Public Opinion Quarterly, Vol. 72, No. 2, Summer 2008, pp. 216–240

ESTIMATING CAUSAL EFFECTS OF BALLOT ORDER FROM A RANDOMIZED NATURAL EXPERIMENT

THE CALIFORNIA ALPHABET LOTTERY, 1978–2002

DANIEL E. HO
KOSUKE IMAI

Abstract Randomized natural experiments provide social scientists with rare opportunities to draw credible causal inferences in real-world settings. We capitalize on such a unique experiment to examine how the name order of candidates on ballots affects election outcomes. Since 1975, California has randomized the ballot order for statewide offices with a complex alphabet lottery. Adapting statistical techniques to this lottery and addressing methodological problems of conventional approaches, our analysis of statewide elections from 1978 to 2002 reveals that, in general elections, ballot order significantly impacts only minor party candidates, with no detectable effects on major party candidates. These results contradict previous research, finding large effects in general elections for major party candidates. In primaries, however, we show that

DANIEL E. HO is with the Stanford Law School, 559 Nathan Abbott Way, Stanford CA 94305, USA. KOSUKE IMAI is with the Department of Politics, Princeton University, Princeton NJ 08544, USA. An earlier version of this article is available as Ho and Imai (2004). We thank the editor, four anonymous referees, Jim Alt, Ian Ayres, Larry Bartels, Barry Burden, Jamie Druckman, Marty Gilens, Brian Jacob, Dale Jorgenson, Gary King, Al Klevorick, Jon Krosnick, Hanna Lee, John Londregan, Becky Morton, Kevin Quinn, Donald Rubin, Jas Sekhon, Sarah Sled, Jim Stock, Liz Stuart, and Stephen F. Williams for helpful comments. Joe Falencki, Claudia Ornelas, and Tim Shapiro provided excellent research assistance. We are also grateful to Janice Atkinson at the Sonoma County Registrar of Voters, Gail Pellerin at the Santa Cruz County Registrar of Voters, Genevieve Troka at the California State Archives, and Karin MacDonald of the California Statewide Data Base at the University of California, Berkeley for their kind and resourceful help in collecting California's randomized alphabets, election returns, and registration data. Finally, we thank seminar participants at Harvard University, Princeton University, and Yale Law School for stimulating discussions. Address correspondence to Daniel E. Ho (e-mail: dho@law.stanford.edu) or Kosuke Imai (e-mail: kimai@princeton.edu).
Research support was provided in part by the National Science Foundation (SES–0550873), the Committee on Research in the Humanities and Social Sciences at Princeton University, the Institute for Quantitative Social Science, the Project on Justice, Welfare and Economics, and the Center for American Political Studies at Harvard University, and the Center for Law, Economics, and Organization at Yale Law School.

doi:10.1093/poq/nfn018 Advanced Access publication May 21, 2008

being listed first benefits everyone. Major party candidates generally gain one to three percentage points, while minor party candidates may double their vote shares. In all elections, the largest effects are for nonpartisan races, where candidates in first position gain three percentage points.

Introduction

For decades, scholars have attempted to assess the effects of ballot forms on elections—an effort that has intensified considerably since *Bush v. Gore*. Ballot reform has significant policy implications, with the Help America Vote Act of 2002 authorizing almost 4 billion dollars to reform efforts. One particular research agenda, spanning five decades and dozens of books and articles, examines the causal effect of name order on ballots. Scholars have worried that seemingly minor rules of election administration may have major unintended, or possibly intended, consequences on election outcomes.

In this article, we assess ballot order effects by analyzing a unique *randomized natural experiment* conducted in California statewide elections from 1978 to 2002. Since 1975, California elections law has mandated that the ballot order for statewide offices be physically randomized—after being "shaken vigorously," alphabet letters would be drawn from a lottery container to determine the order of candidates (Cal. Elec. Code § 13112(a) (2003)). This randomized alphabet determines the ballot order for the first district, which is then systematically rotated throughout the rest of the districts. This alphabet lottery offers a series of ideal randomized natural experiments allowing us to assess effects across candidates, parties, and elections in actual elections. Ho and Imai (2006) discuss the statistical issues that arise when the randomization of treatment is followed by systematic rotation, using California's 2003 recall election. This article extends that analysis to a much wider range of California statewide elections.

Several studies have claimed to find large and statistically significant effects for major candidates running for major offices. Krosnick, Miller, and Tichy (2003), for example, highlight the "most interesting finding" (p. 67) that being listed first in the 2000 presidential election in California statistically significantly increased Bush's vote shares by 9.5 percentage points compared to being last. Related studies similarly find that in Ohio, most major candidates for the US Senate and House (though not the Presidency) exhibited large and statistically significant ballot order effects in 1992 and 2000 general elections (Miller and Krosnick (1998, Tables 2 and 3); Krosnick, Miller, and Tichy (2003, Table 4.2)). These results have led some to conclude that "name order could affect the outcome of a close election – even in a major, highly salient race" (Krosnick, Miller, and Tichy 2003, 68).[1] Other studies find negligible ballot order effects (e.g., Gold 1952; Darcy 1986), with one study concluding that "there

1. See also Jon A. Krosnick, *In the Voting Booth, Bias Starts at the Top*, N.Y. Times (November 4, 2006) ("[E]ven in well-publicized major national races [for candidates such as Clinton in 1996 and Bush in 2000], being listed first can help.").

is no evidence that there is a ballot position advantage in general elections" (Bagley 1966, 649). More recently, Ho, and Imai (2006) find no detectable effects for major candidates in the highly publicized 2003 California gubernatorial recall. Given the conflicting findings, we directly assess here whether ballot order affects major candidates in general elections, contrasting estimates with primary results.

Part of the source of the disagreement in the extant literature may be methodological. Our analysis improves the previous approaches at least in four ways. First, while some scholars rely on observational data, where name order is not randomized and possibly confounded, others have used laboratory experiments that may lack external validity. Randomized natural experiments overcome these challenges, providing exceptional opportunities to draw credible inferences in real-world settings (Meyer 1995; Rosenzweig and Wolpin 2000). Second and most importantly, the unique feature of the California alphabet lottery is that only one randomization is performed for each election—ballot order is not randomized for each district. Thus, per-candidate analyses of ballot order effects may be confounded by observed and unobserved district characteristics. To overcome this problem, we identify robust patterns across a total of 473 candidates (in 80 races from 13 general elections and 8 primary elections) by examining a much larger data set than those analyzed previously. Third, our analysis employs a nonparametric approach which avoids conventional, but restrictive, parametric assumptions (e.g., constant additive effects and homoskedasticity) and directly accounts for California's nonclassical randomization. Ho and Imai (2006, sec. 4.3) show that under such nonclassical randomization, standard parametric analyses produce confidence intervals that are too narrow. Finally, we show that the exaggerated ballot effects for major candidates found in the previous literature stem in part from the problem of multiple testing.

When these methodological issues are appropriately addressed, estimated effects for major party candidates in general elections are negligible. In general elections, ballot order substantially impacts minor candidates, but has inconclusive effects on major candidates. In primaries, however, being listed first significantly increases vote shares for all candidates: major party candidates generally gain two percentage points of the total party vote, while minor party candidates may increase their vote shares by 50 percent of their baseline vote. In fact, primary effects are so substantial that ballot order might have changed the winner in as many as 12 percent of all primary races examined.[2]

We find the largest overall effect for nonpartisan races, where candidates in first position gain two percentage points on average. We observe no detectable differences in effects across types of offices for general elections, although effects appear to be somewhat larger for major offices in primaries. Our results are largely consistent with (1) a simple cognitive cost model of voting, where

2. This finding about primary races is consistent with results of Democratic primaries in New York City by Koppell and Steen (2004).

ballot order effects are due to cognitive costs of processing each candidate, and (2) partisan cue theory, where party labels convey information to uninformed voters (e.g., Schaffner and Streb 2002; Snyder and Ting 2002). In closer races and when party labels are not available, as in nonpartisan races, or not informative, as in party primaries, voter decisions are more likely to be influenced by ballot order.

Elections and Ballot Order

Social scientists have rediscovered the importance of ballot design since the days of counting chads in Florida (Niemi and Herrnson 2003). Recent studies have ranged from examining the causal effects of the butterfly ballot (Wand et al. 2001), forms of voting equipment (Tomz and Van Houweling 2003), absentee ballots (Imai and King 2004), partisan labels (Ansolabehere et al. 2003), and the ballot order of candidates. Current interest in ballot order is rooted in a half century of research investigating the causal effect of the order in which candidates appear on ballots (e.g., Gold 1952; Bain and Hecock 1957; Scott 1972; Darcy 1986; Darcy and McAllister 1990; Miller and Krosnick 1998; Krosnick, Miller, and Tichy 2003; Koppell and Steen 2004; Ho and Imai 2006). Research extends beyond the United States, with studies in Australia (MacKerras 1970), Britain (Bagley 1966), Spain (Lijphart and Pintor 1988), and Ireland (Robson and Walsh 1973).

Beyond the academic literature, practical implications abound. Dozens of US court decisions[3] and the drafting of electoral statutes in all 50 states[4] rely on a version of the claim that vote shares will accrue to a candidate solely for being listed first on the ballot. And electoral jurisdictions have proposed remedying ballot order effects by instituting some form of rotation or randomization. At the heart of these reform efforts is an assumption of ballot order effects.

We build on the theoretical propositions scholars have developed about ballot order effects and derive implications from a simple cognitive cost model of voting. Psychological theory offers a hypothesis of "primacy effects," whereby voters are satisfied by finding reasons to support rather than oppose a candidate (Miller and Krosnick 1998, 293–95). In contrast, scholars have proposed that candidates listed last should benefit from a "recency effect" (Bain and Hecock 1957), as these candidates are freshest in the minds of voters, or even that candidates toward the middle of the ballot should be advantaged (Bagley 1966). Alternatively, ballot order effects may exist because ballot order is *informative* in many states where major party candidates are listed earlier on the ballot.

We posit a simple decision-theoretic cognitive costs model of voting. Voters are assumed to be sincere and to maximize the benefit associated with

3. See, e.g., *Bradley v. Perrodin*, 106 Cal. App. 4th 1153 (Cal. Ct. App. 2003), *Gould v. Grubb*, 14 Cal. 3d 661 (Cal. 1975); *Mann v. Powell*, 333 F. Supp. 1261 (D. Ill. 1969).

4. See, e.g., Ohio Rev. Code Ann. Section 3505.03 (Anderson 2003); N.M. Stat. Ann. Section 1-10-8.1 (2003).

each candidate subject to costs of voting. Voters incur some nonzero cost to processing the information about each candidate in the order that they are printed. The result from such a simple model is that a voter will choose a candidate without reading the remainder of the ballot if the perceived marginal benefit of subsequent candidates, discounted by the probability of the pivotal vote, exceeds the cognitive cost of processing the merits of an additional candidate. Such a model can be considered a decomposition of the cost component of the canonical decision-theoretic voting model of Riker and Ordeshook (1968), and is related to behavioral formalizations of confirmatory bias (Rabin and Schrag 1999) and anchoring effects (Ariely, Loewenstein, and Prelec 2003).

This simple model also clarifies an observable implication of ballot order effects; cognitive costs are larger when less information exists about candidates in a race and when more candidates are running.[5] This suggests that ballot order effects are larger for elections with many candidates, for minor than for major candidates, for off-year than on-year elections, for lesser known offices, and for ballots containing less information such as partisan cues. This model excludes the possibility of recency and middle effects, since it assumes that there are positive marginal costs as voters read down the ballot. The model also excludes the possibility that the ballot position is informative, because ballot order effects are solely driven by the cost of processing ballot information.

The California Alphabet Lottery

In this section, we first describe the procedure of the California alphabet lottery as mandated by state election law. Second, we conduct statistical tests to show that the alphabets used for the elections in the past 20 years are indeed randomly ordered, a crucial identification assumption of our analysis.

LOTTERY PROCEDURE

California election ballots are printed in column-vertical format, depicting the name, party, and occupation of all candidates. Until 1975, California election law mandated that incumbents appear first on the ballot in the majority of statewide elections (Scott 1972, 365). In 1975, the California Supreme Court struck down the provision that reserved the first ballot position to incumbents, and held as unconstitutional, on equal protection grounds, ballot forms that present candidate names in alphabetical order (*Gould v. Grubb*, 14 Cal. 3d 661 (Cal. 1975)). The decision relied prominently on studies and testimonies by Bain and Hecock (1957) and Scott (1972). Scott (1972, 376) investigated the effect of ballot order using ballot rotations in 10 nonincumbent California races. While providing only point estimates of the ballot order effect, the study concluded that "one can attribute at least a five percentage point increase in the

5. We assume that the perceived differences in the probability of a pivotal vote across different races are negligible because the absolute magnitude of such probability is small (Gelman, King, and Boscarding 1998).

first listed candidate's vote total to a positional bias," a figure that has often been quoted by the Secretary of State since.

In response to that decision, the California legislature passed an alphabet randomization procedure to determine the ballot order of candidates.[6] The randomization applies to US Presidency and Senate races, as well as statewide races for Governor, Lieutenant Governor, Secretary of State, Controller, Treasurer, Attorney General, Insurance Commissioner, and Superintendent of Public Instruction. The law spells out in precise detail the procedure for drawing a "randomized alphabet":

> Each letter of the alphabet shall be written on a separate slip of paper, each of which shall be folded and inserted into a capsule. Each capsule shall be opaque and of uniform weight, color, size, shape, and texture. The capsules shall be placed in a container, which shall be shaken vigorously in order to mix the capsules thoroughly. The container then shall be opened and the capsules removed at random one at a time. As each is removed, it shall be opened and the letter on the slip of paper read aloud and written down. The resulting random order of letters constitutes the randomized alphabet, which is to be used in the same manner as the conventional alphabet in determining the order of all candidates in all elections. For example, if two candidates with the surnames Campbell and Carlson are running for the same office, their order on the ballot will depend on the order in which the letters M and R were drawn in the randomized alphabet drawing [Cal. Elec. Code § 13112(a) (2003)].

The container used in the drawing is in the same style as once used in one of the official state lotteries. The code further mandates that the drawing be open to public inspection and advance notice be given to the media, the representative of local election officials, and party chairmen (Cal. Elec. Code § 13112(c) (2003)). The explicit procedures defined in the law are designed to ensure accurate implementation of randomization. California election officials appear to have taken this duty seriously. The Secretary of State, in charge of the randomization, maintains two designated "random alpha persons" who draw the letters from a lottery bin. When asked about the process, officials insist that "it's the law" to randomize.[7]

Equally important to our estimation strategy, California elections law mandates that the randomized ballot order is rotated through the 80 assembly districts for all statewide candidates,

> the Secretary of State shall arrange the names of the candidates for the office in accordance with the randomized alphabet . . . for the First Assembly District. Thereafter, for each succeeding Assembly district, the name appearing first in the last preceding Assembly district shall be placed last, the order of the other names remaining unchanged [Cal. Elec. Code § 13111(c) (2003)].

6. The provision was added under Assembly Bill 1961, 1975–1976 Regular Session of the California Assembly, as Stats 1975, ch. 1211, Sections 16 and 17.

7. Telephone interview with Melissa Warren, Elections Officer at Office of Secretary of State, August 15, 2003.

The rotation is not implemented randomly, which, unlike previous analyses (but cf. Ho and Imai 2006), we take explicitly into account. The procedure nonetheless provides substantial variation of the ballot order, enabling the estimation of candidate-specific ballot order effects. Further, the ordering of Assembly Districts is *not random*, a property that we explicitly address in our analysis. The California Constitution provides that (a) districts be numbered from north to south, (b) the population be "reasonably equal" across districts, (c) all districts be contiguous, and (d) geographical subregions be respected to the extent possible (Cal. Const., art. XXI, § 1). Every 10 years following the census (here 1982, 1992, and 2002), districts are adjusted in state legislative reapportionment. The randomization-rotation procedure has remained virtually unchanged since 1975, allowing us to examine ballot order for a large number of elections from the past 25 years.

One concern about the California alphabet lottery is that the randomized alphabet may induce behavioral changes of candidates, making it difficult to isolate the direct effects of ballot order. For example, candidates listed last on the ballot in a particular assembly district might campaign more intensely in that district, in fear of ballot order effects. Or candidates might be chosen to assure a higher ballot order in favorable districts (Masterman 1964). However, such behavior seems unlikely given that the randomized alphabet is drawn very late in the game. All but write-in candidates must have declared candidacy and have been certified by the time that the drawing of a randomized alphabet takes place, and even sample (nonrandomized) ballots are printed before the drawing. Only minor adjustments, such as removal of a candidate from the ballot in the case of a death, occur after the drawing.

VERIFYING ALPHABET RANDOMIZATION

Given anecdotal evidence of manipulation of ballot order in other states (e.g., Darcy and McAllister 1990), we test for accurate implementation of randomization (see Imai 2005). Table 1 displays randomized alphabets for 23 California statewide elections since 1982. We conduct a rank test under the null hypothesis that the alphabet is completely randomized. We compare the relative positions of all possible pairs of letters by calculating the mean absolute rank differences of paired letters across elections, $\frac{1}{325}\sum_{i=1}^{26}\sum_{j\neq i}^{26}|\frac{1}{23}\sum_{k=1}^{23}\{R(L_{ik}) - R(L_{jk})\}|$, where $R(L_{ik})$ denotes a rank or position of the ith letter of the alphabet on the randomized list of the kth election. This statistic averages the relative positions of two distinct letters over 23 elections and all such possible pairs.

The resulting sample statistic for the data in table 1 is 2.07, representing the average absolute difference in the relative positions of all possible pairs of distinct letters. Under the null hypothesis of complete randomization, the distribution of this statistic can be calculated *exactly* by considering all possible lists of alphabet which are equally likely (Fisher 1935; Rosenbaum 2002; Ho and Imai 2006). Since there are 26! such lists for each election, we approximate

Table 1. Randomized Alphabets Used for the California Statewide Elections Since 1982

Year	Election	Randomized alphabet																									
1982	Primary	S	C	X	D	Q	G	W	R	V	Y	U	A	N	H	L	P	B	K	J	I	E	T	O	M	F	Z
	General	L	S	N	D	X	A	M	W	V	T	O	F	I	B	K	Y	U	P	E	Q	C	J	Z	H	R	G
1983	Consolidated	L	C	P	K	I	A	U	G	Z	O	N	B	X	D	W	H	E	M	F	V	R	S	T	Y	Q	J
1984	Primary	W	M	F	B	Q	Y	T	D	J	U	O	V	I	K	R	H	S	N	P	C	A	E	L	Z	G	X
	General	V	W	I	H	R	Q	G	J	O	M	T	S	Y	C	A	F	U	X	K	B	P	E	Z	N	D	L
1986	General	Q	N	H	U	B	J	E	G	M	V	L	W	X	C	K	O	F	D	Z	R	Y	I	T	S	P	A
1988	Primary	W	O	K	N	Q	A	V	T	H	J	F	Z	L	B	U	D	Y	M	I	R	G	C	E	S	X	P
	General	S	W	F	M	K	J	U	Y	A	T	V	G	O	N	Q	B	D	E	P	L	Z	C	I	X	R	H
1990	Primary	E	J	B	Y	Q	F	K	M	O	V	X	L	N	Z	C	W	A	P	R	D	G	T	H	I	S	U
	General	W	F	C	L	D	I	N	J	H	V	K	O	S	A	R	E	Q	B	T	M	Y	U	G	Z	X	P
1992	Primary	U	R	F	A	J	C	D	N	M	K	P	Z	Y	X	G	W	O	H	E	B	I	S	V	L	Q	T
	General	F	Y	U	A	J	S	B	Z	G	O	E	Q	R	L	I	M	H	V	N	T	P	D	K	X	C	W
1994	Primary	K	J	H	G	A	M	I	Q	U	N	C	Z	S	W	V	R	P	Y	B	L	O	T	D	F	E	X
	General	V	I	A	E	M	S	O	K	L	B	G	N	W	Y	D	P	U	F	Z	Q	J	X	C	R	H	T
1996	Primary	G	E	F	C	Y	P	D	B	Z	I	V	A	U	S	M	L	H	K	N	T	O	J	Q	R	X	W
	General	J	Y	E	P	A	U	S	Q	B	H	T	R	K	N	L	X	F	D	O	G	M	W	I	Z	C	V
1998	Primary	L	W	U	J	X	K	C	N	D	O	Q	A	P	T	Z	R	Y	F	E	V	B	H	G	I	M	S
	General	W	K	D	N	V	A	G	P	Y	C	Z	I	S	T	L	J	X	Q	O	F	H	R	B	U	M	E
2000	Primary	O	P	C	Y	I	H	X	Z	V	R	S	Q	E	K	L	G	D	W	J	U	T	M	B	F	A	N
	General	I	T	F	G	J	S	W	R	N	M	K	U	Y	L	D	C	Q	A	H	X	O	E	B	V	P	Z
2002	Primary	W	I	Z	C	O	M	A	Q	U	K	X	E	B	Y	N	P	T	R	L	V	S	J	H	D	F	G
	General	H	M	V	P	E	B	Q	U	G	N	D	K	X	Z	J	A	W	Y	C	O	S	F	I	T	R	L
2003	Recall	R	W	Q	O	J	M	V	A	H	B	S	G	Z	X	N	T	C	I	E	K	U	P	D	Y	F	L

this statistic by Monte Carlo simulation, drawing and calculating the test statistic for 10,000 lists of 23 randomized alphabets. The resulting two-tailed p-value (comparing the observed statistic with its randomization distribution) is 0.30; we cannot reject the null of complete randomization. Conducting similar tests for rank differences between even and odd letters, and letters in the top and bottom half of the alphabet, yields p-values of 0.54 and 0.60, respectively. There is little evidence that election officials in California have incorrectly randomized the ballot order.

Causal Effects of Ballot Order

With the aid of the California State Archives and the Statewide Data Base at UC Berkeley, we coded election returns data by 80 assembly districts for a total of 80 statewide races (44 primary races and 36 general races), going back to 1978. Table 2 lists all the races examined in this article. These include 13 general elections and eight primaries for 10 statewide offices, yielding a total of 473 candidates analyzed ($n = 37{,}840$). Using official randomized alphabets and ballots, we reconstructed the ballot order for each of these races in each district.

While our data provide a nearly ideal test of ballot effects, it is also limited in several ways. First, since California publicizes the randomization, voting behavior may differ from jurisdictions where voters are unaware of the assignment of ballot order. If California voters adapt to counteract randomization, that should bias our effects downwards. Nonetheless, even where voters are aware of randomized cues, such cues may still play an important role (cf. Tversky and Kahneman 1974; Ariely, Loewenstein, and Prelec 2003). Second, our data set consists of only statewide races, which may provide little information about the effects in smaller, local races. To the degree that cognitive costs are greater in local races, our estimates provide a lower bound. Lastly, our data set consists of a relatively small number of observed outcomes for each ballot position, as there are only 80 Assembly Districts.

Below, we describe our analysis of the California alphabet lottery and present results. We first place our analysis in a formal statistical framework of causal inference. Second, we describe our estimation strategy and interpret identification assumptions. Third, we present estimates and effects conditional on parties, offices, elections, number of candidates, and incumbency to test implications of our simple cognitive cost model. Finally, we compare the effects to the margins of victory to assess potential substantive impact if candidate names were ordered differently.

IDENTIFYING CAUSAL EFFECTS OF BALLOT ORDER

In the majority of experimental studies, researchers assign treatment to units that are randomly selected with equal probability. In contrast, the unique

Table 2. Number of Candidates Running in All Races Examined by Candidate Vote Share in 80 Assembly Districts ($n = 37{,}840$)

Election		President	Senate	Governor	Lt. Gov.	Atty Genl	Controller	Ins. Comm.	Sec. State	Treasurer	Supt Educ
1978	General	–		5							
1980	General	7	5	–	–	–	–	–	–	–	–
1982	General	–	5	5							
	Primary	–	19	20							
1984	General	5	–	–	–	–	–	–	–	–	–
1986	General	–	5	5							
	Primary	–	20	9							
1988	General	5	5	–	–	–	–	–	–	–	–
	Primary		6	–	–	–	–	–	–	–	–
1990	General	–	–	5							
	Primary	–	–	19							
1992	General	6	5, 5[a]	–	–		–	–	–	–	–
1994	General	–	6	5							
	Primary	–		12							
1996	General	8	–	–	–	–	–	–	–	–	–
1998	General	–	7	7	7	5	7	6	7	6	2
	Primary	–	13	17	13	10	7	8	8	9	5
2000	General	7	7	–	–	–	–	–	–	–	–
	Primary	23	15	–	–	–	–	–	–	–	–
2002	General	–	–	6	7	5	5	6	7	6	2
	Primary	–	–	11	8	6	10	11	13	7	4

NOTE.—"–" indicates that no election was held for that office in a particular year. Blank cells represent races where election returns data were not available by assembly districts. The number of candidates in this table differs slightly from total number of candidates analyzed because of several uncontested party primaries.

[a]There were two senatorial elections in 1992 both of which had five candidates running.

feature of the California alphabet lottery is that the randomization applies only to the first district and treatment for other districts is systematically determined thereafter by rotation. We call this procedure "systematically randomized treatment assignment." The name, *systematic*, stems from the fact that randomization-rotation directly resembles systematic sampling in sampling theory (e.g., Cochran 1977, ch. 8). We can therefore adapt well-known results from this literature to account for rotation.

Following the literature, we estimate candidate-specific effects. Suppose there are J candidates, and for the sake of simplicity, 80 is assumed to be divisible by J. Let Z_j denote the randomized variable representing the ballot order in the first district for candidate j. For reasons that will become apparent, we focus on the effect of being in the first position compared with the rest of the positions. We use T_{jk} to denote the indicator variable whether candidate j in district k is listed first. Under the systematically randomized treatment assignment, T_{jk} is a deterministic function of Z_j; formally $T_{jk} = \mathbf{1}[\{(Z_j + k - 2) \bmod J\} = 0]$, where $\mathbf{1}(\cdot)$ is the indicator function and $a \bmod b$ represents the remainder of the division of a by b. Note that only the ballot position in the first district, Z_j, not the ballot position in each district, i.e., T_{jk}, is randomized.

Our analysis is based on the widely used potential outcomes framework for causal inference (Holland 1986). Accordingly, $Y_{jk}(1)$ denotes the potential vote share for candidate j in district k when she is listed first. Similarly, $Y_{jk}(0)$ is the potential vote share when not listed first. Under this setting, we can identify the average ballot order effect of being at the first position (compared to the rest of the positions) for each candidate from the observed data with uniformly fewer assumptions than regression approaches commonly used in the literature. Specifically, the average treatment effect of being in the first position for candidate j, i.e., $\tau_j \equiv \frac{1}{80}\sum_{k=1}^{80}\{Y_{jk}(1) - Y_{jk}(0)\}$, can be estimated without bias. To see this more formally, define the observed vote share as $Y_{jk} \equiv T_{jk}Y_{jk}(1) + (1 - T_{jk})Y_{jk}(0)$. Then, our nonparametric estimator is given by $\hat{\tau} \equiv \frac{J}{80}\{\sum_{k=1}^{80} T_{jk}Y_{jk} - (1 - T_{jk})Y_{jk}/(J-1)\}$. Noting the fact that $E_{Z_j}[T_{jk}] = 1/J$, we have $E_{Z_j}(\hat{\tau}) = \tau$. Thus, $\hat{\tau}$ is an unbiased estimator of τ. Appendix A2 of Ho and Imai (2004) empirically verifies this result by examining the balance of observable covariates from Census and registration data.

Although an unbiased estimate of the average ballot effect is readily available, its variance is not. This is because systematically randomized assignment, unlike completely randomized assignment, involves only one randomization. To address this problem, Ho and Imai (2006) adopt randomization inference and show that ignoring rotation underestimates standard errors. Unfortunately, this approach only works for races with a large number of candidates. As an alternative solution, we thus apply an auxiliary variable approach from the systematic sampling literature (e.g., Zinger 1980; Wolter 1984), detailed in Appendix A1 (see the supplementary data online at http://pubopq.oxfordjournals.org/).

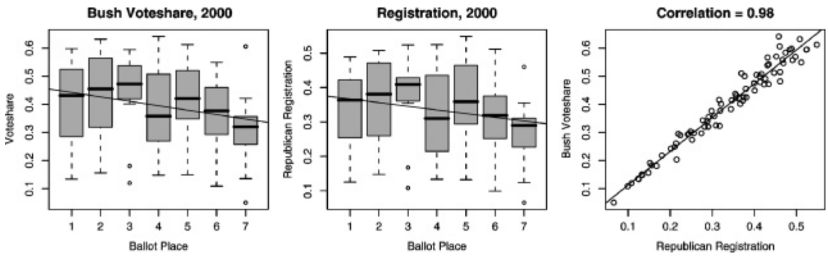

Figure 1. These Panels Show That Trend in Bush Vote Shares is Confounded with Republican Vote Shares.

HOW OUR APPROACH DIFFERS: ILLUSTRATION WITH 2000 PRESIDENTIAL ELECTION

To illustrate how our approach differs from extant approaches, we analyze the 2000 presidential election, previously examined by Krosnick, Miller, and Tichy (2003) (KMT). We compare our approach with KMT as it represents influential, state-of-the-art work, and is applied to a small subset of our data. KMT employs an approach proposed by Miller and Krosnick (1998), regressing vote shares on ballot order, and highlights as the "most interesting finding" a statistically significant effect of nine percentage points for Bush (p. 67). If true, the finding for Bush is daunting because "even in the highly-publicized and hotly-contested presidential race, name order mattered" (p. 52). KMT concludes that ballot order affects *both major and minor candidates in general elections*.

At the outset, we replicate KMT's results, shown in the second column of table 3.[8] The left panel of figure 1 displays boxplots of Bush's vote shares for each ballot place (with seven candidates). Interestingly, the vote share appears to decrease almost monotonically in ballot places, as illustrated by a fitted line from a linear regression.

At first blush, figure 1 provides strong evidence for large effects. But the second panel shows that *Republican registration rates in 2000 produce an almost identical boxplot* though registration rates, measured before the election, should not be affected by ballot order. Of course, Republican registration rates and vote shares, in turn, are highly correlated (0.98), as depicted by the third panel. Thus, the large ballot order effect for Bush appears entirely an artifact of partisanship (measured by registration rates).

This brings up the first crucial methodological point. With a single randomization for a single major candidate, ballot order can be highly confounded with observed and unobserved district characteristics. If the order were randomized

8. See left two columns of table 4.2 of KMT.

in *each district*, the correlation between ballot order and registration rates (and any other covariate) should be zero. But systematic randomization yields only *one* randomization. Combined with nonrandom district order, this can spell disaster for conventional approaches. To assess effects for major candidates in general elections, we require more data (i.e., more candidates, races, and hence randomizations), which we proceed to present below.

To further address the potential for confounding in any single randomization, we use the average gain of first place versus other positions instead of first versus last (see final two columns of table 3).This has the advantage of using all the data, thereby yielding more precise estimates, while also reducing the influence of a small number of confounded districts. For example, the estimated nine percentage point difference for Bush between first and last positions, reduces to roughly one percentage point using all districts. Similarly, for Gore comparing first to last yields larger negative effects than comparing first to the rest.

Second, conventional regression frameworks impose strong (and unnecessary) assumptions of constant ballot order effects (e.g., the difference between the first and second positions is the same as that between the fifth and sixth) and homoskedasticity. Table 3 shows that results differ considerably when using the appropriate standard error for reported point estimates. Using our nonparametric method, the statistical significance for Bush vanishes. Conversely, while KMT reports no significant effect for Browne, a minor candidate, the nonparametric method (using first versus rest) suggests distinguishable effects. Both point estimates and standard errors may differ between parametric and nonparametric methods. Linear regression suggests a point estimate of 2.19 for Gore, but this reverses sign with a -4.47 difference in means. Such sensitivity to parametric assumptions militates in favor of nonparametric methods.

Third, conventional approaches ignore multiple testing. KMT, for example, conducts separate tests for each candidate. Ignoring the multiplicity of hypothesis tests is prone to false discoveries (i.e., Type I error) beyond the level of the test. If test statistics are independent, for example, the probability of one false discovery with $\alpha = 10\%$ is 0.52 ($\approx 1 - 0.9^7$). Although accounting for multiple testing is mandatory in some contexts (e.g., medical journals and FDA studies), the problem has been largely ignored in the social sciences. We use standard methods developed by Benjamini and Hochberg (1995) to control the false discovery rate.[9] Asterisks of table 3 denote statistical significance accounting for multiple testing, showing that even under KMT's own parametric models, significant effects for Bush and Phillips vanish. With our approach, we find statistically significant results only for the two most minor candidates.

9. Benjamini and Yekutieli (2001) show that this procedure is valid even when test statistics have positive dependency.

Table 3. Reanalysis of the 2000 Presidential Election

	KMT (2003)	Parametric		Nonparametric			
	First versus Last	First versus Last		First versus Last		First versus Rest	
Candidates	Table 4.2	Linear	Quadratic	Random	Systematic	Random	Systematic
Gore	−4.47	2.19	2.20	−4.47		−3.62	
53.4%	(≥0.1)	(0.65)	(0.66)	(0.46)	(0.42)	(0.40)	(0.32)
Bush	**9.45**	**9.48**	**9.48**	9.40		0.97	
41.7%	(<0.1)	(0.06)	(0.06)	(0.14)	(0.12)	(0.85)	(0.83)
Nader	0.03	0.45	0.45	0.03		−0.44	
3.8%	(≥0.1)	(0.44)	(0.44)	(0.96)	(0.95)	(0.28)	(0.18)
Browne	0.09	0.04	0.04	**0.09**		**0.07**	
0.4%	(≥0.1)	(0.38)	(0.38)	(0.06)	(0.04)	(0.08)	(0.05)
Buchanan	0.06	0.02	0.02	0.06		0.00	
0.4%	(≥0.1)	(0.65)	(0.64)	(0.47)	(0.12)	(0.98)	(0.96)
Phillips	**0.11**	**0.05**	**0.05**	**0.11**		**0.08**	
0.1%	(<0.1)	(0.08)	(0.08)	(<0.01***)	(<0.01***)	(0.01**)	(<0.01**)
Hagelin	**0.06**	**0.05**	**0.05**	**0.06**		**0.06**	
0.1%	(<0.01)	(<0.01***)	(<0.01***)	(<0.01***)	(<0.01***)	(<0.01***)	(<0.01***)

NOTE.—For each candidate whose overall vote share is given below his name, the first row represents the estimated average treatment effects and the second row gives its p-values based on the normal test in parentheses. The figures in bold are statistically significant at the 90 percent level without taking into account multiple testing. The asterisks indicate statistical significance with respect to the false discovery rate of multiple testing, based on the procedure of Benjamini and Hochberg (1995): *, **, and *** indicate significance levels of 0.1, 0.05, and 0.01, respectively. The table shows that the point estimates, of Krosnick, Miller, and Tichy (2003, Table 4.2) are based on the nonparametric difference-in-means estimates, while statistical significance is determined by linear regression. For nonparametric methods, we compare the first and last positions as well as the first and the rest of the positions. For each of the nonparametric estimates, p-values based on the random list assumption ("random") and our proposed variance calculation ("systematic") are reported. The results indicate that there is little ballot effect among major candidates once parametric assumptions are relaxed and multiple testing is taken into consideration.

Last, KMT specifically is internally inconsistent in reporting results. While KMT's point estimates are apparently the difference in means between first and last positions (see fifth column of Table 3)—the nonparametric approach we recommend—statistical significance appears based on linear regressions (see third column).

In sum, next to the variance problem described in the previous subsection, previous analyses face distinct methodological challenges: (1) confounding due to rotation, (2) strong parametric assumptions, (3) multiple testing, and (4) internal inconsistency in reporting significance. When any one of these problems is addressed, detectable effects are limited to minor candidates. We now show that this is a robust pattern across all elections.

OVERALL RESULTS FROM 1978–2002

We now present results across a large set of elections. We report effects by party, office, and type of election. Although we investigated effects of other positions, we confine ourselves to the primary robust effect of first position. We start by presenting results for the 1998 and 2000 elections, and then summarize effects for all elections considering each race as a repeated experiment.

The top panels of figure 2 present estimates for the average gain (percentage points of the total vote) of all candidates in the 1998 and 2000 general elections, with major party and nonpartisan candidates in the left panel and minor party candidates in the right panel. Vertical bars indicate estimated 95 percent confidence intervals, using the minimum MSE variance estimator (see Appendix A1, available online as supplementary data at http://pubopq.oxfordjournals.org/). For 28 of 68 candidates, intervals are positive and do not intersect zero. Accounting for multiple testing, 27 of 28 candidates remain statistically significant. The median gain was roughly 0.2 percentage points. *All* statistically significantly positive effects (with multiple testing procedure) stem from minor party and nonpartisan candidates, as seen by the fact that confidence intervals for Democratic and Republican candidates in the top left panel largely overlap with zero. Third party candidates have a median gain score of 0.2 percentage points, compared to a median *loss* of 0.4 for major party candidates.

The bottom panel of figure 2 presents comparable estimates for 1998 and 2000 primaries. Effect magnitude (albeit measured as proportion of party vote share) is substantially larger than in general elections. For 74 out of the 128 candidates, confidence intervals are positive and do not include zero. Accounting for multiple testing, 72 of these 74 results remain significant at the 5 percent level. In primaries, ballot order affects major and minor party candidates alike, with a median ballot effect of roughly 1.6 percentage points, and a striking range of gains across candidates. This result is consistent with the analysis of New York City primary elections by Koppell and Steen (2004).

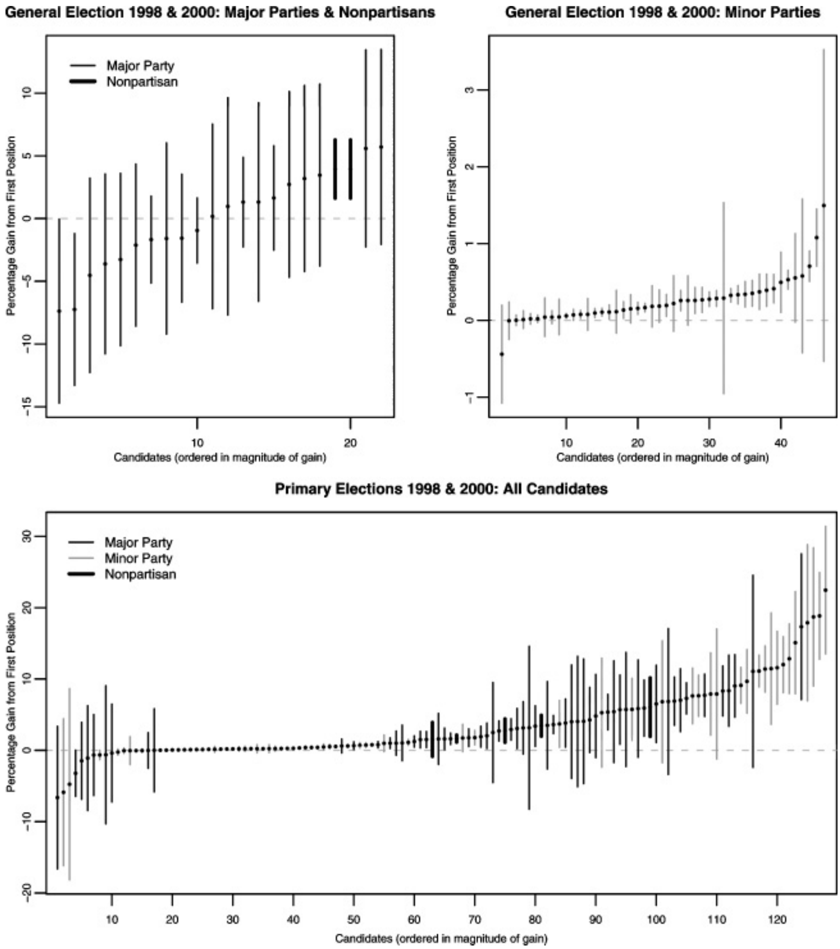

Figure 2. Candidate-Specific Average Gain Due to Being Listed in First Position on Ballots for 1998 and 2000 Elections. The top panels show results for general elections, and the bottom panel displays those for primary elections. Circles indicate point estimates for each candidate, and vertical bars represent estimated 95 percent confidence intervals. In general elections, only minor party and nonpartisan candidates are affected by the ballot order. In primary elections, all candidates are affected.

Table 4 summarizes effects across all races from 1978 to 1992.[10] The general patterns of the 1998 and 2000 elections hold across all elections studied. In

10. In cases where multiple candidates from the same party or multiple nonpartisan candidates ran, such as in primaries or nonpartisan elections, averages of those candidate-specific point estimates and standard errors are used to obtain an estimate for each race, and these estimates are in turn averaged across elections with the number of candidates in each race as weights.

Table 4. Party-Specific Average Causal Effects (Percentage Points) of Being Listed in First Position on Ballots Using All Races from 1978 to 2002 ($n = 37{,}840$)

	General		Primary	
	ATE	SE	ATE	SE
Democratic	0.05	0.46	1.89	0.32
Republican	−0.06	0.53	2.16	0.46
American independent	0.16	0.02	2.33	0.15
Green	0.56	0.17	3.15	1.16
Libertarian	0.23	0.02	6.59	1.42
Natural law	0.31	0.06	0.40	0.08
Peace and freedom	0.28	0.03	6.31	0.53
Reform	0.25	0.07	4.11	1.56
Nonpartisan	1.95	0.30	3.44	0.78

NOTE.—ATE and SE represent the average causal effects and their standard errors, respectively. Each candidate-specific effect is averaged over different races to obtain the overall average effect for each party. In general elections, only minor party and nonpartisan candidates are affected by the ballot order. In primaries, however, the candidates of all parties are affected. The largest effects are found for nonpartisan candidates.

general elections, major party candidates exhibit no discernible ballot order effect, while the effect on minor party candidates is substantial given that their initial vote shares are small. Minor party candidates typically gain roughly 0.2 to 0.6 percentage points.

Because cognitive costs are highest when races are close and when party labels are uninformative, ballot effects should be most pronounced for nonpartisan races, independent candidates, and primary races. These predictions bear out consistently. First, independent and nonpartisan candidates exhibit statistically significant gains even in general elections when listed first. When the office itself is nonpartisan, candidates gain roughly two percentage points in the general election. More information about candidates may be conveyed in races where at least some candidates are partisans (see also Miller and Krosnick 1998). That said, the only nonpartisan office in our data set is that of Superintendent of Education, so we cannot determine whether larger cognitive biases might stem from lack of partisan labels, lower prominence of the office, or both.

Second, in primaries, where the least information is conveyed by party affiliation and where cognitive costs are greatest, ballot order affects all candidates. Both Democrats and Republicans gain roughly one to two percentage points of the party vote when in first position. Since the number of candidates is generally much larger in primaries, with, for example, five Republican and six Democratic candidates running for the gubernatorial party nomination in 1998, this does not mean that the effect is confined to minor candidates in the

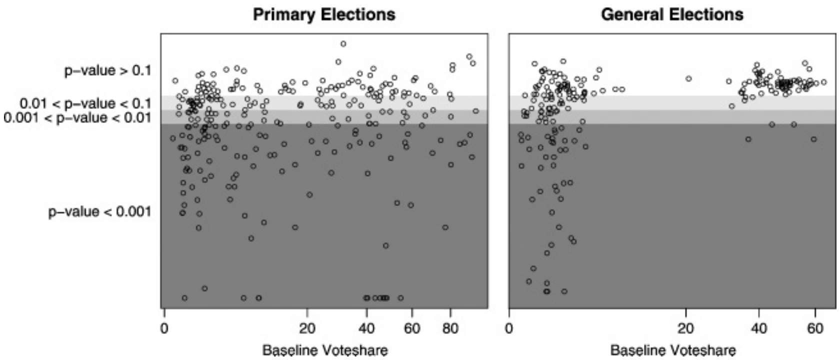

Figure 3. p-Values on y-Axis Against Vote Shares on x-Axis for All Candidates.
p-values and voteshares are transformed by logistic and square root transformations, respectively, for visualization. The figure is based on the data from all statewide elections listed in table 2.

major parties. On the contrary, many of the major Democratic and Republican candidates are affected by ballot order (e.g., Michael Huffington (1994), Barbara Boxer (1998), Dianne Feinstein (2000), Gary Mendoza (2002)). In the race for the Republican nomination for Lieutenant Governor in 1998, the average effect for Tim Leslie, who won the nomination by 10 percentage points, is 11 percentage points (SE = 6.8), and the effect on the runner-up, Richard Mountjoy, was 9 percentage points (SE = 2.2).

Minor party candidates in primaries receive average gains of several percentage points, with Libertarian and Reform party candidates exhibiting the largest relative gains. Nonpartisan candidates gain roughly two to six percentage points when listed first, which does not differ appreciably from nonpartisan gains in general elections or gains by other candidates in primaries. Given that partisan labels are relatively uninformative in primaries, where there are often multiple party candidates running, this result is consistent with our cognitive cost model.

To summarize the major distinctions between primaries and general elections and between major and minor candidates, figure 3 plots (logit transformed) p-values for all candidates against vote shares on the x-axis (square root transformed). This figure conclusively shows that for general elections, significant effects are limited to minor candidates, whereas effects exist across the board in primaries. These results contrast sharply with Miller and Krosnick (1998) and Krosnick, Miller, and Tichy (2003), which find large effects for major candidates for the US Presidency, Senate, and House.

Table 5 presents estimated average gains broken down by office and party, respectively. In both general and primary elections, no discernible patterns emerge with respect to the prominence of the office, or to the order in which

Table 5. Average Gain (Percentage Points) Due to Being Listed in First Position on Ballots using All Races from 1978 to 2002 ($n = 37{,}840$)

	General elections									
Party	President	Senate	Governor	Lt. Gov.	Atty Genl	Controller	Ins. Comm.	Sec. State	Treasurer	Supt. Educ.
Democrat	1.1	0.7	0.2	−1.1	−0.7	−1.9	0.2	−3.0	0.4	
	(1.0)	(0.7)	(1.0)	(3.0)	(1.4)	(2.0)	(1.5)	(2.8)	(1.7)	
Republican	−0.8	−0.6	1.5	2.2	−0.7	−5.0	1.5	2.6	−2.0	
	(1.2)	(0.9)	(1.1)	(2.7)	(1.6)	(2.1)	(2.3)	(3.0)	(2.3)	
Amer. Indep.	0.1	0.2	0.1	0.1	0.3	0.1	0.1	0.4	0.2	
	(0.0)	(0.0)	(0.0)	(0.1)	(0.1)	(0.1)	(0.0)	(0.1)	(0.0)	
Green	0.1	0.9	0.4	1.0	0.8	−0.5	0.2	1.3	0.5	
	(0.4)	(0.5)	(0.4)	(0.7)	(0.4)	(0.4)	(0.3)	(0.8)	(0.3)	
Libertarian	0.0	0.2	0.4	0.2	0.3	0.0	0.5	0.6	0.2	
	(0.0)	(0.0)	(0.0)	(0.1)	(0.0)	(0.1)	(0.1)	(0.1)	(0.1)	
Natural Law	0.0	0.1	0.1	0.1		0.2	0.6	0.7	0.5	
	(0.0)	(0.0)	(0.0)	(0.1)		(0.2)	(0.3)	(0.1)	(0.3)	
Peace & Frdm	0.1	0.4	0.2	1.1	0.1	0.3	0.2	0.5	0.0	
	(0.0)	(0.1)	(0.0)	(0.2)	(0.1)	(0.1)	(0.2)	(0.2)	(0.1)	
Reform	0.3	0.1		0.3		0.1		0.3		
	(0.3)	(0.0)		(0.0)		(0.1)		(0.1)		
Nonpartisan	0.4		0.1							4.0
	(0.4)		(0.4)							(0.6)

Continued

Table 5. Continued

	Primary elections									
Party	President	Senate	Governor	Lt. Gov.	Atty Genl	Controller	Ins. Comm.	Sec. State	Treasurer	Supt. Educ.
Democrat	1.6	1.5	0.6	5.6	4.6	3.3	3.6	2.4	7.1	
	(2.5)	(0.5)	(0.5)	(2.8)	(2.0)	(1.0)	(1.3)	(1.1)	(1.8)	
Republican	−0.9	2.8	0.6	5.4	4.8	2.1	3.2	2.8	3.1	
	(1.6)	(1.0)	(0.4)	(2.7)	(1.8)	(1.0)	(1.0)	(1.3)	(1.5)	
Amer. Indep.	0.0	0.2	8.6	0.4	0.4	0.2	0.0	0.8	0.1	
	(0.0)	(0.0)	(0.6)	(0.1)	(0.1)	(0.2)	(0.1)	(0.1)	(0.1)	
Green	0.9	4.6	−0.2	−0.6		6.2				
	(0.8)	(2.8)	(0.2)	(0.3)		(0.9)				
Libertarian	17.9	0.5	0.2	0.2	0.4	0.2	0.2	0.7	0.0	
	(4.0)	(0.1)	(0.1)	(0.3)	(0.1)	(0.1)	(0.3)	(0.2)	(0.2)	
Natural Law	0.1	0.2	0.1			0.1	0.5	1.1	1.0	
	(0.0)	(0.0)	(0.0)			(0.1)	(0.2)	(0.2)	(0.6)	
Peace & Frdm		3.1	8.2	11.5	9.8	0.1	8.2	5.4	0.2	
		(0.7)	(0.8)	(3.3)	(2.0)	(0.2)	(3.3)	(1.1)	(0.2)	
Reform	5.2	5.8		0.5		0.5		0.6		
	(3.3)	(1.6)		(0.2)		(0.1)		(0.1)		
Nonpartisan										3.4
										(0.8)

NOTE.—Standard errors are in parentheses. As in Table 4, all candidate-specific effects are averaged over different elections to obtain the overall average effect for each office and party. In general elections, no discernible patterns emerge with respect to the prominence of the office, or to the order in which the office appears on the ballot. In primary elections, ballot order effects are sometimes larger for major offices. In both cases, nonpartisan candidates for the Superintendent of Education are significantly affected by ballot order.

the office appears on the ballot. The only exception is the Superintendent of Education, which is a nonpartisan race. This suggests that cognitive costs are constant across offices.

Appendix A2 (see the supplementary data online at http://pubopq.oxfordjournals.org/) presents several other conditional effects to further test implications of our model. First, one might expect ballot order effects to be smaller in nonincumbent races, since incumbency may act as an informational cue to voters, and since the pivotal vote probability is larger in open races. Incumbents are denoted on California ballots, which provide current employment descriptions for all candidates. While we find few differences for incumbent and open races in general elections, in primaries open seat races appear to be associated with larger ballot order effects (see table 6). Second, we test the degree to which ballot order effects are driven by small uninformed groups of voters who turn out only for the prominent races. We do this by examining on-year versus off-year (or midterm) elections. Since contested offices differ in on-year and off-year elections with the exception of US Senate elections, we examine Senate results. Effects for on-year elections are generally larger (see table 7): Democratic candidates in on-year general elections gain roughly two percentage points when listed first, exhibiting no gains in off-year elections.

Third, we investigate the magnitude of ballot order effects conditional on the number of candidates. This should distinguish the cognitive cost model from behavioral models positing that the first position solves a coordination problem between voters (e.g., Forsythe et al. 1993; Mebane 2000). The cognitive cost model implies monotonically increasing ballot effects in the number of candidates (albeit offset by the increased likelihood of being a pivotal voter), while the latter provides an unclear prediction when the number of candidates is greater than two. We find that ballot order effects roughly increase monotonically in the number of candidates, lending further credence to the cognitive cost model (see table 8).

Lastly, our results suggest little evidence for recency or middle effects, thereby sharply rejecting such models. A simple cognitive cost model thereby appears to perform relatively well in explaining variation in ballot order effects.

MARGIN OF VICTORY AND BALLOT ORDER EFFECT

To assess potential substantive effects, figure 4 plots estimates for the second highest vote-getter of each race against the margin of victory (the difference in vote shares between the winner and the second highest vote-getter). Thick confidence intervals indicate that they include or exceed the margin of victory. Naturally, the substantive effect of ballot order on election outcomes depends on the closeness of contests. In general elections, as suggested by our previous results, we find no conclusive evidence of ballot order effects on major candidates. In contrast, ballot order effects were significantly positive and possibly greater than the margin of victory in 7 of 59 primary races. Ballot order might

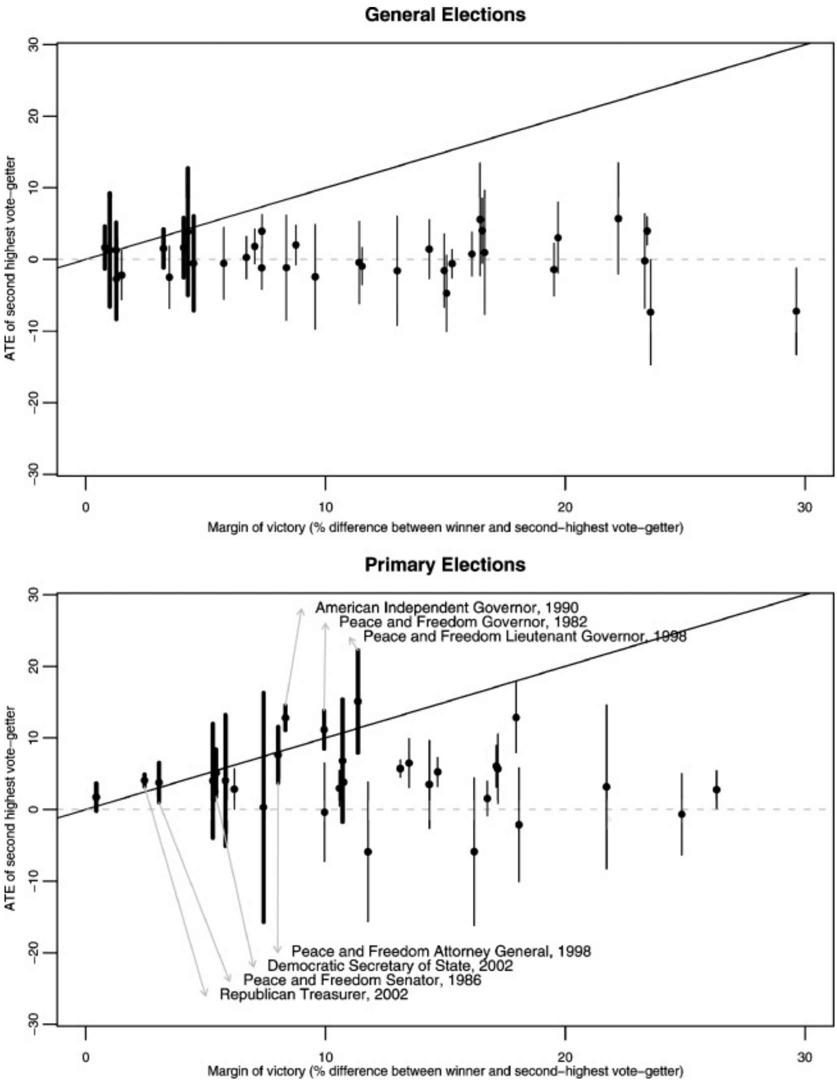

Figure 4. Comparison of Estimated Average Ballot Order Effect for Second-Highest Vote-Getter and Margins of Victory from 1978 to 2002.

The top panel shows general elections, and the bottom panel represents the primary elections. Dots indicate the point estimate for the (absolute) average ballot order effect, whereas vertical bars represent 95 percent confidence intervals. The 45° lines represent the instances where the ballot order effect equals the margin of victory. Thicker intervals indicate the races where the margin of victory is included in or below the 95 percent confidence interval. The figure implies that the outcomes of four primaries might have been different if the candidates were listed differently on ballots.

potentially have changed the winner of the Democratic primary for the office of Secretary of State in 2002 *if the second place candidate had been listed first*. This is not implausible, as many jurisdictions explicitly mandate that one candidate be listed first on all ballots.

Concluding Remarks

Our analysis of the California alphabet lottery from 1978 to 2002 places the study of ballot order effects on solid empirical ground and reconciles many of the findings in the field. In general elections, few effects exist for major candidates, contrasting with Miller and Krosnick (1998) and Krosnick, Miller and Tichy (2003). In primary elections, robust effects exist across the board (see also Koppell and Steen 2004). These results are largely consistent with a model of cognitive costs of voting, as we detect the largest effects when voters lack substantial information about candidates. Ballot order matters, though not as widely as believed by some, but widely enough to affect ultimate election outcomes in a large proportion of primaries.

Methodologically, our use of a randomized natural experiment avoids external validity problems of laboratory experiments and potential biases of observational studies. Free from financial, ethical, and other practical constraints of field experiments, randomized natural experiments provide a promising way to make causal inferences. While such experiments provide rare opportunities for research, they are not without limitations. Finely tuned statistical methods are required to adjust for nonclassical randomization.

Our results also have considerable implications for electoral administration, suggesting that arbitrary ballot format (determined by partisan administrators in many states) may be shaping outcomes of primary elections. Randomization can drastically reduce such biases—and methodology, in turn, may inform the fair and effective design of electoral administration.

Finally, although we analyze a wide range of offices in both general and primary elections, our inferences are limited to statewide races in California. Additional research is required to investigate whether our conclusions hold in other situations.

Supplementary Data

Supplementary data are available online at http://pubopq.oxfordjournals.org/.

References

Ansolabehere, Stephen, Shigeo Hirano, Jim Snyder, and Michiko Ueda. 2003. "Voting Cues and the Incumbency Advantage: Non-Partisan and Partisan Elections to the Minnesota State Legislature, 1950–1988." Technical Report.

Ariely, Dan, George Loewenstein, and Drazen Prelec. 2003. "Coherent Arbitrariness: Stable Demand Curves without Stable Preferences." *Quarterly Journal of Economics* 118:73–105.

Bagley, C. R. 1966. "Does Candidates' Position on the Ballot Paper Influence Voters' Choice? A Study of the 1959 and 1964 British General Elections." *Parliamentary Affairs* 19:162–74.

Bain, Henry M., and Donald S. Hecock. 1957. *Ballot Position and Voter's Choice*. Detroit, MI: Wayne State University.

Benjamini, Yoav, and Daniel Yekutieli. 2001. "The Control of the False Discovery Rate in Multiple Testing under Dependency." *Annals of Statistics* 29:1165–88.

Benjamini, Yoav, and Yosef, Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society, Series B* 57:289–300.

Cochran, William G. 1977. *Sampling Techniques*. 3rd ed. New York: John Wiley and Sons.

Darcy, R. 1986. "Position Effects with Party Column Ballots." *Western Political Quarterly* 39:648–62.

Darcy, R., Ian McAllister. 1990. "Ballot Position Effects." *Electoral Studies* 9:5–17.

Fisher, Ronald A. 1935. *The Design of Experiments*. London: Oliver and Boyd.

Forsythe, R., R. B. Myerson, T. A. Rietz, and R. J, Weber. 1993. "An Experiment on Coordination in Multi-Candidate Elections: The Importance of Polls and Election Histories." *Social Choice and Welfare* 10:223–47.

Gelman, Andrew, Gary King, and John W. Boscarding. 1998. "Estimating the Probability of Events that Have Never Occurred: When Is Your Vote Decisive." *Journal of the American Statistical Association* 93:1–9.

Gold, David. 1952. "A Note on the "Rationality" of Anthropologists in Voting for Officers." *American Sociological Review* 17:99–101.

Ho, Daniel E., and Kosuke Imai. 2004. The Impact of Partisan Electoral Regulation: Ballot Effects from the California Alphabet Lottery, 1978–2002. Princeton Law & Public Affairs Paper No. 04-001. Technical Report, available at SSRN http://ssrn.com/abstract=496863.

Ho, Daniel E., and Kosuke Imai. 2006. "Randomization Inference with Natural Experiments: An Analysis of Ballot Effects in the 2003 California Recall Election." *Journal of the American Statistical Association* 101:888–900.

Holland, Paul W. 1986. "Statistics and Causal Inference (with Discussion)." *Journal of the American Statistical Association* 81:945–60.

Imai, Kosuke. 2005. "Do Get-Out-The-Vote Calls Reduce Turnout?: The Importance of Statistical Methods for Field Experiments." *American Political Science Review* 99:283–300.

Imai, Kosuke, and Gary King. 2004. "Did Illegal Overseas Absentee Ballots Decide the 2000 U.S. Presidential Election." *Perspectives on Politics* 2:537–49.

Koppell, Jonathan G. S., and Jennifer A. Steen. 2004. "The Effects of Ballot Position on Election Outcomes." *Journal of Politics* 66:267–81.

Lijphart, Arend, and Rafael Lopez Pintor. 1988. "Alphabetic Bias in Partisan Elections: Patterns of Voting for the Spanish Senate, 1982 and 1986." *Electoral Studies* 7:225–31.

MacKerras, Malcolm H. 1970. "Preference Voting and the 'Donkey Vote'." *Politics: The Journal of the Australasian Political Studies Association* 5:69–76.

Madow, William G., and Lillian H, Madow. 1944. "On the Theory of Systematic Sampling, I." *Annals of Mathematical Statistics* 15:1–24.

Masterman, C. J. 1964. "The Effect of the Donkey Vote on the House of Representatives." *Australian Journal of Politics and History* 10:221–25.

Mebane, Walter R. Jr. 2000. "Coordination, Moderation and Institutional Balancing in American Presidential and House Elections." *American Political Science Review* 94:480–512.

Meyer, Bruce D. 1995. "Natural and Quasi-Experiments in Economics." *Journal of Business & Economic Statistics* 13:151–61.

Miller, Joanne M., and Jon A. Krosnick. 1998. "The Impact of Candidate Name Order on Election Outcomes." *Public Opinion Quarterly* 62:291–330.

Niemi, Richard G., and Paul S. Herrnson. 2003. "Beyond the Butterfly: The Complexity of U.S. Ballots." *Perspectives on Politics* 1:317–26.

Rabin, Matthew, and Joel L, Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics* 114:37–82.

Riker, William H., and Peter C. Ordeshook. 1968. "A Theory of the Calculus of Voting." *American Political Science Review* 62:25–42.

Robson, Christopher, and Brendan Walsh. 1973. "The Importance of Positional Voting Bias in the Irish General Election of 1973." *Political Studies* 22:191–203.

Rosenbaum, Paul R. 2002. "Covariance adjustment in randomized experiments and observational studies (with Discussion)." *Statistical Science* 17:286–327.

Rosenzweig, Mark R., Kenneth I. Wolpin. 2000. "Natural 'Natural Experiments' in Economics." *Journal of Economic Literature* 38:827–74.

Schaffner, Brian F., and Matthew J. Streb. 2002. "The Partisan Heuristic in Low-Information Elections." *Public Opinion Quarterly* 66:559–81.

Scott, James W. 1972. "California Ballot Position Statutes: An Unconstitutional Advantage for Incumbents." *Southern California Law Review* 45:365–95.

Snyder, James M., and Michael M. Ting. 2002. "An Informational Rationale for Political Parties." *American Journal of Political Science* 46:364–78.

Tomz, Michael, Robert P. Van Houweling. 2003. "How Does Voting Equipment Affect the Racial Gap in Voided Ballots." *American Journal of Political Science* 47:46–60.

Tversky, Amos, and Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185:1124–31.

Wand, Jonathan N., Kenneth W. Shotts, Jasjeet S. Sekhon, Walter R. Mebane Jr., Michael C. Herron, and Henry Brady. 2001. "The Butterfly Did It: The Aberrant Vote for Buchanan in Palm Beach County, Florida." *American Political Science Review* 95:793–810.

Wolter, Kirk M. 1984. "An Investigation of Some Estimators of Variance for Systematic Sampling." *Journal of the American Statistical Association* 79:781–90.

Zinger, A. 1980. "Variance Estimation in Partially Systematic Sampling." *Journal of the American Statistical Association* 75:206–11.