



**Post-Election Risk-Limiting Audit Pilot Program 2011-2012
Semi-Annual Progress Report to the Election Assistance Commission
Reporting Period Close: November 30, 2011**

Summary

The Secretary of State (SOS) Post-Election Risk-Limiting Audit Pilot Program is designed to test the effectiveness and efficiency of risk-limiting audits. As of November 30, 2011, audits have been conducted in 4 counties following elections held between March and November 2011. All 4 audits successfully confirmed official election results by reviewing a small number of ballots (i.e., a few dozen to a few hundred ballots) cast in each contest audited. By contrast, California law requires elections officials to hand tally 100% of the ballots from 1% of all precincts after each election, an exercise that typically requires hand counting thousands of ballots.

Despite the high number of ballots hand tallied for the 1% manual tally, the pilot project team's analysis showed this statutorily-mandated manual tally to be ineffective and inefficient at confirming election results. The post-election risk-limiting audits were able to confirm with 90% confidence that election results were correct after hand counting very few randomly selected ballots. By contrast the 1% manual tally in the same elections gave very little statistical proof that the election outcomes were correctly calculated by the voting system.

Twenty counties have volunteered to participate in the program and 8 audits were conducted in 2011. Of the 8 audits, 2 were conducted in spring 2011 prior to the grant award, 2 were conducted during the current progress reporting period and 4 were conducted in December 2011). Since 6 counties participated after the grant award, the SOS will conduct audits in the up to 14 additional counties following the June and November 2012 elections.

The \$230,000 two-year grant from the federal Election Assistance Commission (EAC) helps fund:

- 1) Audits of election results following live elections in 20 California counties;
- 2) Detailed analyses of the efficacy of risk-limiting audits and recommendations on modifications needed to make current voting systems auditable; and
- 3) Creation of auditing tools for elections officials. The pilot program team has developed draft audit rules for selecting the initial sample size and for determining when enough ballots have been audited, methods for ballot-level audits, and user-friendly web-based tools and procedures for conducting and reporting on risk-limiting audits. (A preliminary version is available at statistics.berkeley.edu/~stark/Vote/auditTools.htm)

Risk-limiting audits are audits based on modern statistical principles. The number of ballots initially reviewed in a risk-limiting audit varies based on the margin of victory. The audit escalates – potentially to a full hand count of every ballot cast – if significant differences between the hand tally and the voting system tally are found. Risk-limiting audits are efficient when conducted at the “ballot level,” meaning individual ballots from the entire voting jurisdiction are subject to the random draw and the audit. This contrasts with California’s statutorily-mandated 1% manual tally, where only the precincts are subject to the random draw from across the entire election jurisdiction, not the ballots themselves. Put another way, risk-limiting audits generally involve hand counting fewer ballots overall, but those ballots come from across the entire voting jurisdiction, whereas the 1% manual tally generally involves significantly more ballots but only from specific areas of the voting jurisdiction.

Now that a variety of methods for risk-limiting audits have been tested, the SOS believes efficient and effective election auditing requires auditing at the ballot level. Therefore the team will develop standards, procedures and tools for such audits as part of this project. This will:

- 1) Help California and other states develop new, more robust and informative election auditing laws,
- 2) Inform the design of next generation voting systems,
- 3) Provide election auditing best practices and procedures that can be used by many jurisdictions in the U.S. using a broad variety of voting systems; and
- 4) Build public confidence that if there are errors in election results, those errors will be caught and corrected.

Due to delays with the state budget revision process, the SOS was unable to draw down grant money for the period of May 23, 2011, to November 30, 2011, for the Post-Election Risk-Limiting Audit Pilot Program. However, the SOS proceeded with the program despite this delay. On December 21, 2011, the SOS has confirmed that the California Department of Finance completed processing the necessary documents to enable the SOS to begin drawing down EAC grant funds for the pilot program. The expenses incurred during this reporting period are included in the accompanying fiscal report and will be drawn down prior to the end of the next reporting period.

Background

Why audit?

State and federal voting system testing and certification help ensure voting systems used in the U.S. can count ballots accurately and securely, while protecting voter privacy. But front-end regulation and testing isn’t enough. How can the public know whether voting

systems actually got the job done right, unless the election results are audited after the election?

History has shown that election fraud is not theoretical. Computer experts have demonstrated that voting systems can be hacked. But even setting aside the chance of voter fraud or tampering, no voting system – no machine – can operate to perfection. Neither can humans. Machines misinterpret ballots, people mis-mark ballots. Errors happen, and auditing determines whether those errors matter – in other words, whether a full hand count would show a different winner.

What is a “risk-limiting audit?”

Risk-limiting audits can determine with precision how much auditing is necessary to confirm election results with a high confidence. For the pilot program, the model is set to provide 90% confidence that the audit will require a full hand count of ballots if the official results (i.e. winners) tallied by the voting system are in fact wrong. The size of the initial sample of ballots depends on the margin of victory in the contest: the narrower the margin, the larger the initial sample. For a higher confidence level of, for example, 99%, that any voting system tallying errors that led to the wrong winner being declared would be caught, more ballots would be audited. The audit can stop after the initial sample if auditors find no errors or only statistically insignificant errors. An “error” is any difference between the voting system and a human eye interpretation of a vote.

Risk-limiting audits can lead to a full hand count of ballots to determine that the voting system tallying software correctly declared the election winner. If the initial sample uncovers significant errors, then the audit escalates, and additional ballots are selected for hand counting. Escalation continues until the error rate falls below certain level, a level which indicates statistically that a full hand count would confirm the winner. If significant errors persist during escalation, then the audit can lead to a full hand count of ballots to confirm (or overturn) results.

Risk-limiting audits sample individual ballots spread across a jurisdiction – a method that is dramatically more effective and efficient than hand counting 100% of the ballots from significantly fewer randomly selected precincts and increases the likelihood that any voting system tallying errors will be caught.

California’s 1% manual tally law dates back to the 1960s and requires elections officials to hand count the votes for all contests on all ballots from 1% of all precincts statewide. In a regular election year, counties must count tens of thousands of ballots as part of the required 1% manual tally, yet doing so gives very little proof that the voting system tallied ballots correctly and got the right winners. With risk-limiting audits, individual contests or groups of contests on the same ballot can be audited and the winners confirmed by looking at relatively few individual ballots. For the simplified version of the risk-limiting audit model, called the “Super-Simple” audit, which has been used for some of the audits in this pilot project, the size of the initial sample is the number 4.8

divided by the margin of victory. If elections officials expect to see a few errors then a slightly larger sample is drawn:

Margin	Equation	Initial Sample Size (in ballots)
50%	4.8/.5	10
40%	4.8/.4	12
30%	4.8/.3	16
20%	4.8/.2	24
10%	4.8/.1	48
5%	4.8/.05	96
2%	4.8/.02	240
1%	4.8/.01	480
.5%	4.8/.005	960

Professor Philip B. Stark (Stark) from the University of California, Berkeley, has been a pioneer in developing risk-limiting election audit methods. His methodology has since been endorsed by Common Cause, the League of Women Voters, the American Statistical Association, Verified Voting, and other groups involved in election integrity.

Progress

During this first phase of the project the pilot program team:

- 1) Gathered a mix of 20 urban and rural counties, including at least one county representing each voting system in use in California, to participate;
- 2) Entered into a contract with the University of California (UC) to allow Stark to serve as lead researcher for the pilot program;
- 3) Convened an advisory panel and established a webpage for the pilot program;
- 4) Conducted audits in 4 counties before the November 30, 2011, close of the EAC's semi-annual reporting period (eight audits total as of the date of this report); and
- 5) Developed and tested draft web-based tools, instructions and methods for conducting risk-limiting audits.

Counties

The 20 counties below volunteered to participate in the pilot program. Two of these counties participated in pilot audits prior to the grant award. The results of their participation are included in this report, even though the expenses related to these counties are being absorbed by the counties and the Secretary of State.

County	Election Date	Audit Date
Alameda	Nov. 8, 2011	Dec. 5, 2011
Alpine	TBD	TBD
Colusa	Nov. 6, 2012	TBD
El Dorado	Nov. 6, 2012	TBD
Humboldt	Nov. 8, 2011	Dec. 16-19, 2011
Madera	TBD	TBD
Marin	TBD	TBD
Merced	Nov. 8, 2011	Dec. 12, 2011
Monterey*	May 3, 2011	May 6, 2011
Napa	Jun. 5, 2012	TBD
Orange*	Mar. 8, 2011	Mar. 14, 2011
Sacramento	TBD	TBD
San Diego	TBD	TBD
San Francisco	TBD	TBD
San Luis Obispo	Aug. 30, 2011	Sept. 12, 2011
Santa Cruz	TBD	TBD
Stanislaus	Nov. 8, 2011	Dec. 2, 2011
Sutter	Nov. 6, 2012	TBD
Ventura	Nov. 8, 2011	Nov. 29, 2011
Yolo	TBD	TBD
Yuba	TBD	TBD
TBD*	TBD	TBD

*Monterey and Orange Counties volunteered to serve as pilot counties prior to the award being issued. Both are willing to participate again in 2012.

University of California

After the EAC grant was awarded, the SOS contracted with UC to engage Stark as lead researcher for the project.

Advisory Panel

At the outset of the pilot program, the SOS established an advisory panel of the following experts, advocates, and community activists in the field of election auditing and reform:

Dean Logan
Registrar-Recorder/County Clerk, Los Angeles County

Pam Smith
President, Verified Voting

Joseph Lorenzo Hall
Postdoctoral Research Fellow, New York University Department of
Media, Culture and Communication

Hovav Shacham
Assistant Professor, University of California, San Diego, Department of
Computer Science and Engineering

Mark Halvorson
Director and Founder, Citizens for Election Integrity Minnesota

Susannah Goodman
Director, Common Cause National Campaign for Election Reform

Conducting Ballot-Level Risk-Limiting Audits Using a Parallel Scan and Tally

For most election audits, the results of a hand tally are compared to the results recorded by the voting system. For California's 1% manual tally, elections officials hand tally entire precincts of ballots and compare those hand tally totals to the machine-tallied totals generated by the voting system.

In order for risk-limiting audits to be efficient, they must be conducted at the individual ballot level, not at the precinct level. A ballot-level audit compares the result tallied by the voting system for a given ballot to a hand tally of the same ballot. To conduct a risk-limiting audit at the ballot level, two things are necessary: 1) the voting system must have a cast vote record (CVR) for each ballot. A CVR is a line of data that shows how the votes on a given ballot were actually tallied by the voting system; and 2) elections officials must be able to match a CVR to the corresponding physical ballot, which requires keeping ballots and CVRs in the same identical order.

The pilot project team conducted a series of conference calls with voting system vendors to determine the capabilities of existing voting systems. Through these calls and discussions with participating counties, the team determined that none of the voting

systems in use in California is capable of exporting CVRs that can be associated with corresponding physical ballots.

For this reason, the team has conducted some of the audits for this pilot program by means of a parallel scan and tally of the votes. A parallel scan and tally is a second tally of the ballots, using commercial-off-the-shelf (COTS) scanners and open source tally software, which was developed during spring and summer 2011 for the pilot program.

County elections officials scanned the ballots using a COTS scanner and either marked the ballots or kept the ballots in order to permit each physical ballot to be paired with its scanned ballot image. This method allowed auditing the interpretation of individual ballots rather than auditing vote subtotals for entire precincts. Making individual ballots auditable – i.e., creating auditable “batches” of one ballot each – brings incredible efficiency, as described above. The hand counting work load for a ballot level audit can be smaller than the workload of a precinct level audit by a factor of 1,000 or more. Since the parallel tally for each audit showed the same results (winners and losers) as the official voting system, the audit was able to confirm the official results *transitively* (i.e., If $A = B$, and B is correct, then A is correct).

Conducting “Ballot Polling” Audits

Stark developed a new risk-limiting audit model for the pilot program, called a “Ballot Polling” audit. The advantage of the Ballot Polling audit model is that elections officials need only the overall election results – not precinct subtotals or individual ballot results (CVRs). The Ballot Polling audit model becomes inordinately time consuming if the margin of victory is very small. The hand count workload for Ballot Polling audit grows rapidly as the margin shrinks. An audit that checks a sampling of individual ballots against the voting system’s results for the same ballots is more efficient, but requires knowing how the voting system interpreted those ballots.

To perform a Ballot Polling audit, physical ballots are selected at random and interpreted by human eye. This selection continues until the number of ballots in the sample is sufficiently high – and the spread between winners and losers is sufficiently similar to the official results – to give strong statistical evidence that a full hand count would show the same winners and losers as the voting system.

Ballot Polling makes its own statistical assessment of who won directly from a random sample of ballots compared to the overall election result totals. This kind of auditing was developed and tested for the first time during the pilot program in the Monterey County audit (see below). Ballot Polling audits may be an excellent way to efficiently confirm large contests that cross two or more counties as well as contests on the statewide ballot, provided the margins of victory in the contests audited are relatively large.

Web-Based Tools and Instructions

The pilot program team, led by Stark, has honed the audit models and developed a set of web-based tools (statistics.berkeley.edu/~stark/Vote/auditTools.htm) and instructions (attached) designed for elections officials to use to conduct risk-limiting audits. The tools explain how the audits work and show the math that the tools implement, so elections officials and the public can understand risk-limiting audits. These tools are continually being refined and improved as a part of the pilot program.

Audits Conducted Through November 30, 2011

The pilot program team conducted successful risk-limiting audits in 4 counties prior to the November 30, 2011, close of the EAC's semi-annual reporting period. For each audit, the team worked with participating counties and voting system vendors to plan the audits. Stark traveled to each county and provided on-site assistance to jurisdictions carrying out the audits, including performing all computations and helping with the manual tally as required.

In some counties, the team successfully conducted simultaneous audits of several contests at one time. The simultaneous audits proved very efficient where the contests audited overlapped completely (or almost completely) in jurisdiction. The team found that unless there are large margins in the contests to be audited, conducting audits of each contest separately was more efficient if the jurisdictions among the contests did not overlap. The team plans to develop a web tool that will allow elections officials to enter contest data to determine whether it is more efficient to audit two or more contests simultaneously or separately.

The time it took to conduct the audits was minimal – a few minutes to a few hours – compared to the time it takes to conduct the 1% manual tally. However, counties participating in the pilot audits had to spend a considerable amount of time scanning the ballots in preparation for the audits. Each county used a regular scanner, rather than a set of high speed scanners. For the 2012 audits, the team plans to help counties rent high speed scanners to minimize the time spent scanning ballots in preparation for the audits.

1. Orange County: March 14, 2011

The first pilot audit followed an election in Orange County, California. The election was March 8, 2011, and the audit took place on March 14, 2011. The contest audited was a special election for San Clemente Measure A, Playa del Norte Commercial Development Project. There were 17,823 ballots cast, with 42.8% voting Yes and 57.2% voting No. Orange County uses Hart BallotNow v. 3.3.11 and the Hart eSlate v. 4.2.13 for polling place voting.

This audit was conducted as follows:

1) Ballots cast on eSlate direct recording electronic (DRE) voting machines were audited at the DRE level by randomly selecting eSlate DRE machines. Auditors hand counted the voter-verified paper trails (VVPATs) from the selected DREs and matched the totals against the DRE generated totals.

2) Ballots cast on paper (i.e., vote by mail ballots and polling place ballots cast using paper instead of a DRE) were audited at the ballot level, because the Hart system was able to generate a cast vote records (CVR) for each paper ballot and county elections officials kept the paper ballots in the order they were fed into the Hart system.

The initial sample size was:

1) 12 randomly selected Hart eSlate machines for a total of 446 ballots;

2) 21 individual paper ballots

Overall, 467 of the 17,823 ballots cast were manually reviewed and tallied for this audit. No errors were found, meaning the hand tally of these ballots matched the machine tally of these same ballots exactly.

The hand counting burden for this audit was relatively high. The 467 ballots hand tallied represented about 2.5% of all ballots cast. This was because part of the audit had to be conducted at the DRE machine level, not the ballot level. The Hart eSlate DREs did not have the capability to produce CVRs that could be associated with each VVPAT, so entire DRE machines had to be selected and all ballots cast on the DREs hand counted.

If all of the votes had been cast on paper, the entire audit could have been conducted at the ballot level, which would have required the hand-counting of only roughly 33 ballots, about one-tenth of one percent (0.10%) of all ballots cast.

2. Monterey County: May 6, 2011

The second pilot audit followed an election in Monterey County, California. The election was May 3, 2011, and the audit took place on May 6, 2011. The contest audited was a special all-mail election for Monterey Peninsula Water Management District Director, Division 1. Monterey uses the Sequoia Optech 400-C/WinETP v. 1.12.4 voting system with the Sequoia AVC Edge Model II v. 5.0.24 for accessibility in polling places. Two candidates, Brenda Lewis and Thomas M. Mancini, were on the ballot, along with write-in candidates. There were 2,111 ballots cast in all. The reported totals were 1,353 votes for Lewis, 742 for Mancini, and 13 for various write-in candidates. The remaining 3 ballots were recorded as undervotes or overvotes, and as a result, those ballots were voided in the official count. According to the voting system results, Lewis received 64% of the valid votes, while Mancini received 35% of valid votes.

The audit was a “Ballot Polling” audit (see description under “Background” above) which relied only on comparing the margin of victory in the overall election results to the margin of victory in a hand tally of a sample of randomly selected ballots.

The sample size was 89 ballots. The Ballot Polling audit was designed to ensure that if Lewis had at least 64% of the vote, there was at most a 1% chance that the audit would lead to a full hand count (i.e., 99% confidence level). The audit took about 90 minutes, including the time Stark spent explaining the audit procedure to public observers. Public observers helped roll the dice used to select ballots at random and had an opportunity during the audit to confirm that their interpretation of ballots agreed with the auditors’ hand tally.

3. San Luis Obispo County: September 12, 2011

The third pilot audit followed a special election in San Luis Obispo County, California. The election was August 30, 2011, and the audit took place on September 12, 2011. Both contests on the ballot, City of San Luis Obispo Measures A-11 and B-11, were audited. San Luis Obispo County uses Premier AccuVote-OS v. 2.0.12 with AutoMARKs for accessibility.

This was the first simultaneous risk-limiting audit of two contests for the pilot program. The method Stark used was the “Super-Simple” model (described above under “Background”), because it uses a relatively simple, easy-to-understand mathematical formula* to determine the initial sample size and confirm election outcomes. The audit involved a random sample of just 16 ballots, and was finished in one hour, confirming the winner of each measure. There were 10,689 ballots cast in the election, and the narrower of the two margins of victory in the contests was 45%.

Stark and San Luis Obispo County elections officials performed the audit in front of 10 public observers. The public was able to see, hear, and compare the manual tally results for each ballot audited against the cast vote record for the ballot. All 16 ballots matched the votes tallied from the cast vote records, so no escalation was required.

The audit was extremely efficient because Stark conducted a ballot-level audit – i.e., selected individual ballots for the sample – instead of whole precincts. Again, the sample size for this audit was 16 audit units – regardless of whether those audit units are individual ballots, small batches, or entire precincts. The audit confirmed the winner with 90% confidence (10% risk limit). Public notice and observation are built into the audit process. The model calculates the size of the initial audit sample based on the margin of victory.** The smaller the margin of victory is, the larger the initial sample must be.

The audit stops after the initial sample if auditors find zero or only statistically insignificant differences between the voting system interpretation and a human eye

interpretation of the votes on the ballots. The number of differences permissible depends on the margin and the size of the initial sample. As discussed above, risk-limiting audits can lead to a full hand count, if enough statistically significant differences are found.

*The drawback to the “Super-Simple” model is that the formula is slightly less efficient (i.e. leads to a slightly larger initial sample size) than more complex and precise models.

**The margin used for risk-limiting audits is the “diluted margin of victory,” which is the margin in votes divided by the number of ballots cast in the contest(s) audited. This is slightly different from the usual “margin of victory” which is the number of votes between the winner and loser divided by the valid votes cast. The diluted margin is used to ensure all ballots cast are subject to being audited, not just those ballots deemed by the voting system to be valid votes (i.e., an overvote or undervote would be a “ballot cast” by a voter, but it would not be a “valid vote” because a vote could not be recorded for any specific candidate or contest).

4. Ventura County: November 29, 2011

The fourth pilot audit followed an election in Ventura County, California. The election was November 8, 2011, and the audit took place on November 29, 2011. The contest audited was the City of San Buenaventura City Council, for which there were three at-large seats to be filled. Ventura uses the Sequoia Optech 400-C/WinETP v. 1.12.4 voting system, the Sequoia AVC Edge Model II v. 5.0.24 for accessibility, and the Sequoia OptechInsight APX K2.10 HPX K1.42 in polling places.

This was the first multi-winner contests to be audited using the Super-Simple method. The audit was successful: The election outcome was confirmed by looking at just 90 individual ballots.

Stark developed a draft set of web-based auditing tools and tested the web tools for the first time in the Ventura audit. There were 11 candidates in this vote-for-three contest. Official results showed the winners to be Cheryl Heitmann with 7,090 votes, Carl E. Morehouse with 6,793 votes, and Christy Weir with 6,515 votes. The runner up was Kenneth M. Cozzens, with 5,564 votes. There were 17,376 ballots cast in all.

To prepare for the audit, Ventura County staff scanned all of the paper ballots* cast in the election to produce digital images. The digital images were processed using ballot tally software developed for the pilot program. The software created a CVR for each ballot and tallied the votes on the digital images of the ballots. Ventura County staff kept the ballots in the physical order in which they were scanned so the CVRs could be associated with the paper ballots they represented. The ballots were organized into batches of a maximum of 50 before scanning, to make it easier to find individual ballots.

The initial sample size was 90 individual ballots. The 90 ballots were retrieved and compared to the CVRs. All ballots matched their CVRs exactly, so the audit stopped and the election outcomes were confirmed with 90% confidence (10% risk limit).

*Three ballots could not be located. Since the ballots could have contained votes for the loser, auditors treated the ballots as such, slightly narrowing the margin of victory calculation for the audit, which affects the initial sample size calculation. This ensured that the initial sample size took into account that the ballots may have been cast for the loser.