

Momentum and Social Learning in Presidential Primaries *

Abstract

This paper investigates social learning in sequential voting systems. In the econometric model, voters are uncertain over candidate quality, and voters in late states attempt to infer the information held by those in early states from voting returns. Candidates experience momentum effects when their performance in early states exceeds expectations. The empirical application uses daily polling data from the 2004 presidential primary. Our empirical strategy involves comparing support for candidates among late voters before and after the release of voting returns from early states. We find that Kerry benefited from surprising wins in early states and took votes away from Dean, who held a strong lead prior to the beginning of the primary season. Due to these momentum effects, early voters had up to 5 times the influence of late voters in the selection of candidates, and this disproportionate influence helps to explain the cross-state distribution of advertising expenditures. Finally, we use the estimated model to conduct two counterfactual experiments, showing that the election would have been more competitive under a simultaneous election and that electoral outcomes are sensitive to the order of voting.

Brian Knight
Department of Economics
Brown University
Providence RI 02912
Brian.Knight@brown.edu

Nathan Schiff
Sauder School of Business
University of British Columbia
Vancouver, BC V6T 1Z2
Nathan.Schiff@sauder.ubc.ca

*The authors received helpful comments from two anonymous referees, Nageeb Ali, Eric Brunner, Andy Foster, Alan Gerber, Rob McMillan and seminar participants at Yale, Johns Hopkins, University of Pennsylvania (Wharton), Northwestern, Toronto, Brown, the NBER political economy meetings, and the 2009 AEA meetings. Data on advertising were obtained from a project of the University of Wisconsin Advertising Project includes media tracking data from TNSMI/Campaign Media Analysis Group in Washington, D.C. The University of Wisconsin Advertising Project was sponsored by a grant from The Pew Charitable Trusts. The opinions expressed in this article are those of the authors and do not necessarily reflect the views of the University of Wisconsin Advertising Project or The Pew Charitable Trusts.

1 Introduction

While voting occurs simultaneously in many elections, voters choose sequentially in other cases, such as in roll-call voting in legislatures and in general elections for many federal offices prior to 1872. The most widely discussed example of a sequential election, however, is the presidential primary. As shown in Figure 1, the 2004 Democratic primary season began with the Iowa caucus on January 19th, followed by the New Hampshire primary on January 27th, and then mini-Super Tuesday on February 3rd, when voting occurred in the states of Arizona, Delaware, Missouri, New Mexico, North Dakota, Oklahoma, and South Carolina. The primary season continued with various elections in March, April, and May before concluding on June 8th. In 2008 the schedule became increasingly front-loaded with Nevada scheduled between Iowa and New Hampshire and, perhaps more importantly, many states moving their primaries to February 5th.

When considering such changes in the primary schedule, one naturally wonders whether or not the order of voting matters. That is, do outcomes of primaries depend upon the sequencing of states? Relatedly, do sequential, relative to simultaneous, systems lead to different outcomes in terms of the selection of candidates? And, if so, why? In our view, as well as the view of others, the key distinction is that sequential, relative to simultaneous, elections provide late voters with an opportunity to learn about the desirability of the various candidates from the behavior of early voters.¹ This opportunity for late voters to learn from early voting returns can in turn lead to momentum effects, defined as a positive effect of candidate performance in early states on candidate performance in later states.

While conventional wisdom holds that such momentum effects are important in sequential elections, any econometric attempt to identify their existence and measure their magnitude faces several challenges. First, what is the informational content of voting returns from early states? Do the absolute returns matter or should results be measured relative to voter expectations regarding candidate performance? If returns should be gauged relative to expectations, how can these expectations be measured? Second, how should researchers account for unmeasured candidate characteristics? The fact that eventual winners tend to do well in early states has often been interpreted as evidence of momentum effects. But success in both early and late states could simply reflect underlying candidate strength, which is often unobserved by the econometrician. Said

¹ Bikhchandani, Hirshleifer, and Welch (1992) mention presidential primaries as an example of sequential decision making and potential observational learning.

differently, winners in early states might have won the overall primary even with a simultaneous primary system under which momentum effects play no role. Third, how do voters weigh the voting returns from different states? For example, how should voters in states third in the sequence, such as those in South Carolina in 2004, weigh the returns from Iowa, the first state, relative to those from New Hampshire, the second state. A similar question is how do voters account for the fact that voters in states earlier in the sequence might also condition on returns from even earlier states? More concretely, when attempting to learn about the desirability of candidates from voting returns in Iowa and New Hampshire, how do voters in South Carolina account for the fact that, before casting their ballots, voters in New Hampshire may have also conditioned their decisions on voting returns in Iowa?

In this paper we attempt to overcome these econometric challenges through the development of a simple discrete choice econometric model of voting and social learning. In the model voters are uncertain about candidate quality, which is valued by all voters regardless of their ideology. Voters gather information about quality during the campaign, and voters in late states attempt to uncover the information of early voters from voting returns in these states. In the context of this model we show that candidates benefit from momentum effects when their performance in early states exceeds expectations.

In order to estimate the degree of social learning in sequential elections, we examine daily polling data in the 2004 Democratic primaries. Our empirical strategy involves comparing support for candidates among late voters before and after the release of voting returns from early states. To the extent that social learning is important, unexpected strength in voting returns from early states should lead to improved candidate evaluation by voters in late states in the daily polling data. Our estimates demonstrate substantial momentum effects. Using the estimated model, we next calculate the degree of any disproportionate influence for early voters and examine its implications for the distribution of campaign resources across states.

Finally, we use the model to simulate electoral outcomes under counterfactual electoral systems, an area where the scope for real-world randomized experimentation is obviously quite limited. In particular, we show that the race would have been much tighter under a simultaneous system due to the lack of social learning. We also simulate the election under alternative calendars and show that electoral outcomes are sensitive to the order of voting.

The paper proceeds as follows. Section 2 provides an overview of the relevant theoretical and

empirical literature. Section 3 lays out the basic theoretical and econometric model of momentum in primaries. Section 4 describes our empirical application, section 5 calculates voting weights, section 6 describes the counterfactual simulations, and section 7 describes possible extensions and summarizes our key findings.

2 Literature Review

Welch (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Banerjee (1992) provide the first formal analysis of social learning. Agents choose actions sequentially and are uncertain about the correct action, which depends upon the state of the world. If agents are sufficiently unsure about the true state of the world, then they may ignore their private signals and simply follow the actions of others.² Smith and Sorensen (2000) extend this model in two important ways, allowing for a continuum of signals and for heterogeneity in preferences across agents, and these two extensions will also be important ingredients of our model.

One of the key contributions of our paper is providing micro-foundations for empirical analyses of social learning in terms of building an econometric model of social learning from a theoretical model of social learning. Of course, a key question in developing this application is whether these theoretical social learning results extend to the context of sequential elections, with the main distinction being that voters make a social choice, and individual payoffs thus depend upon the actions of all agents.³ Dekel and Piccione (2000) show that every equilibrium of the simultaneous game is also an equilibrium of the sequential game, regardless of the sequence. In their model,

² This social learning framework has been applied in a variety of empirical settings. Welch (2000), for example, studies herding among security analysts. For a general overview of social learning in finance, see Devenow and Welch (1996). In development economics, social learning has been shown to play a key role in the choice of technology, such as in Foster and Rosenzweig (1995) and Munshi (2004). Cai, Chen, and Fang (2009) conduct a field experiment in which the top selling dishes are posted in restaurant menus and find that these postings are influential for orders and especially so for infrequent customers. Finally, Glaeser and Sacerdote (2007) provide a social learning explanation for aggregation reversals, where an individual relationship, such as income and ideology, is reversed at some level of aggregation, such as the state-level. For a more comprehensive overview of the social learning literature, see the survey by Sushil, Hirshleifer, and Welch (1998) and Chamley (2004).

³ Several authors have suggested alternative theoretical reasons for momentum effects. Callander (2007) proposes a model where every voter gains utility from both conforming, defined as supporting the eventual winner, and voting informatively, defined as supporting the best candidate based on their belief about the true state of the world. As the number of voters increases, the conforming component of utility dominates the information-based component and herding results, propelling the leading candidate to victory. On the candidate side, Klumpp and Polborn (2006) specify a model in which an early primary victory increases the likelihood of victory for one candidate and creates an asymmetry in campaign spending that furthers this advantage. Finally, Strumpf (2002) discusses a countervailing force to momentum. In particular, a candidate who is expected to win several of the last elections can credibly commit to not dropping out of the race even if he is trailing early. This effect, which favors later winners, thus moves in the opposite direction of momentum, which favors early winners.

strategic voters condition on being pivotal and hence behave as if they know that all other voters are evenly divided between the two candidates. Thus there is a symmetry between early and late voters, and it does not matter which candidate is supported by the early voters. Note, however, that this result does not demonstrate an equivalence between simultaneous and sequential elections; due to multiplicity, there are equilibria of the sequential game that are not equilibria of the simultaneous game. Indeed, Ali and Kartik (2010) construct an equilibrium in sincere voting in which voters do condition on the history of voting.⁴

Most of the empirical work on momentum in presidential primaries has come from the political science literature.⁵ Bartels (1987 and 1988) uses data from the National Election Study (NES) to predict the dynamics of the 1984 Democratic Primary. He shows that simple ratings of candidates do not fit the dynamics as well as do models that include measures of candidate viability. This model using viability suggests that candidate Gary Hart's surprising early victories convinced later voters of his viability.

The most closely related research is a series of papers that use state-level voting returns from multiple presidential primaries and examine the impact of returns from early states on overall primary outcomes. Adkins and Dowdle (2001) use data from 1980 to 1996 and find that New Hampshire plays a key role in determining the ordinal ranking of candidate finishes. Using data from 1976 to 2000 and a richer set of control variables, Steger, Dowdle, and Adkins (2004) also find an important role for New Hampshire in determining the eventual nominee. Using data through 2004, Steger (2008) confirms the New Hampshire effect but also documents an important role for Iowa.

Relative to these studies using state-level voting returns, our study offers two contributions to the literature. First, our empirical strategy is based upon the use of daily polling data, which allows us to compare the preferences of late voters before and after the release of returns from early states. The existing literature, which has effectively examined the correlation between early returns and

⁴ In related work, McKelvey and Ordeshook (1985) show that momentum effects can be generated even under simultaneous elections if polling data or endorsements are released during the period leading up to the election.

⁵ There have also been experimental tests for momentum effects. Morton and Williams (1999) consider a model with three candidates, liberal, moderate, and conservative. Voters do not observe candidate ideology but can potentially learn about ideology from past voting. Partisan voters (liberal or conservative) are risk averse and thus would rather vote for the moderate if they believe that only the moderate and the opposing candidate have a chance of winning. The authors test this hypotheses in a laboratory setting and find that later voters do use the early results and that a sequential election increases the likelihood of victory for moderate, unknown candidates. In addition, Battaglini, Morton, and Palfrey (2007) test predictions of the sequential voting model of Battaglini (2005), which incorporates costly voting and endogenous turnout.

late returns, may suffer from the problem of unobserved candidate characteristics discussed in the introduction. Second, we build an empirical model from micro-foundations. This provides us a framework for measuring expectations and thus the informational content associated with returns from early states. This framework also helps in the development of weights that voters place on the returns from states voting earlier in the sequence. Finally, the micro-foundations allow us to conduct several counterfactual policy experiments that are not possible using the methods in the existing literature.

3 Theoretical framework

This section lays out our basic theoretical and econometric framework for measuring momentum effects in sequential elections, and the notation here roughly follows Chamley (2004). Given our empirical motivations, we keep things simple and make the assumptions necessary to generate a tractable empirical model. Many of these assumptions, however, will be discussed and relaxed in the empirical section to follow.

3.1 Setup

Consider a set of states (s) choosing between candidates ($c = 0, 1, \dots, C$) in a sequential election, where the order of voting is taken as given. We allow for the possibility that multiple states may vote on the same day; in particular, let Ω_t be the set of states voting on date t and let $N_t \geq 1$ be the size of this set.

Voter i residing in state s is assumed to receive the following payoff from candidate c winning the election:

$$u_{cis} = q_c + \eta_{cs} + \nu_{cis} \tag{1}$$

where q_c represents the quality of candidate c , η_{cs} represents a state-specific preference for candidate c , and ν_{cis} represents an individual preference for candidate c and is assumed to be distributed type-I extreme value and independently across both candidates and voters. Note that quality is valued equally by all voters and can be interpreted as a fundamental positive characteristic, such as competence or integrity. We normalize utility from the baseline candidate to be zero for all voters ($u_{0is} = 0$). While underlying state preferences are assumed to be stable, or time-independent, there is uncertainty and expectations may evolve during the election, as described below.

We assume the following information structure. Voters know their own state-level preference (η_{cs}) but not those in other states. Voters do, however, know the distribution from which these state-level preferences are drawn. In particular, we assume that state-level preferences are normally distributed [$\eta_{cs} \sim N(0, \sigma_\eta^2)$] and independently across states. We further assume that voters are uncertain over candidate quality and are Bayesian. In particular, initial ($t = 1$) priors over candidate quality (q_c) are assumed to be normally distributed with a candidate-specific mean μ_{c1} and a variance σ_1^2 that is common across candidates. Under the assumptions to follow, the posterior distribution will be normal as well. Before going to the polls, all voters in state s receive a noisy signal (θ_{cs}) over the quality of candidate c :

$$\theta_{cs} = q_c + \varepsilon_{cs} \tag{2}$$

where the noise in the signal is assumed to be normally distributed [$\varepsilon_{cs} \sim N(0, \sigma_\varepsilon^2)$] and independently across states. These signals can be interpreted in a variety of ways, including town hall meetings with candidates, media coverage of candidate debates and appearances within the state, endorsements of candidates by either local media outlets or local politicians, political advertising on local television channels, etc. We assume that this signal is common across all voters within a state.⁶ Finally, we assume that the signal is unobserved by voters in other states.⁷

Given the state-level signal (θ_{cs}), expected utility for voter i in state s from candidate c winning can be written as follows:

$$E(u_{cis} | \theta_{cs}, \eta_{cs}, \nu_{cis}) = E(q_c | \theta_{cs}) + \eta_{cs} + \nu_{cis} \tag{3}$$

Finally, regarding voter behavior, we assume sincere voting. In particular, given the information available, voter i in state s at time t supports the candidate who provides the voter with the

⁶ We feel that this assumption of a common signal within states is reasonable given the role of the mass media in modern elections. However, some campaign messages, such as mailings, can be targeted to individual voters, suggesting an alternative formulation that would allow for voters within the same state to receive independent signals. This formulation implies that, in the absence of heterogeneity in state-level preferences ($\sigma_\eta^2 = 0$), quality is perfectly revealed by voting returns from states with large populations. Thus, voters will learn only from returns in the first state and will ignore both their private signals and voting returns from other states thereafter. We view this feature of a model with individual-level signals as both unattractive and unrealistic and thus focus on the case of state-level signals. One could also consider a hybrid model with both individual-level and state-level signals. While this formulation would overcome the problem of perfect revelation of quality after voting in the first state, as described above, it is not clear how the variance in these two signals, which is a key parameter of interest in the empirical analysis to follow, would be separately identified.

⁷ Later in the paper we discuss the implications of an alternative model in which signals are observed by all voters on a national basis.

highest level of expected utility.⁸ This behavioral assumption is myopic in the sense that voters do not account for how their vote will influence later voters. We also abstract from several other commonly studied motives, such as bandwagon effects, under which voters have a preference for supporting the eventual winner. Bandwagon effects can also generate momentum effects that are unrelated to social learning, and this issue of alternative explanations for any measured momentum effects will be discussed more completely below in the empirical application.

3.2 Voting behavior

Then, for voters in state s observing a signal over quality (θ_{cs}) and with a prior given by (μ_{ct}, σ_t^2) , private updating over quality is given by:

$$E(q_c|\theta_{cs}) = \alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct} \quad (4)$$

where the weight on the signal is given by:

$$\alpha_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\varepsilon^2} \quad (5)$$

Reflecting well-known results in the literature on Bayesian learning, voters thus place more weight on their private signal (θ_{cs}) the higher is the variance in the prior over quality (σ_t^2) and the lower is the degree of noise in the signal (σ_ε^2).

Plugging equation (4) into equation (3), we have that:

$$E(u_{cis}|\theta_{cs}, \eta_{cs}, \nu_{cis}) = \alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct} + \eta_{cs} + \nu_{cis} \quad (6)$$

Then, using the fact that ν_{cis} is distributed type-I extreme value, we can write the probability that voter i supports candidate c conditional on the state-level variables (the signal θ_{cs} and unobservables η_{cs}) as follows:

$$\Pr(E(u_{cis}|\theta_{cs}, \eta_{cs}, \nu_{cis}) > E(u_{dis}|\theta_{ds}, \eta_{ds}, \nu_{dis}); \forall d \neq c) = \frac{\exp(\alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct} + \eta_{cs})}{\sum_{d=0}^C \exp(\alpha_t \theta_{ds} + (1 - \alpha_t) \mu_{dt} + \eta_{ds})} \quad (7)$$

⁸ These strategies are similar to the sincere voting strategies analyzed in Ali and Kartik (2010).

Under the assumption of a continuum of voters, state-level vote shares are equal to these voting probabilities.⁹ That is, the vote shares for candidate c relative to the baseline candidate 0 are given by:

$$v_{cst}/v_{0st} = \frac{\exp(\alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct} + \eta_{cs})}{\exp(\alpha_t \theta_{0s} + (1 - \alpha_t) \mu_{0t} + \eta_{0s})} \quad (8)$$

Finally, taking logs and using the normalization that the utility from the baseline candidate equals zero, vote shares in state s voting at time t can be described as follows:

$$\ln(v_{cst}/v_{0st}) = \eta_{cs} + \alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct}. \quad (9)$$

Thus, the log-odds ratio can be expressed as a linear combination of state-level preferences (η_{cs}), the signal (θ_{cs}) received by voters in state s , and the mean of the quality distribution (μ_{ct}) prior to the realization of the signal, where the relative weight on the latter two terms depends upon the parameter α_t . As will be seen below, this expression for aggregate voting returns provides the key link between the individual-level voting data and the aggregate returns in the econometric formulation, and the linearity will be a particularly attractive feature in the analysis of social learning from early voting returns.

3.3 Social learning and momentum

From the perspective of measuring momentum, the key question is then how voters in late states update their beliefs over quality upon observing vote shares in early states (i.e. $E(q_c | v_{cst}, v_{0st})$). Given that state-level preferences (η_{cs}) are unobserved by voters in other states, signals (θ_{cs}) cannot be inferred directly from vote shares in equation 9. However, rearranging equation 9 and using the fact that $\theta_{cs} = q_c + \varepsilon_{cs}$ we can say that transformed vote shares provide a noisy signal of quality:

$$\frac{\ln(v_{cst}/v_{0st}) - (1 - \alpha_t) \mu_{ct}}{\alpha_t} = q_c + \frac{\eta_{cs}}{\alpha_t} + \varepsilon_{cs} \quad (10)$$

where the noise in the voting signal includes the noise in the quality signal (ε_{cs}) but also the noise due to the unobserved state preferences (η_{cs}/α_t); the combined variance of the noise in the

⁹ In a theoretical model with a continuum of voters, sincere voting is an equilibrium since the behavior of individual voters has no effect on overall vote shares and hence the behavior of later voters. While this equilibrium may not be unique, it does help to motivate our assumption of sincere voting.

voting signal thus equals $(\sigma_\eta^2/\alpha_t^2) + \sigma_\varepsilon^2$. Given $N_t \geq 1$ such signals, the posterior distribution is also normal and can thus be characterized by its first two moments:

$$\mu_{ct+1} = \beta_t \left[\frac{1}{N_t} \sum_{s \in \Omega_t} \frac{\ln(v_{cst}/v_{0st}) - (1 - \alpha_t)\mu_{ct}}{\alpha_t} \right] + (1 - \beta_t)\mu_{ct} \quad (11)$$

$$\frac{1}{\sigma_{t+1}^2} = \frac{1}{\sigma_t^2} + \frac{N_t}{(\sigma_\eta^2/\alpha_t^2) + \sigma_\varepsilon^2} \quad (12)$$

where the weight on the voting signals is given by:

$$\beta_t = \frac{N_t \sigma_t^2}{N_t \sigma_t^2 + (\sigma_\eta^2/\alpha_t^2) + \sigma_\varepsilon^2} \quad (13)$$

Before describing the evolution of the mean of the belief distribution, we note that the precision of the posterior, defined as the inverse of the variance ($1/\sigma_{t+1}^2$), is increasing in the number of states (N_t) voting at time t along with the degree of precision in these voting returns $[(\sigma_\eta^2/\alpha_t^2) + \sigma_\varepsilon^2]^{-1}$. To provide further interpretation of this social learning rule, it is useful to re-write equation (11) as follows:

$$\mu_{ct+1} - \mu_{ct} = \frac{\beta_t/N_t}{\alpha_t} \sum_{s \in \Omega_t} [\ln(v_{cst}/v_{0st}) - \mu_{ct}] \quad (14)$$

Social learning ($\mu_{ct+1} - \mu_{ct}$) thus depends upon the surprises in voting returns, defined as the deviations in vote shares from expectations over candidate performance.¹⁰ Interestingly, this learning rule implies that candidates who do not win the primary in state s can still benefit from momentum effects so long as they perform well relative to expectations. At the same time, candidates who win primaries may actually experience reverse momentum effects in the event that their margin of victory is smaller than expected.

To provide a sense of the degree of social learning, note that the effect of increasing vote shares in one state on the mean of the posterior distribution of candidate quality can be expressed as

¹⁰ An important implicit aspect of the above formulation is that expectations over electoral outcomes, as captured by μ_{ct} in equation 14, depend upon national, but not state-specific, factors. This result follows from our assumption of unobserved state preferences. We make this assumption for three reasons. First, if state-level preferences are perfectly observed then private signals can be inferred from voting returns, and preferences of early states would be assumed to have no disproportionate impact. Given this, we would be assuming away one of the most important controversies surrounding the primary system. Second, with perfectly observed preferences, we have that $\alpha_t = \beta_t$ for dates on which a single state votes ($N_t = 1$). In this case, the updating rule takes the form $\mu_{ct+1} - \mu_{ct} = \ln(v_{cst}/v_{0st}) - \eta_{cs} - \mu_{ct}$, and momentum effects would be effectively assumed, rather than estimated. Third, it may be difficult for voters to know the state preferences of most other states given that, with the exception of Iowa and New Hampshire, most polls of other states were reported infrequently, if at all. Even with Iowa and New Hampshire there could be error in the polls, and not all voters in other states may diligently follow this polling, and it is thus far from clear that all voters could know the preferences of other states. Finally, we should note that later in the paper we relax the assumption of completely unknown state preferences by using an alternative specification where state preferences are partially observed.

follows:

$$\frac{\partial \mu_{ct+1}}{\partial \ln(v_{cst}/v_{0st})} = \frac{\beta_t/N_t}{\alpha_t} = \frac{\sigma_t^2 + \sigma_\varepsilon^2}{N_t \sigma_t^2 + (\sigma_\eta^2/\alpha_t^2) + \sigma_\varepsilon^2} \quad (15)$$

This parameter shows the ratio of the weight placed on a public signal, which is an indirect measure of the private signal, to the weight placed on a private signal that is received directly. Voters do not observe preferences of voters in other states and this ratio is thus bounded between zero and one. When there is no heterogeneity in state preferences ($\sigma_\eta^2 = 0$) and for the special case of a single election ($N_t = 1$), public and private weights are equal ($\beta_t = \alpha_t$), and the ratio is one. As the degree of heterogeneity in state preferences increases, the private signal becomes more difficult to infer and the ratio converges to zero. As with updating from a private signal, the degree of social learning is decreasing in the degree of noise in the signal (σ_ε^2) and increasing in the variance of the prior (σ_t^2).

4 Empirical Application

Our empirical application focuses on the 2004 Democratic primary. During the months leading up to the primary season, Howard Dean, governor of Vermont, held a substantial lead in opinion polls. After his third place finish in the Iowa caucuses, however, Dean soon lost that lead in opinion polls to the Iowa winner, senator John Kerry of Massachusetts, and was forced to withdraw after a disappointing performance in Wisconsin. Kerry continued his success in Iowa with a win in New Hampshire and with strong performances in all of the subsequent states. The only serious challenge to Kerry after Iowa came from John Edwards, a senator from North Carolina, who came in a surprisingly strong second in Iowa and proceeded to win in South Carolina and Oklahoma. Edwards was forced to withdraw, however, on March 3, the day after a string of second-place finishes to Kerry on Super Tuesday.

4.1 Data

To measure the degree of social learning in the 2004 primaries, we examine reactions of voters in daily opinion polls to candidate performance in primaries. We focus on the campaigns of the three major candidates: Dean, Edwards, and Kerry, where Kerry is considered the baseline candidate.¹¹ Individual-level data are taken from the National Annenberg Election Survey 2004,

¹¹ Another candidate, Wesley Clark, was considered viable in the months leading up to the primary season. He chose, however, to not participate in the Iowa caucuses and subsequently fell out of serious contention. Given that

which conducted interviews on a daily basis in a rolling cross-section design beginning on October 7, 2003 and continuing through the general election in November 2004.¹² Given our focus on the primary season, we use voting intentions for 4,084 respondents who identify as likely Democratic primary voters between October 7, 2003 and March 2, 2004, the day before Edwards withdrew from the race.¹³ To be clear, these are respondents living in states that have not yet held their primaries. Voters living in states that have already voted are not asked their voting intentions in the survey and are thus excluded from this part of our analysis. Finally, as will be described below, we aim to estimate the state-specific preference parameters (η_{cs}) and thus insufficient data forces us to also delete respondents from DC and seven small states.¹⁴ These individual-level data are then merged with state-aggregate vote shares from the 2004 primary season as reported on the website www.cnn.com.

While the econometric analysis to follow uses voting returns from all states, we next highlight our identification strategy in Figures 2-4 for the case of Iowa. As shown, Dean had a substantial and stable lead over Kerry and Edwards during the month preceding the Iowa primary. Dean under-performed in Iowa relative to expectations, as captured by pre-Iowa polling levels, and voters in the Annenberg survey updated appropriately.¹⁵ Kerry, by contrast, outperformed expectations in Iowa, and survey respondents updated accordingly. Edwards also outperformed his pre-Iowa polling numbers and his polling numbers did increase following Iowa. After a few days, however, his support fell back to pre-Iowa levels.¹⁶

we do not have a model of candidate campaign strategies and the possible negative signals sent by non-participation, we felt it best to exclude him from the analysis. Another candidate, Richard Gephardt, polled well prior to Iowa but withdrew from the race shortly thereafter.

¹² Sampling was done using random digit dialing to create a national cross-section where each observation is a unique individual. To deal with differences in the population distribution of area codes the NAES proportioned the first eight digits of the phone number to match the national incidence. We also include a separate over-sample of New Hampshire voters covering the period surrounding the primary (1/8/2004-2/3/2004).

¹³ We assume that respondents report their voting intentions truthfully. Due to any momentum effects, however, respondents may have incentives to misreport their preferences in polling data to dampen expectations for their preferred candidate. To address this issue, we test for the accuracy of these voting intentions using residents from New Hampshire, a state that was over-sampled in the Annenberg survey, during the period after the Iowa caucuses but before the New Hampshire primary. According to these voting intentions, 33 percent planned to vote for Dean, 52 percent for Kerry, and 15 percent for Edwards. These intentions match the actual three-candidate vote shares surprisingly well. In particular, Dean received 34 percent, Kerry received 50 percent, and Edwards received 16 percent. While we cannot conduct similar tests for other states due to small sample sizes, the close match between intentions and voting returns in New Hampshire suggests that respondents do report truthfully.

¹⁴ Specifically, we dropped AR, ID, KS, ND, SD, VT, and WY.

¹⁵ The “Dean scream” will be discussed in detail below.

¹⁶ These patterns are similar to those in prices from the Iowa Electronic Market, in which market participants purchased contracts that pay \$1 in the event that Kerry, for example, is the party’s nominee in the general election,

4.2 Empirical Model

As noted above, our empirical strategy involves comparing support for candidates among late voters before and after the release of voting returns from early states. In our econometric specification, we assume that voters receive their signals just prior to the date of their primary. That is, these respondents from late states have not yet observed their private signals.¹⁷ In this case, their voting intentions can thus be summarized as follows:

$$\Pr(E(u_{cis}|\eta_{cs}, \nu_{cis}) > E(u_{dis}|\eta_{ds}, \nu_{dis}); \forall d \neq c) = \frac{\exp(\eta_{cs} + \mu_{ct})}{\sum_{d=0}^C \exp(\eta_{ds} + \mu_{dt})} \quad (16)$$

To better understand our empirical strategy for estimating the parameters governing the learning process, it is useful to first note that voter updating over quality can be summarized by the weight on private signals, the weight on public signals, updating over the mean, and updating over the variance as follows:

$$\alpha_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\varepsilon^2} \quad (17)$$

$$\beta_t = \frac{N_t \sigma_t^2}{N_t \sigma_t^2 + (\sigma_\eta^2 / \alpha_t^2) + \sigma_\varepsilon^2} \quad (18)$$

$$\mu_{ct+1} - \mu_{ct} = \frac{\beta_t / N_t}{\alpha_t} \sum_{s \in \Omega_t} [\ln(v_{cst} / v_{0st}) - \mu_{ct}] \quad (19)$$

$$\frac{1}{\sigma_{t+1}^2} = \frac{1}{\sigma_t^2} + \frac{N_t}{(\sigma_\eta^2 / \alpha_t^2) + \sigma_\varepsilon^2} \quad (20)$$

As seen, with information regarding the initial priors (μ_{c1}, σ_1^2) along with the parameters σ_ε^2 and σ_η^2 , one can compute the weight on the private signal in the first period (α_1) and, with this weight in hand, one can then compute the weight placed upon the public voting signals in the first period (β_1). Then, with the entire set of first-period values $(\mu_{c1}, \sigma_1^2, \alpha_1, \beta_1)$, along with information on first-period voting returns, we can successively compute the second-period values $(\mu_{c2}, \sigma_2^2, \alpha_2, \beta_2)$.

and the price of this contract can thus be interpreted as the probability that a given candidate wins the nomination (Wolfers and Zitzewitz, 2004). We choose to focus on polling data, rather than these prices from prediction markets, for two reasons. First, the mapping from voting in primaries to the probability of nomination, as provided by the prediction market prices, is confounded by the presence of superdelegates, which are not pledged to the winning candidate, as well as the possibility that no candidate wins a majority of the delegates, in which case the nominee is chosen through a bargaining process at the party convention. Second, the daily polling data, but not prediction market data, include additional measures of candidate quality, and we will make use of these in our discussion of alternative explanations to follow.

¹⁷ We also considered an alternative formulation under which all voters receive their signals prior to the start of the primary season. For two reasons, we choose to not pursue this alternative formulation. First, identification of the key parameters is more complex since expected quality is now a weighted average of signals and priors. Second, as will be seen below, our analysis demonstrates that respondents who have already voted are better informed when responding to auxiliary questions regarding candidate characteristics than respondents who have not yet voted.

With these second-period values, along with information on second-period voting returns, we can then successively compute the third-period values $(\mu_{c3}, \sigma_3^2, \alpha_3, \beta_3)$, etc. Each day in which there is one or more elections is the start of a new period. There are thus ten periods in our sample, please see Table 4 for periods, states, and election dates.

Thus, it should be clear that the key parameters to be estimated are the initial priors (μ_{c1}, σ_1^2) along with the variance in state-level preferences (σ_η^2) and the degree of noise in the signal (σ_ε^2) .¹⁸ These key parameters are estimated via a two-step approach. In the first step, we use the pre-Iowa polls to estimate the initial conditions. In particular, for the case of $t = 1$, we have that:

$$\Pr(E(u_{cis}|\eta_{cs}, \nu_{cis}) > E(u_{dis}|\eta_{ds}, \nu_{dis}); \forall d \neq c) = \frac{\exp(\eta_{cs} + \mu_{c1})}{\sum_{d=0}^C \exp(\eta_{ds} + \mu_{d1})} \quad (21)$$

We estimate the state-level preference parameters (η_{cs}) , which are normalized to sum to zero and which can be used to calculate σ_η^2 , along with a constant term, which provides an estimate of μ_{c1} . In the second step, we use reactions of voters in post-Iowa opinion polls to the revelation of voting returns in all states voting prior to March 2, the effective end of the primary season, in order to estimate the key parameters $(\sigma_\varepsilon^2, \sigma_1^2)$ governing the social learning process. Given the two-stage estimation approach, conventional confidence intervals will not reflect the uncertainty associated with using generated regressors in the second stage. We address this issue by computing bootstrap confidence intervals.¹⁹

The key social learning parameters are identified by voter responses (in the NAES data) to the release of voting returns in others states. If voters are unresponsive to the release of such information, this suggests an absence of social learning, and the variance in the initial prior (σ_1^2) will have a small estimate or the variance in the degree of noise in the signal (σ_ε^2) will have a large estimate. If voters are responsive to voting returns, by contrast, then the variance in the initial prior (σ_1^2) will have a large estimate or the variance in the degree of noise in the signal (σ_ε^2) will have a small estimate.

¹⁸ Note that we are assuming that the variance in the noise in the signal (σ_ε^2) is stable throughout our sample period. If more information is released in the early part or the later part of the campaign, then the variance in the noise in the signal should depend upon the time period; that is $\sigma_\varepsilon^2(t)$. In this case, our estimate can be considered as the average variance in the noise in the signal across the primary season.

¹⁹ In particular, we draw 100 samples with replacement from the underlying sample. In some replications, an insufficient number of cases were drawn to allow for identification of the specific state fixed effects, and we thus exclude such states from the analysis in these bootstrap samples. Due to this variability in the number of states in a replication we report the median coefficients from our bootstrap samples in Tables 1 and 2.

4.3 Baseline Results

Table 1 provides the results from the first-step of the estimation procedure. As shown in columns 1 and 2, the coefficient on the candidate-specific constant term demonstrates Dean’s substantial lead over Kerry and Kerry’s lead over Edwards prior to the commencement of the primary season. As noted above, this coefficient can be interpreted as the mean of the initial prior (μ_{c1}), and this variable will play a key role in the updating rule given by equation (19). The substantial variation in the state-specific coefficients demonstrates the significant diversity in preferences for the candidates across states. As shown, there are strong regional effects with Kerry holding a substantial advantage in his home state of Massachusetts, and Edwards enjoying a corresponding strong advantage in the South, with statistically significant advantages over Kerry in North Carolina and South Carolina. This advantage likely reflects the fact that Edwards was the only candidate of the three from the South. This issue of regional advantages will be considered below in an alternative specification, which relaxes the assumption that such regional advantages are unobserved by voters in later states.

Table 2 provides estimates of the other key parameters. The degree of heterogeneity in state-level preferences (σ_η^2) is calculated by taking the cross-state and cross-candidate variance in the coefficients on the state dummy variables as reported in Table 1. As shown, the key learning parameters, the initial prior (σ_1^2) and the degree of noise in the signal (σ_ε^2), are both positive and statistically significant.

To provide additional interpretation of these results, we next present in Figures 5 and 6 the key dynamics of the model as implied by these estimated parameters and the aggregate returns. As shown in the upper panel of Figure 5, voters learned about the quality of the candidates relative to one another from the returns in other states. Prior to the first primary, voters viewed Dean as the highest quality candidate followed by Kerry and Edwards, reflecting the pattern of coefficients on the candidate indicator variables in Table 1. Following Kerry’s win in Iowa, Kerry pulled ahead of Dean in terms of mean quality ratings. Although Kerry defeated Edwards in Iowa, voters updated favorably over Edwards relative to Kerry, reflecting the fact that candidates can benefit, even relative to first place finishers, from surprisingly strong second place finishers. On the other hand, although Edwards defeated Dean in Iowa, voters still evaluated Dean and Edwards roughly equally. This in turn reflects the fact that voters also placed some weight on their beliefs prior to

voting in Iowa, and these priors were strongly in favor of Dean relative to Edwards. Following New Hampshire and mini-Super Tuesday, Kerry held a strong advantage, and Dean never recovered from his weak performances in Iowa.

As voters learned from election results, the degree of variance in the beliefs over candidate quality (σ_t^2) fell substantially over the primary season, as depicted in the bottom panel of Figure 5. Prior to the Iowa caucus, the variance in this distribution was around 3.5, reflecting the estimated parameter in Table 2, but falls to around 0.5 by March 2, or Super Tuesday. Thus, voters learn a substantial amount over the course of the campaign about candidate quality purely from the release of voting returns in other states.

To provide further interpretation of these results, Figure 6 plots the implied weights on the private signals observed by voters (α_t), the weights placed upon aggregate vote shares after scaling by the number of primaries (β_t/N_t), and the ratio of these two, which is the key social learning parameter ($\beta_t/\alpha_t N_t$). As shown, voters place less weight on their prior than on the private signal at the beginning of the sample period. This in turn reflects the fact that the estimated degree of noise in the signal is less than the estimated degree of variance in the initial prior ($\sigma_\epsilon^2 < \sigma_1^2$) and that the weight on the private signal can be shown to be inversely related to the ratio of these parameters (i.e. $\alpha_t = (1 + \sigma_\epsilon^2/\sigma_t^2)^{-1}$). As implied by the model, the weight placed upon the private signal falls during the primary season and, by Super Tuesday, voters place roughly 75 percent weight on their priors and only 25 percent on their private signals.

In terms of the weights on public signals, late voters initially place roughly 60 percent weight on these signals and 40 percent on their priors. As more primary returns come in, voters place less weight on voting returns and more weight on their prior and by the end of the sample voters are largely unmoved by primaries held in other states. As this weight falls to zero, the ratio of these two weights ($\beta_t/\alpha_t N_t$) also falls from its starting point of 75 percent to around 10 percent. This is due to the fact that, as states place less weight on their private signal (α_t), the variance of the noise in the public signal ($\frac{\sigma_y^2}{\alpha_t^2} + \sigma_\epsilon^2$) increases and voting returns are thus less informative about quality.

In summary, our estimated model demonstrates that voters in late states placed significant weight on Kerry's early victories. It is the deviations from expectations that matter, however, and Edwards benefited relative to Kerry from a surprisingly strong second-place finish in Iowa. While Dean came in third place in Iowa, he benefited from strong voter beliefs regarding his quality

prior to Iowa and was able to remain viable. At the same time that voters shifted their relative evaluations of candidate quality, they became increasingly confident in these evaluations, and voters in late states thus placed less weight on both their private signals as well as from returns in other states. Taken together, these results demonstrate significant momentum effects as reflected in the effect of early returns on the choices of late voters.

4.4 Additional specifications

As noted above, the baseline model assumes that voters observe their own state-level preferences but not those in other states. As an alternative to this assumption of perfect observability, we consider and estimate a specification in which state-level preferences consist of both an unobserved component (η_{cs}) and an observed component (X_{cs}), such as geography, which could capture advantages enjoyed by politicians campaigning in their home states. Then, aggregate voting returns can be written as follows:

$$\ln(v_{cst}/v_{0st}) = \eta_{cs} + \gamma X_{cs} + \alpha_t \theta_{cs} + (1 - \alpha_t) \mu_{ct}. \quad (22)$$

where γ is a weight, or vector of weights, on observed preferences that will be estimated. It is then straightforward to show that the social learning rule is adjusted for these observed characteristics as follows:

$$\mu_{ct+1} - \mu_{ct} = \frac{\beta_t/N_t}{\alpha_t} \sum_{s \in \Omega_t} [\ln(v_{cst}/v_{0st}) - \gamma X_{cs} - \mu_{ct}] \quad (23)$$

Thus, voters in late states incorporate these observed state-level characteristics into their expectations of candidate performance, and, in our example of geography, returns showing that a candidate performed well in his home state, even relative to national expectations over candidate performance (μ_{ct}), do not necessarily lead to momentum effects.

To operationalize this specification, we incorporate into X_{cs} a measure of the distance between state s and the home state of candidate c , where the measure is relative to the distance between state s and Kerry's home state of Massachusetts. In this way we allow for voters to incorporate perceived regional advantages into their expectations of candidate performance. For example, voters may expect John Edwards to outperform national polling trends in his home state of North Carolina. After the first step, or pre-Iowa, analysis, we regress the estimated fixed effects on this distance measure and use the residuals from this regression as an estimate of unobserved preferences (η_{cs}). As shown at the bottom of columns 3 and 4 of Table 1, distance has a negative and

statistically significant effect on voting intentions. After accounting for this observed dimension of preferences, the regional advantages enjoyed by candidates are diminished although the home state advantage enjoyed by Kerry and Edwards remains. As shown in Table 2, the estimated variance of unobserved preferences (σ_η^2) is reduced in this model, reflecting the assumption that some component of preferences are observed by voters in other states. The other key parameters are qualitatively similar to those in column 1.

The second specification relaxes the assumption that underlying voter preferences are stable over the campaign. Trends in candidate-specific preferences could of course confound the estimation of social learning effects. To address this issue, we estimate a model with a candidate-specific trend (γ_c) in preferences. Then, aggregate voting returns are adjusted as follows:

$$\ln(v_{cst}/v_{0st}) = \eta_{cs} + \gamma_c\tau + \alpha_t\theta_{cs} + (1 - \alpha_t)\mu_{ct}. \quad (24)$$

where τ indexes the interview date and is normalized to equal zero on the date of the Iowa primary. The social learning rule then becomes:

$$\mu_{ct+1} - \mu_{ct} = \frac{\beta_t/N_t}{\alpha_t} \sum_{s \in \Omega_t} [\ln(v_{cst}/v_{0st}) - \gamma_c\tau - \mu_{ct}] \quad (25)$$

Thus, voters in late states incorporate these trends into their expectations of candidate performance. As shown in columns 5 and 6 of Table 1, the pre-Iowa trends tended to favor Dean and Edwards, while Kerry was disadvantaged. Thus, at the time of the Iowa primary, the mean evaluation of Dean and Edwards are higher than are those in the baseline specification. This is reflected in the first row of Table 1. As shown in Table 2, however, the key social learning parameters are similar to those in the baseline specification.

4.5 Alternative Explanations

In this section, we address three alternative explanations for our baseline findings: departures from sincere voting, national information, and campaign finance and persuasion.

4.5.1 Departures from Sincere Voting

In the baseline model, we have assumed sincere voting, under which voters support the candidate that provides the highest expected utility level. We have thus abstracted from strategic voting, several forms of which provide alternative, non-learning explanations for our documented momentum results. The first form of strategic voting involves electability considerations associated with

the general election. For example, consider a voter who prefers Dean over Kerry as president but prefers either over Bush. This voter may learn that Kerry is more popular among other voters after Iowa and thus has a better chance of defeating Bush in the general election; this voter may thus switch to Kerry after Iowa, and this electability-driven switching provides an alternative explanation for our results.²⁰ The second form of strategic voting involves concerns over wasted votes with more than two candidates. Consider a voter, for example, who ranks the candidates as Edwards first, Dean second, and Kerry last. This voter may view Edwards as not viable before Iowa and support Dean instead. After Iowa, Edwards becomes viable and this voter may shift from Dean to Edwards. Again, this viability-driven switching provides an alternative explanation for our results. The third form of strategic voting involves bandwagon effects, under which voters have conforming preferences and thus vote for the candidate expected to win. Again, bandwagon effects may mimic social learning and thus provide an alternative explanation for our results.

Importantly, all three of these alternative explanations for our measured momentum effects involve learning about the *preferences of other voters* rather than underlying candidate quality, as is emphasized by our social learning model. Thus, to distinguish between strategic voting explanations and social learning explanations, we examine the dynamics of measures of candidate quality, which we proxy by auxiliary questions in which respondents evaluated candidates on a 1-10 scale for the following candidate characteristics: favorability, cares about people like me, inspiring, strong leader, trustworthy, shares my values, knowledgeable, and reckless. These characteristics can be interpreted as measures of candidate quality given that they are arguably traits that would be valued by all voters regardless of ideology.

More concretely, for each of these quality proxies, we run the following regression:

$$\text{quality}_{itc} = \delta_c + \kappa \times \mu_{ct} + \xi_{itc} \tag{26}$$

where μ_{ct} is the mean candidate quality at time t , as implied by our estimated model, and reflected in Figure 5a, and quality_{itc} represents responses to the auxiliary questions regarding candidates, relative to Kerry, described above. Under our assumption of sincere voting, we would expect $\kappa > 0$,

²⁰ Note that electability considerations are not necessarily inconsistent with our results if primary voters believe that the probability of winning in the general election is increasing in candidate quality. This could be the case, for example, if independent voters are pivotal in the general election and place significant weight on quality. Then primary voters driven by electability considerations will still value quality although the rationale for this preference is somewhat different.

whereas, as argued above, there should be no link ($\kappa = 0$) under strategic voting explanations given that voters only learn about the preferences of other voters from early returns.

As shown in Table 3a, there is a positive relationship between estimated mean evaluations of candidate quality (μ_{ct}) and respondent evaluations of candidate quality, providing support for our social learning story. The first six measures have the expected positive coefficients under a social learning story, whereas the coefficient associated with the measure of “reckless” has the expected negative sign. The final two measures, however, are statistically insignificant, likely reflecting the reduced sample sizes.²¹ Thus, this analysis demonstrates that, in addition to voting intentions, perceptions of candidate quality also respond to the release of voting returns from early states and thus provides support for our social learning interpretation of our baseline findings.

4.5.2 National Information

A key assumption in our model is that signals are only observed within the state. If all signals are observed nationally, by contrast, then our model does not predict social learning since early voters do not hold an informational advantage. A key question, however, is whether our estimator would mistakenly detect social learning due to purely national signals. To examine this issue, consider a model with national information, under which all voters, both early and late, update their beliefs upon receiving the signal available at time t (θ_{ct}) as follows:

$$\mu_{ct+1} = \alpha_t \theta_{ct} + (1 - \alpha_t) \mu_{ct} \quad (27)$$

and voting returns in state s at time t can thus be summarized by:

$$\ln(v_{cst}/v_{ost}) = \eta_{cs} + \alpha_t \theta_{ct} + (1 - \alpha_t) \mu_{ct}. \quad (28)$$

Averaging (28) across all states voting at time t , solving for $\alpha_t \theta_{ct}$, and substituting into (27), we have that:

$$\mu_{ct+1} - \mu_{ct} = \frac{1}{N_t} \sum_{s \in \Omega_t} [\ln(v_{cst}/v_{ost}) - \eta_{cs} - \mu_{ct}] \quad (29)$$

Note that equation (29) is similar to equation (14) and thus also predicts that updating by early voters will be correlated with voting returns. The key issue is that any information released between periods t and $t + 1$ will be reflected in voting returns at the end of t and in the updating

²¹ All respondents were queried as to candidate favorability but were then randomly queried as to the additional traits.

at the beginning of $t + 1$. So, it is clear that our estimator may detect momentum effects even in the absence of social learning due to the presence of national information.

We address this alternative explanation in two ways. Our first test examines whether any shift in voting intentions in the daily polling data seem to correspond with key election dates. While our baseline estimator uses daily polling data in conjunction with information on the dates of primaries and their associated electoral outcomes, this alternative analysis uses only the daily polling data in order to detect the most likely date for a break in support for candidates. If our results are driven solely by nationally-available information, there is no reason that a shift in support for candidates should correspond with key election dates and may even occur before such events. According to our social learning model, by contrast, a shift in support for candidates should occur on or just following key election dates.

More concretely, using methods analogous to those from the econometrics literature on unknown structural breaks, we first locate the single break in support for which the model best fits the data.²² That is, we estimate the parameters from following model for all possible break dates j and choose the date that maximizes the likelihood:

$$\Pr(\text{prefer } c) = \frac{\exp[\eta_{cs} + \mu_c + 1(\tau > j)\mu_{cj}]}{\sum_{d=0}^C \exp[\eta_{ds} + \mu_d + 1(\tau > j)\mu_{dj}]} \quad (30)$$

where η_{cs} is the state-candidate fixed effect, τ indexes the date of the interview, μ_c is a time-invariant candidate-specific constant, $1(\tau > j)$ indicates that the interview occurred after date j , and μ_{cj} represents the shift in support for candidate c occurring after date j . Relative to equation (16), which formed the basis for our baseline estimator, this alternative model also includes state fixed effects but allows for a more flexible pattern in mean support for the candidates, which is no longer required to follow the updating rule in equation (19) and thus does not incorporate information on the timing and outcomes of elections.

To estimate the most likely break date, we use interviews between January 2 and February 17, the day before Dean dropped out of the race.²³ As shown in Figure 7, the likelihood is clearly maximized by choosing January 19, the date of the Iowa caucus, as the break date.²⁴ That is, the

²² For an overview of this literature, see Hansen (2001).

²³ For the purposes of this analysis, it is helpful to not have time gaps in the data. No interviews occurred on January 1, and some dates prior to January 1 had a very small number of interviews.

²⁴ Following standard techniques, we trim 5 percent of the observations on each side and thus allow for all possible break dates between January 4 and February 15, for a total of 45 days.

model best fits the data by allowing for break in support between January 19 and January 20.²⁵ This evidence, which finds a break in support just after, and not before, a key election date, thus supports our social learning interpretation of the baseline results over this interpretation based upon nationally available information.

This reaction in polling data to the Iowa outcome is potentially confounded by Dean’s reaction, which was dubbed the “Dean scream” and was televised extensively. For a variety of reasons, we feel that this reaction does not fully explain these documented shifts in support. First, this media coverage would not arguably have occurred had Dean fared better in Iowa. That is, the Iowa outcome and Dean’s reaction to that outcome are not necessarily independent events. Second, the votes that Dean lost went to Kerry and Edwards, the winners in Iowa, rather than the losers, including Gephardt and Lieberman. Thus, even if Dean’s vote loss was due to Dean’s reaction, the reallocation of those votes is consistent with our story of momentum associated with social learning from aggregate returns. Third, the television coverage was concentrated in the days following the election, whereas our finding is that the break in support occurred on the election day itself.²⁶ Fourth, when we exclude Dean from the analysis and focus on the choice between Kerry and Edwards, we find a break on January 22, which is just 3 days after Iowa. Thus, our results are not entirely driven by Dean himself. Finally, when we take the Iowa break as given and allow for a second break in support, we find another shift in support on January 27, which is the date of the New Hampshire primary.²⁷

Our second test examines the degree to which respondents who have already voted are better informed than respondents who have not yet voted.²⁸ If information is available nationally and processed in the same manner in every state, as outlined above, then those who have not yet voted and those who have already voted should be equally well-informed. If information is available on a state-level basis, by contrast, then, among voters interviewed on the same date, those who have already voted have received their signals and should thus be better informed. To measure the degree to which voters are informed about candidates, we next examine whether respondents

²⁵ The associated coefficients capturing the size of the break (μ_{cj}) are jointly significant at conventional levels.

²⁶ In the first four days after the Iowa caucuses, the Dean scream was televised 633 times (Boston Globe, January 29, 2004).

²⁷ These two sets of results are available upon request from the authors.

²⁸ Note we did not use the first set of voters in our baseline analysis since they were not asked about their primary voting intentions.

answered correctly or incorrectly three factual questions involving candidate characteristics. For example, respondents were asked which candidate was the son of a mill worker, with the correct answer being Edwards.

As shown in Table 3b, which reports coefficients from a linear probability model with interview date fixed effects, those who have already voted were 20 percentage points more likely to know that Edwards was the son of a mill worker. Similarly, those who have already voted were more likely to know that Dean was the governor of Vermont. Third, those who have already voted are also more likely to know that Edwards was a trial lawyer. As shown in the fourth column, those who have already voted on average answered 0.74 more questions (out of three) correctly. One potential limitation of this analysis is that early voters, such as those from Iowa, may be better informed in general than late voters. To address this concern, we next estimate a model with state fixed effects and these results are presented in the final four columns. While the key coefficient in two of the four specifications is no longer statistically significant, due in part to a loss in power associated with estimating a large number of additional parameters, we still find evidence that those who have already voted answered more questions correctly in total. Thus, this test documents that those who have already voted are better informed and supports our social learning interpretation of the baseline results over an interpretation based upon nationally available information.

4.5.3 Campaign Finance and Persuasion

A final possible alternative explanation for our results involve campaign contributions and expenditures. Suppose that surprising wins in early states lead to an increase in campaign contributions from influence-motivated contributors—those who want to alter the platform of the likely winner of the election.²⁹ Suppose further that these contributions are used to fund political advertisements that persuade late voters to support the candidate sponsoring the advertisement. In this case, our results, which document an increase in support among late voters for candidates that outperform expectations in early primaries, could be driven by this alternative explanation.

We address this alternative explanation by examining the timing of both campaign contributions and expenditures in the 2004 primaries. We first compute the share of contributions, as measured in dollars, flowing to each of the three candidates during the key primary season, using

²⁹ Grossman and Helpman (2001) distinguish between the influence motive for contributing and the electoral motive. According to the latter motive, candidates have fixed platforms and contributors give to enhance the electoral success of their preferred candidate.

data from the Federal Election Commission. For consistency with the structural break analysis described above, we use data beginning on January 1, 2004. We also compute the share of advertising expenditures, as measured in the Wisconsin Advertising Data, by each of the three candidates during this same time period.³⁰

As shown in Figure 8, which focuses on campaign contributions, it is indeed the case that contributions to Kerry, relative to Dean, spiked following his surprising victory in Iowa on January 19. As shown in Figure 9, there is also a spike in advertising expenditures by Kerry, relative to Dean, following his victory in Iowa. Importantly, however, this spike did not occur until around January 26, a full 7 days after the Iowa caucus.³¹ As shown previously in Figure 7, however, the most likely date for a shift in support occurred just after the Iowa election. Given this discrepancy in the timing of the shift in campaign expenditures and support for the candidates, the increase in campaign expenditures cannot fully explain our baseline results.

To summarize, we have addressed and ruled out three alternative interpretations of our baseline results. We should emphasize that this analysis only demonstrates that our results cannot be explained solely by these alternative interpretations. It is of course possible that multiple factors shape voter perceptions of candidates. For example, it could be that both national and state-level information play an important role in sequential elections. Even if other motives are present, however, this analysis does provide strong support for the presence of social learning in this election.

5 Implied Voting Weights and the Allocation of Campaign Resources

Due to these documented momentum effects, early voters may have a disproportionate influence over the selection of candidates. In this section, we attempt to quantify any such influence by measuring the weights, or number of votes, afforded to voters from different states during the primary season. Any over-weighting of early voters associated with sequential voting represents a deviation from the democratic ideal of “one person, one vote.” While this property of sequential

³⁰ In this data set researchers recorded all of the political advertisements in 100 media markets from February 2003 to November 2004, coding characteristics for each ad including candidate, date, time of day, length, cost, tone, and other features. The cost of each ad was estimated based on the timeslot and geographic market.

³¹ This lag presumably reflects the delay in processing contributions and in placing new advertisements in the field.

voting has been frequently discussed in policy debates over the design of the primary system, there is little evidence on its magnitude. After measuring these weights, we then examine the consequences of any disproportionate influence for the distribution of campaign resources across states.

5.1 Measuring Voting Weights

We first use the estimated model to explicitly calculate the voting weights associated with sequential voting in the 2004 primary. Our measure of voting weights is based upon the effect of changes in state-level preferences (η_{cs}) on national vote shares:

$$\omega_{cst} = \frac{1}{N_s} \sum_r \frac{\partial v_{crt}}{\partial \eta_{cs}} = \frac{1}{N_s} \left[\frac{\partial v_{cst}}{\partial \eta_{cs}} + \sum_{r \neq s} \frac{\partial v_{crt}}{\partial v_{cst}} \frac{\partial v_{cst}}{\partial \eta_{cs}} \right] \quad (31)$$

where N_s is the number of states.³² As shown, these weights depend upon two effects. The direct effect ($\partial v_{cst}/\partial \eta_{cs}$) captures the change in candidate vote shares in the home state associated with a shift in preferences, and the second term captures the indirect multiplier effect, which equals zero in the absence of social learning. Even with social learning, however, the second term equals zero for voters from states voting in the final period ($t = 10$). For earlier states, by contrast, the multiplier effects will be positive in the presence of social learning.

We compute this derivative for each state voting on each of the primary dates ($t = 1, 2, \dots, 10$) and present in Figure 10 the average effect by time period.³³ We normalize the weight for the final time period to be equal to one. The plotted measures can thus be interpreted as the number of votes afforded to a representative voter from a state voting at time t . As shown, preferences of voters from the state of Iowa, the first state to vote, have roughly five times the influence relative to those individuals voting on Super Tuesday ($t = 10$), and voters from New Hampshire have over four times the influence. For comparison purposes, we also plot voting weights in the absence of social learning, those associated with only the direct effect ($\partial v_{cst}/\partial \eta_{cs}$). These may vary across states

³² In order to focus on differences in political power due to the timing of elections, the calculation of these weights abstracts from population differences across states. Given that delegates are allocated roughly proportional to population, however, incorporating state population would not significantly change the analysis.

³³ Specifically, we increase state-level preferences by one unit and re-compute vote shares for that state, as expressed in equation 9. In order to predict vote shares for subsequent states, we re-compute the posterior mean quality, as expressed in equation 11, and ultimately vote shares, as expressed in equation 9. Note that simulating these vote shares requires explicit measures of the voting signal (θ_{cs}), which can be backed out of equation 9 with information over state-level preferences (η_{cs}).

due to the fact that preferences enter vote shares in a non-linear manner. As shown, however, there are not large differences across states. Taken together, these results confirm the often-held notion that early states have a disproportionate influence over the selection of candidates in sequential primary systems due to momentum effects and thus represents a significant departure from “one person, one vote.”³⁴

5.2 Distribution of Campaign Resources

We next examine how this disproportionate influence of early voters affects candidates incentives in terms of the allocation of campaign resources, as measured by advertising expenditures. To motivate our empirical exercise, we extend our theoretical model by examining a candidate charged with allocating a fixed budget among each of the states, where A_{cs} denotes the advertising expenditure in state s . In developing this simple model of the allocation of campaign expenditures, we make four key assumptions. First, the budget is fixed and does not depend upon election outcomes.³⁵ Second, candidates commit to an allocation of campaign expenditures in advance of the primary season.³⁶ Third, candidates choose an allocation across states in order to maximize their overall vote share. Fourth, campaign spending in state s is productive via its shifting of the preferences of voters.³⁷ In particular, we assume that preferences depend upon spending as follows:

$$\eta_{cs} = \eta_{cs}^0 + f(A_{cs}) \tag{32}$$

where η_{cs}^0 is the baseline support for candidate c in state s and $f(A_{cs})$ is an increasing and concave function normalized so that $f(0) = 0$.

Given these assumptions and recalling the definition of ω_{cst} , the optimal allocation of campaign

³⁴ We also investigated a measure based upon the effect of changes in state-level information (θ_{cs}) on candidate vote shares. We find that the information held by Iowa voters has roughly 20 times the influence of information held by Super Tuesday voters. These information-based weights are substantially larger than the preference-based weights, as described above, given that voters in late states place less weight on their own signal. This underweighting of signals has a direct effect in the calculation of voting weights but also has an indirect effect as late vote shares are thus a noisier signal of quality.

³⁵ This assumption rules out the possibility that campaign contributions and hence budgets may change following election outcomes in early states.

³⁶ This assumption rules out the possibility that candidates may adjust their allocations in response to these election outcomes.

³⁷ This assumption rules out informative advertising, under which voters learn about candidate attributes, such as quality. See Prat (2006) for a review of these models.

spending can be characterized as follows:

$$\omega_{cst} f'(A_{cs}) = \lambda \tag{33}$$

where λ represents the Lagrange multiplier on the candidate budget constraint. Thus, as weights increase, $f'(A_{cs})$ must decrease, which requires that A_{cs} increase. In short, the model predicts that advertising is increasing in preference weights.

To test this prediction, we use the Wisconsin Advertising Data to examine the geographic distribution of campaign spending, as measured by the exposure (minutes) of residents to candidate advertisements. An important consideration is that these data are measured at the market (DMA) level. DMA's do not cover the entire United States and also often cross state boundaries. To convert these market-based measures to state-based measures, we take a weighted average across markets, where the weights are based upon the number of residents of state s living in the market relative to the number of residents in state s covered by the Wisconsin Advertising data.

As shown in the bottom panel of Figure 10, advertising exposure in state s , aggregated across the three candidates (Dean, Kerry, and Edwards), is increasing in our estimated preference weights. Thus, the prediction of this simple model of the allocation of campaign resources is supported, and this pattern demonstrates that candidates appear to respond to the momentum effects documented in this paper.

6 Counterfactual simulations

One advantage of our empirical model is our ability to conduct counterfactual experiments. We examine two counterfactual scenarios: simultaneous voting and alternative ordering of states under a sequential system.

6.1 Simultaneous primary

We first consider an election in which every state votes in a simultaneous national primary on January 19, 2004. In this case, vote shares in state s can be summarized as follows:

$$\ln(v_{cs1}/v_{0s1}) = \eta_{cs} + \alpha_1 \theta_{cs} + (1 - \alpha_1) \mu_{c1} \tag{34}$$

Accordingly, behavior in states voting after Iowa may be altered, relative to the sequential voting returns, for two reasons. First, all voters use the pre-Iowa prior (μ_{c1}). Given that voter priors favored Howard Dean in the days leading up to Iowa, we thus expect that he would have performed better in a simultaneous national primary. Second, at the time Iowa voted, voters were less certain in their evaluations of candidate quality and thus placed more weight on their private signals ($\alpha_1 > \alpha_t$ for $t > 1$). Thus, signals will be amplified in a simultaneous primary, and this second effect could benefit any of the three candidates depending upon the distribution of the realized signals.

Table 4 provides the key results from the actual sequential primary and the counterfactual simultaneous primary based upon the baseline estimated coefficients in tables 1 and 2. As noted above, Dean dropped out of the race following the Wisconsin primary and we thus cannot calculate counterfactual Dean vote shares for states thereafter.³⁸ We thus run two counterfactual simultaneous primaries: one in which Dean is included but in which states after Wisconsin do not vote and one in which Dean is excluded but all states vote. As shown in the Table 4, the results from the counterfactual three-candidate simultaneous election demonstrate that the election would have been much closer, with Dean winning in Michigan, Washington, Maine, and Nevada and Edwards winning in Oklahoma, Tennessee, Virginia, and Wisconsin. While Kerry would have won a plurality with 40 percent of delegates, he did not win a majority. Similarly, the two-candidate simultaneous primary would have been much closer with Edwards winning nine states. We do not wish to over-emphasize the predictive nature of the results from this simulation for specific states, such as the surprising finding that Edwards would have defeated Kerry in Massachusetts.³⁹

Rather, we hope that the heightened competition under a counterfactual simultaneous election helps to further reinforce our finding that Kerry benefited from the sequential primary system.

³⁸ Of course, under a national primary, he would have been on the ballot in every state. But we abstract from that issue given our inability to measure the signals for these states. Moreover, without a model of candidate exit, it is difficult to predict how Dean would have performed in subsequent states following his decision to drop out of the race.

³⁹ This prediction of a win by Edwards in Massachusetts under a national primary reflects the fact that Edwards did better than expected from the perspective of the econometrician given Kerry's home state advantage and the state of the race going into Super Tuesday. This in turn implies that voters in Massachusetts received a positive signal regarding Edwards relative to Kerry and this signal is amplified when considering the significantly higher weight placed upon signals at the beginning of the primary season, relative to the end of the primary season.

6.2 Alternative sequential schedules

Our second counterfactual election involves changes in the voting order under a sequential schedule. In the context of our model, there are two reasons why the sequence of elections can affect voting outcomes: 1) early signals have more weight and 2) preferences of early states have more weight. For these reasons, the voting outcome may be fragile or sensitive to the order of voting. To investigate this issue, we randomly generated alternative voting sequences, holding constant the number of states voting on each date and assigning each state the same signal it received in the actual sequence. Again, we consider a two-candidate election, in which all states are included, and a three-candidate election, in which only states voting prior to and including Wisconsin are included. As shown in Table 5, Kerry continues to win a plurality of states in most cases under both the sequential two-candidate and sequential three-candidate elections. Using delegate weights, however, the counterfactual sequential elections are somewhat more competitive, with Edwards winning 11 percent of the two-candidate sequential election schedules. Again, while Kerry still wins most of these counterfactual elections, there are a sizeable number of cases in which Edwards would have won. This is surprising given the wide margin by which Kerry won the actual election and highlights the sensitivity of electoral outcomes to the sequencing of states.

7 Conclusion

Given our goal to develop a tractable empirical framework, we have kept the model simple and have thus abstracted from many relevant features of electoral politics in the United States. We thus view this model as a first step in a larger research agenda and plan to extend the environment in a variety of ways in subsequent work. A first possible extension involves the media. While the process through which voters observe signals was taken as exogenous here, one could introduce a media outlet that reports election results in early states to voters in late states. Second, one could model candidate entry and exit, which we have taken as given in this paper. Candidate exit would presumably depend upon the degree of social learning, which may limit the ability of trailing candidates to make up lost ground in late states. A third extension would be to examine whether candidates alter their platforms in favor of issues that are most important to voters in early states. Whether or not such strategies are effective presumably depends upon whether or not voters in later states condition on such candidate behavior when analyzing voting returns

from early states. Finally, one could conduct a welfare analysis of simultaneous versus sequential elections. On the one hand, voters in later states have more information under a sequential system and thus presumably make better choices. On the other hand, signals in early states are overweighted and state preferences may be misinterpreted as signals of quality; outliers in these early signals or early state preferences could lead to selection of a lower quality candidate.

In summary, we have developed and estimated a simple model of voter behavior under sequential elections. In the model, voters are uncertain about candidate quality, and voters in late states attempt to infer private information held by early voters from voting returns in early states. Candidates experience momentum effects when their performance in early states exceeds voter expectations. Our empirical application focuses on the 2004 Democratic primary. We find that Kerry benefited substantially from surprising wins in early states and took votes away from Dean, who stumbled in early states after holding strong leads in polling data prior to the primary season. Early states have up to five times the influence of late states, and candidates respond to these differences by funneling campaign expenditures into early states. Finally, we simulate the election under a number of counterfactual primary systems and show that the race would have been much tighter under a simultaneous system and that electoral outcomes are sensitive to the order of voting. While these results are specific to the 2004 primary, we feel that they are informative more generally in the debate over the design of electoral systems in the United States and elsewhere and also hope that the methods developed in this paper will be applied in other settings as well.

References

- R.E. Adkins and A.J. Dowdle. How Important Are Iowa and New Hampshire to Winning Post-Reform Presidential Nominations? *Political Research Quarterly*, 54(2):431–444, 2001.
- S.N. Ali and N. Kartik. Observational Learning with Collective Preferences. *working paper, U. Columbia*, 2010.
- A.V. Banerjee. A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- L.M. Bartels. Candidate Choice and the Dynamics of the Presidential Nominating Process. *American Journal of Political Science*, 31(1):1–30, 1987.
- L.M. Bartels. *Presidential Primaries and the Dynamics of Public Choice*. Princeton University Press, 1988.
- M. Battaglini. Sequential Voting with Abstention. *Games and Economic Behavior*, 51(2):445–63, 2005.
- M. Battaglini, R. Morton, and T. Palfrey. Efficiency, equity, and timing of voting mechanisms. *American Political Science Review*, 101(03):409–424, 2007.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992, 1992.
- H. Cai, Y. Chen, and H. Fang. Observational Learning: Evidence from a Randomized Natural Field Experiment. *The American Economic Review*, 99(3):864–882, 2009.
- S. Callander. Bandwagons and momentum in sequential voting. *Review of Economic Studies*, 74(3):653–684, 2007.
- C. Chamley. *Rational herds*. Cambridge University Press Cambridge, 2004.
- E. Dekel and M. Piccione. Sequential Voting Procedures in Symmetric Binary Elections. *Journal of Political Economy*, 108(1):34, 2000.
- A. Devenow and I. Welch. Rational Herding in Financial Economics. *European Economic Review*, 40(3-5):603–615, 1996.
- A.D. Foster and M.R. Rosenzweig. Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *The Journal of Political Economy*, 103(6):1176–1209, 1995.
- E.L. Glaeser and B. Sacerdote. Aggregation Reversals and the Social Formation of Beliefs. *NBER Working Papers*, 2007.
- G. Grossman and E. Helpman. *Special Interest Politics, 2001*. Cambridge: MIT Press.
- B.E. Hansen. The new econometrics of structural change: Dating breaks in US labor productivity. *Journal of Economic perspectives*, 15(4):117–128, 2001.

- T. Klumpp and M.K. Polborn. Primaries and the New Hampshire Effect. *Journal of Public Economics*, 90(6-7):1073–1114, 2006.
- R.D. McKelvey and P.C. Ordeshook. Elections with limited information: A fulfilled expectations model using contemporaneous poll and endorsement data as information sources* 1. *Journal of Economic Theory*, 36(1):55–85, 1985.
- R.B. Morton and K.C. Williams. Information asymmetries and simultaneous versus sequential voting. *American Political Science Review*, 93(1):51–67, 1999.
- K. Munshi. Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution. *Journal of Development Economics*, 73(1):185–215, 2004.
- A. Prat. Rational voters and political advertising. *The Oxford handbook of political economy*, page 50, 2006.
- L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- W.P. Steger. Forecasting the presidential primary vote: Viability, ideology and momentum. *International Journal of Forecasting*, 24(2):193–208, 2008.
- W.P. Steger, A.J. Dowdle, and R.E. Adkins. The New Hampshire effect in presidential nominations. *Political Research Quarterly*, 57(3):375, 2004.
- K.S. Strumpf. Strategic Competition in Sequential Election Contests. *Public Choice*, 111(3):377–397, 2002.
- B. Sushil, D. Hirshleifer, and I. Welch. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*, pages 151–170, 1998.
- I. Welch. Sequential Sales, Learning, and Cascades. *The Journal of Finance*, 47(2):695–732, 1992.
- I. Welch. Herding among security analysts. *Journal of Financial Economics*, 58(3):369–396, 2000.
- J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.