in Wisconsin last week. Surely those Wisconsinites knew something, and if so many of them voted for Kerry, then he must be a decent candidate."

— Duncan Watts in *Slate* Magazine

1 Introduction

Many elections take place over time. The most prominent example lies at the heart of the American presidential selection process: the primaries. A series of elections by which a party nominates its candidate for the general presidential election, the primaries are held sequentially across states over a few months. On a smaller scale, but also explicitly sequential, are the roll-call voting mechanisms used by city councils and Congressional bodies. A more subtle example is the general U.S. presidential election itself, where the early closing of polls in some states introduces a temporal element into voting.

In contrast, most theoretical models of voting are static. The distinction between simultaneous and sequential elections is not just of theoretical interest, but also relevant to policy. It is often suggested that in sequential elections, voters condition their choices on the acts of prior voters. Such history dependence is believed to result in *momentum* effects: the very fact that a particular alternative is leading in initial voting rounds may induce some later voters to select it who would have otherwise voted differently. Moreover, voting behavior in primaries suggests that candidates are judged by how they perform relative to expectations: a surprisingly good performance in an early primary may generate more momentum than an anticipated victory (Popkin, 1991).

The beliefs in momentum and performance relative to expectations have shaped electoral policy and strategy. Some U.S. states aim to hold their primaries early in the process,[1] campaign funds and media attention are disproportionately devoted to initial primaries,[2] and candidates strategically attempt to gain surprise victories. For example, Hamilton Jordan, who would become Jimmy Carter's White House Chief of Staff, outlined the importance of a surprise victory in New Hampshire in a memorandum two years before Carter's campaign: "...a strong surprise in New Hampshire should be our goal, which would have a tremendous impact on successive primaries..." This assessment is supported by the analysis of Bartels (1988), who simulates Carter's national popularity prior to his victory in New Hampshire and predicts that Carter would have lost the 1976 nomination to George Wallace had the primaries been held

[1]For instance, since 1977, New Hampshire law has stated that its primary is to be the first in the nation. As a result, the state has had to move its primary, originally in March, earlier in the year to remain the first. New Hampshire's primary was held on February 20 in 1996, February 1 in 2000, and then January 27 in 2004 to compete with front-loaded primaries in other states.

[2]Bartels (1988) and Gurian (1986) document how a substantial share of the media's attention and candidates' campaign resources are devoted to New Hampshire even though that state accounts for merely 4 out of 538 total electoral votes.

simultaneously in all states.

The 1984 Democratic primaries featuring Gary Hart and Walter Mondale exemplifies that voters judge candidates' performances relative to expectations. Expected to have support in Iowa, Mondale's victory with close to 50% of the caucus votes in this state was largely overshadowed by Gary Hart's ability to garner about 17% of the votes. Even though Hart's performance in Iowa was far inferior to Mondale's on an absolute scale, Hart had performed much better than expected. This garnered him momentum, and Hart proceeded to win the following primaries in Vermont and New Hampshire (Bartels, 1988).

This paper provides a positive theory of voter behavior in sequential elections, generating dynamics with momentum effects. Specifically, we consider a sequential version of a canonical election environment (Feddersen and Pesendorfer, 1996, 1997). There are two candidates, and a finite population of voters who vote in an exogenously fixed sequence, each observing the entire history of prior votes. The candidate who receives the majority of votes wins the election. There are two kinds of voters: Neutrals and Partisans. Neutrals desire to elect the "correct" or better candidate, which depends on the realization of an unknown state variable. Partisans, on the other hand, wish to elect their exogenously preferred candidate regardless of the state. Whether a voter is Neutral or Partisan is her private information, and each voter receives a private binary signal that contains some information about the unknown state.

In this setting, we formalize an informational theory of voting that is built on the simple but powerful logic that if (some) initial voters use their private information in deciding how to vote, the voting history provides useful information to later voters. We show that a simple form of history-dependent strategies constitutes fully rational behavior. This equilibrium generates rich momentum effects where a leading candidate is judged relative to his expected partisan support, and some voting histories can cause future voters to entirely ignore their private information, with Neutrals simply joining "bandwagons" for one of the candidates.

The notion of rational bandwagons is reminiscent of the the herd-cascade literature initiated by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). However, the strategic issues involved are very different. In standard herd-cascade models, a player's payoffs do not directly depend on the actions of others; the only externality is informational. Thus, agents have no strategic reason to consider the impact of their actions on future agents' choices, and optimal behavior for any agent is necessarily that of a backward-looking Bayesian. In contrast, an election is a game with payoff interdependencies: a voter's payoffs are determined by the collective outcome, which can depend on the votes of those before and after her. A strategic voter determines her optimal action by conditioning on being pivotal. A vote can be pivotal in two distinct ways in sequential elections. First, just as in simultaneous voting, it can affect the electoral outcome in the event that all other votes are tied. Second, and unique to sequential voting, a vote can reveal some information to future voters, and therefore induce different (distributions over) future votes. Therefore, sequential voting inherently presents *forward-looking* incentives to players, and the possibility of bandwagons or information cascades in our environment is distinct from the standard herd-cascade literature.

The main contribution of this paper is to characterize a Perfect Bayesian equilibrium that generates rational herding in a sequential election. In the equilibrium we characterize, *Posterior-Based Voting* (PBV), a voter uses all available information—her prior, signal, and the observed history of votes—to form her expectation of which candidate is better for her, and votes for this candidate. Importantly, this behavior does not require a voter to condition on being pivotal. That is, when playing PBV, each voter does in fact behave like a backward-looking Bayesian of the standard herding model. We show that in spite of the forward-looking incentives in our model, this is optimal behavior for a strategic voter who conditions on being pivotal (given that others are playing PBV). Moreover, for generic parameters, PBV is a strict equilibrium; thus, it does not rely on the choice of how to resolve indifferences among players, and is robust to small perturbations of the model.

Though simple, PBV generates rich dynamics. Play is inherently dependent on history, and voters update their beliefs about which candidate is better based upon the voting history. Accordingly, PBV generates momentum effects, where the prospects for a candidate can ebb and flow during the course of an election as a function of the voting history. The possibility of partisanship in our model implies that in the PBV equilibium, the performance of a candidate is judged relative to his expected partisan support. Thus, even though a candidate may be trailing in the absolute vote lead on the equilibrium path, he may have momentum because he has outperformed expectations. In large elections, bandwagons for a candidate form with very high probability, where Neutral voters vote for a candidate independent on their private information. Indeed, a bandwagon can even form on a candidate who is trailing in the election.

The essence of why PBV is an equilibrium is that when other voters follow PBV, if a deviation were to change the outcome of the election, the subsequent profile of votes does not generate enough public information to outweigh a voter's posterior that is computed from her observed history and private signal. In fact, this logic is so pervasive that PBV is a (generically strict) equilibrium for *all* anonymous and monotonic voting rules, and not simply majority rule. Momentum can therefore be the outcome of equilibrium behavior independent of the voting rule. Interestingly, we find that if voters play PBV, a large class of voting rules are asymptotically equivalent in terms of the electoral outcomes they induce. Furthermore, PBV is robust to uncertainties over the size of the voting population and remains an equilibrium regardless of the beliefs voters have about the final population size. This is attractive in the context of large elections.

The notion that momentum can be the outcome of a fully rational sequential voting model is novel.[3] Some of the early informal literature in political science invokes assumptions such as a psychological desire to vote for the eventual winner to explain momentum (e.g. Berelson, Lazarsfeld, and McPhee, 1954); this has recently been formalized by Callander (2004a) in

[3]To our knowledge, Fey (2000) and Wit (1997) were the first to formally examine the possibility of rational herding in sequential voting, and both authors concluded that under "reasonable restrictions", herds could not arise. We postpone a a detailed comparison to Section 4.4.

a model with an infinite population of voters, each of whom is partially motivated to elect the right candidate, and partially motivated to simply vote for the eventual winner. While his analysis is innovative and presents interesting comparisons with analogous simultaneous elections (Callander, 2004b), we believe that it is important to establish a non-behaviorally-based benchmark model of momentum, especially because the empirical evidence on whether voters possess conformist motivations is inconclusive.[4]

In an important contribution, Dekel and Piccione (2000) study a model of strategic voting and show that under some assumptions, a class of symmetric equilibria of their simultaneous voting game remain equilibria of the corresponding sequential voting game. Naturally, such equilibria feature history-independent voting. This result has sometimes been misinterpreted as implying that when a voter conditions on being pivotal in a sequential voting environment, the observed history of votes is irrelevant to her voting decision. However, this is only true when other voters are playing history-independent strategies. In the PBV equilibrium we derive, strategic voters *do* learn from previous voters *and* they act fully rationally, voting as if they are pivotal. This illustrates the more general possibility that when other voters play history-dependent strategies, a voter learns from the history how she can be pivotal, and this may be useful for her voting decision. Our findings therefore do not contradict Dekel and Piccione (2000); rather, we identify an appealing equilibrium of the sequential election that has no counterpart in simultaneous elections. While Dekel and Piccione (2000) has important normative implications of whether sequential elections *can* aggregate information efficiently in certain environments, our main interest is to provide a positive description of behavior that can account for the rich momentum effects exhibited in real world elections.

Though the informational model here contributes to the understanding of electoral dynamics, we do not wish to suggest that it captures the whole story. In our model, we abstract from many important institutional details such as campaign finance and media attention.[5] Having said that, it is difficult to shed light on why both financial and media resources are devoted to the first few elections without making specific assumptions about voting behavior. Insofar as the purpose of many elections—especially primaries—is to provide and aggregate information, we believe that these institutional details should be embedded in an informational model similar to the one here. Our work may be viewed as a first step towards understanding the role that these institutions play in voting behavior.

The plan for the remainder of the paper is as follows. Section 2 lays out the model, and Section 3 derives the main results about PBV strategies and equilibrium. We discuss various implications and extensions of our analysis in Section 4. Section 5 concludes. All formal proofs are deferred to the Appendix.

[4]See the discussion in Bartels (1988, pp. 108–112). Kenney and Rice (1994) test the strength of various explanations for momentum, including both the preference for conformity theory and an informational theory similar to the one proposed here. Studying the 1988 Republican primaries, they find that voters acted in ways that are consistent with both theories.

[5]Klumpp and Polborn (2006) offer a model of momentum in this vein, examining the behavior of candidates who must choose how to allocate a fixed pie of campaign funds across states.

2 Model

We consider a voting game with a finite population of n voters, where n is odd. Voters vote for one of two candidates, L or R, in a fixed sequential order, one at a time. We label the voters $1, \ldots, n$, where without loss of generality, a lower numbered voter votes earlier in the sequence. Each voter observes the entire history of votes when it is her turn to vote. The winner of the election, denoted $W \in \{L, R\}$, is selected by simple majority rule. The state of the world, $\omega \in \{L, R\}$, is unknown, but individuals share a common prior over the possible states, and $\pi > \frac{1}{2}$ is the ex-ante probability of state L. Before voting, each voter i receives a private signal, $s_i \in \{l, r\}$, drawn from a Bernoulli distribution with precision γ (i.e., $\Pr(s_i = l|\omega = L) = \Pr(s_i = r|\omega = R) = \gamma$), with $\gamma > \pi$.[6] Individual signals are drawn independently conditional on the state.

In addition to being privately informed about her signal, a voter also has private information about her preferences: she is either an L-partisan (L_p), a Neutral (N), or an R-partisan (R_p). We denote this *preference type* of voter i by t_i. Each voter's preference type is drawn independently from the same distribution, which assigns probability $\tau_L > 0$ to preference type L_p, probability $\tau_R > 0$ to R_p, and probability $\tau_N > 0$ to the Neutral type, N, where $\tau_L + \tau_R + \tau_N = 1$. The preference ordering over candidates is state dependent for Neutrals, but state independent for Partisans. Specifically, payoffs for voter i are defined by the function $u(t_i, W, \omega)$ as follows:

$$\begin{aligned}
u(L_p, W, \omega) &= \mathbf{1}_{\{W=L\}} \text{ for } \omega \in \{L, R\} \\
u(R_p, W, \omega) &= \mathbf{1}_{\{W=R\}} \text{ for } \omega \in \{L, R\} \\
u(N, L, L) &= u(N, R, R) = 1 \\
u(N, L, R) &= u(N, R, L) = 0
\end{aligned}$$

Therefore, a voter of preference type C_p ($C \in \{L, R\}$) is a Partisan for candidate C, and desires this candidate to be elected regardless of the state of the world. A Neutral voter, on the other hand, would like to elect candidate $C \in \{L, R\}$ if and only if that candidate is the better one, i.e. if the state $\omega = C$. Note that each voter cares about her individual vote only instrumentally, through it's influence on the winner of the election.

We now clarify the role of two modeling choices.

Partisans. The Partisan types here are analogous to those in a number of papers in the literature, such as Feddersen and Pesendorfer (1996, 1997) and Feddersen and Sandroni (2006). Nevertheless, Partisans are not necessary for the existence of a history-dependent equilibrium. We analyze the game without Partisans—pure common value elections—in Section 4.4 and show that Posterior-Based Voting remains an equilibrium of that game. While some of the sequential voting literature has restricted attention to the case of pure common values, we believe that

[6] This implies that any individual's signal is more informative than the prior. Our analysis will carry over with obvious changes to cases where the signal precision is asymmetric across states of the world.

the presence of Partisanship is relevant both theoretically and in practice. Partisans introduce private values into the electoral setting, an element that is important in real-world elections.[7] Although we have formalized Partisans as those without a common value element to their preferences whatsoever, this is only for expositional convenience. All that is necessary for our results is that, under complete information, an L-Partisan needs at least three net signals in favor of candidate R to prefer electing candidate R over L (and analogously for an R-Partisan), independently of the population size. In contrast, a Neutral needs only one net signal in favor of a candidate to prefer that candidate being elected. Partisans, therefore, may have a common value component to their preferences; it is simply that their preferences (or priors) are biased in favor of one candidate, although not necessarily to the extent that sufficient information cannot change their views. The presence of Partisans allows our model to generate rich momentum effects, including the feature that a candidate's performance is judged relative to expectations in a non-trivial way.

Information Structure. The binary information structure chosen here is the canonical focus in both the voting literature (e.g. Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1996) and the social learning literature (e.g. Bikhchandani, Hirshleifer, and Welch, 1992). Due to its prevalence and the complexities of studying history-dependent equilibria in sequential voting, this important benchmark is our focus here. Nonetheless, the issue of richer information structures is certainly important, and we discuss this in Section 4.5.

Denote by $G(\pi, \gamma, \tau_L, \tau_R; n)$ the sequential voting game defined above with prior π, signal precision γ, preference type parameters τ_L and τ_R, and n voters. Throughout the subsequent analysis, we use the term *equilibrium* to mean a (weak) Perfect Bayesian equilibrium of this game (Fudenberg and Tirole, 1991). Let $h^i \in \{L, R\}^{i-1}$ be the realized history of votes when it is voter i's turn to act; denote $h^1 = \phi$. A pure strategy for voter i is a map $v_i : \{L_p, N, R_p\} \times \{L, R\}^{i-1} \times \{l, r\} \to \{L, R\}$. We say that a voter i *votes informatively* following a history h^i if $v_i(N, h^i, l) = L$ and $v_i(N, h^i, r) = R$. The posterior probability that voter i places on state L is denoted by $\mu_i(h^i, s_i)$.

3 Posterior-Based Voting

3.1 Definition and Dynamics

We begin the analysis by introducing *Posterior-Based Voting* (PBV) and characterizing its induced dynamics. This characterization allows us to demonstrate that such behavior is an equilibrium in Section 3.2. Let $\mathbf{v} = (v_1, \ldots, v_n)$ denote a strategy profile and $\mathbf{v}^i = (v_1, \ldots, v_{i-1})$ denote a profile of strategies for all players preceding i.

Definition 1. *A strategy profile, $\mathbf{v}$, satisfies (or is) Posterior-Based Voting (PBV) if for every*

[7]For example, in the context of presidential primaries, while some voters in a party hope to nominate the party candidate who is more electable in the general presidential election, there are others who may not be so sophisticated and simply wish to select a particular party candidate without considering the general election.

voter i, type t_i, history h^i, signal s_i, and for any $W, W' \in \{L, R\}$,

1. $\mathbb{E}_\omega[u(t_i, W, \omega)|h^i, s_i; \mathbf{v}^i] > \mathbb{E}_\omega[u(t_i, W', \omega)|h^i, s_i; \mathbf{v}^i] \Rightarrow v_i(t_i, h^i, s_i) = W$
2. $\mathbb{E}_\omega[u(t_i, L, \omega)|h^i, s_i; \mathbf{v}^i] = \mathbb{E}_\omega[u(t_i, R, \omega)|h^i, s_i; \mathbf{v}^i] \Rightarrow \begin{cases} v_i(t_i, h^i, l) = L \\ v_i\left(t_i, h^i, r\right) = R \end{cases}$

PBV is a property of a strategy *profile*. A PBV strategy consequently refers to a strategy for a player that is part of a PBV profile.

The first part of the definition requires that given the history of votes and her private signal, if a voter believes that electing candidate L (R) will yield strictly higher utility than electing candidate R (L), then she votes for candidate L (R). In other words, in a PBV profile, each voter updates her beliefs about candidates using all currently available information (taking as given the strategies of previous voters), and then votes for the candidate she currently believes to be best for her. Observe that this behavior coincides with rational behavior in a standard herding environment, and does not explicitly account for the payoff interdependency of an election. Since Partisan voters have a preference ordering over candidates that is independent of the state of the world, the definition immediately implies that Partisans vote for their preferred candidate in a PBV profile, independent of signal and history. Whenever a Neutral voter's posterior is $\mu_i(h^i, s_i) \neq \frac{1}{2}$, she votes for the candidate she believes to be strictly better.

Part two of the definition is a tie-breaking rule.[8] It requires that when a Neutral voter has posterior $\mu_i(h^i, s_i) = \frac{1}{2}$, she vote informatively. In doing so, she reveals her signal to future voters. While we choose this tie-breaking rule to facilitate exposition, it does not play a significant role in our analysis. Any choice of how to break ties only matters for a non-generic constellation of parameters $(\pi, \gamma, \tau_L, \tau_R)$. For a generic tuple, $(\pi, \gamma, \tau_L, \tau_R)$, when PBV is played, it will never be the case that there is a Neutral voter with posterior $\mu_i(h^i, s_i) = \frac{1}{2}$. We defer a formal discussion of this point to Remark 1 in the Appendix.

The above discussion implies that the behavior of voter i in the PBV profile can be summarized as follows:

$$\begin{aligned} v_i(L_p, h^i, s_i) &= L \\ v_i(R_p, h^i, s_i) &= R \\ v_i(N, h^i, s_i) &= \begin{cases} L & \text{if } \mu_i(h^i, s_i) > \frac{1}{2} \text{ or } \{\mu_i(h^i, s_i) = \frac{1}{2} \text{ and } s_i = l\} \\ R & \text{if } \mu_i(h^i, s_i) < \frac{1}{2} \text{ or } \{\mu_i(h^i, s_i) = \frac{1}{2} \text{ and } s_i = r\} \end{cases} \end{aligned}$$

PBV is sophisticated insofar as voters infer as much as possible from the past history, taking into account the strategies of preceding players. However, the construction of PBV is nevertheless myopic. Since voters are influenced by the voting history in the PBV profile, a strategic voter who conditions on being pivotal should account for how her vote affects the

[8]Note that the tie from the standpoint of PBV does not imply that the voter is strategically indifferent between her choices, because the indifference here does not condition on being pivotal.

decisions of those after her. The requirement of PBV, on the other hand, is simply that a voter cast her vote for the candidate she would wish to select based on what she currently knows. This distinction between myopic and strategic reasoning makes it unclear, *a priori*, that PBV can be an equilibrium. We take up that question in Section 3.2, showing that it indeed is; for the moment, however, the focus is on characterizing the dynamics induced by PBV.

To provide intuition, we begin with an informal description of PBV behavior. As noted previously, the behavior of Partisans is trivial, so we focus on Neutrals. Voter 1, if Neutral, votes informatively, since signals are more informative than the prior. All subsequent Neutral voters face a simple Bayesian inference problem: conditional on the observed history and their private signal, what is the probability that the state is L? Consider $i > 1$ and a history $\tilde{h}^i$ where all preceding Neutrals (are assumed to) have voted informatively, with $\tilde{h}^i$ containing k votes for L and $i-k-1$ votes for R. For any history, we define the *public likelihood ratio*, $\lambda\left(h^i\right)$, as the ratio of the public belief that the state is L versus state R after the history h^i: $\lambda\left(h^i\right) = \frac{\Pr\left(\omega=L|h^i\right)}{\Pr\left(\omega=R|h^i\right)}$. By Bayes Rule, for the history $\tilde{h}^i$,

$$\lambda\left(\tilde{h}^i\right) = \frac{\pi}{(1-\pi)}\left(\frac{\tau_L+\tau_N\gamma}{\tau_L+\tau_N(1-\gamma)}\right)^k\left(\frac{\tau_R+\tau_N(1-\gamma)}{\tau_R+\tau_N\gamma}\right)^{i-k-1} \tag{1}$$

The public likelihood ratio above captures how informative the history $\tilde{h}^i$ is, given the postulated behavior about preceding voters. Since $\gamma > 1-\gamma$, the ratio is strictly increasing in k, i.e. seeing a greater number of votes for L strictly raises Neutral voter i's belief that L is the better candidate. Partisanship makes the public history noisy: when $\tau_L \simeq \tau_R \simeq 0$, the ratio is close to its maximum, which is informationally equivalent to voter i having observed k signal l's and $i-k-1$ signal r's. On the other hand, when $\tau_L \simeq \tau_R \simeq \frac{1}{2}$, the ratio is approximately $\frac{\pi}{1-\pi}$, reflecting a public history only slightly more informative than the prior. The Neutral voter i combines the information from the public history, h^i, with that of her private signal, s_i, to determine her posterior belief, $\mu_i\left(h^i, s_i\right)$, that the state is L. By Bayes Rule,

$$\frac{\mu_i\left(h^i, s_i\right)}{1-\mu_i\left(h^i, s_i\right)} = \lambda\left(h^i\right)\frac{\Pr\left(s_i|\omega=L\right)}{\Pr\left(s_i|\omega=R\right)}$$

Since the signal precision of s_i is γ, it follows that for $\lambda\left(h^i\right) \in \left[\frac{1-\gamma}{\gamma}, \frac{\gamma}{1-\gamma}\right]$, an l signal translates into a posterior belief no less than $\frac{1}{2}$ that the state is L and an r signal translates into a posterior belief no less than $\frac{1}{2}$ that the state is R. However, for $\lambda\left(h^i\right) > \frac{\gamma}{1-\gamma}$, both signals generate posterior beliefs strictly greater than $\frac{1}{2}$ that the state is L, and thus in PBV, a Neutral voter i would vote uninformatively for candidate L. Similarly, for $\lambda\left(h^i\right) < \frac{1-\gamma}{\gamma}$, voter i's posterior favors R regardless of her private signal, and a Neutral i thus votes uninformatively for candidate R.

In sum, PBV prescribes the following behavior: Neutrals vote informatively until the public likelihood ratio, $\lambda\left(h^i\right)$, no longer lies in $\left[\frac{1-\gamma}{\gamma}, \frac{\gamma}{1-\gamma}\right]$; when this happens, all future Neutrals vote uninformatively for the candidate favored by the posterior—they herd. There are several points to be emphasized about the nature of these herds. First, even after a herd begins for a candidate, Partisans continue to vote for their preferred candidate. Thus, it is always possible to see votes contrary to the herd, and any such contrarian vote is correctly inferred by future voters as having come from a Partisan. Second, once herding begins, the public likelihood ratio remains fixed because all voting is uninformative. Therefore, once a herd begins, its length does not influence the beliefs held by subsequent voters; independent of the population size, the (private) posterior belief of any voter lies within $\left(\frac{(1-\gamma)^3}{\gamma^3+(1-\gamma)^3}, \frac{\gamma^3}{\gamma^3+(1-\gamma)^3}\right)$ and is therefore bounded away from 0 and 1. Third, at any history where the winner of the election is yet undecided, a herd forming on, say, candidate L, does not immediately imply a victory for L. This is because if all subsequent voters are R-partisans (an event of positive probability), R will in fact be elected. Fourth, it is possible for a herd to form on a candidate who is *trailing* because the informational content of the voting history is not limited to merely whether a candidate is leading, but also how that candidate is performing relative to the ex-ante distribution of private preferences. This is illustrated by the following example.

Example 1. *Let $n \geq 6$, $\pi = \frac{2}{3}$, $\gamma = \frac{3}{4}$, $\tau_L = 0.1$, and $\tau_R = 0.45$. Suppose voters play the PBV profile, and voter* 6 *observes the history $h^6 = (R, R, R, L, L)$. Straightforward calculations show that given this history, voters 1 through 5 must have voted informatively if Neutral. Due to the relatively small partisan support for candidate L, the two votes for L from voters* 4 *are* 5 *are sufficient to overturn the impact of the 3 votes for R from voters* 1 *through* 3 *in terms of influencing voter 6's posterior. In fact, even if voter* 6 *receives an r signal, she believes that candidate L is the better candidate. Thus, regardless of signal, voter* 6*'s vote is uninformative, and by induction, all future Neutrals vote for candidate L. The bandwagon on L has formed even though L is trailing in the election.*

The discussion so far of dynamics was couched in terms of beliefs. For our equilibrium analysis, we need a characterization in terms of the voting history. It is convenient to use two state variables that summarize the impact of history on behavior. For any history, h^i, the *vote lead* for candidate L, $\Delta\left(h^i\right)$, is defined recursively as follows:

$$\Delta(h^1) = 0; \text{ for all } i > 1, \ \Delta(h^i) = \Delta\left(h^{i-1}\right) + \left(\mathbf{1}_{\{v_{i-1}=L\}} - \mathbf{1}_{\{v_{i-1}=R\}}\right) \tag{2}$$

The second state variable, called the *phase*, summarizes whether learning is ongoing in the system (denoted phase 0), or has terminated in a herd for one of the candidates (denoted phase L or R). The phase mapping is thus $\Psi : h^i \to \{L, 0, R\}$, and defined by the following transition mapping:

$$\Psi\left(h^1\right) = 0; \text{ for all } i > 1, \ \Psi\left(h^i\right) = \begin{cases} \Psi\left(h^{i-1}\right) & \text{if } \Psi\left(h^{i-1}\right) \in \{L, R\} \\ L & \text{if } \Psi\left(h^{i-1}\right) = 0 \text{ and } \Delta\left(h^i\right) = n_L\left(i\right) \\ R & \text{if } \Psi\left(h^{i-1}\right) = 0 \text{ and } \Delta\left(h^i\right) = -n_R\left(i\right) \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Note that herding phases, $\Psi \in \{L, R\}$, are absorbing. The sequences $n_L(i)$ and $n_R(i)$ in the phase map equation (3) are determined by explicitly considering posteriors, corresponding to our earlier discussion of the public likelihood ratio. For example, assuming that all prior Neutrals voted informatively, $n_L(i)$ is the smallest vote lead for candidate L such that at a history h^i with $\Delta(h^i) = n_L(i)$, the public history in favor of L outweighs a private signal r. Therefore, the threshold $n_L(i)$ is the unique integer less than or equal to $i-1$ that solves

$$\Pr\left(\omega = L|\Delta\left(h^i\right) = n_L(i) - 2, s_i = r\right) \leq \frac{1}{2} < \Pr\left(\omega = L|\Delta\left(h^i\right) = n_L(i), s_i = r\right) \quad (4)$$

If it is the case that a history h^i with $\Delta(h^i) = i - 1$ is outweighed by signal r, we set $n_L(i) = i$. Similarly, the threshold $n_R(i)$ is the unique integer less than or equal to i that solves

$$\Pr\left(\omega = L|\Delta\left(h^i\right) = -n_R(i) + 2, s_i = l\right) \geq \frac{1}{2} > \Pr(\omega = L|\Delta\left(h^i\right) = -n_R(i), s_i = l) \quad (5)$$

where again, implicitly, it is assumed that all prior Neutrals voted informatively. If it is the case that a signal l outweighs even that history h^i where $\Delta(h^i) = -(i-1)$, we set $n_R(i) = i$.

We summarize with the following characterization result (all proofs are in the Appendix).

Proposition 1. *Every game $G(\pi, \gamma, \tau_{L,}\tau_R; n)$ has a unique PBV strategy profile. For each $i \leq n$, there exist thresholds, $n_L(i) \leq i$ and $n_R(i) \leq i$, such that if voters play PBV in the game $G(\pi, \gamma, \tau_{L,}\tau_R; n)$, then a Neutral voter i votes*

1. *informatively if $\Psi(h^i) = 0$;*
2. *uninformatively for $C \in \{L, R\}$ if $\Psi(h^i) = C$,*

   where Ψ is as defined in (3). The thresholds $n_L(i)$ and $n_R(i)$ are independent of the population size, n.

When voters play the PBV strategy profile, a herd develops if and only if there is a history h^i such that $\Psi(h^i) \neq 0$. To study how likely this is, assume without loss of generality that the true state is R. Fixing play according to PBV, the realized path of play is governed by the draw of preference-types and signals. Consequently, the public likelihood ratios can be viewed as a stochastic process, which we denote $\langle\lambda_i\rangle$, where each λ_i is the public likelihood ratio when it is voter i's turn to act. It is well-known (e.g. Smith and Sorensen, 2000) that this stochastic process is a martingale conditional on the true state, R. By the Martingale Convergence Theorem, the process $\langle\lambda_i\rangle$ converges almost surely to a random variable, λ_∞. Since PBV is informative and the public likelihood is bounded away from 0 and ∞ so long as $\Psi = 0$, convergence of the public likelihood ratio requires that $\Psi \in \{L, R\}$ in the limit, i.e. herds eventually occur with probability 1. This intuition underlies the following result for our finite voter game.

Theorem 1. *For every* $(\pi, \gamma, \tau_L, \tau_R)$ *and for every* $\varepsilon > 0$*, there exists* $\overline{n} < \infty$ *such that for all* $n > \overline{n}$*, if voters play PBV, then* $\Pr[\Psi(h^n) \neq 0$ *in* $G(\pi, \gamma, \tau_L, \tau_R; n)] > 1 - \varepsilon$*.*

3.2 PBV Equilibrium

In this section, we establish that PBV is an equilibrium of the sequential voting game. In fact, we prove a stronger result. Say that the *election is undecided* at history h^i if both candidates still have a chance to win the election given the history h^i. An equilibrium is *strict* if conditional on others following their equilibrium strategies, it is uniquely optimal for a voter to follow her equilibrium strategy at any undecided history.[9] We show that not only is PBV an equilibrium, but moreover, it is generically a strict equilibrium.

Theorem 2. *The PBV strategy profile is an equilibrium, and generically, is strict.*

There are two points to emphasize about PBV equilibrium. Strictness for generic parameters implies that its existence does not rely upon how voter indifference is resolved when the election remains undecided. Given that others are playing PBV, a strategic voter follows PBV not because she is indifferent between or powerless to change the outcome, but rather because deviations yield strictly worse expected payoffs. Strictness also implies that the equilibrium is robust to small perturbations of the model. Second, because Partisan voters always vote for their preferred candidates in the PBV equilibrium, every information set is reached with positive probability. Therefore, off-the-equilibrium-path beliefs play no role in our analysis, and PBV is a Sequential Equilibrium (Kreps and Wilson, 1982).

It may be surprising that PBV is an equilibrium even within our simple model, for at least two reasons. First, PBV is the generalization of of *sincere voting* (Austen-Smith and Banks, 1996) when moving from simultaneous to sequential voting in incomplete information environments, because PBV prescribes voting for the candidate currently thought to be better, without conditioning on being pivotal. However, as is well-known from the analysis of Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996), sincere behavior is not generically an equilibrium in the simultaneous voting version of our model. Secondly, and related, behavior is myopic in the PBV profile, whereas strategic optimality must account for the forward-looking incentives in the game. In particular, a Neutral voter in the learning phase of the election faces the following tradeoff when deciding to vote informatively, as prescribed by PBV. On the one hand, the public history in the learning phase is not sufficiently informative so as to overturn her private signal received. The benefit of voting informatively, then, is voting in the direction favored by her current posterior. On the other hand, the cost of voting informatively is that it may push future voters towards herding and/or push the election towards being decided,

[9] This definition is non-standard, but is the appropriate modification of the usual definition for voting games. Usually, a strict equilibrium of a game is one in which a deviation to any other strategy makes a player strictly worse off (Fudenberg and Tirole, 1991). A sequential voting game (with $n \geq 3$) cannot possess any strict equilibria in this sense, because after any history where a candidate has captured sufficiently many votes to win the election, all actions yield identical payoffs. That is, only histories where the election remains undecided are strategically relevant to voters. Our definition of strictness restricts attention to these histories.

suppressing valuable information possessed by later Neutrals. Given the myriad of ways in which a voter could be pivotal, resolving this tradeoff for every possible vote lead and for each private signal that a Neutral voter may face in the learning phase is *a priori* quite complicated.

In light of these considerations, the remainder of this section provides a detailed sketch of why PBV is an equilibrium. Readers more interested in robustness and other issues can skip forward to Section 4.

Sketch of the equilibrium argument. For simplicity, we consider a symmetric level of partisanship, $\tau_L = \tau_R$. Symmetry induces thresholds $n_L(j)$ and $n_R(j)$ that do not vary across voters; instead, they can be denoted simply as constants $n_L > 0$ and $n_R > 0$. We also restrict attention here to generic parameters of the game, where it can be shown that when others follow PBV, a voter is never indifferent at any undecided history. (We clarify the details concerning genericity in the Appendix, where our formal results and proofs also deal with the general case in which τ_L can differ from τ_R.)

To see that PBV is an equilibrium, consider a voter acting when the election is undecided, and assume that all others players are playing PBV. Note that such a voter is pivotal with strictly positive probability. We will argue that the following three kind of behaviors are uniquely optimal, depending on the voter's preference type and the phase of the election:

(i) Partisans vote for their preferred candidate;

(ii) In the herding phase for L (R), Neutrals vote uninformatively for L (R);

(iii) In the learning phase, Neutrals vote informatively.

Point (i) above is a consequence a monotonicity property of PBV: when others play PBV, a voter can never increase the probability of a future L vote, say, by herself voting R rather than L.[10] Turning to point (ii), consider the incentives for a Neutral voter i in a herding phase. Since all subsequent voting after voter i is completely uninformative, conditioning on being pivotal does not change i's posterior beliefs about the state of the world. Consequently, i strictly prefers to vote on the basis of her current posterior, $\mu_i(h^i, s_i)$. By construction of the phase map, a herding phase on candidate L (analogously for R) implies that i's posterior is strictly higher on L than R regardless of her private signal. Therefore, it is strictly optimal for Neutral i to vote uninformatively for L.

It remains to establish optimality of (iii): a Neutral voter i votes informatively in the learning phase, when $\Psi\left(h^i\right) = 0$. As previously noted, the incentive compatibility constraint for i concerns the tradeoff between myopically optimal behavior and inducing informative behavior from future Neturals. Lemma 5 in the Appendix shows that if the incentive constraints are satisfied (strictly) at histories h^i in which voting for L or R immediately starts a herd and/or

[10]This monotonicity does not hold in an arbitrary strategy profile. Thus, eliminating weakly-dominated strategies is not sufficient to guarantee that Partisans vote for their preferred candidate, unlike the case of a simultaneous election. It is in fact possible to construct sequential voting equilibria (in undominated strategies) where a Partisan votes against her preferred candidate.

ends the election, then the incentive constraints are satisfied (strictly) at all other histories in the learning phase. Therefore, the trade-off is most stark for the Neutral voter i who faces $\Delta\left(h^i\right) = n_L - 1$ and $s_i = l$ (or $\Delta\left(h^i\right) = -n_R + 1$ and $s_i = r$), since voting for candidate L (or R) in this situation ends social learning altogether.

Accordingly, we need only to show it is strictly optimal for i to vote informatively even when she immediately triggers a herd. Consider the case where voter i receives signal l and faces undecided history h^i where $\Psi\left(h^i\right) = 0$ and $\Delta\left(h^i\right) = n_L - 1$. (The argument is analogous for the case where $\Delta\left(h^i\right) = -n_R + 1$ and $s_i = r$.) Assuming that others are playing PBV, the impact of i's vote can be assessed for any vector of realizations of preference types and signals among the future voters. The set of type-signal realizations in which voter i is pivotal, denoted as Piv_i, consist of all those vectors where $v_i = L$ results in L winning the election and $v_i = R$ results in R winning.[11] Since h^i is undecided, $Piv_i \neq \emptyset$. Denote the set of type-signal profiles where $v_i = R$ results in a herd for candidate $C \in \{L, R\}$ as ξ^C; let the set of type-signal profiles where no herd forms after $v_i = R$ be denoted $\tilde{\xi}$. The proof is completed by showing that i's posterior conditional on being pivotal and on each of the three mutually exclusive and exhaustive events, ξ^L, ξ^R, or $\tilde{\xi}$, is greater than $\frac{1}{2}$.[12]

Consider the event $\xi^R \cap Piv_i$. Because voting is uninformative once the cascade begins, $\Pr\left(\omega = L|\xi^R, Piv_i, h^i, l\right) = \Pr\left(\omega = L|\xi^R, h^i, l\right)$. Since $\Delta\left(h^i\right) = n_L - 1$, the event ξ^R can happen only if following $v_i = R$, candidate R subsequently gains a lead of n_R votes in the learning phase. This requires that after i's vote, R receive an additional $n_L + n_R - 2$ votes over L in the learning phase. Since L has a vote lead of $n_L - 1$ prior to i's vote, conditioning on ξ^R in effect reveals a *net* total of $n_R - 1$ votes for R in the learning phase (not counting i's vote since that is a deviation). By definition of n_R, i's belief given her own signal $s_i = l$ and $n_R - 1$ votes for R in the learning phase is strictly in favor of L. In other words, $\Pr\left(\omega = L|\xi^R, Piv_i, h^i, l\right) > \frac{1}{2}$.

Now consider the events $\xi^L \cap Piv_i$ and $\tilde{\xi} \cap Piv_i$: amongst voters $i+1, \ldots n$, candidate R can receive at most $n_L + n_R - 2$ votes over candidate L; otherwise, an R-cascade would start. By the same logic as before, it follows that $\Pr\left(\omega = L|\tilde{\xi}, Piv_i, h^i, l\right) > \frac{1}{2}$.

Therefore, conditional on her observed history, signal, and being pivotal, voter i believes candidate L to be the better candidate with probability strictly greater than $\frac{1}{2}$. Since voting L leads to a strictly higher probability of L winning the election, it is strictly optimal for voter i to vote L, even though such a choice ends the learning phase.

[11]In general, there are strategy profiles where i can be pivotal in a way that $v_i = L$ results in R winning, whereas $v_i = R$ results in L winning. This is not possible in PBV because as we noted earlier, PBV features a weak monotonicity of subsequent votes in i's vote.

[12]The simplification that $\tau_R = \tau_L$ ensures that the posterior conditional on a herd forming for a candidate is invariant to *when* the herd begins.

4 Discussion

4.1 Other Equilibria

The previous section demonstrated that history-dependent behavior in the form of PBV is an equilibrium of the sequential voting game, and moreover, such behavior engenders momentum, with bandwagons forming almost surely in large elections. It is natural to ask whether the model possesses other equilibria. Given the complexity of the strategy space in sequential voting games, a characterization of all equilibria appears unfeasible. However, there is one other equilibrium that has been identified in the literature, with which we would like to contrast PBV.

Dekel and Piccione (2000) have shown that a class of sequential voting games possess history-independent equilibria that are equivalent in outcomes to symmetric equilibria of otherwise identical simultaneous voting games. Their main insight applies to the model considered here and is as follows. For any parameter set $(\pi, \gamma, \tau_L, \tau_R, n)$, the simultaneous voting analog of our model possesses a symmetric equilibrium where each Partisan voter votes for her preferred candidate, and each Neutral votes for a candidate on the basis of a signal-dependent probability. Turning to the sequential voting game, consider a strategy profile where, independent of history, every voter plays the same strategy as in the above construction. The insight of Dekel and Piccione (2000) is that because all voters are acting independently of history, the events in which a voter is pivotal is identical in both the simultaneous and sequential games; therefore, since the profile is an equilibrium of the simultaneous game, it is also an equilibrium of the sequential game. Using the approach of Feddersen and Pesendorfer (1997), it can be shown that this equilibrium achieves *full information equivalence*, aggregating information efficiently in large elections.

Though the existence of the history-independent equilibrium is an important theoretical benchmark for efficiency, its descriptive implications appear at odds with behavior in real dynamic elections. While the history-independent equilibrium hinges on it being common knowledge that all voting behavior is unaffected by the history of votes,[13] it is generally accepted in practice that the prior performance of candidates influences subsequent voting behavior. Bartels (1988) and Popkin (1991) argue that voters keep careful track of how candidates have performed relative to expectations when deliberating how to vote, and that the information provided to voters during the primaries is little more than *horse-race* statistics that describe candidates' performance in preceding states. Since the comparison of history-independent versus history-dependent behavior is ultimately a difficult empirical question, this suggests that it may be counterfactual to focus theoretically on a history-independent equilibrium.[14]

[13] Were a Neutral voter to consider the possibility that her vote may influence that of future voters, her incentives to play according to the history-independent equilibrium can dissipate.

[14] Experimental evidence can be brought to bear on the issue, but we are hesitant to draw conclusions from existing work on sequential voting. The setup of experiments such as those of Morton and Williams (1999) and Battaglini, Morton, and Palfrey (2005) differ importantly from the model developed here. The former authors consider an election with three options; the latter authors consider an election with only three voters, but where

The myopic nature of PBV also makes it appealing, given the mixed evidence on the strategic sophistication of voters. In the history-independent equilibrium of a large voting game, Neutral voters will often be voting against the candidate they believe to be better given currently available information. To the extent that some voters do not condition on hypothetical events of being pivotal and draw the requisite inferences, an equilibrium that accommodates both myopic and strategic play—if one exists (cf. Section 4.5)—is attractive.

4.2 Population Uncertainty

As argued by, for example, Myerson (1998, 2000), it may be unrealistic to assume that in large elections, each voter knows exactly how many other voters there are in the game. While some models of simultaneous elections (e.g. Myerson, 1998, 2000; Feddersen and Pesendorfer, 1996) account for this possibility, existing models of sequential elections assume that the size of the electorate is commonly known. In principle, population uncertainty can play a significant role in voting behavior in sequential elections: based on the history, a voter can update her beliefs about how many others are participating, and at least, set a lower bound on the number of other voters. In spite of the complexities introduced by population uncertainty, we show that PBV remains an equilibrium of the sequential voting game.

We define an election with population uncertainty as follows. Suppose there is a countably infinite set of available voters, indexed by $i = 1, 2, \ldots$. Nature first draws the size of the electorate, n, according to probability measure ν with (possibly unbounded) support on the natural numbers. The draw is unobserved by any agent. A voter is selected to vote if and only if her index i is no larger than n. All those who are selected to vote do so sequentially, in a roll-call order, only observing the history of votes, and receiving no information about the numbers of voters to follow. The rest of the game, in terms of preferences, information, and how outcomes are determined, remains as before. We define a game with population uncertainty as $G(\pi, \gamma, \tau_L, \tau_R; \nu)$.

The structure of population uncertainty formulated above is general, encompassing the kinds of population uncertainty that have been considered in simultaneous elections, such as Poisson distributions (Myerson, 1998) and binomial distributions (Feddersen and Pesendorfer, 1996). Nevertheless, the logic for why PBV remains an equilibrium is straightforward. Consider the decision faced by a voter i who observes history h^i, signal s_i, and assumes that all other voters who are selected to vote will behave according to PBV. Voter i can assess her best response by conditioning on every possible realization of population size that is weakly larger than i. By Theorem 2, for every such realization, voter i would wish to vote according to PBV. Since the behavior prescribed by PBV is independent of the population size, the result follows.

Proposition 2. *For every game $G(\pi, \gamma, \tau_L, \tau_R; \nu)$, the PBV strategy profile is an equilibrium.*

voters face a cost of voting, building on the theoretical work of Battaglini (2005). The treatment of Hung and Plott (2001) is the closest to our model, and they find evidence of herding in almost 40 percent of experimental rounds.

It is interesting to note that in sequential voting games with a sufficient degree of population uncertainty, there generally cannot exist a symmetric equilibrium in which all voters who are selected to vote play the same strategy. The reason is that by learning about (lower bounds on) the population size, the incentives for voters at different stages of the election are quite different. In particular, a symmetric equilibrium of the simultaneous election with population uncertainty does not generally remain an equilibrium of the corresponding sequential election. A thorough examination of sequential elections with population uncertainty is left for future research, including issues such as whether all equilibria generally feature history dependence, or if there are also equilibria where behavior varies by voter position without depending on the voting history.

4.3 Other Voting Rules

Generally, equilibria of elections are sensitive to the choice of voting rule, and this has been illustrated in the case of simultaneous elections by the important contributions of Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998). Holding fixed a strategy profile, changing the voting rule affects the profiles of votes in which one is pivotal, and can therefore change one's posterior conditional on being pivotal. Surprisingly, however, PBV is an equilibrium in our sequential voting game for any voting rule, for all parameters. Furthermore, a large class of voting rules are asymptotically equivalent in terms of the electoral outcome they induce.

We study the class of voting rules termed *q-rules*, where if the fraction of votes for L strictly exceeds some number $q \in [0,1]$, then L wins the election. Let $G(\pi, \gamma, \tau_L, \tau_R; n, q)$ denote the sequential voting game with parameters $(\pi, \gamma, \tau_L, \tau_R; n)$ where votes are aggregated according to the q-rule. Since the PBV strategy profile is defined independently of the voting rule, for two different rules q and q', PBV generates identical behavior in the games $G(\pi, \gamma, \tau_L, \tau_R; n, q)$ and $G(\pi, \gamma, \tau_L, \tau_R; n, q')$. In fact, with minor modifications to the proof of Theorem 2, the following can be shown.

Theorem 3. *PBV is an equilibrium for any q-rule, and generically, is a strict equilibrium.*

While it is straightforward to see that changing the voting rule leaves the incentives of Partisan voters or Neutrals in the herding phases unchanged for any q-rule, the incentives of Neutral voters to vote informatively in the learning phase would seem to be affected because the vote profiles for which one is pivotal differs across voting rules. However, the crucial point about PBV is that even if a voter's deviation in the learning phase changes the outcome, the subsequent profile of votes never generates enough public information to overturn one's posterior based on the private signal and available public history, regardless of the voting rule. This is because the thresholds for herding in PBV are independent of the voting rule, and the information that can be extracted from future voters' actions is determined by these thresholds.

Theorem 3 raises the question of how changes in the voting rule affect the electoral outcome when voters vote according to PBV. When voters play PBV, different rules may yield different (distributions over) electoral outcomes. However, the following result shows that voting rules can be partitioned into three classes such that all rules within any class are asymptotically *ex-ante* equivalent: they elect the same winner with probability approaching 1 in large voting games.

Theorem 4. *Fix any parameters $\pi, \gamma, \tau_L, \tau_R$, and assume that for any n, q, voters play PBV in the game $G(\cdot; n, q)$. For any $\varepsilon > 0$, there exists $\bar{n}$ such that for all $n > \bar{n}$,*

(a) $|Pr(L \text{ wins in } G(\cdot; n, q) - Pr(L \text{ wins in } G(\cdot; n, q')| < \varepsilon$ for all $q, q' \in (\tau_L, 1 - \tau_R)$;

(b) $\Pr(L \text{ wins in } G(\cdot; n, q)) > 1 - \varepsilon$ for all $q \in [0, \tau_L)$;

(c) $\Pr(L \text{ wins in } G(\cdot; n, q)) < \varepsilon$ for all $q \in (1 - \tau_R, 1]$.

Parts (b) and (c) of Theorem 4 are not surprising given the presence of Partisans: in PBV, the probability with which any voter votes L is at least τ_L and at most $(1 - \tau_R)$. Therefore, in any sufficiently large voting game, L wins with probability approaching 1 if $q < \tau_L$ and loses with probability approaching 1 if $q > (1 - \tau_R)$. The important result is part (a) of Theorem 4: all "interior" voting rules—where outcomes are not determined asymptotically by Partisanship alone—are nevertheless asymptotically equivalent. The intuition relies on the fact that for all such voting rules, once Neutrals begin herding on a candidate, that candidate wins with probability approaching one in large electorates. Therefore, all of these interior rules are asymptotically equivalent once a herd begins on a particular candidate. Asymptotic ex-ante equivalence of these voting rules then follows from the observation that the probability of a herd beginning on a particular candidate is independent of the voting rule (since PBV behavior is defined independently of the voting rule), and by Theorem 1, this probability approaches 1 in large games.[15]

It is interesting to contrast Theorems 3 and 4 with the message of Dekel and Piccione (2000). By showing that (responsive) equilibria of a simultaneous election are outcome-equivalent to equilibria of a voting game with any timing structure, Dekel and Piccione (2000) have demonstrated that strategic behavior can be unaffected by the timing of a voting game. In these history-independent equilibria, conditioning on being pivotal negates any usefulness from observing the history of votes; necessarily such equilibria are sensitive to the choice of voting rule. Theorem 3 demonstrates that the opposite effect can prevail in sequential environments: regardless of the voting rule, if others vote according to PBV, conditioning on being pivotal does not contain more payoff-relevant information than the public history and one's private signal. That PBV is an equilibrium across all the voting rules and renders all interior voting rules asymptotically equivalent in sequential frameworks illustrates the striking difference between static and dynamic elections.

[15] In fact, this argument shows that Theorem 4 can be strengthened to an *ex-post* statement: under PBV, for almost all realized profiles of type-signal vectors, all "interior" voting rules are asymptotically equivalent, and all "extremal" voting rules result in one of the candidates winning with probability approaching 1 in large elections.

4.4 Pure Common Value Elections

We have thus far studied a model where some voters are Neutral, and others are Partisan; in contrast, most of the prior papers on sequential voting games have considered pure common value environments, where every voter is Neutral (e.g. Fey, 2000; Wit, 1997; Callander, 2004a). Before comparing our results with this earlier research, we should emphasize two points. First, although we have formalized Partisans as preferring a candidate regardless of the state of the world, this was purely for expositional purposes. It leaves our analysis unchanged to model Partisans and Neutrals as sharing the same state-dependent rankings over candidates, but simply specifying that Partisan preferences (or priors) be sufficiently biased in favor of one candidate, while remaining responsive to the state of the world. Consequently, even if all voters—Neutrals and Partisans—would agree on the preferred candidate for almost all profiles of publicly observable signals in large elections, our results apply. Second, some element of private values is surely an important aspect of elections in practice, and modeling this aspect adds richness to the theory.

In any case, fix a triple (π, γ, n), i.e. a prior, signal precision, and population size. Our main result extends as follows.

Proposition 3. *Assume $\tau_L = \tau_R = 0$. Then there is an equilibrium where all voters use PBV. In this equilibrium, a transition from the learning phase to the herding phase occurs when $\Delta\left(h^i\right) \in \{1, -2\}$.*

While PBV remains an equilibrium of the pure common value election, its properties differ in one important respect from that of PBV when τ_L and τ_R are strictly positive. Since Partisans vote for their preferred candidates regardless of history in the PBV equilibrium, after a herd begins, there is positive probability that any future vote may be contrary to the herd. On the other hand, in the PBV equilibrium of the pure common values election, once a herd begins (necessarily on the candidate who is leading), every vote thereafter is for the leading candidate. Consequently, a vote for the losing candidate is an off-the-equilibrium-path action after a herd has begun. This implies that some of the beliefs that sustain the equilibrium are necessarily off-the-equilibrium-path beliefs, and a theory of "reasonable" beliefs now becomes necessary: if voters see a vote going against a herd, how should they interpret it, and given their interpretation, would voters still wish to herd?

A natural place to begin would be to investigate the implications of standard beliefs-based refinements for signaling games such as the *Intuitive Criterion*, *D1* (Cho and Kreps, 1987), or *Divinity* (Banks and Sobel, 1987). However, none of these refinements have bite in this environment. To see why, consider the even stronger refinement criterion of *Never a Weak Best Response* (Kohlberg and Mertens, 1986). If future voters interpret a deviation from a herd as being equally likely to come from a voter with signal $s_i = l$ as from a voter with signal $s_i = r$, then future voters should not update their beliefs at all based on i's vote, and hence it is a weak best response for voter i to deviate from the herd regardless of her signal.

Given this, the belief that a deviation is equally likely to come from either signal-type of voter i survives Never a Weak Best Response, which is the strongest of standard dominance-based belief refinements.

In contrast, the aforementioned papers that consider common value sequential voting impose the following belief restriction: if a voter i votes for R once an L-herd has begun, it must be believed that $s_i = r$; similarly, if i votes for L once an R-herd has begun, it must be believed that $s_i = l$. Under this belief restriction—henceforth referred to as *Perpetual Revelation*—Fey (2000) and Wit (1997) show that because of the signaling motive inherent in common value sequential voting, at least one voter with a signal that opposes the herd would always wish to deviate out of the herd and reveal her signal to future voters. That is, in the pure common value setting, Perpetual Revelation is sufficient to halt momentum by inducing anti-herding for at least one voter.

Without recourse to standard belief-based refinements, it is *ad hoc* to impose direct restrictions on off-path beliefs in the pure common values game. The off-path belief of ignoring a deviation from the herd, which supports the PBV equilibrium of Proposition 3, is a limit of the beliefs when the possibility of Partisanship is strictly positive. It is not our contention that this limiting belief is the only reasonable or sensible off-path belief in the pure common value game, but rather, that precluding it—by a priori requiring Perpetual Revelation, for example—lacks justification.

4.5 Richer Information Structures

The binary signal structure we have studied is the canonical framework for both social learning models and the information aggregation approach to elections. In this setting, we have have been able to identify a history-dependent equilibrium in the form of PBV. Although the binary signal structure is an essential benchmark, it is important to understand the extent to which the results here generalize to richer informational structure. The challenge, however, stems from the payoff interdependencies inherent to elections. Due to the forward looking incentives in sequential voting, richer information structures tremendously complicate the analysis of inferences to be drawn from the myriad of ways a voter may be pivotal in a history-dependent strategy profile.[16] This contrasts with standard models without payoff interdependencies, where even a continuum of signals presents no difficulty to equilibrium characterization, adding challenges only to outcome dynamics (Smith and Sorensen, 2000).

Therefore, we leave the study of momentum effects with richer information structures to future research. We simply note by example that PBV will not always be an equilibrium, illustrating the need to consider more complex forms of history-dependent strategies in richer environments.

Example 2. *Consider a sequential voting game with* 3 *voters where* $\pi \simeq \frac{1}{2}$, $\tau_L \simeq 0$, *and* $\tau_R = \frac{1}{2}$.

[16] The results of Dekel and Piccione (2000) can still be applied to yield history-independent equilibria in some cases.

A Neutral voter can be one of two types : a Guru (G), who obtains a perfectly informative signal, or a Follower (F), who obtains a signal of precision $\gamma = \frac{3}{5}$*. Conditional on being Neutral, a voter is a Guru with probability* $\tau_G = \frac{4}{5}$ *and a Follower with probability* $\frac{1}{5}$*. Consider a PBV strategy profile, and suppose that voter* 1 *has voted R, and voter* 2 *is a Neutral Follower with signal r. Given the history and private signal, voter* 2*'s posterior favors R, and PBV therefore prescribes that she vote for R. By following PBV, she immediately determines the winner of the election as R. Voter 2's expected payoff from voting R is therefore approximately* $\frac{8}{11}$*. On the other hand, if she deviates from PBV and votes for L, she makes voter 3 pivotal. If voter* 3 *is a Neutral Follower (and follows PBV), voter 3 will vote L uninformatively; therefore, by voting L, voter* 2 *knows that R will win if and only if voter 3 is an R-Partisan or a Neutral Guru who observes an r signal. Consequently, voter 2's expected payoff from voting L computes as* $\frac{87}{110}$*, which is strictly greater than her payoffs from voting R. It is optimal for her to deviate from PBV following voter 1's vote for R.*

5 Conclusion

This paper has proposed an informational theory of momentum and herd behavior in sequential voting environments. Our model is that of a binary election where a proportion of the voters seek to elect the better candidate, and the remainder have partisan preferences. The central result is that there exists a generically strict equilibrium that leads to herding with high probability in large elections. In this equilibrium, voters learn from the voting history , and use this information to update on optimally to cast their vote. The behavior exhibited in this equilibrium can explain why voters rationally judge candidates relative to expectations.

Our results raise various issues that deserve further study, some of which we have already mentioned in Section 4. We conclude by highlighting some others.

We touched briefly on efficiency comparisons of simultaneous and sequential voting in Section 4.1, using the insights of Dekel and Piccione (2000). However, real-world sequential voting mechanisms may feature informational benefits that have not been considered here, nor in other models. For example, when thinking about presidential primaries, one natural point of departure is that candidates face greater constraints in campaigning across states that hold primaries on the same day than across states whose primaries are sequenced. This can be modeled formally as a constraint on the informativeness of signals that are obtained by voters who vote simultaneously. If sequencing can increase the informativeness of signals, this provides a rationale for greater information aggregation in sequential voting mechanisms, even in the presence of momentum effects.

Our focus in this paper has been on elections with only two options. Given the nature of the candidate winnowing process in the U.S. presidential primaries, it is important to understand the dynamics of sequential voting with more than two candidates. We are currently exploring this idea.

We have also restricted attention here to an environment where voting is entirely sequential, one voter at a time. Though there are elections of this form—for example, roll-call voting mechanisms used in city councils and legislatures—there are many dynamic elections, such as the primaries, that feature a mixture of simultaneous and sequential voting. To what extent such games possess history-dependent equilibria with similar qualitative features is a significant question for future research.

This paper has abstracted away from the role of institutions, and concentrated on voters as being the sole players. Certainly, in practice, there are other forces involved in dynamic elections, many of which are strategic in nature, such as the media, campaign finance contributors, and so forth. By examining the potential for sequential voting alone to create momentum, our model provides a benchmark to understand the impact of different institutions on sequential elections.

Finally, we note that social learning with payoff and information externalities can arise more generally outside the confines of voting. Many environments of economic interest—such as dynamic coordination games, timing of investments, or network choice—feature sequential decision-making, private information that has social value, and payoff interdependencies (cf. Dasgupta, 2000; Neeman and Orosel, 1999). Our results contribute to a better understanding of social learning when there are incentives to reveal or distort one's information to successors.

A Proofs

A.1 Proofs for Section 3.1

We begin with preliminaries that formally construct the thresholds $n_L(i)$ and $n_R(i)$ for each constellation of parameters $(\pi, \gamma, \tau_L, \tau_R)$, and each index i. Define the functions

$$f(\tau_L, \tau_R) \equiv \frac{\tau_L + (1 - \tau_L - \tau_R)\gamma}{\tau_L + (1 - \tau_L - \tau_R)(1 - \gamma)}$$

where the domain is $\tau_L, \tau_R \in \left[0, \frac{1}{2}\right)$. It is straightforward to verify that f strictly exceeds 1 over its domain.

For each positive integer i and any integer k where $|k| < i$ and $i - k$ is odd, define the function $g_i(k) = (f(\tau_L, \tau_R))^k \left(\frac{f(\tau_L,\tau_R)}{f(\tau_R,\tau_L)}\right)^{\frac{i-k-1}{2}}$. Note that for a history h^i where $\Delta\left(h^i\right) = k$ and all prior Neutrals voted informatively and Partisans voted for their preferred candidates, $g_i(k) = \frac{\Pr\left(h^i|\omega=L\right)}{\Pr\left(h^i|\omega=R\right)}$; thus $g_i(k) = \left(\frac{1-\pi}{\pi}\right)\lambda\left(h^i\right)$, as defined in equation (1) in the text. Plainly, $g_i(k)$ is strictly increasing in k.

For a given $(\pi, \gamma, \tau_L, \tau_R)$, define $\{n_L(i)\}_{i=1}^{\infty}$ as follows. For all i such that $g_i(i-1) \leq \frac{(1-\pi)\gamma}{\pi(1-\gamma)}$, set $n_L(i) = i$. If $g_i(i-1) > \frac{(1-\pi)\gamma}{\pi(1-\gamma)}$, we shall set $n_L(i)$ to be the unique integer that solves:

$$g_i(n_L(i) - 2) \leq \frac{(1-\pi)\gamma}{\pi(1-\gamma)} < g_i(n_L(i)) \tag{6}$$

Since $g_i(-(i-1))$ is strictly less than $\frac{(1-\pi)\gamma}{\pi(1-\gamma)}$, and $g_i(k)$ is strictly increasing in k, a unique solution exists to (6).

Similarly, we define $\{n_R(i)\}_{i=1}^{\infty}$ as follows. For all i such that $g_i(-(i-1)) \geq \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$, set $n_R(i) = i$. If $g_i(-(i-1)) < \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$, set $n_R(i)$ to be the unique integer that solves:

$$g_i(-n_R(i) + 2) \geq \frac{(1-\pi)(1-\gamma)}{\pi\gamma} > g_i(-n_R(i)) \tag{7}$$

As before, since $g_i(k)$ is strictly increasing in k, and $g_i(i-1) = (f(\tau_L, \tau_R))^{i-1} \geq 1 > \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$, a unique solution exists to (7).

We use these values of $n_L(i)$ and $n_R(i)$ to define $\Psi(\cdot)$ as in equation (3) from the text, and turn to the proof of Proposition 1.

Proposition 1 on pp. 10

Proof. The claim is obviously true for Voter 1 as $\Psi\left(h^1\right) = 0 \in (-n_R(1), n_L(1))$, and by construction, a PBV strategy involves a Neutral Voter 1 voting informatively. To proceed by induction, assume that the claim about behavior is true for all Neutral voter $j < i$.

Case 1: $\Psi\left(h^i\right) = 0$: All preceding neutrals have voted informatively. It is straightforward to see that the posterior $\mu\left(h^i, s_i\right) = \mu\left(\tilde{h}^i, s_i\right)$ if $\Delta\left(h^i\right) = \Delta\left(\tilde{h}^i\right)$ and $\Psi\left(h^i\right) = \Psi\left(\tilde{h}^i\right) = 0$

(i.e., so long as all preceding neutrals have voted informatively, only vote lead matters, and not the actual sequence). Thus, we can define $\tilde{\mu}_i(\Delta, s_i) = \mu(h^i, s_i)$ where $\Delta = \Delta(h^i)$. By Bayes' rule,

$$\tilde{\mu}_i(\Delta, l) = \frac{\pi\gamma g_i(\Delta)}{\pi\gamma g_i(\Delta) + (1-\pi)(1-\gamma)}$$

Simple manipulation shows that $\tilde{\mu}_i(\Delta, l) \geq \frac{1}{2} \Leftrightarrow g_i(\Delta) \geq \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$. This latter inequality holds since by hypothesis, $\Psi(h^i) = 0$, and therefore, $\Delta \geq -n_R(i) + 1$. If $\tilde{\mu}_i(\Delta, l) > \frac{1}{2}$, then Condition 1 of the PBV definition requires that Neutral voter i vote L given $s_i = l$; if $\tilde{\mu}_i(\Delta, l) = \frac{1}{2}$, then Condition 2 of the PBV definition requires that Neutral voter i vote L given $s_i = l$.

Similarly, using Bayes' rule,

$$\tilde{\mu}_i(\Delta, r) = \frac{\pi(1-\gamma) g_i(\Delta)}{\pi(1-\gamma) g_i(\Delta) + (1-\pi)\gamma}$$

Simple manipulation shows that $\tilde{\mu}_i(\Delta, r) \leq \frac{1}{2} \Leftrightarrow g_i(\Delta) \leq \frac{(1-\pi)\gamma}{\pi(1-\gamma)}$. The latter inequality holds since by hypothesis, $\Psi(h^i) = 0$, and therefore, $\Delta \leq n_L - 1$. If $\tilde{\mu}_i(\Delta, r) < \frac{1}{2}$, then Condition 1 of the PBV definition requires that Neutral voter i vote R given $s_i = r$; if $\tilde{\mu}_i(\Delta, r)$ then Condition 2 of the PBV definition requires that Neutral voter i vote R given $s_i = r$.

Case 2: $\Psi(h^i) = L$. Then all Neutrals who voted prior to the first time Ψ took on the value L voted informatively, whereas no voter voted informatively thereafter. Let $j \leq i$ be such that $\Psi(h^j) = L$ and $\Psi(h^{j-1}) = 0$; therefore, $\Delta(h^j) = n_L(j)$. Then, $\mu(h^j, s_j) = \tilde{\mu}_j(n_L(j), s_j)$. Since all voting after that of $(j-1)$ is uninformative, $\mu(h^i, s_i) = \mu(h^j, s_i) = \tilde{\mu}_j(n_L(j), s_i)$. A simple variant of the argument in Case 1 implies that $\tilde{\mu}_j(n_L(j), l) > \frac{1}{2}$, and therefore Condition 1 of the PBV definition requires that Neutral voter i vote L given $s_i = l$. Consider now $s_i = r$. Since $g_j(n_L(j)) > \frac{(1-\pi)\gamma}{\pi(1-\gamma)}$, it follows that $\tilde{\mu}_j(n_L(j), r) > \frac{1}{2}$, and therefore Condition 1 of the PBV definition requires that Neutral voter i vote L even following $s_i = r$.

Case 3: $\Psi(h^i) = R$. Then all Neutrals who voted prior to the first time Ψ took on the value R voted informatively, whereas no voter voted informatively thereafter. Let $j \leq i$ be such that $\Psi(h^j) = R$ and $\Psi(h^{j-1}) = 0$; therefore, $\Delta(h^j) = -n_R(j)$. Then, $\mu(h^j, s_j) = \tilde{\mu}_j(-n_R(j), s_j)$. Since all voting after that of $(j-1)$ is uninformative, $\mu(h^i, s_i) = \mu(h^j, s_i) = \tilde{\mu}_j(-n_R(j), s_i)$. A simple variant of the argument in Case 1 implies that $\tilde{\mu}_j(-n_R(j), r) < \frac{1}{2}$, and therefore Condition (1) of the PBV definition requires that Neutral voter i vote R given $s_i = r$. Consider now $s_i = l$. Since $g_j(-n_R(j)) < \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$, it follows that $\tilde{\mu}_j(-n_R(j), l) < \frac{1}{2}$, and therefore Condition 1 of the PBV definition requires that Neutral voter i vote R even following $s_i = l$. □

Remark 1. *As promised in the text, we argue here that the tie-breaking Condition (2) of the PBV definition only matters for a non-generic set of parameters* $(\pi, \gamma, \tau_L, \tau_R)$. *Observe that from the proof of Proposition 1, the posterior of voter* i *having observed a history* h^i *and private signal* s_i *is* $\frac{1}{2}$ *if and only if* $\Psi(h^i) = 0$ *and* $g_i(\Delta(h^i)) \in \left\{\frac{(1-\pi)(1-\gamma)}{\pi\gamma}, \frac{(1-\pi)\gamma}{\pi(1-\gamma)}\right\}$. *For* *any particular* (π, γ, τ_L)*, this occurs for at most a countable collection of* τ_R*. Therefore, for a given* (π, γ)*, the set*

$$\Gamma_{\pi,\gamma} \equiv \left\{ (\tau_L, \tau_R) \in \left(0, \frac{1}{2}\right)^2 : g_i(\Delta) \in \left\{ \frac{(1-\pi)(1-\gamma)}{\pi\gamma}, \frac{(1-\pi)\gamma}{\pi(1-\gamma)} \right\} \textit{ for some } i \in \mathbb{Z}^+ \textit{ and } |\Delta| \leq i \right\}$$

is isomorphic to a 1-dimensional set. Thus, the need for tie-breaking arises only for a set of parameters $(\pi, \gamma, \tau_L, \tau_R)$ *of (Lebesgue) measure* 0.

Theorem 1 on pp. 11

Proof. The proof consists of two steps: first, we show that there must almost surely be a herd in the limit as the population size $n \to \infty$; second, we show that this implies the finite population statement of the Theorem. Assume without loss of generality that the true state is R. (If the true state is L, one proceeds identically, but using the inverse of the likelihood ratio λ_i).

Step 1: As discussed in the text, by the Martingale Convergence Theorem for non-negative random variables (Billingsley, 1995, pp. 468–469), $\lambda_i \overset{a.s.}{\to} \lambda_\infty$ with *Support* $(\lambda_\infty) \subseteq [0, \infty)$. Define $\bar{\Lambda} \equiv [0, \underline{b}] \cup [\bar{b}, \infty)$ and $\Lambda = [0, \underline{b}) \cup (\bar{b}, \infty)$, where $\underline{b}$ (resp. $\bar{b}$) is the likelihood ratio such that the associated public belief that the state is L causes the posterior upon observing signal l (resp. r) to be exactly $\frac{1}{2}$. Note that by their definitions, $\underline{b} < \frac{1}{2} < \bar{b}$. To prove that there must almost surely be a herd in the limit, it needs to be shown that eventually $\langle \lambda_i \rangle \in \Lambda$ almost surely.[17]

We claim that *Support* $(\lambda_\infty) \subseteq \bar{\Lambda}$. To prove this, fix some $x \notin \bar{\Lambda}$ and suppose towards contradiction that $x \in$ *Support* (λ_∞). Since voting is informative when $\lambda_i = x$, the probability of each vote is continuous in the likelihood ratio around x. Moreover, the updating process on the likelihood ratio following each vote is also continuous around x. Thus, Theorem B.2 of Smith and Sorensen (2000) applies, implying that for both possible votes, either (i) the probability of the vote is 0 when the likelihood ratio is x; or (ii) the updated likelihood ratio following the vote remains x. Since voting is informative, neither of these two is true—contradiction.

The argument is completed by showing that $\Pr(\lambda_\infty \in \Lambda) = 1$. Suppose not, towards contradiction. Then since *Support* $(\lambda_\infty) \subseteq \bar{\Lambda}$, it must be that $\Pr\left(\lambda_\infty \in \left\{\underline{b}, \bar{b}\right\}\right) > 0$. Without loss of generality, assume $\Pr(\lambda_\infty = \underline{b}) > 0$; the argument is analogous if $\Pr\left(\lambda_\infty = \bar{b}\right) > 0$. Observe that if $\lambda_m < \underline{b}$ for some m, then by definition of $\underline{b}$ and PBV, $\lambda_{m+1} = \lambda_m < \underline{b}$ and this sequence of public likelihood ratios converges to a point less than $\underline{b}$. Thus $\Pr(\lambda_\infty = \underline{b}) > 0$ requires that for any $\varepsilon > 0$, eventually $\langle \lambda_i \rangle \in [\underline{b}, \underline{b} + \varepsilon)$ with strictly positive probability. But notice that by the definition of $\underline{b}$, if $\lambda_i = \underline{b}$ then voter i votes informatively under PBV and thus if $\lambda_i = \underline{b}$, either $\lambda_{i+1} < \underline{b}$ (if $v_i = R$) or $\lambda_{i+1} = \frac{1}{2}$ (if $v_i = L$). By continuity of the updating process in the public likelihood ratio on the set $\left[\underline{b}, \bar{b}\right]$, it follows that if $\varepsilon > 0$ is chosen small

[17] To be clear, when we say that $\langle \lambda_i \rangle$ eventually lies (or does not lie) in some set S almost surely, we mean that with probability one there exists some $k < \infty$ that for all $i > k$, $\lambda_k \in (\notin)\ S$.

enough, then $\lambda_i \in [\underline{b}, \underline{b}+\varepsilon)$ implies that $\lambda_{i+1} \notin [\underline{b}, \underline{b}+\varepsilon)$. This contradicts the requirement that for any $\varepsilon > 0$, eventually $\langle \lambda_i \rangle \in [\underline{b}, \underline{b}+\varepsilon)$ with strictly positive probability.

Step 2: Since $\lambda_i \overset{a.s.}{\to} \lambda_\infty$, λ_i converges in probability to λ_∞, i.e. for any $\delta, \eta > 0$, there exists $\overline{n} < \infty$ such that for all $n > \overline{n}$, $\Pr(|\lambda_n - \lambda_\infty| \geq \delta) < \eta$. Since $\Pr(\lambda_\infty \in \Lambda) = 1$, for any $\varepsilon > 0$, we can pick $\delta > 0$ small enough such that $\Pr\left(\lambda_\infty \in [0, \underline{b} - \delta) \cup \left(\bar{b} + \delta, \infty\right)\right) > 1 - \frac{\varepsilon}{2}$. Pick $\eta = \frac{\varepsilon}{2}$. With these choices of δ and η, the previous statement implies that there exists $\overline{n} < \infty$ such that for all $n > \overline{n}$, $\Pr(\lambda_n \in \Lambda) > 1 - \varepsilon$, which proves the theorem. □

A.2 Proofs for Section 3.2

The proof for Theorem 2 follows the logic of the argument laid out in the text. We need various intermediate steps to prove the results. Throughout, to prove that PBV is an equilibrium, we assume that the relevant history is undecided since all actions at a decided history yield the same payoffs. We first prove that following PBF is strictly optimal for Partisan voters (conditional on others playing PBV strategies).

Definition 2. *(Winning Prob.) For a history h^i, let $P\left(\Psi\left(h^i\right), \Delta\left(h^i\right), n-i+1, \omega\right)$ be the probability with which L wins given the phase $\Psi\left(h^i\right)$, the vote lead $\Delta\left(h^i\right)$, the number of voters who have not yet voted $(n-i+1)$, and the true state is ω.*

Note that once $\Psi\left(h^i\right) \in \{L, R\}$, all players are voting uninformatively, and therefore, $P\left(\Psi\left(h^i\right), \cdot\right)$ is independent of state. For the subsequent results, let K denote $n-i$, and Δ denote $\Delta\left(h^i\right)$.

Lemma 1. *For all h^i, $P\left(\Psi\left(h^i, L\right), \Delta+1, K, \omega\right) \geq P\left(\Psi\left(h^i, R\right), \Delta-1, K, \omega\right)$ for all $\omega \in \{L, R\}$. The inequality is strict if $K > \Delta - 1$.*

Proof. Consider any realized profile of preference types and signals of the remaining K voters given true state ω (conditional on the state, this realization is independent of previous voters' types/signals/votes). In this profile, whenever a voter i votes for L given a vote lead $\Delta - 1$, he would also vote L given a vote lead $\Delta + 1$. Thus, if the type-signal profile is such that L wins given an initial lead of $\Delta - 1$, then L would also win given an initial lead of $\Delta + 1$. Since this applies to an arbitrary type-signal profile (of the remaining K voters, given state ω), it follows that $P\left(\Psi\left(h^i, l\right), \Delta+1, K, \omega\right) \geq P\left(\Psi\left(h^i, r\right), \Delta-1, K, \omega\right)$. That the inequality is strict if $K > \Delta - 1$ follows from the fact that with positive probability, the remaining K voters may all be Partisans, with exactly Δ more R-partisans than L-partisans. In such a case, L wins given initial informative vote lead $\Delta + 1$, whereas R wins given initial informative vote lead $\Delta - 1$. □

Lemma 2. *If all other players are playing PBV and the election is undecided at the current history, it is strictly optimal for a Partisan to vote for her preferred candidate.*

Proof. We will begin by showing that an L-partisan always votes L if others are playing PBV strategies. By voting L, an L-partisan's utility is:

$$\mu\left(h^i, s_i\right) P\left(\Psi\left(h^i, L\right), \Delta+1, K, L\right)+\left(1-\mu\left(h^i, s_i\right)\right) P\left(\Psi\left(h^i, L\right), \Delta+1, K, R\right)$$

If she voted R, her utility is

$$\mu\left(h^i, s_i\right) P\left(\Psi\left(h^i, R\right), \Delta-1, K, L\right)+\left(1-\mu\left(h^i, s_i\right)\right) P\left(\Psi\left(h^i, R\right), \Delta-1, N, R\right)$$

It follows from Lemma 1 that the L-partisan voter i strictly prefers to vote L when the election is undecided (i.e. $K>\Delta-1$).

The same arguments apply *mutatis mutandis* to see that R-partisans strictly prefer to vote R when the election is undecided. □

To show that following PBV is optimal for a Neutral voter (conditional on others following PBV strategies), we need to describe the inferences a Neutral voter makes conditioning on being pivotal. As usual, let a profile of type and signal realizations for all other voters apart from i be denoted

$$\left(t_{-i}, s_{-i}\right) \equiv\left(\left(t_1, s_1\right), \ldots,\left(t_{i-1}, s_{i-1}\right),\left(t_{i+1}, s_{i+1}\right), \ldots,\left(t_n, s_n\right)\right)$$

Given that other players are playing PBV, for any realized profile $\left(t_{-i}, s_{-i}\right)$, i's vote deterministically selects a winner because PBV does not involve mixing. For a vote by voter i, $V_i \in\{L, R\}$, denote the winner of the election $x\left(V_i ;\left(t_{-i}, s_{-i}\right)\right) \in\{L, R\}$. Then, denote the event in which voter i is pivotal as $Piv_i=\left\{\left(t_{-i}, s_{-i}\right): x\left(L ;\left(t_{-i}, s_{-i}\right)\right) \neq x\left(R ;\left(t_{-i}, s_{-i}\right)\right)\right\}$. By arguments identical to Lemma 1, for a given profile $\left(t_{-i}, s_{-i}\right)$, if a subsequent voter after i votes L following $V_i=R$, then she would also do so following $V_i=L$. Therefore,

$$Piv_i=\left\{\left(t_{-i}, s_{-i}\right): x\left(L,\left(t_{-i}, s_{-i}\right)\right)=L \text{ and } x\left(R,\left(t_{-i}, s_{-i}\right)\right)=R\right\} \tag{8}$$

Let $U\left(V_i \mid h^i, s_i\right)$ denote a Neutral Voter i's expected utility from action $V \in\{L, R\}$ when she faces a history h^i and has a private signal, s_i. If $\operatorname{Pr}\left(Piv_i \mid h^i, s_i\right)=0$, then no action is sub-optimal for Voter i. If $\operatorname{Pr}\left(Piv_i \mid h^i, s_i\right)>0$, i's vote changes her expected utility if and only if her vote is pivotal. Therefore, in such cases,

$$U\left(V \mid h^i, s_i\right)>U\left(\tilde{V} \mid h^i, s_i\right) \Leftrightarrow U\left(V \mid h^i, s_i, Piv_i\right)>U\left(\tilde{V} \mid h^i, s_i, Piv_i\right) \text{ for } V \neq \tilde{V}$$

It follows from equation (8) that $U\left(L \mid h^i, s_i, Piv_i\right)=\operatorname{Pr}\left(\omega=L \mid h^i, s_i, Piv_i\right)$ and $U\left(R \mid h^i, s_i, Piv_i\right)=1-\operatorname{Pr}\left(\omega=L \mid h^i, s_i, Piv_i\right)$. Therefore, if $\operatorname{Pr}\left(\omega=L \mid h^i, l, Piv_i\right)>\frac{1}{2}$, it is strictly optimal for a Neutral Voter i to vote for L, and if $\operatorname{Pr}\left(\omega=L \mid h^i, l, Piv_i\right)<\frac{1}{2}$, it is strictly optimal for a Neutral Voter i to vote R.

Lemma 3. *If all other players are playing PBV and the election is undecided at the current history, it is strictly optimal for a Partisan to vote for her preferred candidate.*

Proof. Consider a history, h^i, where $\Psi\left(h^i\right) = L$. Since all future Neutral voters vote uninformatively for L, $\Pr\left(\omega = L|h^i, s_i, Piv_i\right) = \Pr\left(\omega = L|h^i, s_i\right)$, which by construction strictly exceeds $\frac{1}{2}$ for all s_i (since $\Psi\left(h^i\right) = L$). Therefore, a Neutral voter strictly prefers to vote L. An analogous argument applies when $\Psi(h^i) = R$. $\square$

Lemma 4. *If all other players are playing PBV and the election is undecided at the current history, h^i, it is (generically, strictly) optimal for a Neutral voter i to vote informatively when $\Psi\left(h^i\right) = 0$ and $\Delta(h^i) \in \{-n_R(i+1)+1, \dots, n_L(i+1)-1\}$.*

The proof proceeds in a series of steps. We shall first use an intermediate lemma (Lemma 5) to show that if the incentive constraints hold for certain voters at certain histories of the learning phase, then they hold for all other possible histories in the learning phase. This simplifies the verification of many incentive constraints to that of a few important constraints. We shall then verify that those constraints also hold in Lemmas 6, 9, and 10.

Lemma 5. *Consider any h^i where $\Psi\left(h^i\right) = 0$ and $\Delta\left(h^i\right) = \Delta$. Then if it is incentive compatible for Neutral Voter $(i+1)$ to vote informatively when $\Delta\left(h^{i+1}\right) \in \{\Delta-1, \Delta+1\}$, then it is incentive compatible for Neutral Voter i to vote informatively when $\Delta\left(h^i\right) = \Delta$. Moreover, if the incentive compatibility condition for Neutral Voter $(i+1)$ holds strictly at least in one of the two cases when $\Delta\left(h^{i+1}\right) \in \{\Delta-1, \Delta+1\}$, then it holds strictly for Neutral Voter i.*

Proof. We prove that it is optimal for i to vote L given signal $s_i = l$; a similar logic holds for optimality of voting R with signal r. It is necessary and sufficient that

$$\begin{array}{l} \tilde{\mu}_i(\Delta, l)\left[P(0, \Delta+1, K, L) - P(0, \Delta-1, K, L)\right] \\ -\left(1-\tilde{\mu}_i(\Delta, l)\right)\left[P(0, \Delta+1, K, R) - P(0, \Delta-1, K, R)\right] \end{array} \geq 0 \tag{9}$$

Define the state-valued functions $p(\cdot)$ and $q(\cdot)$

$$\begin{aligned} p(\omega) &= \begin{cases} \tau_L + (1-\tau_L-\tau_R)\gamma & \text{if } \omega = L \\ \tau_L + (1-\tau_L-\tau_R)(1-\gamma) & \text{if } \omega = R \end{cases} \\ q(\omega) &= \begin{cases} \tau_R + (1-\tau_L-\tau_R)(1-\gamma) & \text{if } \omega = L \\ \tau_R + (1-\tau_L-\tau_R)\gamma & \text{if } \omega = R \end{cases} \end{aligned}$$

Since voter $i+1$ votes informatively if Neutral (because both $\Delta+1$ and $\Delta-1$ are non-herd leads), the probability that $i+1$ votes L and R in state ω is $p(\omega)$ and $q(\omega)$ respectively. Noting the recursive relation

$$P\left(\Psi\left(h^i\right), \Delta, K+1, \omega\right) = p(\omega)\, P\left(\Psi\left(h^i, l\right), \Delta+1, K, \omega\right) + q(\omega)\, P\left(\Psi\left(h^i, r\right), \Delta-1, K, \omega\right)$$

it follows that the above inequality holds if and only if

$$\begin{aligned}0 \;\leq\; & \tilde{\mu}_i(\Delta,l)\left[\begin{array}{l}\left(P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,L\right)-P\left(0,\Delta,K-1,L\right)\right)p\left(L\right)\\+\left(P\left(0,\Delta,K-1,L\right)-P\left(\Psi\left(h^i,r,r\right),\Delta-2,K-1,L\right)\right)q\left(L\right)\end{array}\right]\\&-\left(1-\tilde{\mu}_i(\Delta,l)\right)\left[\begin{array}{l}\left(P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,R\right)-P\left(0,\Delta,K-1,R\right)\right)p\left(R\right)\\+\left(P\left(0,\Delta,K-1,R\right)-P\left(\Psi\left(h^i,r,r\right),\Delta-2,K-1,R\right)\right)q\left(R\right)\end{array}\right]\end{aligned}$$

Dividing by $p(R)\left(1-\tilde{\mu}_i(\Delta,l)\right)$, the above is equivalent to

$$\begin{aligned}0 \;\leq\; & \left(\begin{array}{l}\frac{\tilde{\mu}_i(\Delta,l)}{1-\tilde{\mu}_i(\Delta,l)}\frac{p(L)}{p(R)}\left[P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,L\right)-P\left(0,\Delta,K-1,L\right)\right]\\-\left[P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,R\right)-P\left(0,\Delta,K-1,R\right)\right]\end{array}\right)\\&+\left(\begin{array}{r}\frac{\tilde{\mu}_i(\Delta,l)}{1-\tilde{\mu}_i(\Delta,l)}\frac{q(L)}{p(R)}\left[P\left(0,\Delta,K-1,L\right)-P\left(\Psi\left(h^i,r,r\right),\Delta-2,K-1,L\right)\right]\\-\frac{q(R)}{p(R)}\left[P\left(0,\Delta,K-1,R\right)-P\left(\Psi\left(h^i,r,r\right)\Delta-2,K-1,R\right)\right]\end{array}\right)\end{aligned}$$

We now argue that each of the two lines of the right hand side above is non-negative.

1. Since $\frac{\tilde{\mu}_i(\Delta,l)}{1-\tilde{\mu}_i(\Delta,l)}=\frac{\pi\gamma}{(1-\pi)(1-\gamma)}g_i(\Delta)$ and $\frac{p(L)}{p(R)}=f(\tau_L,\tau_R)$, it follows that

$$\begin{aligned}\frac{\tilde{\mu}_i(\Delta,l)}{1-\tilde{\mu}_i(\Delta,l)}\frac{p(L)}{p(R)} &= \frac{\pi\gamma}{(1-\pi)(1-\gamma)}g_{i+1}(\Delta+1)\\&= \frac{\tilde{\mu}_{i+1}(\Delta+1,l)}{1-\tilde{\mu}_{i+1}(\Delta+1,l)}\end{aligned}$$

   Since IC holds for voter $i+1$ with vote lead $\Delta+1$, observe that if the election is undecided for $i+1$, $\frac{\tilde{\mu}_{i+1}(\Delta+1,l)}{1-\tilde{\mu}_{i+1}(\Delta+1,l)}\geq\frac{P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,R\right)-P(0,\Delta,K-1,R)}{P\left(\Psi\left(h^i,l,l\right),\Delta+2,K-1,L\right)-P(0,\Delta,K-1,L)}$, which proves that the first line of the desired right hand side is non-negative. If the election is decided for $i+1$ with vote lead $\Delta+1$, then the first line of the desired right hand side is exactly 0.

2. Using the previous identities,

$$\begin{aligned}\frac{\tilde{\mu}_i(\Delta,l)}{1-\tilde{\mu}_i(\Delta,l)}\frac{q(L)}{q(R)} &= \frac{\pi\gamma}{(1-\pi)(1-\gamma)}g_{i+1}(\Delta-1)\\&= \frac{\tilde{\mu}_{i+1}(\Delta-1,l)}{1-\tilde{\mu}_{i+1}(\Delta-1,l)}\end{aligned}$$

   Since IC holds for voter $i+1$ with vote lead $\Delta-1$, observe that if the election is undecided for $i-1$, then $\frac{\tilde{\mu}_{i+1}(\Delta-1,l)}{1-\tilde{\mu}_{i+1}(\Delta-1,l)}\frac{q(L)}{q(R)}\geq\frac{P(0,\Delta,K-1,R)-P\left(\Psi\left(h^i,r,r\right),\Delta-2,K-1,R\right)}{P(0,\Delta,K-1,L)-P\left(\Psi\left(h^i,r,r\right),\Delta-2,K-1,L\right)}$, and thus the second line of the desired right hand side is non-negative. If the election is decided for $i+1$ with vote lead $\Delta-1$, then the second line of the desired right hand side is exactly 0.

Observe that if incentive compatibility holds strictly for voter $i+1$ in either one of the two cases, then at least one of the two lines of the right hand side is strictly positive, and consequently inequality (9) must hold strictly. □

By the above Lemma, we are left to only check the incentive conditions for a Neutral voter i with undecided history h^i such that $\Psi(h^i) = 0$, but voter $i+1$ will not vote informatively when Neutral if either $v_i = L$ or $v_i = R$. This possibility can be divided into two cases:

1. either i's vote causes the phase to transition into a herding phase; or
2. i is the final voter ($i = n$) and $\Delta(h^n) = 0$.

Lemma 6 below deals with the latter case; Lemmas 9 and 10 concern the former. (Lemmas 7 and 8 are intermediate steps towards Lemma 9.)

Lemma 6. *If there exists a history h^n such that $\Psi(h^n) = 0$ and $\Delta(h^n) = 0$, then it is incentive compatible for Voter n to vote informatively. For generic parameters of the game, the incentive compatibility conditions hold strictly.*

Proof. Since $\Delta(h^n) = 0$ and n is the final voter, $\Pr(\omega = L|h^n, s_n, Piv_n) = \Pr(\omega = L|h^n, s_n)$. Since $\Psi(h^n) = 0$, $\Pr(\omega = L|h^n, l) \geq \frac{1}{2} \geq \Pr(\omega = L|h^n, r)$. Therefore, voting informatively is incentive compatible. Recall from Remark 1 that $\Pr(\omega = L|h^n, s_n) = \frac{1}{2}$ for some $s_n \in \{l, r\}$ only if $(\pi, \gamma, \tau_L, \tau_R)$ is such that $(\tau_L, \tau_R) \in \Gamma_{\pi,\gamma}$, which is a set of (Lebesgue) measure 0. If $(\tau_L, \tau_R) \notin \Gamma_{\pi,\gamma}$, then given that $\Psi(h^n) = 0$, $\Pr(\omega = L|h^n, l) > \frac{1}{2} > \Pr(\omega = L|h^n, r)$, and therefore for generic parameters, voting informatively is strictly optimal for voter n. $\square$

Lemma 7. *Consider any h^i where $\Psi(h^i) = 0$ and $\Delta(h^i) = \Delta$. Then, $P(\Psi(h^i, l), \Delta+1, K, L) \geq P(\Psi(h^i, l), \Delta+1, K, R)$ and $P(\Psi(h^i, r), \Delta-1, K, L) \geq P(\Psi(h^i, r), \Delta-1, K, R)$ implies $P(0, \Delta, K+1, L) \geq P(0, \Delta, K+1, R)$.*

Proof. Simple manipulations yield

$$\begin{aligned}
& P(0, \Delta, K+1, L) - P(0, \Delta, K+1, R) \\
= & \begin{array}{l} p(L) P(\Psi(h^i, l), \Delta+1, K, L) + q(L) P(\Psi(h^i, r), \Delta-1, K, L) \\ -\left[p(R) P(\Psi(h^i, l), \Delta+1, K, R) + q(R) P(\Psi(h^i, r), \Delta-1, K, R)\right] \end{array} \\
\geq & \begin{array}{l} p(L) P(\Psi(h^i, l), \Delta+1, K, L) + q(L) P(\Psi(h^i, r), \Delta-1, K, L) \\ -\left[p(L) P(\Psi(h^i, l), \Delta+1, K, R) + q(L) P(\Psi(h^i, l), \Delta-1, K, R)\right] \end{array} \\
= & \begin{array}{l} p(L)\left[P(\Psi(h^i, l), \Delta+1, K, L) - P(\Psi(h^i, l), \Delta+1, K, R)\right] \\ +q(L)\left[P(\Psi(h^i, r), \Delta-1, K, L) - P(\Psi(h^i, r), \Delta-1, K, R)\right] \end{array} \\
\geq & \; 0
\end{aligned}$$

where the first inequality uses the fact that $p(L) \geq p(R)$ and $P(\Delta+1, K, R) \geq P(\Delta-1, K, R)$; and the last inequality uses the hypotheses of the Lemma. $\square$

Lemma 8. *For all h^i, $P(\Psi(h^i), \Delta, K+1, L) \geq P(\Psi(h^i), \Delta, K+1, R)$.*

Proof. Base Step: The Claim is true when $K = 0$. To see this, first note that Δ must be even for $P\left(\Psi\left(h^i\right), \Delta, 1, \omega\right)$ to be well-defined. If $\Delta \neq 0$ (hence $|\Delta| \geq 2$), then $P\left(\Psi\left(h^i\right), \Delta, 1, L\right) = P\left(\Psi\left(h^i\right), \Delta, 1, R\right)$. For $\Delta = 0$, we have $P(0, 1, L) = p(L) > p(R) = P(0, 1, R)$.

Inductive Step: For any $K \geq 2$, the desired inequality trivially holds if $\Delta \in \{-n_R, n_L\}$ because $P(n_L, K, L) = P(n_L, K, R)$ and $P(-n_R, K, L) = P(-n_R, K, R)$. So it remains to consider only $\Delta \in \{-n_R + 1, \ldots, n_L - 1\}$. Assume inductively that $P(\Delta + 1, K - 1, L) \geq P(\Delta + 1, K - 1, R)$ and $P(\Delta - 1, K - 1, L) \geq P(\Delta - 1, K - 1, R)$. [The Base Step guaranteed this for $K = 2$.] Using the previous Lemma, it follows that $P(\Delta, K, L) \geq P(\Delta, K, R)$ for all $\Delta \in \{-n_R + 1, \ldots, n_L - 1\}$. $\square$

Lemma 9. *Consider history h^i such that $\Delta\left(h^i\right) = n_L(i+1) - 1$ and $\Psi\left(h^i\right) = 0$. Then if all other voters are playing PBV, and a neutral voter i receives signal r, it is strictly optimal for her to vote R. Analogously, if $\Delta\left(h^i\right) = -n_R(i+1) + 1$, and if all other voters are playing PBV, and a neutral voter i receives signal l, it is strictly optimal to vote L.*

Proof. Consider h^i such that $\Delta\left(h^i\right) = n_L(i+1) - 1 = \Delta$, and $\Psi\left(h^i\right) = 0$. For it to be strictly optimal for the voter to vote informatively, it must be that

$$\begin{aligned} & \tilde{\mu}_i(\Delta, r) P(0, \Delta - 1, K, L) + (1 - \tilde{\mu}(\Delta, r))(1 - P(0, \Delta - 1, K, R)) \\ > \; & \tilde{\mu}_i(\Delta, r) P(L, \Delta + 1, K, L) + (1 - \tilde{\mu}_i(\Delta, r))(1 - P(L, \Delta + 1, K, R)) \end{aligned}$$

which is equivalent to

$$\frac{\tilde{\mu}_i(\Delta, r)}{1 - \tilde{\mu}_i(\Delta, r)} < \frac{P(L, \Delta + 1, K, R) - P(0, \Delta - 1, K, R)}{P(L, \Delta + 1, K, L) - P(0, \Delta - 1, K, L)} \tag{10}$$

By Lemma 8, $P(0, \Delta - 1, K, L) \geq P(0, \Delta - 1, K, R)$, and by definition, $P(L, \Delta + 1, K, R) = P(L, \Delta + 1, K, L)$. Therefore, the right-hand side of (10) is bounded below by 1. Since $\Psi\left(h^i\right) = 0$, $\mu\left(h^i, r\right) = \tilde{\mu}_i(\Delta, r) < \frac{1}{2}$, the left-hand side of (10) is strictly less than 1, establishing the strict inequality. An analogous argument applies to prove the case where $\Delta\left(h^i\right) = -n_R(i+1) + 1$ and $s_i = l$. $\square$

Lemma 10. *Consider history h^i such that $\Delta\left(h^i\right) = n_L(i+1) - 1$ and $\Psi\left(h^i\right) = 0$. Then if all other voters are playing PBV, and a neutral voter i receives signal l, it is optimal for her to vote L. Analogously, if $\Delta\left(h^i\right) = -n_R(i+1) + 1$, and if all other voters are playing PBV, and a neutral voter i receives signal r, it is optimal to vote R. For generic parameters of the game, the optimality is strict.*

Proof. Consider the information set where $\Psi\left(h^i\right) = 0$, $\Delta\left(h^i\right) = n_L(i+1) - 1$, and $s_i = l$. By the discussion in the text (p. 26), it suffices to show that $\Pr\left(\omega = L | h^i, l, Piv_i\right) \geq \frac{1}{2}$. For any i, and for any $k > i$, let ξ_k^Ψ be the set of types $\{(t_j, s_j)\}_{j \neq i}$ that is consistent with history h^i, induces $\left(\Psi\left(h^{k-1}\right), \Psi\left(h^k\right)\right) = (0, \Psi)$ where $\Psi \in \{L, R\}$ after $v_i = R$, and where i's vote is pivotal. Let $K^\Psi = \left\{k > i : \xi_k^\Psi \neq \emptyset\right\}$. Denote by ξ_Δ^0 the set of types $\{(t_j, s_j)\}_{j \neq i}$ that are consistent with

h^i, induces $\Psi\left(h^n\right)=0$ and $\Delta\left(h^{n+1}\right)=\Delta<0$ after $v_i=R$, and where i's vote is pivotal. Let $K_\Delta^0=\left\{\Delta:\xi_\Delta^0\neq\emptyset\right\}$. Observe that since the event $\left(h^i,Piv_i\right)=\cup_\Psi\left(\cup_{k\in K^\Psi}\xi_k^\Psi\right)\cup\left(\cup_{\Delta\in K_\Delta^0}\xi_\Delta^0\right)$, by the definition of conditional probability

$$\begin{aligned}\Pr\left(\omega=L|h^i,l,Piv_i\right) &= \sum_{\Psi\in\{L,R\}}\sum_{k\in K^\Psi}\Pr\left(\xi_k^\Psi|h^i,l,Piv_i\right)\Pr\left(\omega=L|\xi_k^\Psi,l\right)\\ &+\sum_{\Delta\in K_\Delta^0}\Pr\left(\xi_\Delta^0|h^i,l,Piv_i\right)\Pr\left(\omega=L|\xi_\Delta^0,l\right)\end{aligned}$$

We shall argue that $\Pr\left(\omega=L|h^i,l,Piv_i\right)\geq\frac{1}{2}$ by showing that $\Pr\left(\omega=L|\xi_k^L,l\right)>\frac{1}{2}$ for each $k\in K^L$, $\Pr\left(\omega=L|\xi_k^R,l\right)\geq\frac{1}{2}$ for each $k\in K^R$, and $\Pr\left(\omega=L|\xi_\Delta^0,l\right)\geq\frac{1}{2}$ for each $\Delta\in K_\Delta^0$.

Consider $k\in K^L$: ξ_k^L denotes a set of signal-type realizations that induce an L-herd after the vote of voter $(k-1)$ (and meet the aforementioned conditions). Since only votes in the learning phase are informative,

$$\Pr\left(\omega=L|\xi_k^L,l\right)=\Pr\left(\omega=L|l,\Psi\left(h^{k-1}\right)=0,\Delta\left(h^k\right)=n_L\left(k\right)\right)$$

Given that $v_i=R$, the informational content of this event is equivalent to a history $\tilde{h}^{k-1}$ where $\Delta\left(\tilde{h}^{k-1}\right)=n_L\left(k\right)+1$, and all Neutrals are assumed to have voted informatively. Therefore,

$$\Pr\left(\omega=L|\xi_k^L,l\right)=\frac{\pi\gamma g_{k-1}\left(n_L\left(k\right)+1\right)}{\pi\gamma g_{k-1}\left(n_L\left(k\right)+1\right)+\left(1-\pi\right)\left(1-\gamma\right)}$$

Observe that $g_{k-1}(n_L\left(k\right)+1)>g_k\left(n_L\left(k\right)\right)>\frac{(1-\pi)\gamma}{\pi(1-\gamma)}$. Therefore, $\Pr\left(\omega=L|\xi_k^\Psi,l\right)>\frac{\gamma^2}{\gamma^2+(1-\gamma)^2}>\frac{1}{2}$.

Now consider $k\in K^R$: ξ_k^R denotes a set of signal-type realizations that induce an R-herd after the vote of voter $(k-1)$ (and meet the aforementioned conditions). As before, only votes in the learning phase contain information about the state of the world; thus, $\Pr\left(\omega=L|\xi_k^R,l\right)=\Pr\left(\omega=L|l,\Psi\left(h^{k-1}\right)=0,\Delta\left(h^k\right)=-n_R\left(k\right)\right)$. Given that $v_i=R$, the informational content is equivalent to a history $\tilde{h}^{k-1}$ where $\Delta\left(\tilde{h}^{k-1}\right)=-n_R\left(k\right)+1$, and all neutrals are assumed to have voted informatively. Therefore,

$$\Pr\left(\omega=L|\xi_k^R,l\right)=\frac{\pi\gamma g_{k-1}\left(-n_R\left(k\right)+1\right)}{\pi\gamma g_{k-1}\left(-n_R\left(k\right)+1\right)+\left(1-\pi\right)\left(1-\gamma\right)}$$

As by assumption, $\Delta\left(h^{k-1}\right)=-n_R\left(k\right)+1$ and $\Psi\left(h^{k-1}\right)=0$, we have $g_{k-1}\left(-n_R\left(k\right)+1\right)\geq\frac{(1-\pi)(1-\gamma)}{\pi\gamma}$. Therefore, $\Pr\left(\omega=L|\xi_k^R,l\right)\geq\frac{1}{2}$.

Now consider the event $\Delta\in K_\Delta^0$: ξ_Δ^0 denotes a set of signal-type realizations that induce no herd and a final vote lead of $\Delta<0$. Therefore,

$$\Pr\left(\omega=L|\xi_\Delta^0,l\right)=\frac{\pi\gamma g_n\left(\Delta+1\right)}{\pi\gamma g_n\left(\Delta+1\right)+\left(1-\pi\right)}$$

Since $\Delta(h^n) \in \{\Delta-1, \Delta+1\}$and $\Psi(h^n) = 0$, we have $g_n(\Delta+1) \geq \frac{(1-\pi)(1-\gamma)}{\pi\gamma}$, and therefore $\Pr\left(\omega = L|\xi_\Delta^0, l\right) \geq \frac{1}{2}$.

We use the above three facts to deduce that $\Pr\left(\omega = L|h^i, l, Piv_i\right) \geq \frac{1}{2}$: observe that

$$\sum_{\Psi\in\{L,R\},k\in K^\Psi} \Pr\left(\xi_k^\Psi|h^i, l, Piv_i\right) + \sum_{\Delta\in K_\Delta^0} \Pr\left(\xi_\Delta^0|h^i, l, Piv_i\right) = 1$$

Therefore, $\Pr\left(\omega = L|h^i, l, Piv_i\right)$ is a convex combination of numbers that are bounded below by $\frac{1}{2}$.

An analogous argument can be made to ensure optimality at the information set where $\Delta\left(h^i\right) = -n_R(i+1)+1$ and $s_i = r$.

Let us now explain why the incentive conditions hold strictly for generic parameters of the game. Observe that from our arguments above that indifference arises only if there exists some $k \leq n$ and history h^k such that $\Pr\left(\omega = L|h^k, s_k\right) = \frac{1}{2}$. Recall from Remark 1 that this can hold only if $(\pi, \gamma, \tau_L, \tau_R)$ is such that $(\tau_L, \tau_R) \in \Gamma_{\pi,\gamma}$, which is a set of (Lebesgue) measure 0. If $(\tau_L, \tau_R) \notin \Gamma_{\pi,\gamma}$, then for every k, $\Psi\left(h^k\right) = 0$ implies that $\Pr\left(\omega = L|h^k, l\right) > \frac{1}{2} > \Pr\left(\omega = L|h^k, r\right)$. Therefore, for generic parameters of the game, following PBV is strictly optimal for Voter n regardless of history. $\square$

Lemmas 5, 6, 9, and 10 establish Lemma 4: conditional on all others playing according to PBV, it is optimal for Neutrals to vote informatively in the learning phase. Observe that generic parameters of the game yield strict optimality of the incentive conditions in Lemmas 6 and 10, and therefore, by Lemma 5, all the incentive conditions in the learning phase hold strictly generically.

A.3 Proofs for Section 4

Theorem 3 on pp. 16

Proof. The result follows from minor modifications of Theorem 2; in particular, modifying Lemma 6 to $\Psi(h^n) = 0$ and $\Delta(h^n) \in \{\lfloor qn \rfloor, \lceil qn \rceil\}$. $\square$

Theorem 4 on pp. 17

Proof. We shall consider a PBV strategy profile and consider two threshold rules q and q' where $\tau_L < q < q' < 1-\tau_R$. Given a profile of n votes, let S_n denote the total number of votes cast in favor of L.

Pick $\varepsilon > 0$. From Theorem 1, we know that there exists $\bar{k}$ such that for all $k \geq \bar{k}$, $\Pr\left(\Psi\left(h^k\right) = L\right) + \Pr\left(\Psi\left(h^k\right) = R\right) > 1 - \frac{\varepsilon}{2}$. Pick any $k \geq \bar{k}$. By the Weak Law of Large Numbers, for every $\kappa > 0$, $\lim_{n\to\infty} \Pr\left(\left|\frac{S_n}{n} - (1-\tau_R)\right| < \kappa|\Psi\left(h^k\right) = L\right) = 1$ and $\lim_{n\to\infty} \Pr\left(\left|\frac{S_n}{n} - \tau_L\right| < \kappa|\Psi\left(h^k\right) = R\right) = 1$. Pick some $\kappa < \min\{(1-\tau_R) - q', q - \tau_L\}$. There exists some $\bar{n} > k$ such that for all $n \geq \bar{n}$, $\Pr\left(\left|\frac{S_n}{n} - (1-\tau_R)\right| < \kappa|\Psi\left(h^k\right) = L\right) > 1 - \frac{\varepsilon}{2}$

and $\Pr\left(\left|\frac{S_n}{n} - \tau_L\right| < \kappa | \Psi\left(h^k\right) = R\right) > 1 - \frac{\varepsilon}{2}$. Observe that by the choice of κ, $\left|\frac{S_n}{n} - (1 - \tau_R)\right| < \kappa$ implies that L wins under both rules q and q' whereas $\left|\frac{S_n}{n} - \tau_L\right| < \kappa$ implies that L loses under both rules q and q'. For any $n \geq \bar{n}$, we have

$$
\begin{aligned}
& \left|Pr(\text{L wins in } G\left(\pi, \gamma, \tau_L, \tau_R; n, q\right) - Pr(\text{L wins in } G\left(\pi, \gamma, \tau_L, \tau_R; n, q'\right)\right| \\
< & \left|\sum_{x \in \{L,R\}} \Pr\left(\Psi\left(h^k\right) = x\right) \Pr\left(\frac{S_n}{n} \in \left(q, q'\right] | \Psi\left(h^k\right) = x\right)\right| + \frac{\varepsilon}{2} \\
< & \ \varepsilon
\end{aligned}
$$

which proves part (a) of the Theorem.

For part (b), consider any $q < \tau_L$. The probability with which a voter votes L is at least τ_L. Therefore, invoking the Weak Law of Large Numbers, $\lim_{n \to \infty} \Pr\left(\frac{S_n}{n} < q\right) = 0$. The argument is analogous for part (c). □

Proposition 3 on pp. 18

Proof. We shall begin by describing the PBV strategy profile and then proceed to establish that it is an equilibrium. Observe that for $\tau_L = \tau_R = 0$, for every i, $g_i\left(k\right) = \left(\frac{\gamma}{1-\gamma}\right)^k$. As $\frac{\gamma}{1-\gamma} > \frac{(1-\pi)\gamma}{\pi(1-\gamma)} > 1 > \frac{1-\gamma}{\gamma} > \frac{(1-\pi)(1-\gamma)}{\pi\gamma} > \left(\frac{1-\gamma}{\gamma}\right)^2$. Therefore,

$$
\left(n_L\left(i\right), n_R\left(i\right)\right) = \begin{cases} (1,3) & \text{if } i \text{ is even} \\ (2,2) & \text{if } i \text{ is odd} \end{cases}
$$

Therefore, PBV prescribes that if $\Psi\left(h^{i-1}\right) = 0$, and $\Delta\left(h^i\right) \in \{1, -2\}$, then $\Psi\left(h^i\right) \neq 0$. To show that this strategy profile is an equilibrium, we shall consider the incentives in the learning and herding phases separately.

Observe that if $\Psi\left(h^i\right) = L(R)$, voting for $R\,(L)$ occurs with zero-probability on the path of play. We shall consider a belief-restriction that specifies that any off-path vote is simply ignored and does not affect the public belief. Given this belief-restriction, it can be shown using Theorem 2 (particularly Lemmas 3, 4) that PBV is an equilibrium. □

B Supplementary Results on Cut-Point Voting

In this Appendix, we consider the class of *Cut-Point Voting* (CPV) strategy profiles introduced by Callander (2004a). While this class does entail some restrictions, it covers a wide class of strategies, generalizing PBV to permit behavior ranging from fully informative to uninformative voting. We prove that for generic parameters, any equilibrium within this class leads to herding with high probability in large elections. This result is of interest because it suggests that our results are more general beyond the PBV equilibrium.

To define a CPV profile, let $\mu\left(h^i\right) \equiv \Pr\left(\omega = L|h^i\right)$, so that $\mu\left(h^i\right)$ denotes the public belief following history h^i.

Definition 3. *A strategy profile,* $\mathbf{v}$*, is a Cut-Point Voting (CPV) strategy profile if there exist* $0 \leq \mu_* \leq \mu^* \leq 1$ *such that for every voter* i*, history* h^i*, and signal* s_i*,*

$$\begin{aligned} v_i\left(N, h^i, s_i\right) &= \begin{cases} L & \textit{if } \mu\left(h^i\right) > \mu^* \textit{ or } \left\{\mu\left(h^i\right) = \mu_* \textit{ and } s_i = l\right\} \\ R & \textit{if } \mu\left(h^i\right) < \mu_* \textit{ or } \left\{\mu\left(h^i\right) = \mu^* \textit{ and } s_i = r\right\} \end{cases} \\ v_i\left(L_p, h^i, s_i\right) &= L \\ v_i\left(R_p, h^i, s_i\right) &= R \end{aligned}$$

In a CPV strategy profile, Neutrals vote according to their signals alone if and only if the public belief when it is their turn to vote lies within $[\mu_*, \mu^*]$; otherwise, a Neutral votes for one of the candidates independently of her private signal. Denote a CPV profile with belief thresholds μ_* and μ^* as $CPV\left(\mu_*, \mu^*\right)$. These thresholds define the extent to which a CPV profile weighs past history relative to the private signal: $CPV\left(0,1\right)$ corresponds to informative voting (by Neutrals) where history never influences play, whereas $CPV\left(1-\gamma, \gamma\right)$ corresponds to PBV by Neutrals. Similarly, $CPV\left(0,0\right)$ and $CPV\left(1,1\right)$ represent strategy profiles where every Neutral votes uninformatively for candidate L and R respectively. Therefore, CPV captures a variety of behavior for Neutrals.

A CPV equilibrium is an equilibrium whose strategy profile is a CPV profile. While we are unable to derive a tight characterization of what non-PBV but CPV profiles—if any—constitute equilibria, we can nevertheless show that generically, large elections lead to herds with high probability within the class of CPV equilibria.

Theorem 5. *For every* $(\pi, \gamma, \tau_L, \tau_R)$ *such that* $\tau_L \neq \tau_R$*, and for every* $\varepsilon > 0$*, there exists* $\overline{n} < \infty$ *such that for all* $n > \overline{n}$*, if voters play a CPV equilibrium,* $\Pr[$*a herd develops in* $G\left(\pi, \gamma, \tau_L, \tau_R; n\right)] > 1 - \varepsilon$*.*

Proof

We argue through a succession of lemmas that there exist $\overline{\mu}^* < 1$ and $\underline{\mu}_* > 0$ such that when $\tau_L \neq \tau_R$, in a large enough election, a CPV $\left(\mu_*, \mu^*\right)$ is an equilibrium only if $\underline{\mu}_* \leq \mu_* < \mu^* \leq \overline{\mu}^*$. This suffices to prove the Theorem, because then, the arguments of Theorem 1 apply with trivial modifications. Note that in all the lemmas below, it is implicitly assumed when we consider a particular voter's incentives that she is at an undecided history.

For any $CPV(\mu_*, \mu^*)$, we can define threshold sequences $\{\tilde{n}_L(i)\}_{i=i}^{\infty}$ and $\{\tilde{n}_R(i)\}_{i=i}^{\infty}$ similarly to $\{n_L(i)\}_{i=i}^{\infty}$ and $\{n_R(i)\}_{i=i}^{\infty}$, except using the belief threshold μ^* (resp. μ_*) in place of the PBV threshold γ (resp. $1-\gamma$). That is, for all i such that $g_i(i-1) \leq \frac{(1-\pi)\mu^*}{\pi(1-\mu^*)}$, set $\tilde{n}_L(i) = i$. If $g_i(i-1) > \frac{(1-\pi)\mu^*}{\pi(1-\mu^*)}$, we set $\tilde{n}_L(i)$ to be the unique integer that solves $g_i(\tilde{n}_L(i) - 2) \leq \frac{(1-\pi)\mu^*}{\pi(1-\mu^*)} < g_i(\tilde{n}_L(i))$. For all i such that $g_i(-(i-1)) \geq \frac{(1-\pi)(\mu_*)}{\pi(1-\mu_*)}$, set $\tilde{n}_R(i) = i$. If $g_i(-(i-1)) < \frac{(1-\pi)\mu_*}{\pi(1-\mu_*)}$, set $\tilde{n}_R(i)$ to be the unique integer that solves $g_i(-\tilde{n}_R(i) + 2) \geq \frac{(1-\pi)\mu_*}{\pi(1-\mu_*)} > g_i(-\tilde{n}_R(i))$. Given these thresholds sequences $\tilde{n}_L$ and $\tilde{n}_R$, we define the phase mapping $\tilde{\Psi} : h^i \to \{L, 0, R\}$ in the obvious way that extends the PBV phase mapping Ψ. We state without proof the following generalization of Proposition 1.

Proposition 4. *Fix a parameter set $(\pi, \gamma, \tau_L, \tau_R, n)$. For each $i \leq n$, if voters play $CPV(\mu_*, \mu^*)$ in the game $G(\pi, \gamma, \tau_{L,}\tau_R; n)$, there exist sequences $\{\tilde{n}_L(i)\}_{i=i}^{\infty}$ and $\{\tilde{n}_R(i)\}_{i=i}^{\infty}$ satisfying $|\tilde{n}_C(i)| \leq i$ such that a Neutral voter i votes*

1. *informatively if $\tilde{\Psi}\left(h^i\right) = 0$;*
2. *uninformatively for L if $\tilde{\Psi}\left(h^i\right) = L$;*
3. *uninformatively for R if $\tilde{\Psi}\left(h^i\right) = R$;*

   where $\tilde{\Psi}$ is the phase mapping with respect to $\tilde{n}_L$ and $\tilde{n}_R$. The thresholds $\tilde{n}_L(i)$ and $\tilde{n}_R(i)$ do not depend on the population size, n.

Lemma 11. *There exists $\bar{\mu}^* < 1$ and $\underline{\mu}_* > 0$ such that in any $CPV(\mu_*, \mu^*)$,*

1. *if $\mu^* > \bar{\mu}^*$ then $\tilde{n}_L(i) > n_L(i)$ for all i such that $n_L(i) < i$;*
2. *if $\mu_* < \underline{\mu}_*$, then $-\tilde{n}_R(i) < -n_R(i)$ for all i such that $-n_R(i) > -i$.*

Proof. We give the argument for part (1); it is similar for part (2). Define $\bar{\mu}^*$ by the equality $\frac{\bar{\mu}^*}{1-\bar{\mu}^*} = \frac{\gamma}{1-\gamma} f(\tau_L, \tau_R) f(\tau_R, \tau_L)$. It is straightforward to compute from the definition of $g_i(\cdot)$ that for any k (such that $|k| < i$ and $i-k$ is odd), $g_i(k-2) f(\tau_L, \tau_R) f(\tau_R, \tau_L) = g_i(k)$. Suppose $\mu^* > \bar{\mu}^*$ and there is some i with $\tilde{n}_L(i) \leq n_L(i)$. By the definitions of $n_L(i)$ and $\tilde{n}_L(i)$, and the monotonicity of $g_i(k)$ in k,

$$
\begin{aligned}
g_i(n_L(i) - 2) &= g_i(n_L(i)) \left[f(\tau_L, \tau_R) f(\tau_R, \tau_L)\right]^{-1} \\
&\geq g_i(\tilde{n}_L(i)) \left[f(\tau_L, \tau_R) f(\tau_R, \tau_L)\right]^{-1} \\
&> \frac{(1-\pi)\mu^*}{\pi(1-\mu^*)} \left[f(\tau_L, \tau_R) f(\tau_R, \tau_L)\right]^{-1} \\
&> \frac{(1-\pi)\gamma}{\pi(1-\gamma)}
\end{aligned}
$$

contradicting the definition of $n_L(i)$ which requires that $g_i(n_L(i) - 2) \leq \frac{(1-\pi)\gamma}{\pi(1-\gamma)}$. ∎

Lemma 12. *If all Neutral voters play according to a CPV profile, it is uniquely optimal for an L-partisan to vote L and an R-partisan to vote R.*

Proof. This follows from the weak monotonicity imposed by CPV; trivial modifications to the argument in Lemma 2 establish this result. □

Lemma 13. *In a large enough election, CPV* $(0, 1)$ *is not an equilibrium unless* $\tau_L = \tau_R$.

Proof. Suppose all voters play CPV strategy $(0, 1)$. Without loss of generality assume $\tau_L > \tau_R$; the argument is analogous if $\tau_L < \tau_R$. Let $\varsigma_t(n)$ denote the number of voters of preference-type $t \in \{L, R, N\}$ when the electorate size is n. Denote $\tau_N = 1 - \tau_L - \tau_R$. Suppose voter 1 is Neutral and has received signal l. She is pivotal if and only if amongst the other $n - 1$ voters, the number of L votes is exactly equal the number of R votes. Let $\varsigma_{N,s}(n)$ denote the number of Neutrals who have received signal $s \in \{l, r\}$. Under the CPV profile $(0, 1)$, voter 1 is pivotal if and only if $\varsigma_{N,r}(n) - (\varsigma_{N,l}(n) - 1) = \varsigma_L(n) - \varsigma_R(n)$. By the Weak Law of Large Numbers, for any $\varepsilon > 0$ and any $t \in \{L, R, N\}$, $\lim_{n\to\infty} \Pr\left(\left|\frac{\varsigma_t(n)}{n} - \tau_t\right| < \varepsilon\right) = 1$. Consequently, since $\tau_L > \tau_R$, for any $\varepsilon > 0$ and $k > 0$, there exists $\bar{n}$ such that for all $n > \bar{n}$, $\Pr(\varsigma_L(n) - \varsigma_R(n) > k) > 1 - \varepsilon$. Thus, denoting Piv_1 as the set of preference-type and signal realizations where the Neutral voter 1 with $s_i = l$ is pivotal, we have that for any $\varepsilon > 0$ and $k > 0$, there exists $\bar{n}$ such that for all $n > \bar{n}$, $\Pr(\varsigma_{N,r}(n) - \varsigma_{N,l}(n) > k | Piv_1) > 1 - \varepsilon$. Since $\Pr(\omega = L | \varsigma_{N,r}(n), \varsigma_{N,l}(n))$ is strictly decreasing in $\varsigma_{N,r}(n) - \varsigma_{N,l}(n)$, it follows that by considering k large enough in the previous statement, we can make $\Pr(\omega = L | Piv_1) < \frac{1}{2}$ in large enough elections. Consequently, in large enough elections, voter 1 strictly prefers to vote R when she is Neutral and has received $s_i = l$, which is a deviation from the CPV strategy $(0, 1)$. □

Lemma 14. *In a large enough election, CPV* (μ_*, μ^*) *is not an equilibrium if either* $\mu_* > \frac{1}{2}$ *or* $\mu^* < \pi$.

Proof. If $\mu_* > \pi$, then the first voter votes uninformatively for R if Neutral, and consequently, all votes are uninformative. Thus, conditioning on being pivotal adds no new information to any voter. Since $\mu_1(h^1, l) > \pi > \frac{1}{2}$ (recall that $h^1 = \phi$), voter 1 has an incentive to deviate from the CPV strategy and vote L if she is Neutral and receives signal $s_1 = l$.

If $\mu_* \in \left(\frac{1}{2}, \pi\right]$, let h^{k+1} be a history of k consecutive R votes. It is straightforward that for some integer $k \geq 1$, $\mu(h^k) \geq \mu_* > \mu(h^{k+1})$. Since an R-herd has started when it is voter $k+1$'s turn to vote, conditioning on being pivotal adds to information to voter $k+1$. Suppose voter $k+1$ is Neutral and receives $s_{k+1} = l$. Then since an R-herd has started, she is supposed to vote R. But since $\mu_{k+1}(h^{k+1}, l) > \mu(h^k) \geq \mu_* > \frac{1}{2}$, she strictly prefers to vote L.

If $\mu^* < \pi$, the argument is analogous to the case of $\mu_* > \pi$, noting that $\mu_1(h^1, r) < \frac{1}{2}$ because $\gamma > \pi$. □

Lemma 15. *In a large enough election, CPV* $(\mu_*, 1)$ *is not an equilibrium for any* $\mu_* \in (0, \pi]$.

Proof. Suppose $CPV(\mu_*, 1)$ with $\mu_* \in (0, \pi]$ is an equilibrium. Consider a Neutral voter m with signal $s_m = r$ and history h^m such that $\mu(h^m) \geq \mu_*$ but $\mu\left(h^{m+1}\right) < \mu_*$ following $v_m = R$. (To see that such a configuration can arise in a large enough election, consider a sequence of consecutive R votes by all voters.) Voter m is supposed to vote R in the equilibrium. We will show that she strictly prefers a deviation to voting L in a large enough election.

Claim 1: If the true state is R, then following $v_m = L$, the probability of an R-herd converges to 1 as the electorate size $n \to \infty$. *Proof*: Recall that the likelihood ratio stochastic process $\lambda_i \stackrel{a.s.}{\to} \lambda_\infty$ (where the domain can be taken as $i = m+1, m+2, \ldots$). Since voter i votes informatively if and only if $\lambda_i \geq \frac{\mu_*}{1-\mu_*}$, the argument used in proving Theorem 1 allows us to conclude that $Support(\lambda_\infty) \subseteq \left[0, \frac{\mu_*}{1-\mu_*}\right]$ and $\Pr\left(\lambda_\infty = \frac{\mu_*}{1-\mu_*}\right) = 0$. Consequently, there is a herd on R eventually almost surely in state R.

Claim 2: $\Pr(Piv_m | \omega = R)$ converges to 0 as the electorate size $n \to \infty$. *Proof*: To be explicit, we use superscripts to denote the electorate size n, e.g. we write Piv_m^n instead of Piv_m. Denote

$$
\begin{aligned}
X^n &= \{(t_{-m}, s_{-m}) \in Piv_m : \text{ L-herd after $v_m = L$, R-herd after $v_m = R$}\} \\
Y^n &= \{(t_{-m}, s_{-m}) \in Piv_m : \text{ no herd after $v_m = L$, R-herd after $v_m = R$}\} \\
Z^n &= \{(t_{-m}, s_{-m}) \in Piv_m : \text{ R-herd after $v_m = L$ and $v_m = R$}\}
\end{aligned}
$$

We have $Piv_m^n = X^n \cup Y^n \cup Z^n$; hence it suffices to show that $\Pr(X^n) \to 0$, $\Pr(Y^n) \to 0$, and $\Pr(Z^n) \to 0$. That $\Pr(X^n) \to 0$ and $\Pr(Y^n) \to 0$ follows straightforwardly from Claim 1. To show that $\Pr(Z^n) \to 0$, let Ψ_k^n denote the phase after voter k has voted, i.e. when it is voter $k+1$'s turn to vote. For any n, consider the set of $\{(t_j, s_j)\}_{j=m+1}^n$ such that after $v_m = L$, $\Psi_n^n \neq L$; denote this set Ξ^n. Partition this into the sets that induce $\Psi_n^n = 0$ and $\Psi_n^n = R$, denoted $\Xi^{n,0}$ and $\Xi^{n,R}$ respectively. Clearly, $Z^n \subseteq \Xi^{n,R}$. For any ε, for large enough n, regardless of m's vote, $\Pr(\Psi_n^n = 0) < \varepsilon$ by Claim 1, and thus, $\Pr\left(\Xi^{n,0}\right) < \varepsilon$. Now consider any $n' > n$. $Z^{n'} \subseteq \Xi^n$ because if there is a L-herd following $v_m = L$ with electorate size n, there cannot be an R-herd following $v_m = L$ with electorate size n'. Thus, $\Pr\left(Z^{n'}\right) = \Pr\left(\Xi^{n,0}\right)\Pr\left(Z^{n'}|\Xi^{n,0}\right) + \Pr\left(\Xi^{n,R}\right)\Pr\left(Z^{n'}|\Xi_n^R\right) < \varepsilon + \Pr\left(\Xi^{n,R}\right)\Pr\left(Z^{n'}|\Xi_n^R\right)$ for large enough n. We have $\Pr\left(Z^{n'}|\Xi^{n,R}\right) = \dfrac{\Pr\left(Z^{n'} \cap \Xi^{n,R}\right)}{\Pr\left(\Xi^{n,R}\right)}$. It is straightforward to see that $\Pr\left(Z^{n'} \cap \Xi^{n,R}\right) \to 0$ as $n' \to \infty$, using the fact that $\tau_L < 1 - \tau_L$ and invoking the Weak Law of Large Numbers similarly to Lemma 13. Note that $\Pr\left(\Xi^{n,R}\right)$ is bounded away from 0 because if sufficiently many voters immediately after m are R-partisans, then an R-herd will start regardless of m's vote. This proves that $\Pr\left(Z^{n'}\right) \to 0$.

Claim 3: If the true state is L, then following $v_m = L$, the probability that L wins is bounded away from 0 as the electorate size $n \to \infty$. *Proof*: Define $\xi\left(h^i\right) = \frac{\Pr\left(\omega = R | h^i\right)}{\Pr\left(\omega = L | h^i\right)}$; this generates a stochastic process $\langle \xi_i \rangle$ $(i = m+1, m+2, \ldots)$ which is a martingale conditional on state L, and thus $\langle \xi_i \rangle \stackrel{a.s.}{\to} \xi_\infty$. Note that $\xi_{m+1} < \frac{1-\mu_*}{\mu_*}$ since $\mu(h^m) \geq \mu^*$ and $v_m = L$. Since voter

i votes informatively if and only if $\xi_i \in \left(0, \frac{1-\mu_*}{\mu_*}\right]$, the argument used in proving Theorem 1 allows us to conclude that $Support\left(\xi_\infty\right) \subseteq \{0\} \cup \left[\frac{1-\mu_*}{\mu_*}, \infty\right)$ and $\Pr\left(\xi_\infty = \frac{1-\mu_*}{\mu_*}\right) = 0$. Suppose towards contradiction that $0 \notin Support\left(\xi_\infty\right)$. This implies $\mathbb{E}\left[\xi_\infty\right] > \frac{1-\mu_*}{\mu_*}$. By Fatou's Lemma (Billingsley, 1995, p. 209), $\mathbb{E}\left[\xi_\infty\right] \leq \lim_{n\to\infty} \mathbb{E}\left[\xi_n\right]$; since for any $n \geq m+1$, $\mathbb{E}\left[\xi_n\right] = \xi_{m+1}$, we have $\frac{1-\mu_*}{\mu_*} < \mathbb{E}\left[\xi_\infty\right] \leq \xi_{m+1} < \frac{1-\mu_*}{\mu_*}$, a contradiction. Thus, $0 \in Support\left(\xi_\infty\right)$, and it must be that $\Pr\left(\xi_\infty = 0\right) > 0$. The claim follows from the observation that for any history sequence where $\xi_i\left(h^i\right) \to 0$ it must be that $\Delta\left(h^i\right) \to \infty$.

Consider the expected utility for voter m from voting R or L respectively, conditional on being pivotal: $EU_m\left(v_m = R|Piv_m\right) = \Pr\left(\omega = R|Piv_m\right)$ and $EU_m\left(v_m = L|Piv_m\right) = \Pr\left(\omega = L|Piv_m\right)$. Thus, she strictly prefers to vote L if and only if $\Pr\left(\omega = L|Piv_m\right) > \Pr\left(\omega = R|Piv_m\right)$, or equivalently, if and only if $\Pr\left(Piv_m|\omega = L\right) > \Pr\left(Piv_m|\omega = R\right) \frac{1-\mu_m(h^m,r)}{\mu_m(h^m,r)}$. By Claim 2, $\Pr\left(Piv_m|\omega = R\right)$ converges to 0 as electorate grows. On the other hand, $\Pr\left(Piv_m|\omega = L\right)$ is bounded away from 0, because by Claim 3, the probability that L wins following $v_m = L$ is bounded away from 0, whereas if $v_m = R$, a R-herd starts and thus the probability that R wins converges to 1 as the electorate size grows. Therefore, in a large enough election, $\Pr\left(Piv_m|\omega = L\right) > \Pr\left(Piv_m|\omega = R\right) \frac{1-\mu_m(h^m,r)}{\mu_m(h^m,r)}$, and it is strictly optimal for m to vote L following his signal $s_m = r$, which is a deviation from the CPV strategy. □

Lemma 16. *In a large enough election, $CPV\left(0, \mu^*\right)$ is not an equilibrium for any $\mu^* \in [\pi, 1)$.*

Proof. Analogous to Lemma 15, it can be shown here that in a large enough election there is a voter who when Neutral is supposed to vote L with signal l, but strictly prefers to vote R. □

Lemma 17. *In a large enough election, $CPV\left(\mu_*, \mu^*\right)$ is not an equilibrium if $\mu^* \in (\bar{\mu}^*, 1)$ and $\mu_* \in (0, \frac{1}{2}]$.*

Proof. Fix an equilibrium $CPV\left(\mu_*, \mu^*\right)$ with $\mu^* \in (\bar{\mu}^*, 1)$ and $\mu_* \in (0, \frac{1}{2}]$. By Lemma 11, $\tilde{n}_L(i) > n_L(i)$ for all i. Consider a Neutral voter m with signal $s_m = r$ and history h^m such that $\mu\left(h^m\right) \geq \mu_*$ but $\mu\left(h^{m+1}\right) < \mu_*$ following $v_m = R$. (To see that such a configuration can arise in a large enough election, consider a sequence of consecutive R votes by all voters.) Voter m is supposed to vote R in the equilibrium. We will show that she strictly prefers a deviation to voting L in a large enough election.

First, note that by following the argument of Theorem 1, it is straightforward to show that regardless of m's vote, a herd arises with arbitrarily high probability when the electorate size n is sufficiently large. Define X^n, Y^n, and Z^n as in Lemma 15, where n indexes the electorate size. Plainly, $\Pr\left(Y^n\right) \to 0$. The argument of Claim 2 in Lemma 15 implies with obvious modifications that $\Pr\left(Z^n\right) \to 0$. Finally, $\Pr\left(X^n\right) \nrightarrow 0$ because there exists $m' > m$ such that if $v_i = L$ for all $i \in \{m+1, \ldots, m'\}$, then $\Psi^n_{m'} = L$, and $\Pr\left(v_i = L \text{ for all } i \in \{m+1, \ldots m'\}\right) \geq \left(\tau_L\right)^{m'-m} > 0$. Since $Piv^n_m = X^n \cup Y^n \cup Z^n$, we conclude that as $n \to \infty$, $\Pr\left(X^n|Piv^n_m\right) \to 1$, whereas $\Pr\left(Y^n|Piv^n_m\right) \to 0$ and $\Pr\left(Z^n|Piv^n_m\right) \to 0$. Consequently, for any $\varepsilon > 0$, there exists $\bar{n}$ such that for all $n > \bar{n}$,

$$|EU_m(v_m = L|X^n, s_m = r) - EU_m(v_m = L|Piv_m^n, s_m = r)| < \varepsilon$$

and

$$|EU_m(v_m = R|X^n, s_m = r) - EU_m(v_m = R|Piv_m^n, s_m = r)| < \varepsilon$$

Therefore, it suffices to show that for any $n > m$,

$$EU_m(v_m = L|X^n, s_m = r) > EU_m(v_m = R|X^n, s_m = r),$$

or equivalently, $\Pr(\omega = L|X^n, s_m = r) > \Pr(\omega = R|X^n, s_m = r)$. For any $k \in \{m+1, \dots, n\}$, define

$$X_k^n = \left\{(t_{-m}, s_{-m}) \in Piv_m^n : \ \Psi_{k-1}^n = 0 \text{ but } \Psi_k^n = L \text{ after } v_m = L,\ \Psi_n^n \text{ after } v_m = R\right\}$$

Clearly, this requires $\tilde{n}_L(k+1) < k+1$. For $i \neq j$, $X_i^n \cap X_j^n = \emptyset$, but $X^n = \cup_{k=m+1}^n X_k^n$, and thus $\Pr(\omega|X^n, s_m = r) = \cup_{k=m+1}^n \Pr(\omega|X_k^n, s_m = r)\Pr(X_k^n|X^n)$. It therefore suffices to show that for any $k \in \{m+1, \dots, n\}$, $\Pr(\omega = L|X_k^n, s_m = r) > \Pr(\omega = R|X_k^n, s_m = r)$. Given that $v_m = L$, the informational content of X_k^n is equivalent to a history h^{k+1} where $\Delta\left(h^{k+1}\right) = \tilde{n}_L(k+1) - 2$, and all neutrals are assumed to have voted informatively. Therefore,

$$\Pr(\omega = L|X_k^n, s_m = r) = \frac{\pi\gamma g_{k+1}(\tilde{n}_L(k+1) - 2)}{\pi\gamma g_{k+1}(\tilde{n}_L(k+1) - 2) + (1-\pi)(1-\gamma)}$$

Since $\tilde{n}_L(k+1) < k+1$ and $\tilde{n}_L(i) > n_L(i)$ for all i, it must be that $\tilde{n}_L(k+1) - 2 \geq n_L(k+1)$. Consequently,

$$\begin{aligned}\Pr(\omega = L|X_k^n, s_m = r) &\geq \frac{\pi\gamma g_{k+1}(n_L(k+1))}{\pi\gamma g_{k+1}(n_L(k+1)) + (1-\pi)(1-\gamma)} \\ &> \frac{1}{2}\end{aligned}$$

where the second inequality is by the definition of $n_L(k+1)$. □

Lemma 18. *In a large enough election, CPV (μ_*, μ^*) is not an equilibrium if $\mu^* \in [\pi, 1)$ and $\mu_* \in (0, \underline{\mu}_*)$.*

Proof. Analogous to Lemma 17, it can be shown here that in a large enough election there is a voter who when Neutral is supposed to vote L with signal l, but strictly prefers to vote R. □

References

AUSTEN-SMITH, D., AND J. S. BANKS (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *The American Political Science Review*, 90(1), 34–45.

BANERJEE, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797–817.

BANKS, J. S., AND J. SOBEL (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647–661.

BARTELS, L. M. (1988): *Presidential Primaries and the Dynamics of Public Choice*. Princeton University Press, Princeton, NJ.

BATTAGLINI, M. (2005): "Sequential Voting with Abstention," *Games and Economic Behavior*, 51, 445–463.

BATTAGLINI, M., R. MORTON, AND T. PALFREY (2005): "Efficiency, Equity, and Timing in Voting Mechanisms," unpublished manuscript, Department of Politics, NYU.

BERELSON, B. R., P. F. LAZARSFELD, AND W. N. MCPHEE (1954): *Voting*. University of Chicago Press, Chicago, IL.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

BILLINGSLEY, P. (1995): *Probability and Measure*. Wiley, New York, third edn.

CALLANDER, S. (2004a): "Bandwagons and Momentum in Sequential Voting," Kellogg-MEDS, working paper.

——— (2004b): "Vote Timing and Information Aggregation," Kellogg-MEDS, working paper.

CHO, I.-K., AND D. M. KREPS (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–221.

DASGUPTA, A. (2000): "Social Learning with Payoff Complementarities," London School of Economics.

DEKEL, E., AND M. PICCIONE (2000): "Sequential Voting Procedures in Symmetric Binary Elections," *Journal of Political Economy*, 108, 34–55.

FEDDERSEN, T., AND W. PESENDORFER (1996): "The Swing Voter's Curse," *The American Economic Review*, 86(3), 408–424.

——— (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica*, 65, 1029–1058.

——— (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *The American Political Science Review*, 92(1), 23–35.

Feddersen, T., and A. Sandroni (2006): "Ethical Voters and Costly Information Acquisition," *Quarterly Journal of Political Science*, 1, 287–311.

Fey, M. (2000): "Informational Cascades and Sequential Voting," University of Rochester Working Paper.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge, MA.

Gurian, P. (1986): "Resource Allocation Strategies in Presidential Nomination Campaigns," *American Journal of Political Science*, 30(4), 802–821.

Hung, A. A., and C. R. Plott (2001): "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review*, 91(5), 1508–1520.

Kenney, P. J., and T. W. Rice (1994): "The Psychology of Political Momentum," *Political Research Quarterly*, 47, 923–938.

Klumpp, T. A., and M. K. Polborn (2006): "Primaries and the New Hampshire Effect," *Journal of Public Economics*, 31, 1073–1114.

Kohlberg, E., and J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), 1003–1037.

Kreps, D., and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

Morton, R. B., and K. C. Williams (1999): "Information Asymmetries and Simultaneous versus Sequential Voting," *American Political Science Review*, 93(1), 51–67.

Myerson, R. (1998): "Extended Poisson Games and the Condorcet Jury Theorem," *Games and Economic Behavior*, 25, 111–131.

——— (2000): "Large Poisson Games," *Journal of Economic Theory*, 94, 7–45.

Neeman, Z., and G. Orosel (1999): "Herding and the winner's curse in markets with sequential bids," *Journal of Economic Theory*, 85, 91–121.

Popkin, S. L. (1991): *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. The University of Chicago Press, Chicago.

Smith, L., and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2), 371–398.

Wit, J. (1997): "Herding Behavior in a Roll-Call Voting Game," Department of Economics, University of Amsterdam.