EXHIBIT 7
DATE 02/19/2013
HB 466

How Much is Enough? The "Ballot Order Effect" and the Use of Social Science Research in Election Law Disputes

R. Michael Alvarez, Betsy Sinclair, and Richard L. Hasen*

Abstract

Previous empirical research and other related research from survey methodology holds that candidates listed first on an election ballot may gain some measure of advantage from this ballot placement. Using data from the 1998 general election in California, we test whether a candidate's relative position on the ballot has any statistical effect on vote shares. We find little systematic evidence that candidate vote shares benefit from being listed first on the ballot. We show that there is not a primacy ballot order effect (defined as being listed first on the ballot) in every contest, that when the effect exists it is often very small, and that the effect is evenly distributed between primacy and latency (defined as being listed last on the ballot). We consider how courts should balance the concern over ballot order effect against other interests, such as the costs and potential confusion associated with rotation and randomization.

1 Introduction

Conventional wisdom holds that candidates listed first on an election ballot may gain some measure of advantage from this ballot placement. This belief is shared by courts, which have relied on it to change election outcomes. For example, the 2001 mayoral election results in Compton, California were recently overturned and the losing

* R. Michael Alvarez is Professor of Political Science, California Institute of Technology. Betsy Sinclair is a graduate student, Division of Humanities and Social Sciences, California Institute of Technology. Richard L. Hasen is the William H. Hannon Distinguished Professor of Law, Loyola Law School, and co-editor, *Election Law Journal*. Alvarez and Hasen served as experts and consultants for the

incumbent mayor was (temporarily) reinstated because of a supposed failure by the city clerk properly to randomize the candidate names on the ballot. The trial court—eventually overruled—believed that permitting a single candidate to be first necessarily gives that candidate an advantage. How much proof of this "ballot order effect" should be necessary before courts hold that failure to randomize and rotate the names of candidates on the ballot to neutralize the effect denies candidates (or voters) some measure of equal protection of the laws? And how should courts balance the concern over the ballot order effect against other interests, such as the costs and potential confusion associated with rotation and randomization? Is there enough proof of the effect for courts to order a change in the outcome of an election by reason of a failure to rotate or randomize?

Courts have reached widely divergent results on the legal questions. In 1975, the California Supreme Court held that there was enough evidence of the ballot order effect to strike down, on equal protection grounds, a state law requiring incumbents to be listed first on the ballot.[1] In 1999, however, a New Jersey court refused to accept the defendants' concession that the effect existed without an examination of supporting evidence: "If the plaintiffs and defendants believed in the legitimacy of the Easter Bunny, this court would not be required to find as a fact (or through judicial notice) that such a creature truly exists."[2]

Shortly after the New Jersey case, the loser of a mayoral election in the city of Compton, California argued that he lost because the local elections official failed to use the

---

city of Compton in the litigation described in this paper. We thank Jonathan Katz, Geralyn Miller, Jennifer Steen, and Melanie Goodrich for helpful discussions about this paper

[1] Gould v. Grubb, 536 P.2d 1337 (Cal. 1975).

proper "randomized" alphabet as required by state law. The loser claimed his opponent gained an unfair advantage by being listed first and that this error led to his narrow loss. The trial court, relying upon one study of the ballot order effect and the testimony of the study's coauthor, allocated a certain number of votes that had been cast for the winner as instead cast for the loser, and declared the loser the winner of the election. An appellate court later reversed, holding, among other things, that "[w]hile many courts and legislatures have recognized the advantage afforded to candidates whose names are listed first on the ballot, no judicial or statutory authority exists to reverse the results of an election where, due to *unintentional* clerical error, the ballot listed the candidates in the wrong alphabetical order."[3] Even more recently, a New Hampshire court refused to strike down a state law providing that the party which won the most votes in the previous election goes first on the general election.[4]

Social science evidence is not new to the courtroom, nor even to election law cases. It is regularly used in redistricting and campaign finance disputes.[5] Its use acquired the potential to become even more prevalent after the Supreme Court decided *Bush v. Gore*,[6] because the Supreme Court appeared to recognize an equal protection right not to have one person's vote valued more than that of another.[7] Thus, the Ninth Circuit Court

---

[2] New Jersey Conservative Party v. Farmer, 753 A.2d 192, 198 (N.J. Super. 1999).
[3] Bradley v. Perrodin, 131 Cal.Rptr. 2d 402, 417 (App. 2003) (original emphasis). The court also ruled that the Compton city clerk did not err in using the wrong randomized alphabet for the election.
[4] Associated Press, *New Hampshire Judge Upholds Ballot Order*, BOSTON GLOBE, Oct. 18, 2005.
[5] Richard H. Pildes, *The Politics of Race: Quiet Revolution in the South*, 108 HARV. L. REV. 1359 (1995); Richard L. Hasen, *Measuring Overbreadth: Using Empirical Evidence to Determine the Constitutionality of Campaign Finance Laws Targeting Sham Issue Advocacy*, 85 MINN. L. REV. 1773 (2001).
[6] 531 U.S. 98 (2000).
[7] On how Bush v. Gore increases the role of social science in election disputes, *see* Richard L. Hasen, *The Uses, Normative Implications, and Unintended Consequences Of Voting Reform Research in Post*-Bush v.

recently considered – though it ultimately rejected – the argument that a special California gubernatorial recall election should be postponed because of disparities in vote-counting errors allegedly caused by the use of different voting technology in California counties.[8]

This article uses the example of the ballot order effect to argue that courts should be cautious before using generalized social science findings to decide election law cases. Although there seems little question that a ballot order effect exists in some circumstances, there has been little solid social science analysis of the direction and strength of the effect in different electoral contexts. In Part I of the article, we briefly review four previous studies on the topic. As we shall see, there is still dispute over the proper methodology to measure the effect and even whether the effect exists at all in general elections, as opposed to primaries.

In Part II of this paper, using a different methodology from those discussed in Part I, we find little support for the ballot order effect in a 1998 California *general* election. In particular, we find little evidence that candidates listed first on the ballot are necessarily advantaged over those listed last or those listed in the middle of the ballot. When we do find evidence that being first on the ballot helps candidates, we find that the effect is very small.

In Part III of the paper, we discuss the legal implications of the muddled state of social science evidence. We conclude that although there is enough evidence of a ballot order effect to justify a legislative decision to use randomization and rotation, we do not

---

Gore *Equal Protection Challenges*, in RETHINKING THE VOTE 185 (Oxford University Press, Ann Crigler, Marion Just, and Edward McCaffery eds., 2004).

believe there is yet enough evidence for courts to mandate such a change, especially given the costs that are associated with the process of randomization and rotation. We also argue that in cases where the loser of an election alleges a failure to follow state law on ballot order, the appropriate remedy is not to reallocate votes from one candidate to another. We evaluate the trade off between increasing the number of rotations and decreasing the possible ballot order effect.

2 The Existing Literature

Recent examination of the ballot order effect began with Miller and Krosnick's 1998 analysis.[9] The authors began by examining 26 earlier studies on ballot ordering effects. They found that 18 of these studies had significant methodological flaws (though 14 of the 18 did find that candidates listed early did better than candidates listed later on ballots). Six stronger but still imperfect studies produced mixed results. Some found that candidates listed early did better, but others found that candidates listed last did better. Neither of the two remaining studies found any ordering effect.

Miller and Krosnick then undertook their own analysis focused on 1992 elections in the three largest counties in Ohio – Franklin, Cuyahoga and Hamilton– each of which had different procedures for listing candidate names on the ballot. Of the 182 county elections studied by Miller and Krosnick, only 40.7% had a statistically significant ballot order effect (defined as a benefit or detriment for the candidate listed first on the ballot as compared with the candidate listed last), of which 1.6% were negative. Thus, in 39.0% of the cases in their study, the ballot order gave the first-listed candidate an advantage over

[8] Southwest Voter Registration and Education Project v. Shelley, 344 F.3d 882 (9 th Cir.), *rev'd en* banc, 344 F.3d 914 (9th Cir. 2003).

the last-listed candidate. Furthermore, the ballot order effect varied dramatically across the three counties: 72.2% of the cases in Franklin County were statistically significant, 31.3% in Cuyahoga, and 20.0% in Hamilton County; without the negative and significant cases, the heterogeneity across counties is more dramatic, as Cuyahoga's estimate drops to 27.7%.

In general, Miller and Krosnick found little overall difference in the ballot order effect in two-candidate races relative to three-candidate races. They found that 42.7% of the two-candidate elections have statistically significant ballot order effects (without the three negative and significant cases, the percentage drops to 39.6%) relative to 38.4% of the three-candidate races. In Franklin County, the ballot order effect appeared significant in 71% of two-candidate and 73.9% of three-candidate races. However, in Hamilton County the rate was slightly higher for three-candidate races, but the opposite was true in Cuyahoga County.[10]

Koppell and Steen analyzed precinct-level data from the 1998 Democratic primary in New York City.[11] Their results indicate that in the four statewide primaries in their sample, candidates received a statistically significant increase in votes when listed first. For example, Democratic gubernatorial candidates received a vote advantage of 2.3% when in the first ballot position.[12] Koppell and Steen also looked at the local races, finding that 67 of 75 candidates in the first ballot position received more than the

---

[9] Joanne E. Miller & Jon A. Krosnick, *The Impact of Candidate Name Order on Election Outcomes* , 62 PUBLIC OPINION QUARTERLY 291 (1998).

[10] Krosnick has extended his analysis to include the 2000 elections in Ohio, North Dakota, and California (J.A. Krosnick, J.M. Miller, & M.P. Tichy, *An Unrecognized Need For Ballot Reform: Effects of Candidate Name Order, in* Rethinking the Vote: THE POLITICS AND PROSPECTS OF AMERICAN ELECTION REFORM 51 (A.N. Crigler, M.R. Just, & E.J. McCaffery eds., 2004).

[11] Jonathan GS Koppell & Jennifer A. Steen, *The Effects of Ballot Position on Election Outcomes* , 66 JOURNAL OF POLITICS 267 (2004).

expected number of votes, while in 17 of 75 cases the advantage was statistically significant. The magnitude of the effect of being first on the ballot was greater for down-ballot primary races than for the top-of-the-ticket primary races—almost a 4.5% benefit in the state committeewoman race, compared to 2.3% for candidates for governor. Koppell and Steen considered only the advantage that might come to a candidate if she appears first on the ballot — not considering the advantage that might arise if she appears in other positions on the ballot, especially last.

Ho and Imai examined state-wide California primaries and general elections from 1978 through 2002.[13] Employing a sophisticated statistical analysis of these data, Ho and Imai find little systematic evidence that major party candidates in general elections are favored by being first on the ballot. They find that only minor party or nonpartisan candidates are advantaged by being listed first on the ballot. But they do interpret their results as showing a stronger advantage for major party candidates when listed first on the ballot in *primary* elections.[14] Again, as in much of the contemporary literature in this area, Ho and Imai are interested in the effects of being first on the ballot. Like the other recent research papers, they find evidence they claim supports the hypothesis that being listed first on the ballot provides an advantage to those candidates.

---

[12] See Table 1 of their paper.

[13] Daniel E. Ho & Kosuke Imai, *The Impact of Partisan Electoral Regulation: Ballot Effects from the California Alphabet Lottery, 1978-2002*, manuscript, *available at http://imai.princeton.edu/research/alphabet.html. Table 2 of their paper lists the races they concentrate on,* which are most of the presidential, senate and gubernatorial primary and general elections from 1978-2002, and the primary and general elections for the other statewide constitutional offices in 1998 and 2002.

[14] Methodologically, Ho and Imai's technique relies on the assumption of "no interference among units," which here means ignoring the compositional nature of the data, as an increase in one candidate's vote share may decrease the votes for the other candidates. As Ho and Imai state, "the assumption is violated in an analysis that pools candidates, typical in this literature: since the candidate vote shares in one district must sum to 1, a ballot order effect on one candidate necessarily affects the remaining candidates" (page 13). Ho and Imai state that they have reproduced their analysis using another technique that relaxes this assumption, and that those "results are largely consistent with those presented here" (page 14).

In summary, the recent literature has focused on the effect of being first. It is possible, however, that there are multiple locations on the ballot which result in a positive or negative change in vote shares; the first ballot position may in fact be the most preferable position but it is not sufficient to simply compare being first to being located elsewhere unless being located elsewhere is guaranteed of producing a lower vote share. In order to make that assumption it is necessary to evaluate multiple positions on the ballot, especially the effect of being last. Furthermore, voters experience different levels of available information based upon the type of election. In primary elections, such as those studied by Koppell and Steen, voters do not have the differences in candidate partisanship from which to base their voting decision. Ballot order effects are likely to be different across election types. Also, these papers fail to account for the fact that an increase for one candidate in terms of vote share is likely to imply a decrease for another candidate's vote share. The individual candidates' vote shares therefore are not independent. That fact casts doubt on the use of ordinary least squares analysis (OLS), for which independence is crucial. Our analysis, which we detail below, estimates the ballot order effect more efficiently than OLS and does not rely upon the independence assumption.

3 The California Context

The order in which candidate names are listed on a ballot has been a matter of state and local law. Not surprisingly, states and localities have devised many different rules and procedures for determining ballot order. In the past, the question was often resolved through some system of "privilege" – incumbents were listed first, or the candidates of the party controlling the state government were listed first. Elsewhere,

methods like alphabetical order were used. Today, most states and localities use some type of random method to determine where candidate names fall on the ballot.[15]

In California, state courts had held in 1975 that a system, based on alphabetical order or listing incumbents first was unconstitutional due to a supposed 5% ballot order effect among undecided voters. Legislation passed in 1975 (codified in California's Election Code Sections 13111 through 13114) details procedures for choosing the order of candidate names.

The general principles are random ization and rotation. The Secretary of State, before an election, conducts a random drawing of letters of the alphabet.[16] This randomized list is used for all candidate races in the upcoming election. In statewide races, the randomized alphabet list drawn by the Secretary of State determines the order of candidate names on all ballots in Assembly District 1. The order is then rotated, so that in Assembly District 2 the candidate who appeared first in Assembly District 1 is moved to the bottom of the ballot list and the second candidate is moved up to be the first candidate on the ballot in Assembly District 2. This rotation process continues throughout all of the 80 Assembly Districts in California.

All other types of elections in California begin with the same randomized alphabet list. However, the process for rotation differs. Congressional candidates' names rotate by Assembly Districts, but State Senate and Assembly candidate names do not rotate unless the legislative district crosses county lines, in which case county election officials conduct random drawings to determine the order in their own counties. Elections that

---

[15] *See* Ho & Imai, Table 7, for a list of the methods that states use to position candidate names on their ballots.

[16] CAL. ELEC. CODE § 13112 (Deering 2005).

occur throughout a single county begin with the random alphabet list and are rotated based on the county supervisorial districts (for smaller counties) or the Assembly Districts (for larger counties). All other local elections, occurring within but not throughout a county, are randomized but not rotated.

4 A New Methodology for Studying the "Ballot Order Effect"

We study the eight statewide general elections in 1998 across California. We use a new methodology to study the relative impact of candidate name order on the ballot, based on recent advances in the study of multiparty and multicandidate electoral data by Katz and King and Tomz, Tucker and Wittenberg.[17]

We used the 1998 general election for three reasons. First, as 1998 was a California statewide election year, all of the partisan state constitutional offices were contested (7 races) and there was a U.S. Senate election on the ballot. That gives a total of eight races that we can examine across the 80 Assembly Districts in California. Second, the 1998 election returns are readily available in a format suitable for analysis. We have a very detailed data set available to study the impact of candidate order on the ballot, as we employ a census tract-level database with a wide array of control variables. Third, as California's primary process has been in flux in recent years, we decided to focus on a recent general election.

---

[17] Jonathan Katz & Gary King, *A Statistical Model for Multiparty Electoral Data*, 93 AM. POL. SCI. REV. 15 (1999); Michael Tomz, Michael, Joshua A. Tucker & Jason Wittenberg, *An Easy and Accurate Regression Model for Multiparty Electoral Data*, 10 POLITICAL ANALYSIS 66 (2002).

Specifically, we obtained the state wide election data for the 1998 election from the Statewide Database, which collects and archives data from California elections.[18] The database includes the census tract election return file in the 1998 general election. While the data are organized by census tract, each tract is located within an Assembly District. Any census tracts that are not completely contained within a single Assembly District are discarded.

The dependent variables of our analysis are the vote shares in each census tract received by every candidate on the statewide ballot for each contest. Thus, to take the governor's race as an example, we had seven dependent variables, one for the respective percentage of the total vote received by the candidates running as Democrats, Republicans, American Independence Party (AIP), Green, Libertarian, National Law Party, and Peace and Freedom Party. Our goal is to test whether or not the appearance of each party's candidate first or last on the ballot, relative to appearing in the middle of the ballot, produces a marginal increase or decrease in the party candidate's vote share.

We have three methodological improvements to introduce relative to previous work in this area. The first improvement is to incorporate control variables into a multivariate statistical model to account for differences in vote shares across census tracts and assembly districts. The second improvement is that we control for both primacy and latency, where previous literature has focused entirely upon primacy. The third improvement is to utilize a more appropriate statistical framework for estimating the marginal effects of

---

[18] Statewide Database, *available at* http://swdb.berkeley.edu/index.html.

candidate name order using a method appropriate for multiparty data, holding constant the control variables.

In the available tract -level data, we have an important array of control variables for each census tract. Census tracts vary widely in terms of partisan registration, for example, so a detailed analysis of ballot order effects must incorporate those variables which are likely to affect vote shares. While it is true that the alphabet is determined at random, it is not true that the rotation occurs at random -- rotations move across the state in the order of assembly districts. As a consequence, it would be possible for a candidate to be placed first on the ballot in every 6th assembly district, an order which would give her more first-place locations on the coast. California's coast tends to be more heavily Democratic than the interior of the state; as a consequence what might first be observed as a "boost" for that particular candidate's rotation might in actuality occur due to the specific partisan registration associated with that rotation.

First, in the 1998 general election, there were eleven initiatives and propositions included on the statewide ballot. The voting patterns on these eleven initiatives provide an opportunity to include a measure of the tract's ideology. Accordingly, we undertake a factor analysis of each tract's vote on all eleven proposition and initiatives, and the factor score from this analysis is used as a measure of ideology in our multivariate statistical model.[19] Second, we have partisan registration data for each tract: We use construct variables that measure the Democratic, Republican, decline-to-state, and nonpartisan registration in each census tract. Third, we have some demographic information for each tract: we include variables that measure the percentage of Hispanic registered voters in

[19] The results from this factor analysis are available from Sinclair.

each census tract and the percentage of registered female voters in each census tract. We use this set of independent variables as controls in our multivariate statistical model. These are all variables that may affect the percentage of voters who vote for a particular candidate or party. It seems likely that with the addition of these control variables we are more likely to be able to differentiate the percentage of the voting population choosing a candidate based upon voter preferences and the percentage of choosing based upon the candidate's location on the ballot. All census tracts are not homogeneous; therefore, by controlling for the differences that affect electoral outcomes we are more likely to observe the ballot order effect.

As we know the randomized order of candidate names as they were used in Assembly District 1, and the pattern of rotation for the remaining Assembly Districts, it is simple to construct two dummy variables for each Assembly District and each candidate in order to account for the effect of candidate name order on the ballot. We have one dummy variable for each candidate in a race, coded 1 for Assembly Districts where that candidate was first on the ballot, and zero otherwise. We have a second dummy variable for each candidate in a race, coded 1 for Assembly Districts where that candidate was last on the ballot, and zero otherwise. These two dummy variables for each candidate will allow us to estimate the effect of candidate ballot name order across Assembly Districts, holding our control variables constant.[20]

[20] Of course, we are simplifying the problem somewhat here, as we are only examining the effect of being first, last, or at any position between first and last. Our methodology could be used to estimate the effect of all positions in a race; for example, in a race with eight candidates on the ballot, we could estimate *the effect of being first, second, third, (and so on) on vote shares. As we have no theoretical reason that* would indicate that the ballot positions between first and last should be salient focal points for voters, we do not use such a broader specification of the possible impact of ballot position on vote choice; future work could seek to determine if, contrary to our expectations, positions between first and last indeed have some salience as focal points for voters.

Given how we have operationalized these two indicator variables for the candidate's position on the ballot, we can determine if their relative position on the ballot had any statistical effect on their vote shares. The regression coefficients on these two dummy variables allow for the clear differentiation of four different types of ordering effects (here, for simplicity we will call the dummy variable indicating whether or not the particular candidate was first on the ballot "first" and the dummy variable for whether or not the candidate was last on the ballot "last"). If we find a statistically significant and positive coefficient on the "first" dummy, that is evidence for primacy (as this indicates that the candidate received a statistically significant increase in votes where he or she was first on the ballot). If we find a statistically significant and positive coefficient on the "last" dummy, that is evidence for latency (as the candidate received a statistically significant increase in votes where he or she was listed last on the ballot). If we find a statistically significant and negative coefficient for the "first" dummy variable, that is evidence for anti-primacy – here the candidate received fewer votes in precincts where he or she was listed first on the ballot. Finally, if we observe a statistically significant but negative coefficient on the "last" dummy variable, we have evidence for anti-latency, as here the candidate is receiving fewer votes in precincts where he or she is listed last on the ballot. We expect to observe both primacy and anti-latency based upon the survey literature that indicates that respondents are more likely to check the items listed first.[21]

Our statistical approach is the same as that recently developed by Tomz, Tucker and Wittenberg, and is similar to that earlier exposited by Katz and King. Katz and King noted that when estimating models with dependent variables involving multiparty elec-

[21] DON DILLMAN, MAIL AND INTERNET SURVEYS: THE TAILORED DESIGN METHOD (2000).

tion outcomes, important assumptions that underlie ordinary least squares (OLS) regression are untenable and thus OLS regression is likely to produce incorrect results.[22] One concern about OLS is that the percentage vote that each candidate received in a tract is necessarily bounded between 0 and 100. Since OLS assumes that the dependent variable is unbounded and continuous, OLS is likely to produce incorrect results in this setting. The second problematic assumption is the independence of each candidate's vote share. The application of OLS sequentially to each candidate's vote share assumes that these vote shares are independent, when in fact they are certainly not independent as the vote shares must all sum to 100%.[23]

Katz and King develop a methodology that is suitable for data involving three candidates or parties; Tomz, Tucker and Wittenberg have developed a more general procedure for data like ours that includes more than three parties or candidates. The Tomz, Tucker, and Wittenberg approach first transforms candidate vote shares into log-odds ratios, and then uses seemingly unrelated regression (SUR) on a series of J-1 regressions (where J indexes the number of candidates in the particular application).[24] Our use of this methodology to estimate possible ballot order effects is superior to previous studies that have used OLS for two reasons.

First, the use of the log-odds ratios alleviates the problems associated with the bounded nature of this vote share data. The vote shares need to be transformed into log-

---

[22] Katz & King, *supra* note 18.

[23] Similar work in this area is by John E. Jackson, *A Seemingly Unrelated Regression Model for Analyzing Multiparty Elections*, 10 POLITICAL ANALYSIS 49 (2002); James Honaker, Jonathan N. Katz, & Gary King, *A Fast, Easy, and Efficient Estimator for Multiparty Electoral Data*, 10 POLITICAL ANALYSIS 84 (2002); and Nikolai Mikhailov, Richard G. Niemi & David Weimer, *Application of Theil Group Logit Methods to District-Level Vote Shares: Tests of Prospective and Retrospective Voting in the 1991, 1993, and 1997 Polish Elections*, 21 ELECTORAL STUDIES 631 (2002).

odds ratios since it is impossible for any candidate to get less than 0% or more than 100% of the vote. Permitting the dependent variable to extend beyond these bounds could result in coefficient estimates that would predict vote shares outside the realm of possibility – for example, one coefficient could predict a candidate would get over 100% of the vote if the vote shares had not been incorporated as log-odds ratios. Log-odds ratios are appropriate and they constrain the predictions to a feasible interval, helping to insure that we do not produce estimates that are clearly illogical.

Second, the use of SUR allows the error terms across the J-1 regressions to be correlated meaning that we can avoid making the independence assumption made in previous studies. This methodology incorporates the fact that the vote shares of candidates are correlated; for example, an increase in the Democratic candidate's vote share is likely to mirror a decrease in the vote shares the other candidates would have received. This correlation would be reflected in the disturbance terms of the regressions if they were calculated individually. SUR estimates the coefficients by assuming that the error terms are correlated across candidates and thus, by incorporating this additional information in the estimation of coefficients, produces more efficient coefficient estimates.

5 Results from the 1998 California General Election

We use the Tomz, Tucker and Wittenberg estimation procedure, and we thus are estimating J-1 regressions for each contest. In this situation, we normalize by excluding the Republican party from the SUR analysis. On the right-hand side of each of these

---

[24] Arnold Zellner, *An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for*

regressions are variables measuring whether that particular candidate was first on the ballot, last on the ballot, ideology, partisanship, and some demographic attributes of the census tract.

We do not report the full SUR results.[25] Instead, we report a summary table (Table 1) of our results, which shows the number of significant coefficients in each contest. We subclassify these into the number of significant primacy effects (that candidates who appeared first on the ballot in that contest had a significant and positive estimated coefficient on the dummy variable measuring if they were first at the alpha =.05 level), the number of significant latency effects (that candidates who appeared last on the ballot in that contest had a significant and positive estimated coefficient on the dummy variable measuring if they were last); the number of significant antiprimacy effects (candidates appearing first had a significant but negative estimated coefficient on the dummy variable measuring if they appeared first on the ballot); and the number of significant anti-latency effects (candidates appearing last had a significant and negative estimated coefficient on the dummy variable measuring if they appeared last on the ballot). Note that although the number of candidates in Table 1 is inclusive of Republicans, there are no Republican coefficients because all coefficients are normalized against the Republican Party. Note further that it is possible for the same party to exhibit both an effect for primacy and

---

Aggregation Bias, 57 JOURNAL OF THE AMERICAN STASTICAL ASSOCIATION 298, 348-368 (1962).

[25] Available from Sinclair upon request. We include nine independent variables on the right- hand side of each SUR regression; thus for each contest we produce J1 equations of results, or 9 times J1 coefficients and standard errors.

latency. As a consequence the number of significant effects may be greater than the number of candidates for any particular contest.[26]

Table 1 Goes Here

We find in Table 1 little systematic evidence supporting the hypothesis that the primacy phenomenon prevails. We do see evidence in every contest that some candidates appear to have received some marginal benefit by having their names at the top of the ballot, but never for more than 2 of 7 candidates in a contest. On the other hand, in Table 1 we see exactly the same support for the hypothesis of latency – candidates whose names are at the bottom of the ballot win additional votes as often as candidates who have their names at the top of the ballot.

The third and fourth columns of Table 1 provide results for what we call anti-primacy and anti-latency. These are situations where we obtained statistically significant but negative estimates for the candidate's name appearing at the top or bottom of the ballot. Notice here that anti-primacy and anti-latency effects occur more often than primacy and latency. But anti-primacy and anti-latency do not always prevail. In other contests (like Controller) we see few effects from any ballot-positioning.

The summary results presented in Table 1 do not provide support for the hypothesis that when candidates are listed first on the ballot they systematically receive a greater vote share than when listed elsewhere on the ballot. We find evidence for a arbitrary

[26] The Republican Party was chosen as the comparison party because it appears in all contests on the ballot. Furthermore, the advantage of selecting a major party is that a major party is more likely to have greater variance in its vote shares and thus introduce more information into the estimates.

pattern: sometimes being first helps some candidates, but sometimes being last increases a candidate's vote share. In a larger number of cases, being first or last appears to statistically decrease the candidate's vote share, a result that is completely opposite of what we expect from the primacy and latency hypotheses.

In Table 2 we provide the same basic summary of our SUR analysis, but by party. That is, we give the number of partisan candidates for whom we estimate the presence of each effect (primacy, latency, anti-primacy and anti-latency). This breakdown of our multivariate analysis gives us the ability to determine if there are any patterns in primacy or latency for particular parties or types of parties (for example, major relative to minor parties).

Table 2 Goes Here

In brief, no clear pattern emerges from Table 2. Democratic candidates seem slightly more likely to have primacy or latency effects in this election. AIP, Green and Peace and Freedom candidates are also essentially equally likely to be estimated to have experienced one of the four effects. Libertarian and Natural Law candidates appear more likely to have experienced anti-primacy and anti-latency, as were Reform candidates. Again, the conclusion from the presentation of the SUR results as summarized in Table 2 is that there is no clear pattern, by party of the candidate. Again note that as all coefficients are estimated with respect to the Republican Party, no Republican

coefficients themselves are estimable (otherwise the model would not be identified).[27] We next present these coefficients as point estimates surrounded by their 95% confidence intervals. In Figure 1, we present all the primacy and latency coefficients from the 1998 election. The coefficients are sorted from smallest to largest; the vertical line is drawn to emphasize those coefficients whose 95% confidence intervals include zero.

Figure 1 Goes Here

First, note that the majority of the coefficients have confidence intervals below zero, and that many of the estimates have 95% confidence intervals which include the vertical line at zero, which implies that they are not statistically distinguishable from zero. These two points underscore visually the basic result that more often than not primacy and latency effects are either of the incorrect sign or are not significantly different from zero. In some cases, the primacy and latency coefficients are positive and significantly different from zero, as shown in Figure 1.

To give further credence to our claim that no ballot placement produces a systematic positive increase in vote share, we separate the coefficients into the primacy and latency coefficients and replicate the figure above. These coefficients can be seen in Figure 2 and Figure 3. For the primacy coefficients (Figure 2), again no consistent pattern of positive coefficients is seen, as the estimates are visually quite uniformly distributed around zero (roughly equal proportions of estimates are negative, are indistinguishable from zero, and

[27] See technical appendix for details. Furthermore, to ensure that there were no dramatic differences between the Democratic Party and the Republican Party with respect to ballot order we replicated our results using the Democratic Party to normalize the vote shares and found substantively identical results.

are positive). For the latency coefficients (Figure 3), we see some skew in the coefficient estimates toward the left of the graph — reflecting that more of the latency coefficients are negative than zero or positive. We also see a handful of cases where the estimated latency parameters are positive and statistically significant.

Figure 2 and Figure 3 Go Here

To gain some perspective on the magnitude of the marginal effects we selected the Lieutenant Governor's race as a typical example (as it had a mixture of different name order effects) and generated a counterfactual analysis from our SUR results. The counterfactual analysis begins by using the SUR results to predict the expected vote share that the candidate from each party would have received based on our model, with all of the independent variables set to their sample average values and the two dummy variables for candidate ballot position set to 0 (thus assuming the candidate is not first or last). We then change each of the dummy variables to 1, and recompute the vote shares for each party's candidate in the race.[28] The counterfactual estimates of the candidate ballot name order effect for the Lieutenant Governor's race are provided in Table 3.

[28] We produce this simplified counterfactual to isolate only the change from no one being first or last to one specific party being first, to emphasize the possible change for that party alone without any additional complications. Of course it is possible to develop a wide array of alternative counterfactual analyses that while possibly more realistic, would also be more complicated to estimate and to present in this paper. For example, it is possible that were we to compare the outcome with one party being always first to the average of all other parties we would observe a smaller change. Furthermore, it is possible that incorporating the fact that some party must be last would also change the effect. There is no reason to anticipate that other counterfactuals will depict substantive results that are qualitatively distinct from what we report below.

There, we first provide the "baseline vote share" predicted by our SUR model, then we give the estimated change and then vote share for a scenario in which the candidate was listed first on the ballot. Last we give the same estimates for the scenario in which the candidate was listed last on the ballot. The numbers in brackets below each estimate in Table 3 provide the upper and lower bounds on each estimate's 95% confidence interval. Boldfaced entries in the table denote counterfactuals based on statistically significant SUR estimates.

Table 3 Goes Here

There were seven candidates contesting this statewide election in 1998. There are two instances generating a statistically significant estimate that candidates for Lieutenant Governor received more votes when they were first on the ballot: the Democratic and Green party candidates for Lieutenant Governor. In both cases, though, the estimated magnitude of the effect is unlikely to determine the election outcome. The Democratic candidate's vote share would increase from 59.54% to 62.13% when he was first on the ballot (an increase of 2.59%), while the Green party candidate's vote share would increase from 2.71% to 3.21% when first on the ballot (an increase of 0.5%). But, no inference can be made from this analysis about whether one party's projected increase or decrease in vote share would be at the expense of any other particular party's vote share.[29] We estimate that the candidates of two parties, the AIP and Libertarians, receive

[29] We simulate the first differences using Clarify; please see the appendix for details.

a statistically lower vote share when they are first on the ballot. The estimated magnitude of the decreases due to being listed first on the ballot are under 0.55%.

The last two columns in Table 3 provide estimates of the name order effect of appearing last on the ballot. Here, we find two examples (Democratic and Green) of candidates estimated to have received a statistically significant increase in votes when they were last on the ballot, by an estimated margin of 2.21 percentage points for the Democratic party, and an estimated margin of 0.57 percentage points for the Green party.. In four cases, we find that the candidate's vote share would decrease (and the decrease was statistically significant) when listed last: AIP, Reform, Libertarian, and PAF party candidates all get fewer votes when listed last in this counterfactual example. In all cases, the decrease in candidate vote share as the result of being last on the ballot is slight.

In conclusion, our statistical study of the candidate ballot order effect has analyzed a wide variety of different electoral candidates and all of the candidates on the statewide ballot in the 1998 California general election. We employed a multivariate statistical model that controlled for a number of other important variables that are likely to influence the outcomes of these contests (ideology, partisanship, and demographic attributes of census tracts). We also utilized a statistical methodology that is more appropriate for this type of electoral data than have previous studies. We have found little systematic evidence that indicates that candidates are benefited by being listed first on the ballot. Rather, sometimes candidates appear to benefit by being first, other times being first actually decreases their vote shares. Sometimes candidates benefit by being last on the ballot, but sometimes they also do worse if they are last on the ballot.

We also demonstrated that, regardless of the direction of the ballot order effect, the impact of being first or last on the ballot is generally of small magnitude. In the example we presented above, the largest positive increase a candidate received for being first was 2.59% (with a 95% confidence interval ranging from 1.86 to 3.27). Thus, we find little reason to believe that once we control for partisanship, ideology, and demographic factors, ballot order effects (no matter their direction) are potentially large enough to influence anything but a very small fraction of races that are close. Third, we have seen no evidence in our presentation that the candidate order effects–whether being first or last, or whether positive or negative–seem more substantial in any type of California statewide contest, or for any party. The ballot order effects we estimate as significant in this analysis appear to be uniform across contests and parties. That is, our analysis does not indicate that ballot order effects are consistently present (of any direction, or for either being first or last), are more likely for minor party candidates nor are more likely in less salient races.

Future research should consider when the ballot order effect is likely to be most salient. The Koppel & Steen and Imai & Ho papers suggest a stronger ballot order effect in *primary* elections. Thus, primacy may be more salient in primary elections because (1) the partisan voting cue is absent; (2) voters are less likely to have formed opinions about candidates in primary elections; or (3) some other factors. Ballot order effects could also vary by region, education level, voting technology, and the general salience of the election. There might also be different effects in nonpartisan elections. Additional research is needed to predict when candidates will be affected by the ballot order effect.

6 Conclusions: Implications of the Social Science Research on Ballot Order Effect Lawsuits

Our analysis has demonstrated that the "ballot order effect" is not a single effect that may easily be characterized by candidate position, by contest or by party. Furthermore, our estimates indicate that when it exists, the ballot order effect is small. Incorporating those two facts together results in the positive conclusion that California's current rotation and randomization system is, for the most part, not distributing additional vote share to any particular candidate or party in any systematic way, though additional rotations might help eliminate any possible effect. It seems possible that even with a small effect, the additional votes procured by being in the "remainder"–being a candidate appears first or last an additional time–could give a candidate a boost in vote shares. However, additional printing costs would be incurred for additional rotations, and it is not clear that it is worth the expense.

These results differ from those found in the rest of the literature for three principal reasons. First and foremost, we examined whether or not a candidate did better by being listed first, last, or somewhere in-between. We found that sometimes it was the case that candidates actually lost votes by being located first or last. It seems likely that there is a complicated process which determines the focal point on the ballot for each voter and this could differ across elections, races, and ballots. Ballot layout and voting technology are also possible factors. Second, we used a statistical procedure which accounted for the possible relationship between vote shares and candidates within a census tract. Finally, we account for factors which may determine different vote shares (our control variables).

Our analysis leads naturally into a discussion of how courts should utilize these results their resolution of election disputes. The sections below explore settings in which courts may adjudicate disputes and offer our suggestions as to their resolution.

7 How Should Courts Handle Claims that Failure to Randomize and Rotate Violate Equal Protection rights?

There is no question that despite the muddle of existing social science literature, a legislative decision to randomize ballot order and perhaps to provide for rotation is proper. Randomization works like a lottery, allocating whatever ballot order advantage (or disadvantage) there may be in a nondiscriminatory manner. Rotation works to neutralize whatever ballot order advantage (or disadvantage) there may be by insuring that each candidate appears in the same ballot order position the same number of times in the race. Assuming the effect does not vary by census tract, randomization would eliminate any benefit (or cost) to a candidate of being listed first. Legislators might choose to adopt randomization and rotation based upon no more than a hunch about the ballot order effect – they need not justify their decisions to anyone but the voters.

But it is a different question when it comes to judicial intervention into the question of randomization or rotation. Under what circumstances should courts order such changes in a state or locality's ballot order laws?

At first glance, it might appear that there is little cost to a court decision ordering a change in these laws on grounds that failure to randomize and perhaps rotate denies

candidates (or voters) equal protection of the laws under the Fourteenth Amendment of the United States Constitution or under similar provisions of state laws.[30]

A court order to randomize and rotate may have significant costs, however. These costs fit into two categories. First, rotation (but not randomization) increases administrative costs of running elections. Ballots, sample ballots, and related materials must be printed in a variety of styles to track the proper ballot order as the order is rotated in each area. These costs are of course multiplied in jurisdictions that print ballots in multiple languages. Electronic voting may decrease the costs of preparing actual ballots, but the other costs remain the same. Second, both randomization and rotation increase the potential for voter confusion. Randomization may increase confusion, especially on a long ballot, because voters may have a hard time (compared, for example, to an alphabetical listing or a uniform party ordering) finding a preferred candidate's name on the ballot. Rotation compounds the problem. Thus, a candidate in the 135-person race to replace Governor Gray Davis on the 2003 California gubernatorial recall ballot could not campaign with a simple slogan, "Vote for Smith, Number 118 on your ballot."

Courts are not particularly well equipped to make a decision about whether to order randomization and rotation in the face of incomplete evidence on both the need for a change in existing ballot order law and a state's interest in the change. In *Sonneman v. State of Alaska*, for example, the state rejected an equal protection challenge to a decision of the Alaska legislature to end the practice of rotating candidates' names on the ballot.[31] The court held that the burden on the candidates was minor and the state had an important

[30] In Gould v. Grubb, *supra* note 1, for example, the California Supreme Court held that California law giving incumbents the first ballot position violated both the federal and state equal protection clauses.
[31] 969 P.2d 632 (Alaska 1998).

regulatory interest in saving approximately $64,000 per election cycle on the printing of ballots and in reducing voter confusion.

The *Sonneman* decision appears correct because the question depends upon a number of highly contested empirical judgments that are best left to the legislature. For a plaintiff to prevail, she should have to come forward with significant evidence that the ballot order effect is likely to change election outcomes. If the plaintiff can do so, the state should then have to produce real evidence of significant savings (such as monetary costs or elimination of voter confusion) that outweigh making such a change.

It may be that *Bush v. Gore*'s equal protection holding will cause more courts to examine ballot order claims using strict scrutiny.[32] In our view, this would be an error. In the absence of evidence that a voting procedure systematically works against one set of voters or candidates over others, courts should let states and localities continue creating their own rules for ballot order.

8 How Should Courts Handle Election Contests Based on a Failure to Follow State Law on Ballot Order?

Although some of the questions related to ex ante court invalidation of ballot order laws are close, questions related to judicially-mandated ex post changes in the outcome of elections based on evidence of the ballot order effect are not. Courts should not change election results based on a concern about ballot order errors. At most, courts should order new elections using the correct ballot order rules. We use the Compton example to illustrate our point here about the perils of overreliance on social science research in the courtroom.

---

[32] *See* Hasen, *supra* note 9, at 189.

In the Compton case, the trial court judge made an unprecedented ruling. The judge invalidated the results of the contested election and reinstated the losing incumbent (Bradley) as mayor of the City of Compton. The judge based her ruling exclusively on the testimony of one social scientist, whose testimony was based on the Miller and Krosnick research from three Ohio counties.

As we have argued in this article, the evidence provided in the original Miller and Krosnick paper provides inconclusive evidence on whether the primacy effect exists in the three Ohio counties they examined. But even if the Ohio evidence were more convincing, extrapolation to Compton, a jurisdiction with starkly different demographic and political attributes than the three counties in Ohio, would be unjustified.

We have produced our own analysis of statewide elections in California from a single election year. Our statistical analysis uses a methodology more appropriate for the type of data studied here, and we employ a number of politically important control variables, in particular partisanship, ideology, and precinct demographics. We find no clear or compelling evidence that primacy effects dominate, nor for that matter that latency effects dominate. Rather, we find that primacy and latency are not systematically present in California's statewide elections.

Of course, the specific Compton situation did not involve a statewide election. It involved a nonpartisan municipal runoff election, and a heated and highly salient one at that. Municipal elections in California only require initial randomization of candidate names on the ballot; rotation is not required, probably because rotation would be more expensive for already financially-strapped city and county clerks to implement for

municipal elections.[33] But, as rotation is not required for municipal elections in California, it is simply impossible to determine, using past election data, whether the lack of rotation influences California's municipal elections. Thus, as we studied races that were conducted with randomization and rotation, our analysis itself sheds uncertain light on the problem confronting a trial judge in a situation like the Compton case.

What should trial court judges do in such situations? First, they should bear in mind the fact that social science involves estimation. All social science estimates contain a degree of imprecision and uncertainty. Trial court judges must understand what estimation uncertainty means for the proper interpretation of social science research. For example, take our analysis of the Lieutenant Governor's race presented in Table 3 above. The largest estimate for primacy we found was an estimated increase of 2.59% for the Democratic candidate. The 95% confidence interval for that estimate ranges from a low of 1.86% to a high of 3.27%. This confidence interval tells a social scientist two things: first, that she can be more than 95% confident that this estimate is greater than zero (in the terminology of social science, that we have a "statistically significant" estimate); second, that she is 95% confident that the effect that would be expected for the Democratic candidate in this particular race ranges from 1.86% to 3.27%.

Second, attention to context is important. To explain this point requires a simple example. Let's say a Southern California city government manager wanted to produce estimates of the future manufacturing growth in her city, but she finds that manufacturing

[33] For example, one might imagine a situation where municipal elections in the City of Compton could begin with an initial randomization, and then the candidates could be rotated by precinct. This would produce additional expenses for municipal elections, as the city would have to print more sample ballots to match the specific rotations voters would see in the polling places. Rotation might also create substantial voter confusion.

growth data are only available for counties. She could select three counties from Ohio, produce a statistical analysis of manufacturing growth in those counties, and then use the results from that statistical analysis to estimate her city's future manufacturing base. On the other hand, she could select three counties in California, do a statistical analysis, and use the results to project her city's future manufacturing growth.

It is likely that using California data, not Ohio data, is more appropriate for the city manager's decision. Using data from a different state, in a different region of the country (where laws, taxes, and economic development may be vastly different) is likely to add substantial uncertainty to the application of statistical analysis from Ohio to California. This additional uncertainty is difficult, if not impossible, to account for when a researcher uses estimates from one context and applies them to a different context. If only the Ohio data are available, then they should be used — but they should be used with caution, and with the full awareness that there is more uncertainty about the estimates as they are applied to California than is readily apparent from typical measures of statistical uncertainty (such as standard errors). If data from a more similar setting are available, they ought to be employed

As an example, return to the analysis we presented in Table 3. Assume for pur poses of discussion that this Lieutenant Governor race most closely approximates the Compton municipal election context and give the plaintiffs the further benefit of the doubt by assuming that his situation is best approximated by the Democratic candidate results in Table 3, where we estimated that candidates listed first could receive an increase in their vote share of 2.59%. But instead of using the point estimate given in Table 3, a better

counterfactual is to examine the estimated increase in candidate vote share for precincts in the Compton area. This gives us at least the ability to control for the local context, in making our estimate of possible name order effects.

To approximate the Compton context, we produced an estimate of the candidate name order effect for the Democratic candidate in the Lieutenant Governor race, for a typical precinct in the Compton area.[34] This analysis produced a predicted baseline vote share of 77.33% (with a 95% confidence interval of 76.98% to 77.66%) for the Democratic candidate when not listed first or last. Our prediction is that the Democratic candidate would receive a vote share of 79.17% (with a 95% confidence interval of 78.75% to 79.79%) when listed first. We estimate an increase in the Democratic candidate's vote share, if listed first, of 1.86% (95% confidence interval of 1.38 to 2.36), relative to being listed in the middle.

Notice that this predicted effect of being first on the ballot, for a Democratic candidate in a context like that of the 1998 Lieutenant Governor's race in the Compton area, is 0.73% smaller than we estimated statewide. Thus, by taking the Compton area context into account, by acknowledging statistically that the Compton area is much more heavily Democratic and more liberal than the state as a whole, we estimate a substantially smaller name order effect. The Compton mayoral race was decided by 281 votes, and the losing candidate, who was listed last on the ballot throughout the citywide election, received 5,191 votes. Our estimates, based on this specific counterfactual analysis, are

[34] The profile of the typical precinct in the Compton area is 67.19% Democratic registration, 16.24% Republican registration, 11.49% decline-to-state registration, 5.19% other party; 54.71% female, .102 for our ideology measure, and 30.09% Hispanic. We reached this profile by estimating the averages for each variable for all of the precincts in this election that were found in Congressional District 37, State Senate District 28, and Assembly District 55 – the political districts within which Compton was located in 1998.

that were this candidate listed first on the ballot in the Compton area context, he could have received between 61 and 104 additional votes, but nowhere near enough additional votes to throw the election into uncertainty. Of course, it is worth pointing out again that we are assuming a situation quite favorable for the plaintiff in this case — a situation where a Democratic candidate in a context like Compton's actually could have received an increase in vote margin by being first on the ballot. We are also ignoring the fact that in this context, our analysis also estimated that a candidate listed last would have also received a marginal increase in vote share (estimated in Table 3 to be 2.21%, with a 95% confidence interval of 1.47% to 2.9%).

Clearly, social science research is becoming increasingly important in the resolution of election law litigation. Given the prominence of social science research in a number of recent election law cases, an understanding of social science research methodology is becoming more and more necessary for those involved in electoral litigation, and for those who analyze election law. There has been some debate about the use of social science methodology in legal research.[35] We welcome this debate, and hope that our research fuels further discussion about social science research and the law.

Address reprint requests to

Betsy Sinclair
e-mail: betsy@hss.caltech.edu

---

[35] Lee Epstein & Gary King, *The Rules of Inference* , 69 U. C HI. L. R EV. 1 (2002); Lee Epstein & Gary King, *A Reply*, 69 U. CHI. L. REV. 191 (2002); Frank Cross et al., *Above the Rules: A Response to Epstein and King*, 69 U. CHI. L. REV. 135 (2002); Jack Goldsmith & Adrain Vermeule, *Empirical Methodology and Legal Scholarship*, 69 U. CHI. L. REV. 153 (2002).

Table 1: Summary Results from SUR Analysis by 1998 Contest

	Primacy	Latency	Anti-Primacy	Anti-Latency	# of Candidate
Governor	2	0	2	1	7
U.S. Senate	1	2	2	2	7
Lt. Governor	2	3	2	3	7
Sec. of State	1	1	5	5	7
Controller	1	2	0	1	7
Treasurer	2	2	2	3	6
Attorney General	1	1	2	3	5
Insurance Commissioner	2	2	3	2	6
Total	12	13	18	20	52

Table 2: Number of Party Candidates By Name Order Effect, Summary Results from SUR Analysis

	Primacy	Latency	Anti-Primacy	Anti-Latency	# of Party Candidates Across Races
Democrat	4	5	2	2	8
AIP	3	2	3	3	8
Green	1	1	0	1	2
Liber.	1	1	5	5	8
Natural Law	0	0	3	3	6
PAF	3	3	3	4	8
Reform	0	1	2	2	4

Table 3: Lt. Gov. Counterfactual Analysis, Summarized from SUR Analysis of Contest

	Baseline Vote Share	Where Candidate is First		Where Candidate is Last	
Party		Estimated Change	Estimated Vote Share	Estimated Change	Estimated Vote Share
Democrats	59.54	**2.59** **[1.86, 3.27]**	**62.13** **[61.4, 62.81]**	**2.21** **[1.47, 2.9]**	**61.75** **[61.01, 62.44]**
AIP	1.40	**-0.43** **[-0.49,-0.37]**	**0.97** **[0.91, 1.03]**	**-0.06** **[-0.14, 0.01]**	**1.34** **[1.26, 1.41]**
Reform	1.26	-0.01 [-0.06, 0.06]	1.25 [1.2, 1.32]	**-0.28** **[0.33, -0.23]**	**0.98** **[0.93, 1.03]**
Lib.	2.24	**-0.12** **[-0.21, -0.03]**	**2.12** **[2.03, 2.21]**	**-0.22** **[-0.31, -0.13]**	**2.02** **[1.93, 2.11]**
Green	2.71	**0.50** **[0.34, 0.66]**	**3.21** **[3.05, 3.37]**	**0.57** **[0.42, 0.73]**	**3.28** **[3.13, 3.44]**
PAF	1.71	0.00 [-0.01, 0.01]	1.71 [1.7, 1.72]	**-0.06** **[-0.16, 0.04]**	**1.65** **[1.55, 1.75]**

Note: boldfaced entries are counterfactuals based on statistically significant SUR estimates.

Figure 1: All Coefficients

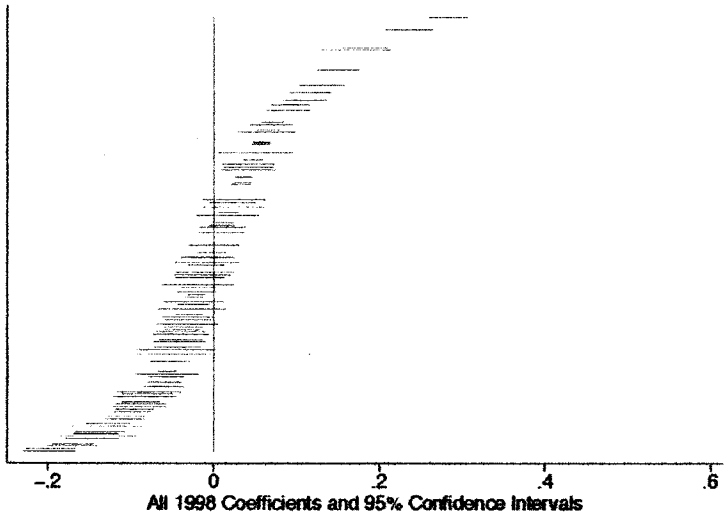

Figure 2: Primacy Coefficients

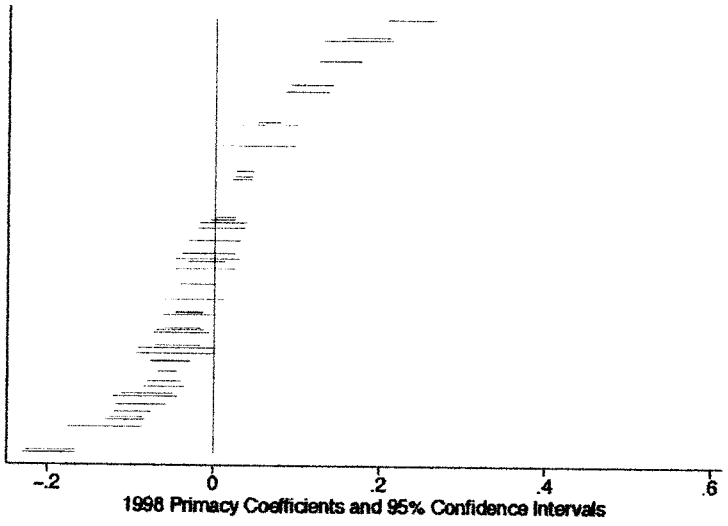

Figure 3: Latency Coefficients

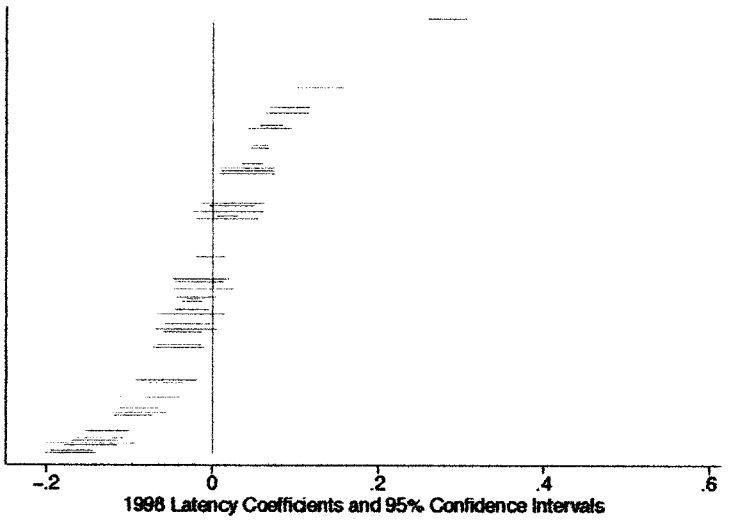

Section: Technical Appendix

This section discusses the technical details of the estimation procedure. It provides several definitions and details the nature of the statistical problem we face in estimating the ballot order effect. We explain why we are able to use Clarify, a software package designed for compositional data. Finally we explain how to interpret both the table of statistically significant coefficients and the counterfactual.

Compositional data is data in which each component is the proportion that falls into the ith category in a list of n categories where the sum of all proportions is 1. Election return data is compositional, since each party has some percentage share of the total vote, which implies that an increase in one party's vote share will likely affect another party's vote share. *Seemingly unrelated regression* is a tool that allows correlation between error-terms across J-1 regressions for J dependent variables. The correlation between the error structure introduces additional information into each regression. *Feasible generalized least squares* is a two-stage process which produces coefficients for seemingly unrelated regression by first estimating the correlation between regressions in the variance-covariance matrix and then weights by the variance-covariance matrix when estimating the coefficients.

In our data we observe, by contest, J election returns for J parties. We take a large party that appears in all contests (the Republican party) to use as the comparison category. We use the Republican Party for two reasons; first because it appears in each contest and second because we need a party with variance in its election returns (this variance will provide the necessary information for our estimates). For each party's vote share, we transform it by taking the log of that party divided by the Republican vote share. Using this log ratio as the dependent variable, we then run SUR using Clarify (which obtains estimates using FGLS). For J-1 parties (every party except the comparison category) this produces coefficient estimates on the ballot location (first and last) indicator variables. We sum the number of coefficients that are statistically significant and indicate whether or not they are positive or negative in the first table. These coefficients describe the impact of being first or last on the particular party of interest with respect to all of the other parties. Note that this is not simply with respect to

the comparison party because the coefficients were calculated from SUR (so that information is "borrowed" from the other party estimates via the error terms).

One disadvantage of summarizing the statistical significance of the coefficients is that the summary fails to describe the magnitude of the effects. In order to provide some insight into the impact ballot order has on election returns, we take the *largest* primacy coefficient we find (in the Lt. Governor's race) and produce a counterfactual. We do this to contradict claims in the literature that there is a large primacy effect. To produce the counterfactual, we fix each control variable at its mean and we fix each indicator variable for all parties as if all parties were last with the exception of the party of interest. We then look at the difference (again with Clarify) between vote share if the party of interest had been first on every single ballot in California and vote share if the party of interest had been last on every single ballot in California. Although we know this counterfactual does not correspond to how California actually rotates the ballots, we perform this calculation to imagine the effects this current rotation is preventing. The difference we observe is the change in vote share that particular party would have received – and each of the other parties would have lost some fraction of their vote share.