Florida 2006: Can Statistics Tell Us Who Won Congressional District-13?

Arlene Ash and John Lamperti

Figure 1. Map of Congressional District 13

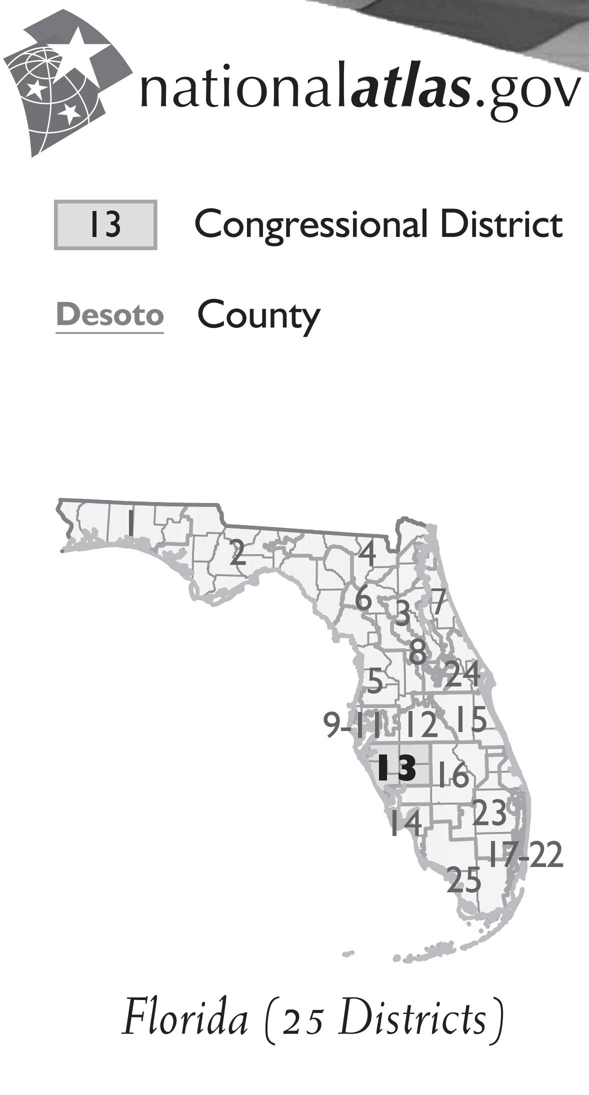

Elections seem simple. People go to the polls. They make choices about one or more contests or issues. The votes are counted. What can go wrong with that?

Unfortunately, many things can go wrong. In the United States, voters are often confronted with bewildering numbers of issues. Ballot choices and designs vary from election to election and district to district—or even within a district. People may have trouble casting the votes they intend. Both machine and human issues affect how votes are recorded and counted. Especially in a close race, the official results may not reflect the actual choices of the voting public.

Florida's 13th Congressional District 2006 Election

The 2006 contest for the U.S. House of Representatives in Florida's District 13 was such a race. The Republican candidate, Vern Buchanan, was declared the winner by just 369 votes, triggering a "mandatory recount." Unsurprisingly, re-querying the same "touch screen" machines that delivered the vote the first time changed nothing. The Democrat, Christine Jennings, challenged the result well into 2008. The problem is not that the race was close. It is that, in Sarasota County, an area of relative Democratic strength, some 18,000 people—almost 15% of those who went to the polls and cast ballots—had no choice recorded for their representative to Congress. A cast ballot with no recorded choice in a race is called an "undervote." The rest of the district contributed about half the total vote, but fewer than 3,000 undervotes. Jennings believes the excess missing votes in Sarasota would have tipped the race to her. Can statistical analysis help evaluate that claim?

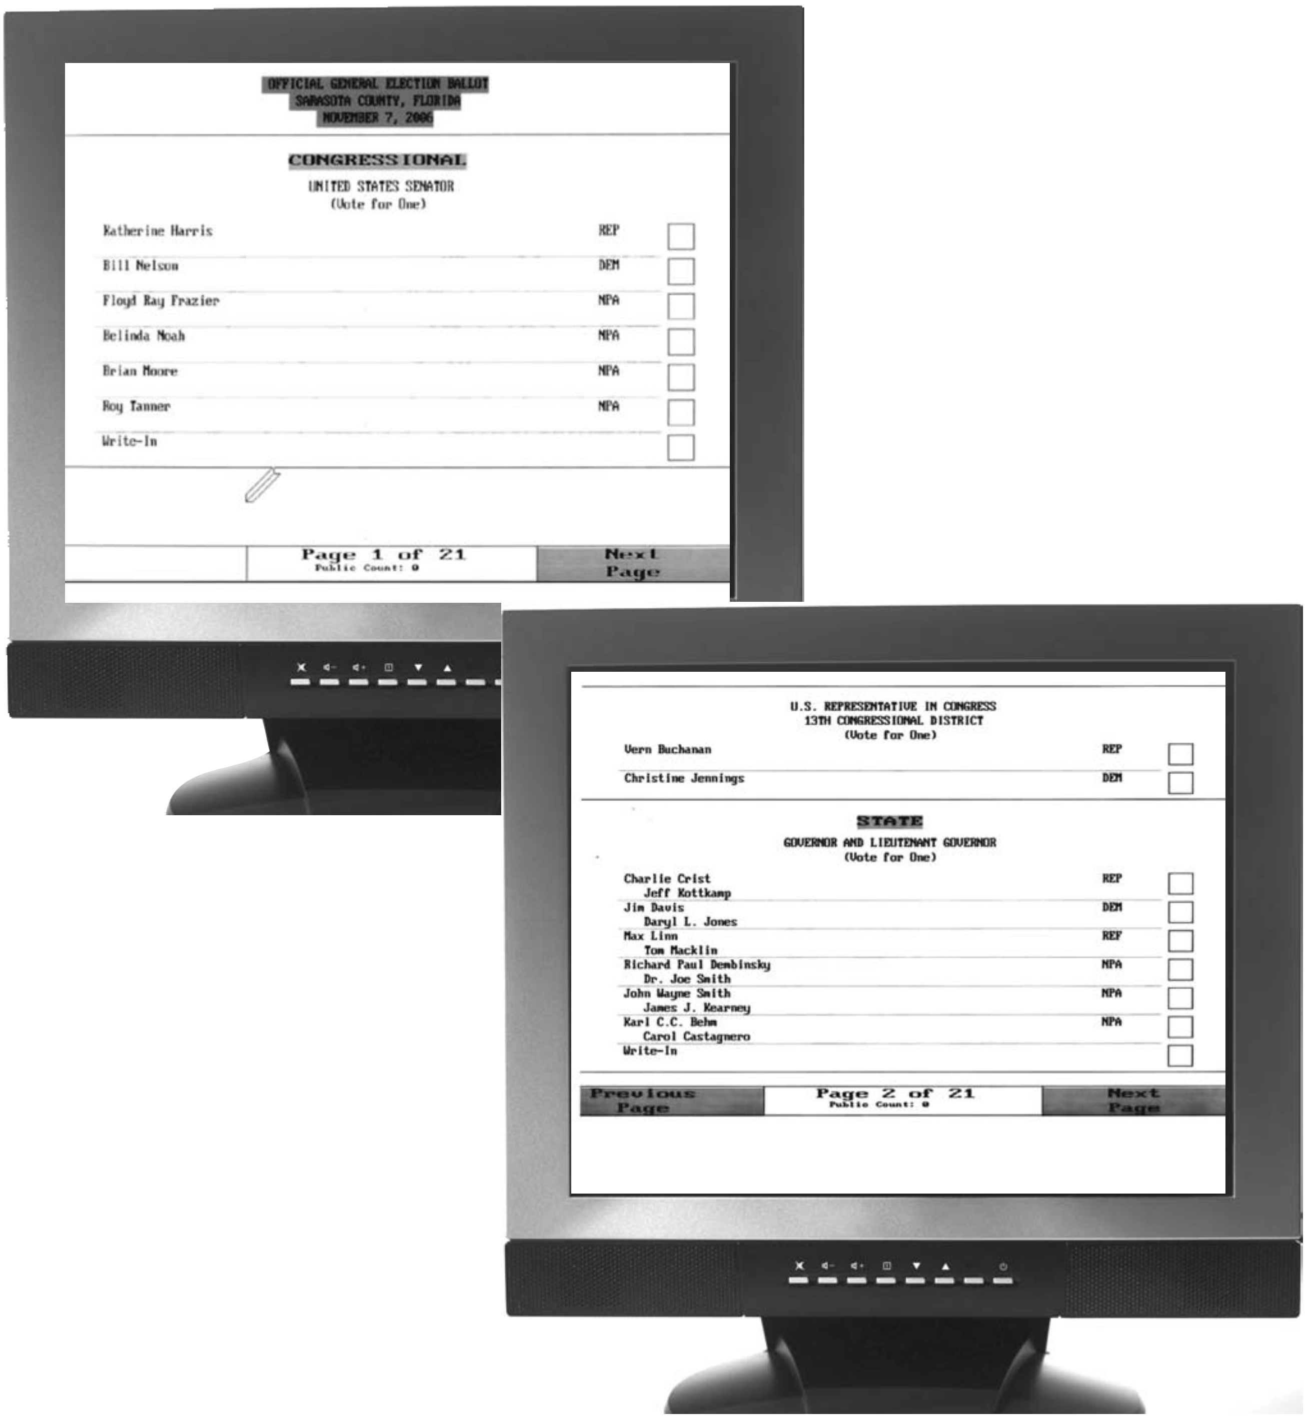

Figure 2. Screen shots of the first two (of 21) pages of the Sarasota County 2006 touch screen ballot

Congressional District 13 (CD-13) is geographically diverse (see Figure 1), including all of Sarasota; all or most of DeSoto, Hardee, and Manatee Counties; and a small part of Charlotte County. About half the district's population (a count of about 370,000 people) is in Sarasota. Manatee has a population of 310,000. DeSoto and Hardee together contribute 65,000 residents. Some issues and candidates are county-specific, so voters in different parts of the district faced different ballots. George Bush received 56% of the entire CD-13 vote in 2004. However, Sarasota County leans Democratic, and, of course, the broader political climate also shifted between 2004 and 2006.

In 2006, all voters in CD-13 participated in the House race plus five statewide elections—for U.S. Senate and four state offices: gubernatorial (for a combined governor/lieutenant governor slate), attorney general, chief financial officer, and commissioner of agriculture. They were also presented with numerous county-specific races and issues. Indeed, each District 13 voter faced a ballot presenting anywhere from 28 to 40 choices. Voting occurred in one of three ways: absentee

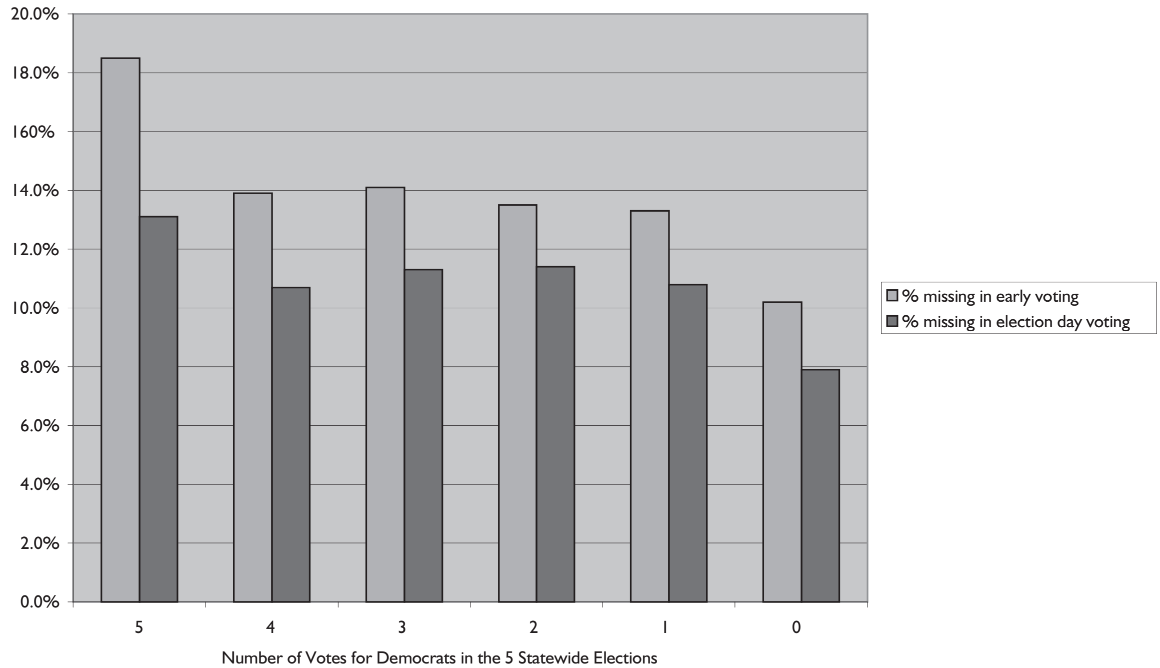

Figure 3. Undervotes in the House race by voting venue and partisanship of other votes among 104,631 ballots with votes recorded in all five statewide contests

ballot, early in-person voting, or traditional Election Day voting. Touch screen voting machines (also known as Direct Recording Electronic, or DRE) were used at all polling stations in Sarasota County for both early and same-day voting. Except for the absentee ballots, the machine totals are the only record of the vote.

What accounts for the 18,000 missing votes for U.S. representative? What would their effect have been?

Undervotes

Undervotes may be intentional—for example, in little-contested local races, where voters have no knowledge or preference. They also may be unintentional—the voters accidentally do not register a vote in a particular race. Finally, they may be entirely "false"—the voters choose, but no choice registers, as with the famous hanging chads of 2000. In well-publicized statewide or national races, undervoting is normally in the 1% to 3% range, with unknown contributions of intentional, unintentional, and false. The campaign for this important, open U.S. House seat had been intense and, by many accounts, dirty. Yet, in Sarasota County, about one out of every seven ballots cast by touch screen recorded no vote in this race. Why?

State officials at first echoed the explanation offered by aides of the declared winner: voters must have abstained due to disgust at the nasty campaign. However, none of the other counties had unusual undervotes in the same race. Manatee County, for example, reported normal undervoting of only 2%. Why would voter disgust stop at the county line? Moreover, the undervote on absentee ballots was low everywhere; only ballots in Sarasota County that had been voted on touch screens displayed abnormally high undervoting.

In Sarasota County, the highest undervote rate occurred in early voting. Thus, the huge undervote in Sarasota was specific to that county, applied to in-person voting but not absentee ballots, and moderated somewhat between early and election-day voting. There is at least one obvious explanation for this pattern—a ballot design (faced by touch-screen voters in Sarasota County only) that made it more difficult to vote for U.S. Representative there than elsewhere in CD-13. Indeed, the Sarasota *Herald-Tribune* cited contacts from "more than 120 Sarasota County voters" reporting problems, mainly with ballot screens that "hid the race or made it hard to verify if they had cast their votes." This alone would hurt Jennings, since Sarasota County voters were more favorable to her than were voters in the other counties.

The ballot design in Sarasota County certainly caused problems. Computer Screen 1 was devoted entirely to Florida's U.S. senatorial race, with seven lines of choices presented, immediately beneath a bright blue banner labeled "Congressional." The undervote rate in this race was normal (that is, low). But Screen 2 presented the House race at the top with only two voting lines and no special banner. The bulk of the page, following a second bright blue banner ("State") listed seven choices on 13 lines for the gubernatorial election. See Figure 2.

Laurin Frisina and three collaborators believe the CD-13 undervote in Sarasota County was due to the ballot screen layout. They point out that abnormally high undervote rates (ranging from 17% to 22%) also were found in the attorney general's race, and just in one part of CD-13: Charlotte County. On that ballot (only), it was the attorney general race with only two candidates that shared a screen with 13 lines of choices for the gubernatorial election.

Other factors likely contributed, as well. For example, there were abnormally slow machine response times that could have led people to "unvote" while trying to ensure their vote registered. This was flagged as a problem by the voting machine supplier the previous August, but not fixed prior to early voting. Furthermore, there are strong patterns in the undervote

Table 1—Florida's CD-13 Race in Sarasota County for All with Votes in Five out of Five Statewide Contests

Number of Democratic Votes in the Other Five Contests	Recorded and Missing Votes in the CD-13 Contest for the U.S. House of Representatives						Proportional Allocation of the Undervote		Change in Buchanan Minus Jennings Tally From Including the Undervotes
	Total # of Ballots	Buchanan	Jennings	No Vote	Recorded % for Buchanan	% Under-vote	Buchanan	Jennings	
Early Voting									
5	10,764	122	8,655	1,987	1.4%	18.5%	28	1959	-1,932
4	2,789	151	2,250	388	6.3%	13.9%	24	364	-339
3	1,170	174	831	165	17.3%	14.1%	29	136	-108
2	1,167	346	664	157	34.3%	13.5%	54	103	-49
1	2,173	1,227	657	289	65.1%	13.3%	188	101	87
0	9,455	8,059	435	961	94.9%	10.2%	912	49	863
Election Day									
5	25,326	468	21,541	3,317	2.1%	13.1%	71	3246	-3,176
4	7,637	561	6,261	815	8.2%	10.7%	67	748	-681
3	3,629	691	2,529	409	21.5%	11.3%	88	321	-233
2	3,847	1,387	2,022	438	40.7%	11.4%	178	260	-82
1	7,305	4,402	2,116	787	67.5%	10.8%	532	255	276
0	29,364	25,676	1,359	2,329	95.0%	7.9%	2212	117	2,095

within Sarasota County (see below), despite all Sarasota voters facing the same ballot. Walter Mebane and David Dill, after extensive study, believe the cause of "the excessive CD-13 undervote rate in Sarasota County is not yet well understood and will not be understood without further investigation." In any case, problems became evident during early voting, eventually leading Sarasota County's supervisor of elections to issue warnings to precinct captains. On election day, undervoting on these machines was lower than in early voting, but still exceeded 10%. This much is beyond dispute.

Consequences of the Undervote

But did it matter that 18,000 Sarasota voters had no recorded votes in the House race? Assuming a normal rate of intended undervotes, the choices of some 15,000 voters were not counted. What inferences can be made about how those votes would have divided between the candidates if they had been recorded? Would they have changed the outcome? There are several ways to tackle this question, and we'll describe perhaps the simplest one. Imagine a group of N voters, with R of them intending to vote for the Republican candidate and D for the Democrat so that R+D = N. Suppose a random group of N-n votes are "lost," creating an undervote. Thus, n votes are actually counted: r Republican votes and d Democratic ones ($d = n - r$). Let's think of these n recorded votes as a random sample taken without replacement from the population of N would-be voters. Of course, we often make inferences from samples to the whole population. Usually, the sample size, n, is a small fraction of the population size, N. Here, we have a very large sample; n is more than 85% as large as N! Never mind, the calculations are the same.

The r Republican votes in the sample are viewed as the result of n "trials," draws without replacement from a population of size N, where the "success" probability is p = R/N, here approximately 1/2. Thus, the expected value of r and its variance are computed in the familiar way:

$$E(r) = n \cdot p;$$
$$Var(r) = np(1-p)\frac{N-n}{N-1} \approx \frac{n(N-n)}{4N}.$$

The multiplier (N–n)/(N–1) is the familiar "finite population correction factor" for sampling without replacement, found in any survey sampling text. It can often be neglected—but not here! Both N–n and n are large, so the distribution of r is nearly normal. In this case, all we need do to estimate the Republican advantage (possibly negative) in the whole population is "inflate" r–d, the Republican advantage in the counted votes, by N/n, the fraction by which the whole population exceeds the counted vote. Thus, a statistically unbiased estimator of R–D is:

$$\text{Estimated } (R-D) = \frac{N}{n}(r-d) = \frac{N}{n}2r-n).$$

The associated standard error is:

$$SE = \sqrt{N(N-n)/n}\ .$$

This translates easily into a 95% confidence interval for R – D:

$$\frac{N}{n}(r-d) - 2\sqrt{N\frac{(N-n)}{n}} \le R - D$$

$$R - D \le \frac{N}{n}(r-d) + 2\sqrt{N\frac{(N-n)}{n}}.$$

How does this result apply to the District 13 election? First, let's imagine that, say, 20,000 nonvoters were randomly chosen from the whole voting population of the district, which was roughly N = 240,000 in 2006. The counted ballots gave Republican Buchanan an edge of 369 votes; that's the value of (*r–d*). By the above formula, the 95% confidence interval for R–D ranges from a low of just more than 100 to a high of nearly 700. Since the interval contains only positive numbers, we conclude with (greater than) 95% confidence that there would not be enough Democratic votes among the missing 20,000 to shift the outcome. Thus, despite the tiny winning margin (less than 1/6 of 1%) and the huge number of missing votes—if the missing votes were distributed just like the whole population—random error due to their loss would not threaten the outcome.

Of course, the missing votes were not chosen randomly from the whole district. For starters, the vast majority came from Sarasota County where Jennings had an advantage. Suppose there was a "normal" intentional undervote of 2.5% among the 120,000 voters in that county, so that only 15,000 (of the 18,000) undervotes were unintentional. Assume the 15,000 uncounted votes were chosen randomly from the county. Would that matter? Indeed, it would! In Sarasota, the recorded votes gave Jennings an edge of 6,833, so *r–d* = –6,833. If R–D now stands for the true Republican advantage among 117,000 would-be voters in Sarasota County, the point estimate for R-D is –7,838, with a 95%-confidence interval ranging from about –8,100 to –7,575. Elsewhere in the district, Buchanan had an advantage of 7,202 votes. If we treat the votes in the other parts of the district as error-free, the estimate indicates a win for Jennings by 636 votes, with a 95% confidence interval for R–D ranging from –898 to –373.

Again, the interval does not cross zero, and so, with more than 95% confidence, we conclude that Jennings should have won. In fact, had we used ±4 SE instead of ±2 SE, the confidence interval still would not include zero; this raises the confidence level to 99.9%. Moreover, in the context of a one-sided question—did Buchanan really get more votes than Jennings?—one-sided confidence bounds could be used, raising the level of certainty even higher.

Refining the Estimate

In making this estimate, we assumed 15,000 unintentional undervoters in Sarasota County differ from those who did vote only in that their votes were not recorded. Can this assumption be tested? Table 1 and Figure 3 are based on "ballot image" data from Walter Mebane that show the sets of choices for the 104,631 Sarasota County ballots with touch screen votes recorded in all five statewide contests. The data are arranged by early versus Election Day voting and by the number of Democrats chosen in the five statewide contests. We'll soon see how useful such data can be.

First, in both early and Election Day balloting, there is a steep gradient associating partisan voting in the other races and the preference of voters—those whose choices were captured—in the House race. For example, in early voting among otherwise "straight ticket" Democrats, only 1.4% of votes for the House race went to Buchanan, as opposed to 94.9% of recorded votes among early voting Republican stalwarts.

Second, it was far easier to "lose" Democratic votes than Republican ones in this race. For example, the straight ticket Democrats had 18% uncounted votes in early voting as opposed to "only" 10% for their early voting Republican counterparts. Understanding what caused these differences is crucial for the legal challenge to this election and for avoiding future voting debacles. For our purposes, we merely note that—in contrast to our previous assumption—not all Sarasota voters were equally at risk for unintentional undervotes. We'll return in a minute to the more refined calculation of the expected effect of the lost votes these data allow.

A third important fact that emerges (see Figure 3) is that the undervote declined substantially within all categories of voters between early voting and Election Day voting. Apparently, many voters were helped by actions taken to mitigate the problems seen in early voting. A study exploring associations between corrective actions taken at individual precincts and undervote rates could be very informative. We do not have such data.

What we do have in the ballot image data leads to a sharper estimate of the likely disposition of most of the missing congressional votes. First, it is hard to imagine that many of the 12,000 voters who expressed a choice in all five statewide races (including commissioner of agriculture and chief financial officer), but had no vote recorded in the House race, intentionally undervoted. Let's suppose they all intended to vote. How would they have voted? A good guess is that the people with missing House votes in each of the 12 strata in Table 1 would have voted in the same proportions as those in the same stratum whose votes were recorded. That is, we perform the same calculations as above, this time within each stratum of Table 1. Then, we sum the estimates of the "full" vote across the strata, leading to a new estimate of R–D that represents the Republican advantage after imputing values for the undervote among these 12,000 people. This calculation suggests Jennings' advantage among these lost votes alone was almost certainly greater than 3,000. It swamps Buchanan's original 369-vote winning margin.

For whatever reasons, it was harder to cast a successful vote for Jennings than for Buchanan in Sarasota County. The higher observed undervote among presumed Democrats means our previous confidence interval calculation was conservative; the conclusion that Jennings was the real winner in CD-13 becomes even surer.

The study by Frisina uses two methods to analyze the CD-13 undervote. Both infer undervoters' choices from their votes for other candidates. One uses precinct-level data from Sarasota County. The other involves matching Sarasota voters with counterparts in Charlotte County. Both show that Jennings was almost certainly the preferred choice among the majority of CD-13 voters.

These different estimates may seem confusing. However, the key point is that all plausible models of what the lost votes would have been point to the same conclusion. Furthermore, the more carefully we examine the data, the more support we see for that conclusion. While poor ballot design may or may not fully account for the Sarasota undervote, it is clear that those missing votes switched the outcome of the congressional race from Jennings to Buchanan.

What Happens Now?

Finally, two questions. How should Florida and other states fix their flawed electoral processes? Requiring a paper record is useful, but not enough, since recounting such a record in District 13 might have simply confirmed that 18,000 Sarasota County voters recorded no choice for their U.S. representative. The paper record, therefore, must at least be confirmed by each voter. We favor paper ballots, plus optical scanners to read them—the method familiar to us all from grading tests and used now for elections in many states. It is relatively inexpensive and foolproof. It does not require new, possibly fragile, technology or big capital investments. It provides an independent check on what is going on inside the machines that tally the votes. Optical scan ballots are also easier to read and less prone to the design problems that disfigured the CD-13 House race. Indeed, optical scanning was used in 2006 in Sarasota County for the absentee ballots and it worked well.

The second question, of course, is what to do about that dubious 2006 election. The statistical evidence shows, beyond any reasonable doubt, that more voters wanted Jennings than Buchanan. However, there is—as yet—no precedent for a court overturning an electoral "count" based on a statistical analysis. We have recommended doing this election over—and doing it right. For the future, statisticians and voting experts should work together to develop guidelines for the appropriate use of statistical evidence to confirm, or overturn, elections. C

Further Reading

Adams, Greg (2001). "Voting Irregularities in Palm Beach, Florida." *CHANCE*, 14:22–24.

Frisina L., Herron M., Honaker J., and Lewis J. (in press). "Ballot Formats, Touchscreens, and Undervotes: A Study of the 2006 Midterm Election in Florida." *Election Law Journal*. Draft at *www.dartmouth.edu/~herron/cd13.pdf*.

Marker, D., Gardenier, J., and Ash, A. (2007). "Statistics Can Help Ensure Accurate Elections." *Amstat News*, (360):2–3. Online at *www.amstat.org/publications/AMSN/index.cfm?fuseaction=pres062007*.

McCarthy, J., Stanislevic, H., Lindeman, M., Ash, A.S., Addona, V., Batcher, M. (2008). "Percentage-Based Versus Power-Based Vote Tabulation Statistical Audits." *The American Statistician*, 62(1):1–6. (A more detailed version is available at *www.verifiedvotingfoundation.org/auditcomparison* as "Percentage-Based Versus SAFE Vote Tabulation Auditing: A Graphic Comparison.")

Mebane, W. Jr., Dill, D.L. (2007). "Factors Associated with the Excessive CD-13 Undervote in the 2006 General Election in Sarasota County, Florida." *www-personal.umich.edu/~wmebane.*

Meyer, Mary C. (2002). "Uncounted Votes: Does Voting Equipment Matter?" *CHANCE*, 15:33–38.

Pynchon, S., and Garber, K. (2008). "Sarasota's Vanished Votes: An Investigation into the Cause of Uncounted Votes in the 2006 Congressional District 13 Race in Sarasota County, Florida." Florida's Fair Election Coalition. *www.floridafairelections.org*.

Wallace, J. (2006). "Political Operatives Gather for Recount." *Herald Tribune*, *www.heraldtribune.com/apps/pbcs.dll/article?AID=/20061111/NEWS/611110643*.

Special Section: District 13 Election, *www.heraldtribune.com/apps/pbcs.dll/section?CATEGORY=NEWS0521&template=ovr2*.

Wolter, K., Jergovic, D., Moore, W., Murphy, J., O'Muirheartaigh, C. (2003). "Reliability of the Uncertified Ballots in the 2000 Presidential Election in Florida." *The American Statistician*, 57(1):1–14.

Statistical Solutions to Election Mysteries

Joseph Lorenzo Hall

Elections in the United States are strange. While other nations have problems with violence at the polls or seemingly insurmountable logistical issues, the problems in our country cluster around complexity. No other country votes so frequently, for so many contests at all levels of government, using dozens of methods to enfranchise all eligible voters. Naturally, such complexity results in frequent errors and a few genuine mysteries.

Arlene Ash and John Lamperti confidently (with greater than 99.9% confidence) conclude that the wrong candidate is currently holding the CD-13 office. This is perhaps the worst possible outcome in an election, with a close second being that there is no discernable winner. Although Ash and Lamperti don't address it, the case of the disputed Florida 2000 presidential election was similar. Florida 2000 received a lot of attention in political science literature. Researchers such as Walter Mebane arrived at similar conclusions, but due to a different mechanism: Instead of mysterious undervotes changing the outcome of the race, the problem in Florida 2000 was with spurious overvotes—where ballots show too many choices recorded for a particular race.

Both of these cases enjoy peculiar features that many election mysteries do not. First, the underlying data in terms of ballot image data and precinct-level vote data could be obtained by using Florida's public records laws. Florida permits some of the highest access to the inner workings of its government via the Florida Public Records Act. In both of these cases, researchers were able to obtain crucial data that would typically not be made publicly available in other states.

Second, when these data were analyzed, researchers found a definitive answer with respect to the disposition of the outcome. Many election mysteries remain mystifying, even after forensic investigation. A case in point is the search for an answer to a different question about the same CD-13 race: What was the cause of the prodigious undervote? As Ash and Lamperti point out, a team of academic computer security experts examined the software that runs the voting machines used in Sarasota's CD-13 race and could not find a software-based cause.

The problem that Ash and Lamperti address is a subset of a more general problem: measuring how confident we are that an election has been decided correctly. In hindsight, one would think mechanisms to ensure election confidence would have been designed into our electoral system, given its fundamentally adversarial nature. Unfortunately, in many cases, the only checks performed on election results are recounts, which can have significant costs and legal barriers and be noninformative.

Part of the answer proposed by Ash and Lamperti is to regularize checking the math behind our elections. This requires two elements: There needs to be something to audit—an audit trail—and there needs to be the appropriate regulatory and procedural infrastructure to conduct election audits. For auditability, voting systems must produce an independent, indelible, and secure record of each ballot voters check for correctness. Fortunately, only a minority of 12 states currently do not require their voting systems to produce such records. However, a recent study by Sarah Everett provides compelling evidence that people don't check these records, and, when they do, they don't notice errors. To improve auditability, we need a combination of voter education about audit record verification and further usability research to make these records easily verifiable.

Unfortunately, despite states overwhelmingly moving toward producing audit records, audits of these records are only performed in one-third of all states, and then they are performed under a wide variety of standards. A white paper authored by Lawrence Norden, Aaron Burstein, Joseph Lorenzo Hall, and Margaret Chen (see *www.brennancenter.org/dynamic/subpages/download_file_50227.pdf*) with the input of a blue-ribbon technical panel highlights this disparity and reviews the various types of post-election audit models in theory and practice. To address this imbalance and inconsistency, it appears we need federal legislation that mandates election audits and audit standards for federal elections.

One solution that comes to my mind, which Ash and Lamperti do not propose, is that of user testing of ballot styles. User testing would involve usability testing each ballot style with a number of actual users to detect strange or unintended behavior. This kind of testing would discover both problems with particular ballot styles and other types of interaction problems, including software bugs. An analogy could be made to the use of focus groups and pilot studies to test survey instruments. Usability testing on this scale, where ballot styles can number in the thousands for certain jurisdiction's primary elections, would be prohibitively expensive in terms of time and resources. Thousands of ballot styles can result when factors including political party, level of election (e.g., federal, state, local), ballot status (official or provisional), and language are crossed. Limited user testing would certainly be less expensive, but it would be much less effective.

Ash and Lamperti propose a less intense, but equally radical, solution to these kinds of mysteries. They advocate allowing elections to be overturned based on statistical evidence. Compared to regularizing post-election audits, this proposal is obviously more complex, involving legal line-drawing standards about when to consider an election suspect based on statistical evidence. For example, is 95% confidence that the election was decided incorrectly enough? 90%? 99.9%? Should the standard be overwhelming statistical evidence, indisputable statistical evidence, or something else? And how will assumptions about undervotes, such as those discussed by Ash and Lamperti, and overvotes be evaluated? Who will do the evaluation? Different assumptions, in some cases, will make a difference. Developing such guidelines for statistical challenges to elections will be difficult, but it might be exactly what judges look for in future litigation involving election mysteries.

Joseph Lorenzo Hall can be reached at *joehall@berkeley.edu*.

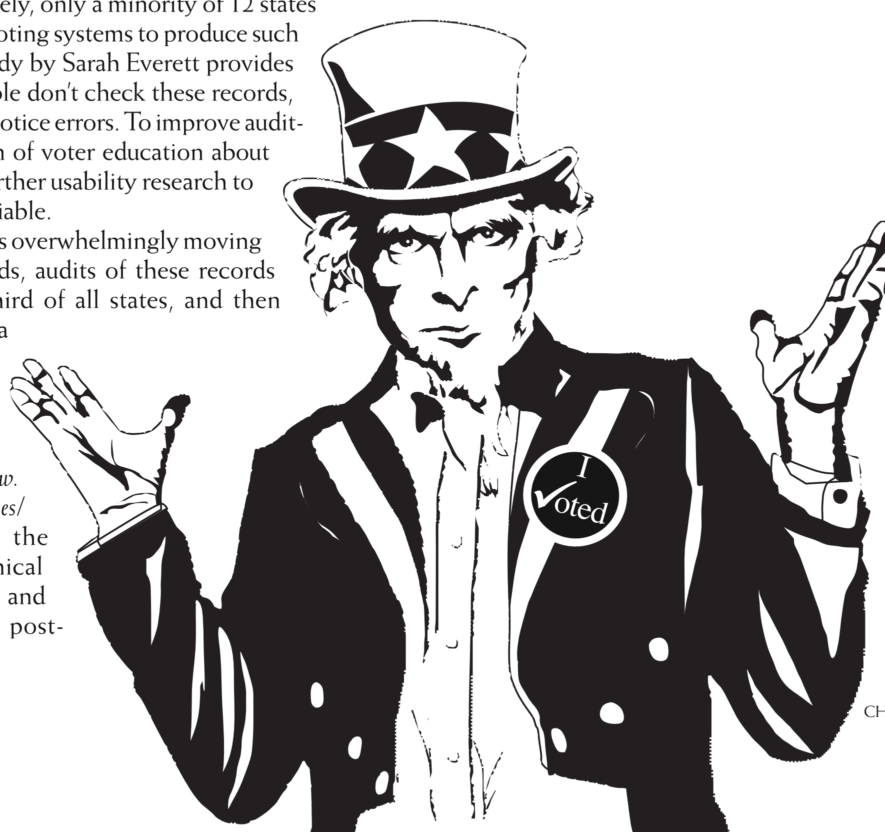

Counting Frustrated Voter Intentions

Walter R. Mebane Jr.

People go to the polls to vote, and then what happens? Recent elections in the United States have seen many cases where voters voted in circumstances that left too many of them doubting whether their votes were counted.

In the 2004 election, this happened not only in Ohio, but in several states that used electronic, touch screen voting technology. In 2006, there were relatively minor problems in various jurisdictions, but initial reports suggested voters' experiences were, in general, better than they had been during 2004, according to a *Washington Post* article by Howard Schneider, Bill Brubaker, and Peter Slevin.

The election for the U.S. House of Representatives in Florida's District 13 helped shatter the illusion of normalcy, reliability, and success. As Arlene Ash and John Lamperti observed, more than 18,000 votes cast on iVotronic touch screen machines in Sarasota County in that race were unaccountably missing (an iVotronic machine like the ones used in Sarasota is described at *www.srqelections.com/ivotronic/ivotronic.htm*).

Ash and Lamperti show that any of several reasonable conjectures about the intentions of the voters who cast these undervotes imply the missing votes are sufficient to have changed the outcome of the election. This finding agrees with the conclusions reached by experts on both sides of one of the lawsuits filed to challenge the outcome based on allegations of defects in the voting machines.

Voter intent is, at first glance, a straightforward idea. Out of a set of candidates or a set of options regarding a ballot initiative, each voter has, at the moment of voting, decided to choose one or has decided not to make a choice. The voter's intention is to have that choice conveyed accurately into the final vote count, or if the voter abstained, the intention is to not have an effect on the final vote count. The voter undertakes some physical gesture—for example, marking on a paper ballot or touching a video screen—with the idea that gesture will ultimately cause the final vote count to be changed, or not, in the way the voter intended.

Nuances come to light when we think about different ways a voter's intentions may be frustrated. Once the voter is at the moment of voting, there are, broadly speaking, two ways things can go wrong. Something can prevent the voter from making the gesture that would express the voter's choice. Or, something can prevent the voter's gesture from having the desired effect on the final vote count. In both cases, there are further important distinctions pertaining to where the difficulty occurs. When a voter is unable to make the appropriate gesture, is that something about the voter or something about the circumstances? Was the Election Day environment the same for all voters, but this particular voter was somehow unable to do the right thing in that setting? Or, were different voters somehow treated differently? When an appropriate gesture does not have the desired effect, is the obstacle something occurring immediately in the voting machine the voter is using or something that happens later in the process, perhaps long after the voter has left the polling place?

Table 1—Sarasota 2006, District 13 Election Day Undervote Rate by Occurrence of Event 18 ("Invalid Vote PEB") on Machine and Event 36 ("Low Battery Lockout") in Precinct

Event 18			
Event 36		No	Yes
No	CD-13 Undervote Rate	13.7%	14.6%
	Total Ballot Count	67,748	9,879
Yes	CD-13 Undervote Rate	14.6%	15.0%
	Total Ballot Count	9,716	1,699

Note: Rates are the proportion of the Election Day ballots in each category that have a CD-13 undervote.

Maybe all of these ways voters' intentions may be frustrated can serve equally well to motivate a what-if exercise designed to see what would have happened had all votes been counted as they were intended. Ash and Lamperti do not try to decide among the several explanations that have been suggested for the excessively high rate of undervoting. But, it may be important to take a stronger stand on this. Suppose one believes the high number of undervotes was the result of some voters being unable to make an appropriate gesture in an environment that was the same for all voters. Someone with such beliefs may be skeptical that these undervotes are unintentional. After all, following the 2000 election debacle in Florida, an elections supervisor stated that the blame for spoiled ballots falls on the voters, wondering, "Where does their stupidity enter into the picture?" Such people seem to believe that would-be voters who fail to solve perceptual or procedural puzzles that all voters have been given to solve do not deserve to have their votes counted. To fully motivate what-if exercises such as the ones Ash and Lamperti carried out, it may be important to demonstrate that the frequency of undervotes varied with circumstances that varied across voters.

So, one can have two attitudes about the claim Laurin Frisina et al. make, that "the exceptionally high Sarasota undervote rate in the 13th Congressional District race was almost certainly caused by the way Sarasota County's electronic voting machines displayed on a single ballot screen for the congressional contest and the Florida gubernatorial race." One view is that, because the ballot's format varied across Florida counties, voters in different counties did face different circumstances. Such a perspective may carry the implication that voters' experiences within each county were homogeneous. One might argue, then, that any variations in the undervote rate among voters within each county must trace back to something about the voters. The cross-county heterogeneity perspective might lead one to think the what-if exercises are well motivated, but the within-county homogeneity perspective might point in the opposite direction.

In fact, different voters in Sarasota faced significantly different circumstances, because different voting machines

performed differently. Recent reports by Susan Pynchon and Kitty Garber document problems that afflicted iVotronic touch screen voting machines not only in Sarasota County, but wherever they were used throughout the state. These reports go beyond previous investigations that considered a limited range of evidence regarding software failures. Pynchon and Garber document extensive problems ranging from low battery errors and power failures to poor security for critical voting machine hardware. They also show that, across the state, undervote rates were higher for many races where iVotronic touch screen machines were used, regardless of the ballot format.

The force of these recent reports is to suggest not only that machines, and not voters, were responsible for excessive undervotes, but that it is possible that security failures allowed vote counts to be altered long after the polls closed. The reports do not demonstrate that manipulations before, during, or after the election definitely occurred, but they do document striking security failures and show that previous investigations were not sufficient to rule out such possibilities.

Using data from Sarasota, one can show that readily measurable problems with the voting machines correlate with significant variations in the frequency of undervotes in the election for the U.S. House of Representatives in District 13. I consider variations in this undervote rate across four categories, defined by two kinds of error conditions. One is an error indicating that an invalid Personalized Electronic Ballot (PEB) was used with the voting machine. PEBs are electronic devices used to conduct all transactions with the iVotronic touch screen machines, including the action of loading the ballot each voter will see and enabling the voter to vote. PEBs are described at *www.srqelections.com/ivotronic/ivotronic.htm* (click on #1). In records produced to show all of the transactions on each voting machine, an "invalid vote PEB" error is denoted as event 18. Walter Mebane and David Dill highlight the relationship between this error and variations in the undervote rate at *www-personal.umich.edu/~wmebane/smachines1.pdf.*

The second kind of error is whether any voting machine in a precinct had a power failure. Such an event for a voting machine is indicated by event 36 ("low battery lockout") in the machine's transaction log. Garber observes that voting machines were often not plugged directly into a wall socket to receive power, but daisy-chained, with one machine plugged into another machine. She and Pynchon also point out that low power conditions or power failure may cause a variety of machine performance failures.

Table 1 shows that on election day in Sarasota, the District 13 undervote rate was lowest (13.7%) on machines not subject to either of the two kinds of error, and the undervote rate was highest (15.0%) on machines on which both kinds of error occurred. Having only the invalid vote PEB error on a machine and having only the low battery lockout error on a machine in the same precinct are each associated with an increase of almost 1% in the undervote rate (to 14.6%). These percentage differences are arguably small, relative to the overall undervote rate, but even they are enough to potentially have had a significant impact on the election outcome. If all four categories of votes shown in Table 1 had had the lowest displayed undervote rate, there would have been 202 fewer undervotes—a number about two-thirds of the margin of victory in the election.

By presenting Table 1, I do not mean to suggest the undervote problem mostly traces to circumstances unrelated to voting machine performance. Especially in view of the wide range of concerns Pynchon and Garber document, Table 1 should be viewed as expressing a lower bound on the share of the undervotes caused by mechanical failures.The final message about undervoting in the 2006 election in Florida is that we still don't know precisely what caused the problem. In Sarasota, more than 18,000 votes effectively vanished into thin air, but, across Florida, the number of mysteriously missing votes is several times that number. Without paper ballots to recount and inspect, and barring purely statistical adjustments, it is difficult to know what can be done practically to remedy the situation in a way that inspires everyone's full confidence. The worst fear is that, as bad as they are, the problems we can see are only a small part of what's really wrong.

Further reading can be found in the supplemental material at *www.amstat.org/publications/chance.*

Walter R. Mebane Jr. can be reached at *wmebane@umich.edu*.