Do Audible Alerts Reduce Undervotes? Evidence From Illinois

Michael G. Miller

Assistant Professor

Department of Political Science

Institute for Legal, Legislative, and Policy Studies

University of Illinois, Springfield

mmill24@uis.edu

Abstract

In November of 2007, Illinois became the only state to require that voters casting optical scan ballots be alerted to undervotes via the emission of an audible beep and ballot "kickback" when they insert their ballot into the optical scanning machine. The 2010 Illinois General Primary was the first election for which the undervote alert was to be implemented. Yet, of 99 counties employing optical scanning technology, 12 did not upgrade their machines in time for the election, citing a shortage of time to do so. The uneven implementation of the Illinois undervote alert presents a unique opportunity to examine whether such an external stimulus is an effective deterrent of aggregate undervoting, within the context of a precinct-level panel study. I find little evidence that the audible alert affected undervote patterns in the 2010 General Primary election.

Acknowledgments: I thank Jennifer Carter for valuable research assistance, and two anonymous reviewers who have improved the paper immensely.

In 2007, the Illinois General Assembly mandated a significant change to election administration in that state. Specifically, Illinois required that by 2010 voters who cast optical-scan ballots were to be alerted whenever those ballots included undervotes in elections for the six statewide constitutional offices: Governor, Lieutenant Governor, Attorney General, Secretary of State, Treasurer, and Comptroller. In practice, this alert is delivered via an audible beep emitted when the voter inserts her ballot into the optical scanner. The beep is accompanied by a ballot "kickback:" the machine feeds the ballot back to the voter, who is told by an election judge that the ballot contains undervotes. Ostensibly, the policy is designed to reduce erroneous undervoting. While few people would likely disagree with the merit of this goal, the law was met with widespread apprehension from Illinois election officials on the grounds that it compromises ballot anonymity. Because the audible beep occurs in the relatively public setting of the scanner, as opposed to inside of the voting booth, critics of the alert have expressed concern that it may allow people nearby to discern a voter's choices.

Despite all the controversy, it remains to be seen whether the undervote alert has affected patterns of aggregate undervotes in Illinois. This question has both practical and academic implications. For one, election administrators should undertake a cost-benefit calculus with full information. Officials in Illinois and other states certainly have an interest in conducting elections fairly and with a minimum of voter error. However, in addition to raising questions surrounding ballot secrecy, the alert policy requires many machines to receive an upgrade, which can impart a significant financial cost. From the standpoint of creating regulations that are generally in the public interest, it is therefore worthwhile to gauge whether the state is receiving a return on its investment.

Determining the effect of the alert policy is also important in an academic context.

The composition of a democratic government should accurately reflect the intent of voters, but a significant percentage of voters typically cast ballots on which a preference is not registered for all contests. Particularly since 2000, political scientists have sought to determine the extent to which voting equipment (e.g., Shocket, Heighberger, and Brown 1992; Tomz and Van Houweling 2003; Sinclair and Alvarez 2004; Buchler, Jarvis, and McNulty 2004) or other institutional arrangements (e.g., Wattenberg 2002; Hayes and McKee 2009) determine ballot "roll-off" (the percentage of ballots with an unrecorded preference in a given race).

Most undervotes are likely intentional in high-profile elections such as presidential races (Knack and Kropf 2003). However, roll-off does increase in elections further down the ballot (Wattenberg, McAllister, and Salvanto 2000), and there is some evidence that an external stimulus can reduce instances of undervoting in elections conducted on direct-recording electronic (DRE) voting machines (e.g., Nichols and Strizek 1995). Yet, the alert mandated in Illinois does not occur while the voter is marking her ballot; as such, it is a different kind of stimulus. Optical-scan voting is widely utilized not only in Illinois, but also in much of the United States. Addressing the question of whether the Illinois policy achieved its goals is therefore a worthwhile endeavor and marks a contribution to political science scholarship.

During the 2010 Illinois General Primary Election, 99 of Illinois's 102 counties employed optical scanning technology exclusively. Of these, 87 counties implemented the alert in time for the election, while 12 did not. The irregular adoption of the alert policy facilitates strong conclusions about the effectiveness of the audible beep. Drawing on results from regression models that calculate panel difference-in-differences, I find little reason to conclude that the alert systematically reduced aggregate undervotes. This paper proceeds in five parts. First, I provide an overview of the policy, including its legislative history and immediate reaction from election officials. I then review previous academic studies of undervoting and ballot roll-off in an effort to determine whether there is a theoretical basis to expect that an alert like the one in Illinois should be expected to achieve its goal. Next, I detail the data sources and methodology employed to conduct the analysis, followed by a description of the results. I close with a brief consideration of the impact of these findings for both political science and policy.

1 Illinois Undervote Alert: Background

Media coverage of failing punch card machines and "hanging chads" following the 2000 presidential election raised public awareness of the importance of both defined voting standards and well-maintained equipment. In response, Congress passed the *Help America Vote Act (HAVA)* of 2002, the main purpose of which was to provide federal funds to the states in order to upgrade old voting systems. Yet the Act also established minimum administrative requirements that were intended to diminish the chances of subsequent events like the Florida recount. As such, one major focus of HAVA was the reduction of "residual" voting such as blank ballots, overvotes, and undervotes. Of these, over-voting was the main target for reduction, since it almost always represents a voter error. For instance, Section 301a of HAVA mandated that voting systems used for federal elections must notify voters who "select votes for more than one candidate for a single office." Moreover, HAVA requires that the systems communicate the consequences of over-voting and that voters be allowed a chance to correct the overvote. Accordingly, since 2006 alerting voters for overvotes has become a nearly ubiquitous practice in the United States.

In November of 2007, the Illinois General Assembly passed an omnibus bill that altered a number of election procedures in that state. Many of the changes, such as adjusting candidate certification procedures and the pay of poll workers, were relatively routine. However, another component of the bill extended the HAVA guideline for overvote alerts to undervotes as well, constituting a sharp departure from previous practice. Prior to 2007, the optical scan systems, which are the primary Election Day voting technology in 100 of 102 Illinois counties, alerted voters only when they submitted a ballot that either contained an overvote or one that registered zero votes for any office. The amended code (10 ILCS 5/24B-16(e)) mandates not only that optical scan ballot machines reject ballots with overvotes, but also that by 2010 they must also "identify when a voter *has not voted for all statewide constitutional offices"* (emphasis mine). In other words, rather than hearing an alarm only when they submitted a ballot on which they failed to vote *at all*, in 2010 Illinois voters began receiving alerts when they failed to vote in *any single* statewide race.

In practice, the procedure for alerting to undervotes on optical scan ballots is the emission of an audible beep when voters insert their ballot into the scanning machine, followed by a "kickback" reversal of the ballot out of the machine and back to the voter. At that point, an election judge will review the ballot and inform the voter that an undervote has been cast. Should they decide to correct their ballots, voters return to the voting booth and repeat the process. To my knowledge, this practice marks Illinois as the only state to require that voters be alerted to the presence of undervotes under the supervision of an election official while departing the polling place.

Given the political environment in Illinois, there is a range of possible rationales for the policy. For instance, it could be that the Assembly truly hoped to improve the democratic process by reducing undervoting, but a more cynical explanation is that the Democratic Assembly in 2007 was merely repaying Republicans for a law the latter had enacted a decade earlier. In early 1997, shortly before returning control of the

Assembly to Democrats after holding it for only two years, the outgoing Republican majority passed a bill (subsequently signed by Republican Governor Jim Edgar) that ended the previous practice of straight-ticket voting in Illinois. The bill was seen by many as Republican "revenge" for the straight-ticket Democratic voting in Cook County that had contributed to Republicans' 1996 election losses (Christian 1997). Presumably, by disallowing voters to register a preference for all offices with a single punch on a voting card, Republican legislators believed that they were increasing the odds that people who may have voted a straight Democratic ballot would either split their ticket (voting for Republicans in at least some elections) or simply abstain from voting in some cases. Since such an abstention provides a net benefit to Republicans, Democrats in 2007 could have seen the undervote alert as a mechanism to bring wayward Democratic voters back into the fold.

Regardless of legislative intent, the law was to take effect in time for the 2010 Illinois General Primary, which took place on February 2. The General Primary featured major-party primary elections for all six of the constitutional offices for which the 2007 omnibus law mandated undervote alerts. Accounting for the fact that both major parties held primaries, the undervote alert affected twelve separate elections; seven of those were contested by two or more candidates. Pre-election polling indicated that the gubernatorial races in both parties were too close to call. The primary campaign also received significant media attention due not only to its early date, but also to the concurrent primary in both parties for the United States Senate. The election for that seat drew national attention since it had formerly been occupied by President Barack Obama, and had played a central role in the impeachment, removal, and conviction of Quinn's predecessor, Rod Blagojevich. Relevant characteristics of each race are contained in Table 1.

Table 1: 2010 Primary Characteristics, By Office and Party

	Governor (Dem. Incumbent)	Lt. Governor (Open)	Attorney General (Dem. Incumbent)
Democrats:			
Candidates	2	6	1
Winner's Vote Share	50.5%	26.0%	100%
Mean Roll-Off	5.0%	17.7%	13.8%
Republicans:			
Candidates	7	6	1
Winner's Vote Share	20.3%	33.9%	100%
Overall Roll-Off	3.1%	11.4%	19.3%
	Secretary of State (Dem. Incumbent)	Comptroller (Open)	Treasurer (Open)
Democrats:			
Candidates	1	3	2
Winner's Vote Share	100%	46.7%	57.9%
Roll-Off	12.8%	18.0%	17.4%
Republicans:			
Candidates	1	3	1
Winner's Vote Share	100%	59.2%	100%
Roll-Off	22.5%	9.1%	17.4%

In the months preceding the election, the undervote alert component of the 2007 election bill provoked a sharp reaction from county election administrators in Illinois. The concerns ranged from practical considerations about implementation to how the alert would affect the voting experience. In November of 2009, Champaign County Clerk Mark Shelden brought suit against the Illinois State Board of Elections, seeking an injunction to stop the undervote alert provision from commencing in the 2010 General Primary Election. Shelden alleged that because the beep is audible to others nearby, the intention of a large number of voters to not vote in a given race would be revealed, thus compromising ballot secrecy.[1]

Shelden's request for an injunction was denied, but the concerns of front-line officials extended beyond both ballot secrecy and the borders of Champaign County. For instance, while Whiteside County adapted its Diebold Accuvote optical scan systems to comply with the law, its county board passed a resolution in March of 2010 (after the primary election) advocating the repeal of the undervote provision on the grounds that its "deleterious effects...include, but are not limited to: increased lines at the polls, voter frustration, wishes of voters intending to undervote to be challenged, and the loss of the right to a secret ballot to voters who intentionally undervote."[2] These concerns were echoed by a relatively large number of other counties in Illinois, many of which passed similar resolutions after the primary had taken place.

Some county officials went even further. For instance, McHenry County Clerk Katherine Schultz openly stated in advance of the 2010 primary election that she would not implement the undervote alert on her Accuvote systems.[3] For other coun-

[1] Retrieved October 10, 2010 from: www.champaigncountyclerk.com/elections/undervote/Shelden%20v%20SBE.pdf

[2] Whiteside County Board March meeting minutes. http://www.whiteside.org/37-county-board-meeting-minutes/view-category/Page-2.html

[3] Associated Press. "Clerks worry about Ill. law aimed at 'undervoting."' http://www.dailyherald.com/story/?id=355898&src=109

ties, implementation posed practical challenges, since the Diebold Accuvote scanners used in 62 counties required an upgrade in order to comply with the policy. Despite the law's passage in late 2007, the State Board of Elections did not approve the upgrade until November of 2009, less than three months prior to the early February primary. As such, eleven other counties cited insufficient time to upgrade their Accuvote systems, and therefore also failed to provide undervote alerts for the primary election.[4]

The uneven implementation of the undervote alert policy presents a unique opportunity to engage a crucial question: Does an undervote alert in the form of an audible beep and ballot kickback reduce undervotes on optical scan ballots? Such an inquiry is relevant not only to policymakers seeking clear and effective election regulations, but also to political scientists and other scholars who seek to understand determinants of voting. For policymakers, defining the existence and magnitude of a relationship is a first step toward a more complete cost-benefit analysis, especially considering the widespread objection of county election administrators and other public officials in Illinois on the grounds that an audible alert might compromise the anonymity of a ballot. The undervote alert law also affords a chance to test prevailing academic theories regarding voter roll-off. For political scientists, who have long studied means of increasing the probability that citizens' preferences are accurately expressed at the election polls, clarifying the relationship between audible alerts and undervotes will mark a contribution to the growing body of literature on how to reduce errant residual voting.

In the next section, I review existing academic studies of undervotes in an effort to determine whether there is a theoretical basis for a hypothesis in support of an effective undervote alert policy.

[4] All counties alerted voters to overvotes and blank ballots.

2 Should an Audible Alert Reduce Undervotes?

Compared to voting turnout, ballot roll-off is an understudied phenomenon in the realm of political behavior. Roll-off occurs when citizens fail to register votes (casting undervotes) in some races, usually after voting for others higher on the ballot. The classic rational choice voting models (e.g. Downs 1957; Riker and Ordeshook 1968) predict that even very small perceived benefits will outweigh the cost of voting when it approaches zero. It stands to reason that the marginal costs of voting in one additional race are extremely low if one is already in the voting booth; accordingly, citizens might be expected to register a preference in all contests. Yet, it is not uncommon for citizens to roll-off at rates of 20% or higher in municipal elections (Bullock and Dunn 1996), judicial contests (Hall 2007), or ballot initiatives (e.g. Kimball and Kropf 2008; Feig 2007; Vanderleeuw and Engstrom 1987). The fact that a large proportion of people expend the resources to go to their polling place and then fail to vote in some elections therefore presents a theoretical puzzle from the standpoint of rational choice political science.

One possible explanation for the roll-off phenomenon is that voters intentionally abstain from registering a preference in a contest that they view as unimportant; they may simply not see a clear benefit to voting in such races. For instance, it seems reasonable to suspect that many citizens see the governor as a more important state office than the treasurer. As such, some voters may be motivated to turn out by a perception that it is important to vote for the former, but then subsequently roll-off when it comes time to vote for the latter. Kimball and Kropf (2008) provide some evidence for the importance of salience in the context of ballot initiative questions, finding that roll-off is higher on initiatives that receive less newspaper coverage.

Another possibility is that the marginal costs of voting in down-ballot elections do not in fact approach zero. Specifically, there are still costs of obtaining information sufficient to decide which candidate to support. In relatively low-information settings when voters know little more than the names of the candidates, the odds that they will vote "incorrectly" increase substantially (Lau and Redlawsk 1997). Wattenberg, McAllister, and Salvanto (2000) find that voters approach ballots in much the same way as a test, abstaining when they lack sufficient knowledge to make a clear decision because they fear logging the "wrong" answer. These findings suggest that the uncertainty resultant from insufficient information about candidates in a given race is a significant cause of voter roll-off, and that the more voters know about a given election, the better able they will be to form opinions and vote with confidence.

Even basic personal information such as incumbency status, occupational background, or endorsements can be useful heuristics on which to base a decision (McDermott 2005; Lau and Redlawsk 2001; Popkin 1991). Perhaps both the most reliable and available piece of information is the partisan affiliation of a candidate, which provides a singular label from which a candidate's issue positions may be gleaned. Indeed, even in a low-visibility election, voters with strong partisan preferences may obtain all of the necessary information simply from the partisan affiliations listed on the ballot. Conversely, non-partisan or weakly partisan voters who face an uncertain decision in such an election are likely to conclude that the best action is to delegate the choice to more informed citizens (Feddersen and Pesendorfer 1996).

Illinois appears to have enacted its undervote alert under the assumption that a large percentage of voters were unwittingly casting incomplete ballots. However, in the context of a primary election–in which all candidates are members of the same party–the costs of acquiring information must be paid *before* the voter enters the booth. The discussion above therefore provides a basis from which to question the theoretical grounding of the state's assumption that a large proportion of undervotes are cast by mistake. To the contrary, the bulk of the evidence suggests that relatively few citizens undervote erroneously. For instance, Knack and Kropf (2003) find that people who fail to vote for a candidate in the presidential election usually do so intentionally. Moreover, previous experimental evaluation of ballots cast on optical scan equipment has shown that voters mistakenly undervote 0.2% of the time, compared to a rate of about 3% of voters who unintentionally cast ballots for the wrong candidate (Herrnson et al. 2006). Finally, while there is evidence that very complex ballot questions encourage roll-off (Reilly and Richey 2010), Bullock and Dunn (1996) find that ballot confusion is not as significant as contest salience and voter fatigue in determining roll-off in a Georgia county election.

Yet well-designed ballots (Kimball and Kropf 2005; Darcy and Schneider 1989) and properly administered elections (Stein *et al.* 2008; Ansolabehere and Stewart 2005) can both make a difference when it comes to reducing undervotes. Shocket, Heighberger, and Brown (1992) report the results of a controlled experiment demonstrating that punch card voting results in significantly higher rates of undervoting than paper and electronic balloting. Moreover, the problems associated with punch cards seem to exacerbate a gap in ballot spoilage rates between blacks and whites (see: Sinclair and Alvarez 2004), a problem that optical scan equipment (Buchler, Jarvis, and McNulty 2004) and electronic voting (Tomz and Van Houweling 2003) both appear to remedy. Nichols and Strizek (1995) show that electronic machines significantly reduced undervoting in Columbus municipal elections, speculating that a specific feature–lights that blink until a vote is cast in a given contest–serve as "a conspicuous reminder that portions of the voting task remain, and some persons may feel obligated to extinguish the lights by completing their ballots."

In short, previous analysis indicates that while much roll-off comes as the result of intentional undervoting, some proportion of undervotes might be reduced with greater attention to voting technology and administration. The results from the Nichols and Strizek (1995) study seem to bode well for the potential efficacy of an alarm designed to reduce unintentional undervoting, and the Illinois policy begs the question of whether an external stimulus might reduce residual voting on optical scanning equipment. Still, there is reason for skepticism. Compared to flashing lights that voters encounter *while* voting, Illinois mandates a different kind of stimulus. Illinois optical-scan voters encounter an audible undervote alarm *after* filling out their ballots, during the process of inserting them into the optical scanner. Since voters are literally halfway to the door of the polling place when they perform this action, the audible beep seems a weaker inducement to cast a vote; voters would have to retrieve their ballot, possibly wait in a short line for an available voting booth, complete the ballot, and re-insert it into the scanner. Given the relatively higher costs associated with the Illinois policy, its efficacy is worth studying.

In the sections below, I describe the data and analysis used to engage the question of whether the Illinois alarm provides a sufficiently compelling stimulant to reduce undervoting.

3 Data and Method

Each election authority in Illinois reports its precinct results to the Illinois State Board of Elections (SBE). The SBE also maintains information on county compliance with the undervote alert as well as the voting technology that each county utilized. There are 102 counties in Illinois; 99 use optical-scan technology to log the ballots of both early and Election Day voters. The audible alert policy pertained to optical scanning equipment, so these 99 counties comprise my initial sampling frame.[5] Aside from

[5]While Cook County utilizes optical-scan ballots for Election Day voting, it employs eSlate DRE machines for early voters in the city of Chicago as well as for many suburban Election Day voters.

vote totals, precinct returns generally include other relevant information, such as the number of registered voters and ballots cast.

Of the 99 optical-scanning counties, 87 implemented the undervote alert while 12 did not. This uneven implementation provides a unique counterfactual opportunity to study the effect of the policy, since voters casting residual votes will either receive the audible stimulus or not, depending solely upon the county in which they live. While the implementation of the treatment was likely not random (see below), for design purposes I follow the language of experiments in designating a "treatment" group of precincts located in counties that upgraded their optical scan machines and a "control" group of precincts in counties that did not implement the alert in time for the primary election. I assume that since no precincts using optical scan technology employed an undervote alert in 2006 (the last statewide "General Primary" election prior to 2010), the mean change in undervote levels in the control group from 2006 to 2010 approximates the mean of what would have occurred in the absence of the undervote alert in the treated counties. This assumption establishes a counterfactual baseline for comparison of undervote trends between the treatment and control groups, and provides the basis for a research design supporting causal inference.

To that end, I compare same-precinct difference-in-differences in aggregate voter roll-off for each major party's primary election. That is, I employ a panel design to compare differences in undervotes between the treatment and control groups through time between the primary elections of 2006, the last before the alert policy was enacted, and 2010, when the undervote alert was implemented in the treated counties. I compare undervotes in these two elections because they are both primary elections

Unfortunately, these votes are not separately tallied and there is no way to determine whether ballots were cast on paper or on machines. I therefore exclude Cook County from the precinct analysis. Kane and Peoria counties use DRE voting technology on Election Day. Since the statutory wording of the undervote alert pertains only to optical scan machines, I also exclude those counties.

during years in which all six statewide offices in Illinois were on the ballot.[6] Analyzing undervote patterns using panel difference-in-differences eliminates time-invariant (or relatively so) precinct-level covariates, such as precinct population, urbanity, or racial composition as potential confounders. Moreover, since the included elections are statewide and appear in the same order on the ballots of every precinct, differences in candidates and the level of contestedness of each race bear on all precincts in the same fashion, reducing the list of potential confounding variables only to those that varied by both precinct *and* time.

Three factors necessarily restrict the precincts included in the panel analysis, compared to those that might be included in a cross-sectional analysis of data from the 2010 election alone. First, the SBE possesses full statewide returns for 5,857 precincts and 5,798 precincts in the 2010 Democratic and Republican primaries, respectively. However, full records from the 2006 election are more difficult to obtain; a small number of counties did not submit election returns for one party or the other in 2006, or otherwise omitted information about ballots cast or registered voters. Second, I exclude all precincts from Sangamon County from the panel analysis, since that county alone altered its voting machines between 2006 and 2010. Third, I exclude precincts that changed boundaries, split into new precincts, or were created between 2006 and 2010, since such precincts cannot be analyzed as a panel.

After adjusting for boundary and equipment changes, and including in the analysis all remaining precincts for which enough data were available to examine roll-off in at least one party's primaries, 3,843 precincts remain for analysis; 3,127 of these are in the treated group.[7] The panel design therefore necessitates a trade; while fewer

[6] There was an election for United States Senate in 2010, but not in 2006. Both parties held contested primaries for the Senate race in 2010.

[7] Roughly 80 fewer precincts are available for panel analysis of Republican primaries than Democratic ones, due to more missing data for Republican contests.

precincts may be utilized, the framework supplies a counterfactual for comparison. As a robustness check, throughout the analysis I therefore also examine county-level summary statistics and panel data utilizing all precincts in Illinois.

One additional consideration remains. While the conception of the two groups of precincts in experimental terms provides a convenient framework, it is important to note that the adoption of the undervote alert was likely not random. Since the Accuvote upgrade became available only about three months prior to the election, it seems reasonable to suspect that counties in which implementation was a more complex endeavor were also less likely to implement the policy. For instance, more populous counties generally have more machines to upgrade, making the process more costly and time-consuming. With the goal of proper inference in mind, it is therefore prudent to compare the distribution of observed covariates between treatment and control groups, in order to determine whether the groups differed in a non-random fashion.[8]

Such covariate imbalance would pose potential problems for causal inference about the effect of the undervote alert since there would be a higher likelihood that the treatment variable is related in some way to other covariates (see: Ho *et al.* 2007a; Sekhon 2009). Table 2 indicates that there are some apparent differences between the treatment and control precincts, particularly with regard to population, racial composition, and partisan preference. Balance should therefore be improved before proceeding with the analysis. To that end, I pre-processed the dataset with a genetic matching algorithm intended to improve the balance of the observed covariates between treatment and control groups (On pre-processing, see Ho *et al.* 2007a; 2007b. On genetic matching and search algorithms, see: Sekhon 2008; Sekhon 2006; Dia-

[8]An examination of county-level census data confirms these suspicions: Control counties contain on average about 100,000 more people.

mond and Sekhon 2006; Mebane and Sekhon 1998; Sekhon and Mebane 1998). The algorithm seeks optimal balance via the construction of matched pairs consisting of one precinct that implemented the alert and one precinct that did not, and supplies weights that can be used in subsequent analysis.

I allow the matching algorithm to operate without restriction on eight theoretically relevant observable covariates measured before the outcome (roll-off) occurred on Election Day. These include measures of registered voters and ballots cast in each precinct in 2010, as well as roll-off in 2006 (as reported in election returns), the racial demographics of each precinct and its level of support for Barack Obama in the 2008 election (as reported in Ansolabehere and Rodden 2011), and whether the precinct is in the so-called suburban "collar counties" that comprise part of the Chicago media market.[9] I also match on the propensity score (the probability of receiving treatment given a set of observed covariates), as calculated with a logistic regression model, since its inclusion generally aids a genetic matching algorithm (see: Diamond and Sekhon 2006). I seek improved balance on all eight of these measures. Since the number of ballots cast and previous primary roll-off differs between the parties' primaries, I conduct two matches, obtaining separate weights for Democratic and Republican primary elections. All "treated" precincts were successfully matched.[10]

Table 2 contains a number of balance measures for all covariates, before and after conducting the matches. Columns 1 through 4 of Table 2 depict sample mean

[9] These are: DuPage, Kane, Lake, McHenry, and Will.

[10] Given that the control group is much smaller than the treatment group, I necessarily match with replacement. The matching algorithm attempts to pair as many "treated" units as possible with a control unit, pruning unmatched controls. For Republican primaries, all 3,044 treated precincts were successfully matched. In Democratic primaries, the algorithm matched all 3,108 precincts. The practical effect of this matching exercise is an oversampling of control counties that compared favorably with the treatment group, which provides subsequent sampling weights for the regression models described below.

Table 2: Mean Values of Observable Covariates and Balance Assessment, Precinct-Level Data

	Means: Unmatched Data		Means: Matched Data		Percent Improvement			
	Treated	Control	Treated	Control	Mean Diff	eQQ Med	eQQ Mean	eQQ Max
Democratic Primaries:								
Propensity Score	0.90	0.47	0.90	0.65	43.13	19.16	13.52	1.75
2010 Registered Voters	635.77	738.66	635.77	634.61	98.88	30.81	33.01	0.00
2010 Ballots Cast	52.69	60.70	52.69	48.75	50.94	33.33	10.87	-31.12
2006 Residual Vote Perc.	0.13	0.14	0.13	0.13	91.18	14.25	11.57	0.00
Perc. Obama Vote	0.50	0.55	0.50	0.50	95.88	26.06	29.31	43.54
Perc. African American	0.04	0.10	0.04	0.03	91.85	48.36	55.44	28.43
Perc. Hispanic	0.05	0.07	0.05	0.05	82.69	24.46	13.65	-14.66
"Collar County"	0.15	0.25	0.15	0.12	73.01	0.00	32.21	0.00
Republican Primaries:								
Propensity Score	0.89	0.48	0.89	0.63	36.08	16.33	12.33	2.52
2010 Registered Voters	636.52	747.47	636.52	644.64	92.68	41.58	44.46	-29.65
2010 Ballots Cast	86.62	73.30	86.62	77.88	34.39	-7.69	-27.44	0.00
2006 Residual Vote Perc.	0.18	0.17	0.18	0.18	81.34	12.80	13.09	50.00
Perc. Obama Vote	0.50	0.54	0.50	0.49	92.44	21.61	22.67	0.00
Perc. African American	0.04	0.08	0.04	0.04	91.99	30.65	32.05	36.84
Perc. Hispanic	0.05	0.07	0.05	0.05	36.00	38.69	5.68	-121.94
"Collar County"	0.15	0.25	0.15	0.15	99.35	0.00	41.34	0.00

differences between treatment and control groups, both before and after matching; entries in the four right-most columns reflect the percentage improvement in a range of differences between the groups. On the whole, the cell entries in Table 2 show that genetic matching results in a dataset with superior overall balance. The absolute value of the mean difference is smaller after matching for all covariates, and for most variables the difference is improved by 40% or more.

Matching should raise confidence that unobserved covariates do not drive any apparent results, and the panel design is beneficial in terms of controlling for time-invariant factors. However, uneven changes in certain other variables through time do still hold the potential to obscure the true effect of the undervote alert. I therefore calculate panel difference-in-differences with OLS regression models of same-precinct roll-off. Compared to basic means testing, the main advantage of this approach is that theoretically relevant covariates can be held constant, isolating the effect of the presence of the alert on roll-off differences. I seek the average treatment effect on the treated, or the effect of the audible undervote alert on undervote (or "roll-off") trends in precincts that adopted the alert policy.

In each major party's primary elections, voters may cast between 0 and 6 votes for constitutional offices–one in each of the six statewide offices. I analyze patterns in roll-off overall, since the policy objective is to reduce under-voting. The dependent variable in my analysis is a measure of "overall" roll-off in a given party's primary for a precinct P, which I express as:

$$R_P = 100\left(\frac{\sum U_{IP}}{6B_P}\right)$$

Where R is roll-off, $\sum U_{IP}$ is the total number of undervotes cast for each I of the six statewide races in precinct P, and B is the total number of ballots cast. Again, there are six possible elections in which a voter may cast a preference. Thus, R_p represents the proportion of possible votes in a given precinct that were recorded as undervotes.

The independent variable of interest for all models is a dichotomous treatment indicator coded 1 if the precinct employed the undervote alert and 0 otherwise; the coefficient for this variable represents the difference-in-differences estimate. If the audible alert was an effective reducer of roll-off, the models should return a negative, statistically significant coefficient on the treatment indicator. As noted, I expect that time-variant factors other than the undervote alert may also affect undervote trends between the treatment and control groups. For instance, roll-off should be lower in areas where candidates target their campaign efforts, since such efforts are likely to reduce information costs and/or raise the salience of the contest for voters. If some campaigns chose to focus their money and/or time on different geographic areas within Illinois in 2010 than in 2006, it seems reasonable to expect a correspondent effect on precinct roll-off.

Unfortunately, there is very little data available at either the precinct or county level that would allow theoretically relevant, time-variant campaign (or any other) effects to be held constant. However, the election returns for nearly all counties contain both the number of registered voters and the number of ballots cast, recorded at the precinct level. As such, I calculate precinct turnout for each party's primary election, measured as the percentage of registered voters who cast ballots for a given party. I assume that since turnout occurs before the act of voting, it can serve as an exogenous proxy for other, unobservable effects such as campaign activity. For instance, high-quality partisan mobilization efforts have been shown to positively affect turnout in the context of a gubernatorial election (Nickerson, Friedrichs, and

King 2006). I therefore include in the model a measure of fluctuation in turnout between 2006 and 2010. Assuming that turnout is driven mainly by the races of highest salience such as those for governor, as turnout rises, the proportion of voters interested in the contests further down the ballot should decrease. In short, if a higher percentage of voters cast ballots, then undervotes should rise as well.

Following these considerations, I calculate panel difference-in-differences with two OLS models–one for each major party primary. I fit two models to data for each major party's primaries. In the first, I present results from a model of unmatched data.[11] The second model is constructed from the matched dataset, applying weights derived from the genetic matching exercise described above. In all cases, the model specification is:

$$\Delta R_P = \delta + \beta_1 X_{1P} + \beta_2 \Delta X_{2P} + \epsilon$$

Where:
ΔR_P is the difference between overall roll-off between 2010 and 2006 in Precinct P
X_1 is a dummy variable coded 1 for the treatment condition (i.e., at the precinct utilized an undervote alert) and 0 for control
ΔX_{2P} is the difference in partisan turnout (the percentage of registered voters in Precinct P who cast ballots in a given primary) between 2010 and 2006

Thus, the β_1 coefficient in all models is the difference-in-differences estimator of same-precinct roll-off changes by treatment condition, holding precinct turnout constant. In the next section, I describe the results of my analysis of undervoting in the 2010 Illinois General Primary.

[11] One conventional method of analysis is a weighted least squares regression of roll-off percentage, where the weighting variable is the number of ballots cast (Vanderleeuw and Engstrom, 1987). The appropriate weighting strategy is less clear in a panel model like this one. However, WLS models applying the mean number of ballots cast in 2006 and 2010 return the same substantive results as the unweighted model of unmatched data described here.

4 Results

In this section, I first describe mean panel difference-in-differences of roll-off trends in Democratic and Republican primary elections for both the treatment and control groups, from 2006 to 2010. I then report the coefficients and standard errors from the regression models described in the previous section, which supply panel difference-in-differences estimates of precinct roll-off, holding turnout constant. I close with a brief consideration of county-level data to assess the robustness of the findings.

4.1 Summary Comparisons

A panel analysis of precinct data facilitates the comparison of trends in primary election roll-off from 2006 to 2010 in precincts for which boundaries remained unchanged between the two cycles. Before commencing with the analysis of the dependent variable (changes in roll-off between the two election cycles), I first examine precinct-level roll-off correlations in the control counties between the 2006 and 2010 primary elections for each party. If roll-off demonstrates substantial correlation in precincts for which there was no intervening treatment, then confidence in the difference-in-differences framework described above should be higher. To that end, Figure 1 depicts bivariate correlations between 2006 and 2010 roll-off in control precincts for primaries of each party. The scatterplots in Figure 1 suggest relatively strong correlations: Indeed, in both Democratic and Republican primaries, the Pearson coefficient is .34, which is statistically significant ($p = .0000$).[12] These correlations indicate consistency in precinct roll-off trends through time, which is necessary for the control group to provide a useful counterfactual.

[12] Without exception, the apparent outliers in Figure 2, with either very high or no roll-off in one or both years, come from precincts in which there are very few (i.e., fewer than five) ballots cast in one or both of the elections.

Figure 1: Mean Ballot Roll-Off Correlation in Illinois "Control" Counties, 2006 and 2010

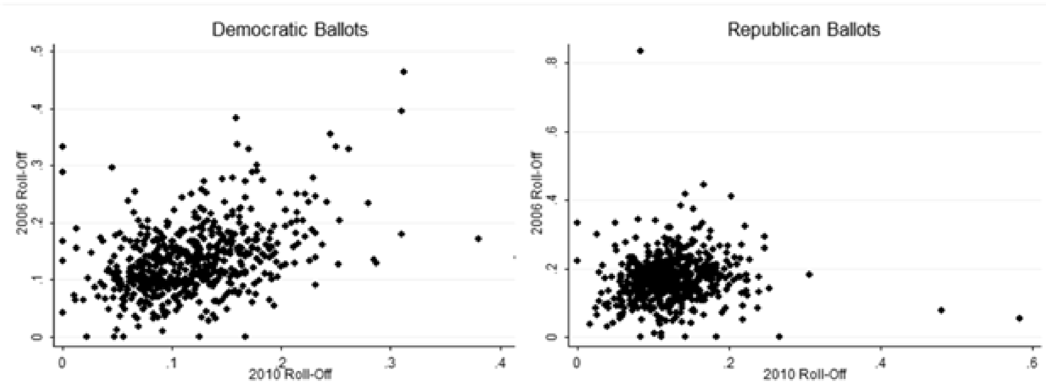

Turning now to summary comparisons with regard to the question of interest, Figure 2 depicts mean same-precinct roll-off differences in all counties for which data is available in both years, and for all precincts that had the same geographic borders in the two elections. Differences are calculated by subtracting 2006 roll-off from 2010 levels; negative values reflect lower roll-off in 2010. I first examine trends in Democratic primary races. Figure 2 suggests that contrary to what might be expected if the alert were effective, roll-off trends are *lower* in 2010 for the control precincts than the treated precincts in five of the six primary elections. In most cases the magnitude of these differences is about one half of one percentage point; however, mean roll-off change was lower in the control group (relative to the treatment group) by about 3 points in the Democratic Comptroller primary, which is a substantively large, statistically significant difference (p=.0000) if a one-tailed test is applied in the direction contrary to the hypothesis. The same is true in the Democratic contest for Secretary of State (p=.0002), although the difference in mean roll-off change is slightly smaller, at one percentage point. The only Democratic primary in which mean roll-off trended significantly lower for the treatment group was the one for Governor, in which treated precincts displayed roll-off that was about 1.5 percentage points lower than the control group (p=.0000).

Figure 2: Mean Same-Precinct Differences in Roll-Off Percentage, 2006 to 2010

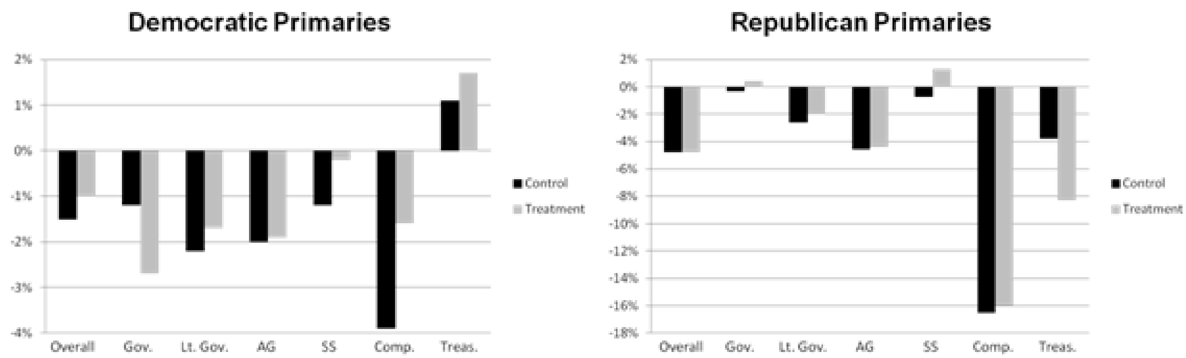

Similarly, Figure 2 shows that treated precincts displayed significantly lower mean roll-off change in only one Republican primary. However, unlike the Democratic races, the significant effect is apparent not in the race at the top of the ballot, but for Treasurer, which appeared last. Treated precincts in the election for Treasurer saw, on average, a mean drop in roll-off of 8.2 percentage points from 2006 to 2010; in control precincts, the difference is 4 points. As noted, the difference-in-differences of 4.2 points in the Republican Treasurer primaries is statistically significant (p=.0000). However, one-tailed tests applied in the opposite direction indicate that roll-off trends significantly *higher* in the treated groups for the Republican gubernatorial primaries (p=.0000), Lieutenant Governor primaries (p=.0487) and those for Secretary of State (p=.0000). The magnitude of the difference-in-differences for the gubernatorial primaries is substantively meaningless (about one-tenth of a percentage point) while in those for Lieutenant Governor and Secretary of State, control precincts saw roll-off drop by about eight tenths and four-tenths of a percentage point more than treatment counties, respectively. In all other Republican primary contests, the clear pattern is that treated precincts apparently exhibited–at best–about the same change in roll-off as control precincts. Overall, changes in Republican primary roll-off between the two election years is almost exactly the same between the treatment and control groups: -4.8 and -5 points, respectively. To sum, Figure 1 shows that the panel design is likely to supply good comparisons of precinct-level roll-off from 2006 to 2010, yet Figure 2 fails to find much of an effect on roll-off trends in optical scanning precincts. I now turn to the results of the controlled analysis described in the previous section.

4.2 Model Results

Table 3 contains coefficients and clustered (by county) standard errors from the OLS regression models used to calculate the difference-in-differences estimators for Democratic and Republican primaries, holding turnout fluctuations constant.[13] I report results from two models for each party's primaries; one is a model of data from the original (unmatched) precinct-level dataset, while the other applies weights obtained from the genetic matching exercise described above.

I first consider the control variable for turnout changes. As expected, all four models return a statistically significant coefficient for that variable; in all cases, the coefficient is also positively signed, indicating that roll-off increases as turnout rises. The effect of a change in turnout appears to be slightly stronger in Democratic primaries than in Republican ones; in the former, a one-point change in turnout is associated with about a one-third point increase in roll-off, while in the latter, the

[13] Clustered robust standard errors do not change the sign or significance of any of the regression coefficients, but do potentially adjust for any non-random error variance.

Table 3: OLS Regression Coefficients: Panel Difference-in-Differences in Roll-Off, Optical Scan Precincts

	Democratic Primaries		Republican Primaries	
	Unmatched	Matched	Unmatched	Matched
Treatment Indicator	0.288	-0.467	-0.047	-0.445
	(0.686)	(.717)	(0.561)	(.747)
Change in Party Turnout	0.324*	0.325*	0.179*	0.187*
	(0.114)	(.119)	(0.058)	(.061)
Constant	-1.265*	-0.510	-4.967*	-4.58*
	(0.475)	(0.510)	(0.420)	(.666)
N	3,802	3,665	3,721	3,596
R-squared	0.061	.059	0.018	.021
RMSE	7.256	7.31	7.522	7.346

Dependent variable is the difference in total roll-off in statewide races from 2006 to 2010.
All percentages measured on a 0-100 scale.
Standard errors are clustered by county.

coefficient for the turnout variable is slightly less than one-fifth of a point. As noted above, the positive relationship between turnout and roll-off is expected, since as the number of voters rises, the mean level of information they possess is likely to decrease. Regardless, the consistency of this result across models lends confidence to the notion of turnout serving as a proxy for unobserved time-variant effects.

I now turn to the relationship of interest, reflected in the coefficient for the treatment dummy variable. If the alert reduces undervoting, then the coefficient for the dummy variable should be negative and significant. However, the results do not present a strong case for an effective undervote alert. The treatment coefficient is negatively signed in both models of data from Republican elections, as well as in the Democratic model of matched data, but is positive in the other Democratic model.

Moreover, the coefficient for the treatment indicator does not achieve–or approach–statistical significance at conventional levels in any of the four models. In short, it is not possible to conclude with any confidence that the alert achieved its objective of reducing the percentage of undervoted races in either Republican or Democratic primary elections. To the contrary, in tandem with the patterns evident in Figure 2, the models presented in Table 3 provide little evidence that the audible alert in Illinois is an effective mechanism for reducing aggregate roll-off in the 2010 General Primary.

4.3 Robustness Check: County-Level Panel Analysis

Having found little effect of the undervote alert using precinct-level panel data in Illinois primary elections, in this section I briefly consider county-level data as a robustness check. Since precincts that split or were created between 2006 and 2010 are not included in the panel analysis above, the examination of county-level data affords the advantage of utilizing all available election data. However, the small size of the control group (12 counties) makes statistical inference difficult. Nonetheless, if the county-level analysis substantively agrees with the findings from the precinct-level one, confidence in the findings from the latter should be higher. To that end, in this section I analyze data both from all counties in the state as well as a subset of only those that only employed optical scanning technology subject to the undervote alert policy.

Figure 3: Mean Ballot Roll-Off in Illinois Counties

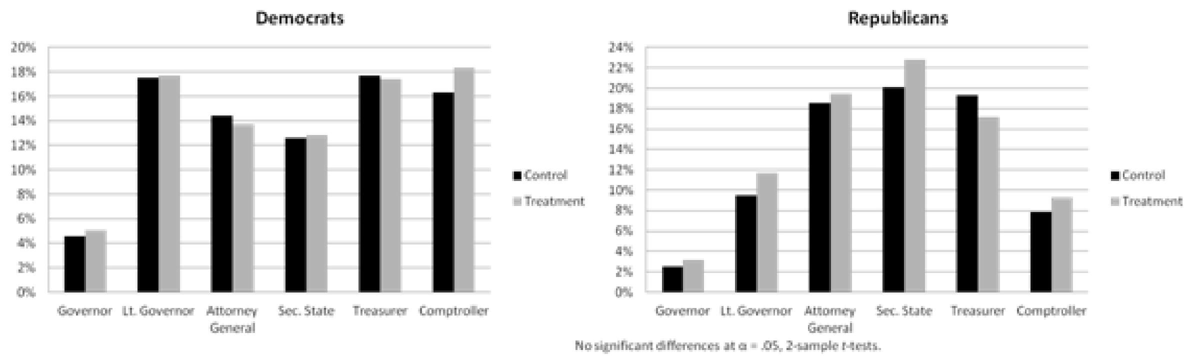

I begin with some basic exploration, comparing 2010 roll-off for each party's primaries at the county level. Figure 3 depicts mean county-level roll-off for both party primaries, by office. Simply, there is no discernible pattern in Figure 3 in support of the undervote alert reducing county-level roll-off. In the 2010 Democratic primaries, mean roll-off levels in the "treated" precincts–those that implemented the alert–and the control precincts are within about one-half of a percentage point of each other in all races, except that for Comptroller, in which the treated group displayed roll-off roughly two points *higher*. In the Republican primaries, roll-off trends higher in the treated group (contrary to expectations if the undervote alert effectively reduces roll-off) for all contests except that for the State Treasurer, in which treated counties displayed roll-off roughly two points lower. However, in no cases are the differences in roll-off for any of the elections depicted in Figure 3 significantly different at conventional levels, using 2-sample t-tests. Thus, Figure 3 yields little evidence on its own that the undervote alert reduced roll-off in the 2010 primary.

An examination of county-level differences in overall roll-off (the number of votes cast on a ballot divided by the number of allowable votes) from 2006 to 2010 yields a split pattern. Figure 4 depicts the mean within-county roll-off differential between the two years, both in all counties and only in those employing optical-scan technology.[14] There are three trends worth mentioning. First, while roll-off declined in both Democratic and Republican primaries, it did so on average by about twice as much in the latter contests. Second, in Democratic primaries, mean roll-off declined more in the control counties than it did in the treated counties. Third, while the "all counties" segmentation of Republican primaries suggests that roll-off declined by about one percentage point more in treated counties, the difference in counties employing optical scanning is roughly half of that amount, and none of the differences apparent in Figure 4 achieve statistical significance at conventional levels. Thus, while there is a suggestive pattern for Republican primaries, Figure 4 on its own does not raise confidence of an effective alert in Democratic primary elections.

Figure 4: County-Level Roll-Off Differences, 2006 to 2010

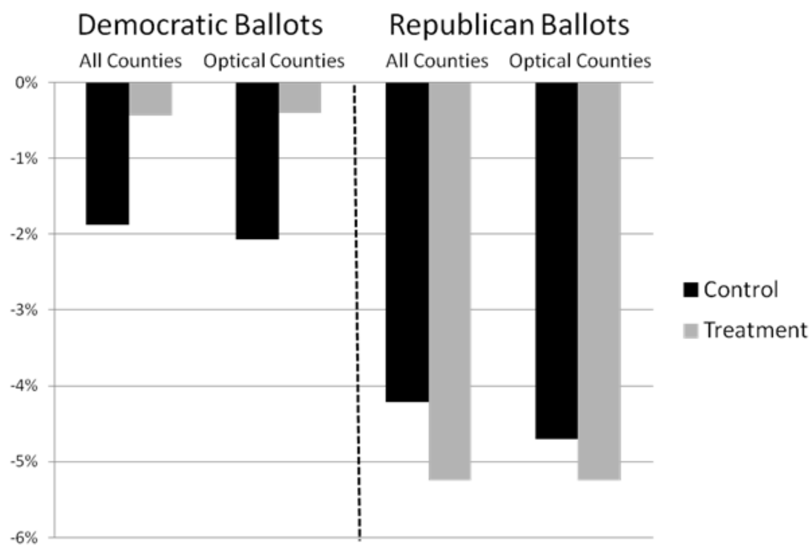

[14]In Figure 2 I group counties that did not employ optical scanning (Cook, Dupage, and Peoria) with 'control' counties, since there was no intervention between 2006 and 2010 in counties employing other technology.

Finally, I replicate the precinct-level analysis presented in Section 4.2 above, using the same model specification and formulation of all variables. As in that section, I present two models for each party's primaries. The first utilizes data from all counties, regardless of their voting technology.[15] The second model examines data only from optical-scan counties, employing weights from a genetic matching algorithm operating without restriction on 19 observable county-level variables measured before the 2010 outcome occurred. These covariates include measures of county population, demographics, socio-economic characteristics, previous political preferences, turnout and other voting behavior, urbanity, language competence, income, and unemployment. As in the precinct analysis, I also match on the propensity score.[16]

Table 4 contains OLS coefficients and robust standard errors from the models of county-level panel roll-off. The results are generally consistent with those of the precinct analysis. Across all four models, fluctuations in turnout are positively and significantly associated with higher levels of roll-off, indicating that as turnout rises, so too does the proportion of people casting undervotes. Moreover, while both models for Republican primaries return negatively signed coefficients, as in the precint-level analysis, none of the treatment indicators achieve statistical significance. To sum, neither mean comparisons nor a panel analysis of optical scanning precincts (holding turnout fluctuations constant) are suggestive of an effective undervote alert in the 2010 Illinois General Primary Election. Subsequent analysis of county-level data employing all precincts in Illinois agrees with these findings and raises confidence in their robustness.

[15]To reiterate, here non-optical scanning counties are included with the "control group," since there was no intervention between 2006 and 2010.

[16]I attempted more than twenty matching algorithms, with and without the propensity score, and without examining outcomes. I proceed with the one that produced the best overall balance. Balance statistics may be obtained from the author.

Table 4: OLS Regression Coefficients and Robust Standard Errors: Panel Difference-in-Differences in Roll-Off, County-Level Data

	Democratic Primaries		Republican Primaries	
	All Counties	Matched Optical Counties	All Counties	Matched Optical Counties
Treatment Indicator	1.05	.868	-.915	-.617
	(.709)	(1.75)	(.666)	(1.14)
Change in Party Turnout	.485*	.484*	.312*	.329*
	(.098)	(.098)	(.070)	(.072)
Constant	-1.30*	-1.12	-4.67*	-4.99*
	(.565)	(1.69)	(.565)	(1.10)
N	102	98	102	98
R-squared	.294	.277	.189	.209
RMSE	3.65	3.82	3.25	3.31

Dependent variable is the difference in total roll-off in statewide races from 2006 to 2010.

All percentages measured on a 0-100 scale.

5 Conclusion

In this paper, I searched for evidence of a relationship between the undervote alert provisions of Illinois's 2007 omnibus elections bill and patterns of aggregate undervoting in the 2010 Illinois General Primary Election. Throughout my analysis, I analyzed same-precinct difference-in-differences of overall roll-off in each party's primary elections, from 2006 to 2010. I first conducted uncontrolled means comparisons, which showed significantly lower roll-off trends in treated precincts in only two of the twelve 2010 statewide primary elections. I then calculated mean panel difference-in-differences, using regression models of Democratic and Republican primaries as well as both matched and unmatched data. In none of the models did aggregate roll-off appear to be lower in precincts that implemented the alert policy. I replicated this finding using county-level panel data.

Political science has long held that undervoting remains a largely intentional act, at least in the more visible races at the top of the ballot such as those affected by Illinois' undervote alert (Knack and Kropf 2003). Yet, previous research has also found that a visual stimulus occurring in the voting booth can significantly diminish undervoting (Nichols and Strizek 1995). In contrast to Nichols and Strizek, who surmised that such stimuli served as effective reminders to cast votes, I find that audible beeps fail to achieve similar results. It could be that audible alerts serve as a less-powerful inducement than flashing lights, and therefore that the Illinois policy is an ineffective means of reducing undervotes. Before moving to broad conclusions either about the impact of the alert on individual voting decisions or its efficacy as public policy, some caveats are therefore required.

With regard to the alert's effects on mass voting behavior, it may be tempting to conclude that the Illinois policy failed because most voters simply intended to cast undervotes, and therefore ignored the audible warning. However, it is important to note that because it employs data aggregated at the precinct and county levels, this paper should be read as a first analysis of the alert's overall effect, and not as an argument about determinants of individual voting behavior. Since I have no data about the citizens who actually voted, here I refrain from drawing conclusions about those citizens' motivations since such conclusions would likely be biased per the ecological fallacy (King 1997). In other words, while this paper provides preliminary conclusions regarding the ineffectiveness of the alert policy in the 2010 primary election, it cannot provide guidance as to *why* that policy might have failed.

Nor should this paper be read as the final and absolute word on the effect of the alert as policy, since the present analysis utilizes data only from primary elections. The 2010 General Primary was the first election for which the alert was in effect, and featured an uneven county-level implementation. Both of these traits make the 2010 primary a useful case in which to examine precinct-level roll-off trends, and the research design described above supports strong conclusions regarding the ineffectiveness of the audible alert *in the primary*. However, the exclusive focus in this paper on primary elections requires an important caveat for readers interested in holistic conclusions about the impact of the policy.

If the "revenge" rationale (described above) for the law is to be believed, then Illinois Democrats passed the alert as a means to reduce roll-off by *general* election voters whose straight-ticket voting prior to the termination of that practice would have aided all Democrats on the ballot. This framework holds that such voters, when reminded of their lack of a preference in a given race, would likely return to the voting booth and mark their ballot for a Democrat. Since Democrats' alleged political motivations were ostensibly to reduce "Democratic" undervotes in general elections, perhaps null effects should be expected in primary elections. It is worth noting however that it is unclear how powerful Democrats' political intentions actually were, since the law that they enacted requires the undervote alert only for elections to statewide offices, and not for state legislative, judicial, or other such "down-ballot" contests.

Nonetheless, the possibility remains that the alert has stronger effects on aggregate undervoting in general elections, which compared to primaries are typically higher-information affairs. If primaries are lower-visibility, lower-salience events in which information is harder to come by, then it seems reasonable to expect that they will attract a particularly dedicated type of voter. Indeed, statewide turnout during the 2010 Illinois General Primary was below 30%–far below the 50% that is typical in a general election. The 30% of Illinois voters who turned out in February 2010 are almost certainly representative of the most committed and information-seeking citizens of that state. Moreover, the models reported above suggest that roll-off rises with turnout. In tandem with the possibility that Democrats intended the alert to be most effective in general elections, the likelihood of a primary electorate less given to errant undervoting suggests that further examination of the alert's efficacy is warranted before concluding that the policy fails to reduce aggregate undervoting in all cases. Future research should therefore re-evaluate the efficacy of the undervote alert in general elections, either by replicating the methods employed in this paper or by employing an alternative research design that exploits the opportunities presented by the irregular implementation of the policy.

It certainly is possible that an examination of data from general elections would reach a different conclusion. That said, the findings reported here yield little optimism regarding the alert's effectiveness in the 2010 primary election. Perhaps the alert was ineffective in that case because it could not overcome natural voter apathy about certain races, be they less-visible or uncontested. It is also possible that the alert failed to reduce undervotes because it demands a more costly action of voters than the immediate quelling of blinking lights. While people who cast errant undervotes may have received a warning, it is not difficult to imagine a voter ignoring the alarm after weighing the costs of re-submitting a ballot against simply continuing the act of leaving the polling place. Both scenarios are likely to be valid in part.

Particularly if similar results are obtained in general elections, the lack of a relationship between the alert and aggregate undervoting has ramifications for both policymakers and political science. In the former arena, it is worthwhile to note that implementation is not a costless endeavor. For instance, the 66 counties that employ Diebold Accuvote machines spent about $135 per machine on upgrades necessary to implement the undervote alert (Guidroz 2010). Particularly given the sizable budget deficit in Illinois in both the 2009 and 2010 fiscal years, the five-figure cash outlay necessary to upgrade even 75 voting machines posed a potential hardship. Moreover, numerous election officials have expressed concern that the alert system may trigger longer wait times and a need for more poll workers, adding indirect implementation costs. Finally, there are real concerns regarding the capacity of undervote alerts to compromise the anonymous ballot. These costs may be bearable if weighed against a benefit of sufficient size. Unfortunately, the present analysis suggests that Illinois is received very little such benefit from its undervote alert policy in the 2010 General Primary Election.

References

[1] Ansolabehere, Stephen, and Jonathan Rodden. 2011. "Illinois Data Files, Version 2." Retrieved October 31, 2011 from: http://hdl.handle.net/1902.1/15845

[2] Buchler,Justin, Matthew Jarvis, and John E. McNulty. 2004. "Punch Card Technology and the Racial Gap in Residual Votes." *Perspectives on Politics* 2(3): 517-524.

[3] Bullock III, Charles S. and Richard E. Dunn. 1996. "Election Roll- Off: A Test of Three Explainations." *Urban Affairs Review* 32 (1): 71-86.

[4] Darcy, R., and Anne Schneider. 1989. "Confusing Ballots, Roll-Off and the Black Vote." *Western Political Quarterly* 42(3): 347-64.

[5] Diamond, Alexis, and Jasjeet S. Sekhon. 2006. "Genetic Matching for Estimating Causal Effects: A General Multivariate Matching Method for Achieving Balance in Observational Studies." Available at: http://sekhon. berkeley. edu/papers/GenMatch. pdf [Accessed March 16, 2009].

[6] Downs, Anthony. 1957. *An Economic Theory of Democracy.* New York: Harper and Brothers.

[7] Feddersen, Timothy J., and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse." *The American Economic Review* 86(3): 408-424.

[8] Feig, Douglas C. 2007. "Race, Roll-Off, and the Straight-Ticket Option." *Politics and Policy* 35(3): 548-568.

[9] Guidroz, John. 2010 (January 9). "Officials Worry About Costs of Catching Undervotes." *The State Journal Register*. Retrieved October 15, 2012 from: http://www.sj-r.com/top-stories/x1672002527/Officials-worried-about-costs-of-catching-undervotes

[10] Hall, Melinda Gann. 2007. "Voting in State Supreme Court Elections: Competition and Context as Democratic Incentives." *Journal of Politics* 69(4): 1147-1159.

[11] Hayes, Danny, and Seth C. McKee. "The Participatory Effects of Redistricting." *The American Journal of Political Science* 53(4): 1006-1023.

[12] Herrnson, Paul S., Benjamin B. Bederson, Richard G. Niemi, Fred Conrad, Michael J. Hanmer, and Michael Traugott. 2006. "The Not So Simple Act of Voting: An Examination of Voter Errors with Electronic Voting." Presented at the *2006 Annual Meeting of the Southern Political Science Association*, Atlanta, Georgia. January 6-8.

[13] Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart. 2007(a). "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(1):199-236.

[14] Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart. 2007(b). "Matchit: Nonparametric Preprocessing for Parametric Causal Inference." *Journal of Statistical Software*. Available at: http://gking.harvard.edu/matchit/.

[15] Kimball, David C., and Martha Kropf. 2008. "Voting Technology, Ballot Measures, and Residual Votes." *American Politics Research* 36(4): 479–509.

[16] Kimball, David C., and Martha Kropf. 2005. "Ballot Design and Unrecorded Votes on Paper-Based Ballots." *Public Opinion Quarterly* 69(4): 508-39.

[17] King, Gary. 1997. *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data.* Princeton, NJ: Princeton University Press.

[18] Knack, Stephen, and Martha Kropf. 2003. "Roll-Off at the Top of the Ballot: Intentional Undervoting in American Presidential Elections." *Politics and Policy* 31(4): 575-594.

[19] Lau, Richard R., and David P. Redlawsk. 1997. "Voting Correctly." *The American Political Science Review* 91(3): 585-598.

[20] Lau, Richard R., and David P. Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45(4): 951-971.

[21] McDermott, Monika L. 2005. "Candidate Occupations and Voter Information Shortcuts." *Journal of Politics* 67(1): 1468-2508.

[22] Nichols, Stephen M., and Gregory A. Strizek. 1995. "Electronic Voting Machines and Ballot Roll-Off." *American Politics Quarterly* 23(3): 300-318.

[23] Nickerson, David W., Ryan D. Friedrichs, and David C. King. 2006. "Partisan Mobilization Campaigns in the Field: Results from a Statewide Turnout Experiment in Michigan." *Political Research Quarterly* 59(1): 85-97.

[24] Reilly, Shauna, and Sean Richey. 2010. "Ballot Question Readability and Roll-Off: The Impact of Language Complexity." *Political Research Quarterly 63*(4): 1-9.

[25] Riker, William H., and Peter C. Ordeshook. 1968. "A Theory of the Calculus of Voting." *The American Political Science Review* 62(1): 25-42.

[26] Sekhon, Jasjeet S. 2006. "Alternative Balance Metrics in Matching Methods for Causal Inference." Available at: http://sekhon.berkeley.edu/papers/SekhonBalanceMetrics.pdf [Accessed March 16, 2009].

[27] Sekhon, Jasjeet S. 2008. "Multivariate and Propensity Score Matching Software with Automated Balance Optimization: The Matching package for R.." Available at: http://sekhon.berkeley.edu/papers/MatchingJSS.pdf [Accessed March 16, 2009].

[28] Sekhon, Jasjeet S. 2009. "Opiates for the Matches: Matching Methods for Causal Inference." Available at: sekhon.berkeley.edu/papers/opiates.pdf [Accessed March 16, 2009].

[29] Sekhon, Jasjeet S., and Walter R. Mebane. 1998. "Genetic Optimization Using Derivatives." *Political Analysis* 7(1): 187-210.

[30] Shocket, Peter A., Neil R. Heighberger, and Clyde Brown. 1992. "The Effect of Voting Technology on Voting Behavior in a Simulated Multi-Candidate City Council Election: A Political Experiment of Ballot Transparency." *Western Political Quarterly* 45 (2): 521-37.

[31] Sinclair, Betsy, and R. Michael Alvarez. 2004. "Who Undervotes, Using Punchcards? Evidence from Los Angeles County." *Political Research Quarterly* 57(1): 15-25.

[32] Stein, Robert M., Greg Vonnahme, Michael Byrne, and Daniel Wallach. 2008. "Voting Technology, Election Administration, and Voter Performance." *Election Law Journal* 7(2): 123-135.

[33] Tomz, Michael, and Robert P. Van Houweling. 2003. "How Does Voting Equipment Affect the Racial Gap in Voided Ballots?" *American Journal of Political Science* 47(1): 46-60.

[34] Vanderleeuw, James M., and Richard L. Engstrom. 1987. "Race, Referendums, and Roll-Off." *Journal of Politics* 49(4): 108 1-92.

[35] Wattenberg, Martin P., Ian McAllister, and Anthony Salvanto. 2000. "How Voting is Like Taking an SAT Test: An Analysis of American Voter roll-off." *American Politics Research.* 28(2): 234-250.