

California Secretary of State
Post-Election Risk-Limiting Audit Pilot Program 2011-2013
Final Report to the United States Election Assistance Commission

Table of Contents

Acknowledgements	
Executive Summary	
Section I: Overview of the Program.....	1
Why Audit Election Results?	
Current Election Audit Law Needs an Overhaul	
What is a “Risk-Limiting” Audit?	
Pilot Program Goals	
Advisory Panel	
Participating Counties	
Section II: The 2011 Pilot Audits.....	8
Conducting Risk-Limiting Post-Election Audits: Comparison Audits	
Conducting Risk-Limiting “Ballot-Polling” Audits of Election Results	
Using a Parallel Scan and Independent Tally	
Auditing Using Small Batches with Corresponding Vote Totals	
Development of Web-Based Tools and Instructions for Elections Officials	
The 2011 Audits	
Section III: The 2012 Pilot Audits.....	19
Section IV: 2013 Project Wrap Up.....	20
Ballot Accounting Best Practices	
Marin County Multi-Contest Audit Completion	
Section V: Recommendations for Modifications to Existing Voting Systems.....	21
Section VI: Cost-Effectiveness: Risk-Limiting Audits Versus the 1% Manual Tally.....	22
Effectiveness	
Cost	
Improving Cost-Effectiveness	
Section VII: Products of the Pilot Program.....	23
Audit Tools	
Step-by-Step Audit Instructions	
Ballot Accounting Best Practices	
OpenCount Software	

Legislation and Regulations to Require Risk-Limiting Audits in Pilots of New Voting Systems

Section VIII Next Steps.....25

- The Legislature Should Revise California’s Election Audit Laws
- The Secretary of State Should Approve Random Selection Methods
- Counties Should Test Cross-Jurisdictional Audits
- California Should Standardize Ballot Accounting Methods

Section IX Conclusion.....27

Appendices

Acknowledgements

The Secretary of State would like to thank the California county elections officials who gave generously of their time to conduct pilots of risk-limiting audits following elections in 2011 and 2012. Thanks are also due to the election experts who served on the advisory panel for the pilot program and who provided valuable insight and creative ideas throughout the program. The Secretary of State would also like to recognize University of California (UC) Berkeley Professor Philip Stark for his groundbreaking work in developing risk-limiting election audits that have become a model in the United States and internationally; UC Berkeley Professor David Wagner for his development of the open source OpenCount ballot tally software used in the pilot program; Verified Voting President Pam Smith for sharing her expertise in election auditing and best practices for ballot accounting; and retired Marin County Registrar of Voters Elaine Ginnold, Yolo County Clerk-Recorder-Assessor Freddie Oakley, and Santa Cruz County Clerk/Registrar of Voters Gail Pellerin for being the first elections officials in the country to step forward to test early versions of Professor Stark's audit methods in 2008 and 2009.

Executive Summary

California's one percent (1%) manual tally law dates back to the 1960s and requires all county elections officials to randomly select 1% of all precincts after each election and hand count *all* of the votes on *all* of the ballots from those precincts. In a regular election year, counties hand count tens of thousands of ballots as part of the 1% manual tally, yet doing so provides little or no statistical evidence that the machine tally found the true winner for each contest on the ballot – and does nothing to correct any erroneous electoral outcomes.

The California Secretary of State (SOS) received a \$230,000 grant from the United States Election Assistance Commission (EAC) under Section 271 of the federal Help America Vote Act to conduct a two-year election audit pilot program during 2011-12 to test new, statistically sound, risk-limiting election audit methods. The California Legislature passed and the Governor signed AB 2023 (Saldaña), Chapter 122, Statutes of 2010, to authorize the SOS to conduct the Post-Election Risk-Limiting Audit Pilot Program.

The SOS partnered with the University of California for the pilot program to allow Berkeley Statistics Professor Philip B. Stark (Stark) – who originally developed the audit methods to be used in the pilot program – to serve as lead researcher. The EAC grant helped fund:

- Fourteen pilot audits of election results following elections in 2011 and 2012;
- Development of web-based tools and procedures for elections officials to use in conducting audits; refinement of statistical methods for conducting risk-limiting audits based on the experiences gained in the pilots; and reimbursement of county elections office costs for conducting the audits with county staff and facilities; and
- Reporting and analysis of the cost and effectiveness of risk-limiting post-election audits compared to the current statutory 1% manual tally law, the development of recommendations for modifications to current voting systems to make auditing easy, and legislation to reform current election audit law.

Eleven counties successfully completed their audits and confirmed the official election results by reviewing a relatively small number of individual ballots (e.g., a few dozen to a few hundred ballots). By contrast, the statutorily-mandated 1% manual tally conducted in the same elections provided little statistical evidence that the election outcomes were correctly tallied by the voting system, despite requiring substantially more ballots to be hand-counted and examined.

Overall, the project team met its goals to develop and test audit methods, procedures and web-based tools for conducting post-election risk-limiting audits. The audit “how-to” materials developed during the pilot program are available to elections officials in California and other jurisdictions across the United States to conduct risk-limiting audits of election results. The adoption of laws and regulations permitting or requiring risk-limiting post-election audits will allow elections officials to use the new audit methods to confirm – or correct – official election results, which will help build public confidence in elections and may reduce the need for voter-requested manual recounts.

Section I: Overview of the Program

The California Secretary of State (SOS) received a grant from the federal Election Assistance Commission (EAC) under Section 271 of the federal Help America Vote Act (HAVA) to conduct a two-year post-election risk-limiting audit pilot program during 2011–12 to test newly developed election auditing methods. The research problem for the project was how to conduct risk-limiting audits following elections, including audits of individual contests, multiple contests simultaneously, and cross-jurisdictional contests.

The SOS partnered with the University of California (UC) Berkeley for the project. UC Berkeley Statistics Professor Philip B. Stark (Stark), who developed the election audit methods used in the program, served as lead researcher. The \$230,000 two-year grant from the EAC helped fund:

- 1) A contract with the University of California (UC) to allow Stark to serve as lead researcher for the pilot program;
- 2) Partnerships with a mix of thirteen urban and rural counties to participate in the program, which included reimbursement of county elections office costs of up to \$5,000 per county for conducting the audits with the help of county staff and facilities;
- 3) Development and testing of risk-limiting audit methods following elections in 2011 and 2012, by which fourteen audits were undertaken, eleven audits were completed successfully, and three audits following the June 2012 election were not completed due to time constraints between the primary and general elections in 2012;
- 4) Development of web-based tools and procedures for elections officials to use in conducting future audits and refinement of Stark’s statistical methods for conducting risk-limiting audits based on the experiences gained in the pilots. The web-based tools allow elections officials to determine initial sample size, select ballots at random in a transparent and reproducible way, determine whether escalation of the audit is necessary based on initial audit results, and report final audit results to the public;
- 5) Reporting and analysis of the cost and effectiveness of risk-limiting post-election audits compared to the current statutory 1% manual tally law; and
- 6) Development of recommendations for modifications to current voting systems to make voting systems more readily auditable and recommendations for legislation to reform current election audit law.

The California Legislature passed and the Governor signed AB 2023 (Saldaña), Chapter 122, Statutes of 2010, to authorize the SOS to conduct the Post-Election Risk-Limiting Audit Pilot Program. See **Appendix A** for a copy of AB 2023.

Why Audit Election Results?

State and federal voting system testing and certification help ensure that voting systems can capture and tally ballots accurately and securely, while protecting voter anonymity. But front-end regulation and testing are not enough. How can the public know whether voting systems actually got the job done right in an election unless the election results are audited after that election?

History has shown 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. Voting machines misinterpret ballots; voters mismark ballots; and sometimes some ballots are not tabulated at all. Errors happen, and auditing can determine whether those errors matter, in other words, whether it is plausible that a full hand count of the ballots would show that a different candidate won the election.

Auditing methods in the banking industry and other sectors have certainly improved over the last 50 years. In the area of elections, modern auditing methods can reduce the number of ballots that need to be audited while increasing the chances of catching and correcting voting system tally errors, discovering security breaches, and reducing the need for full manual recounts. Modern election auditing methods not only ensure the accuracy of election outcomes but also improve public confidence in elections.

Current Election Audit Law Needs an Overhaul

Current California law requires county elections officials to randomly select one percent (1%) of all precincts and manually tally the votes on all of the ballots in those precincts. If a given contest is not captured as part of the 1% sample, the law requires additional precincts to be drawn until all contests are represented in the manual tally. [Elections Code section 15360] The Legislature established California's 1% manual tally in 1965, nearly 50 years ago. By law, the 1% manual tally cannot change election outcomes. To overturn official election results in California, a voter must request and pay for a costly manual recount of the ballots cast in the contest. California law contains no provision for an audit or automatic recount, even in very close contests.

What is a “risk-limiting” audit?

A risk-limiting audit can be thought of as an “intelligent incremental manual tally.” It is a manual tally of randomly selected ballots that stops as soon as it is implausible that a full recount would alter the result. As long as it is statistically plausible that a full recount would overturn the result, the risk-limiting audit continues to examine more ballots. Risk-limiting audits determine precisely how much hand counting is necessary to confirm election results to a given level of confidence. The closer the contest, the more ballots one must examine to have strong evidence – because fewer errors can change the outcome. The higher the desired confidence (e.g., 99%

versus 90%), the more ballots one must examine – because higher confidence requires more evidence.

The pilot program audits provided 90% confidence. That is, if a particular electoral outcome was wrong (that is, if the wrong winner was named), the audit had at least a 90% chance of correcting that outcome. This chance is a worst-case analysis, built on the assumption that the errors were hidden as well as possible. Key to the idea is that we are confirming *outcomes* — who won — rather than exact vote totals. Confirming *exact* vote totals for each candidate or side of a measure always requires a full hand count.

Two approaches to risk-limiting audits were tested in the pilot program: ballot-polling audits and comparison audits. A ballot-polling audit is analogous to an opinion poll. It examines randomly selected ballots until the human eye interpretation of the votes on those ballots gives sufficiently high statistical confidence that a full hand count would confirm the machine results. A ballot-polling audit only requires knowing the overall official result. It does not require precinct level results or more detailed results, which comparison audits require.

In contrast, a comparison audit compares a human interpretation of the votes on randomly selected ballots to the voting system’s interpretation of the votes on those ballots. A comparison audit continues until there is sufficiently high statistical confidence that a full hand count would agree with the machine results, despite any differences between the audit’s manual interpretation of the ballots and the voting system’s interpretation of those ballots.

Comparison audits require detailed information from the voting system. Maximum efficiency requires the voting system to report how it interpreted each physical ballot, in a way that allows the interpretation of any particular ballot to be checked. When that is possible, a comparison audit generally requires examining fewer ballots than a ballot-polling audit. However, if the voting system can only report tallies for large groups of ballots, such as precincts, comparison audits may have little or no advantage over ballot-polling audits, which place much lower demands on the voting system and which have far smaller set-up costs. Both ballot-polling and comparison audits require a ballot manifest that says how many ballots there are in all and how they are stored, e.g., that lists identifiable groups of ballots and says how many there are in each group.

By definition, risk-limiting audits can lead to a full hand count to confirm whether the machine tally of ballots correctly determined the election winner. Indeed, they are guaranteed to have a large chance of progressing to a full hand count whenever that full hand count would show that the machines found the wrong winner or winners. The methods tested under this grant “escalate” gradually, randomly selecting additional ballots for manual inspection until either (i) there is strong statistical evidence that the outcome is correct or (ii) there has been a full hand count, which confirms (or overturns) the machine results.

Efficient risk-limiting audits involve examining individual ballots selected at random from the jurisdiction or the contest being audited. That is dramatically more efficient than hand counting

all the ballots in precincts selected at random. If there is a problem with the results, examining a sample of individual ballots guarantees a greater chance of discovering that problem than examining a much larger number of ballots that come from entire precincts can guarantee.

With risk-limiting audits, individual contests or groups of contests can be audited at the same time, using the same sample of ballots, and the winners of all those contests are confirmed by looking at relatively few individual ballots.

For the simplified version of the risk-limiting audit method, called the “super-simple” audit, which was used for some of the audits in this pilot program, the size of the initial sample for 90% confidence (10% risk limit) is the number 4.8 divided by the “diluted” margin of victory.¹ Table 1 gives examples. The hand count burden is the same regardless of whether the contest is small or large (assuming the audit finds no errors). In other words, the initial sample size for a small town city council race or a statewide measure with an apparent 10% margin of victory would be 48 ballots.

If the initial sample size is an appreciable fraction of the total number of ballots cast, it may be less time-consuming simply to conduct a complete hand count of all the ballots cast in the contest; that depends in part on the number of ballots cast in the contest. For instance, for a margin of 0.1%, the initial sample size is 4,800. If only 20,000 ballots were cast in the contest, it might well be easier to conduct a full hand tally than to select 4,800 ballots at random and inspect them, especially since there is some chance that the audit will need to expand beyond the initial sample. On the other hand, if there were 13 million ballots cast in the contest (roughly the number cast in the 2012 Presidential Election in California), the effort of auditing 4,800 randomly selected ballots would be *de minimis* compared to the effort of conducting a full hand tally or recount.

The risk-limiting audit methods used in the pilot program were first developed and tested by Stark in 2009. Stark has been a pioneer in developing effective election audit methods and his methodology has been endorsed by national organizations, including Common Cause, the League of Women Voters, the American Statistical Association, Verified Voting, and other groups involved in election integrity. Risk-limiting audits were also endorsed by the 2013 Presidential Commission on Election Administration.

¹ If elections officials expect to see errors, it can be efficient logistically to draw a larger initial sample. The “diluted” margin is the number of votes that separates the winner with the fewest votes in a contest from the loser with the most votes in that contest, divided by the total number of ballots cast, including invalid votes such as undervotes and overvotes. This differs from the usual “margin of victory,” which is the number of votes between the winner and loser divided by the number of *valid votes* cast. The diluted margin takes into account the fact that the voting system might have misinterpreted a valid vote as an undervote or overvote, or vice versa. When a collection of contests is audited simultaneously, the diluted margin involves dividing the smallest numerical margin (in votes) by the total number of ballots cast in the largest of the contests.

Table 1: Initial sample sizes for "super-simple" comparison audits

Diluted 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
.1%	4.8/.001	4,800

Pilot Program Goals

The program goals were:

- 1) To develop and test post-election risk-limiting audit methods, audit procedures and web-based audit tools by conducting pilot audits following live elections in California;
- 2) Help California and other states develop more efficient and effective election audit laws;
- 3) Inform the design of next-generation voting systems to ensure that they are efficiently auditable;
- 4) 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
- 5) Build public confidence that if there are errors in election results, those errors will be caught and corrected by audits before the results are final.

Advisory Panel

The project advisory panel was comprised of the following elections officials, experts, and advocates in the field of election reform:

Susannah Goodman
Director, Common Cause National Campaign for Election Reform

Joseph Lorenzo Hall
Chief Technologist, Center for Democracy and Technology

Mark Halvorson
Director and Founder, Citizens for Election Integrity Minnesota

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

Margaret MacAlpine
Independent Security Consultant

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

Pamela Smith
President, Verified Voting

Participating Counties

Twenty counties initially volunteered to participate in the program. Ultimately, 14 pilot audits were undertaken in 13 counties and successfully completed in 11 California counties following elections held during 2011–2012. Eight of the audits were conducted following small local elections held in 2011. Two audits were conducted in small counties, Madera and Napa, following the June 2012 Statewide Presidential Primary Election. Four multi-contest audits were started in Marin, Orange, Santa Cruz, and Yolo counties in July 2012 but had to be terminated before completion due to technical difficulties and time constraints between the primary and general election. The multi-contest audit attempted in Orange County faced an additional challenge: a voter-requested recount in one of the contests on the June 2012 ballot which took precedence over the pilot audit. The Marin County audit was ultimately successfully completed in February 2013.

Below is a chart of the counties and audits conducted for the pilot program.

Table 2: Audits Conducted in the Pilot

County	Election	Audit Completion
Alameda	Nov. 8, 2011	December 2011
Humboldt	Nov. 8, 2011	December 2011
Madera	June 5, 2012	September 2012
Marin	June 5, 2012	February 2013
Merced	Nov. 8, 2011	December 2011
Monterey	May 3, 2011	May 2011
Napa	June 5, 2012	July 2012
Orange*	Mar. 8, 2011	March 2011
Orange	June 5, 2012	Terminated mid-audit due to voter-requested recount of election results.
San Luis Obispo*	Aug. 30, 2011	Sept. 12, 2011
Santa Cruz	June 5, 2012	Terminated mid-audit due to technical delays and the need to prepare for general election.
Stanislaus	Nov. 8, 2011	Dec. 2, 2011
Ventura	Nov. 8, 2011	Nov. 29, 2011
Yolo	June 5, 2012	Terminated mid-audit due to technical delays and the need to prepare for general election.

* Two counties, Orange and San Luis Obispo, participated in the pilot program but conducted their audits prior to the EAC grant award date, so the expenses related to those counties' audits were absorbed by the counties. The remaining audits were funded through the EAC grant.

Four additional counties, Alpine, Colusa, El Dorado, and Sacramento, had volunteered to participate in November 2012 but were unable to due to difficulties the project team experienced with the parallel tally software following the June 2012 election.

Section II: The 2011 Pilot Audits

The project team conducted successful risk-limiting audits in eight counties during 2011. For each audit, the team worked with participating counties and (in some cases) voting system vendors to plan the audits. Stark traveled to all but one of the counties to provide on-site assistance to jurisdictions carrying out the audits, including performing all computations and helping with the random draw and hand count 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. If the jurisdictions among the contests do not overlap substantially, it is more efficient to audit contests separately, drawing samples of ballots from each contest. The web-based audit tools developed as part of the pilot program automates the basic calculations needed for elections officials to enter contest data and determine immediately whether it is more efficient to audit several contests simultaneously or separately.² The tools are transparent and open, allowing anyone to verify that the underlying software correctly implements the methodology.

Conducting Risk-Limiting Post-election Audits: Comparison Audits

As discussed above, comparison audits compare the results of a hand tally of randomly selected ballots or groups of ballots to the results recorded by the voting system for those ballots. California's 1% manual tally is a comparison audit (but not a risk-limiting audit). In the 1% tally, elections officials hand tally 100% of the ballots in 1% of the precincts and compare the results to the subtotals generated by the voting system for those precincts.³

The most efficient risk-limiting audits are those conducted at the individual ballot level, not at the precinct level, which means that each ballot must be subject to the random selection, review and comparison to the voting system's record of the votes on that ballot. Essentially, a ballot-level comparison audit *compares* the voting system's interpretation of a given ballot to a human eye review of the same ballot.

To conduct a risk-limiting *comparison* audit at the ballot level, two things are necessary:

- 1) the voting system must report a cast vote record (CVR) for each ballot. A CVR shows how the marks on a given ballot were actually interpreted as votes by the voting system; and

² The web-based tool was also used for a risk-limiting audit in Arapahoe County, Colorado, funded by an EAC grant to the Colorado Secretary of State under the same grant program.

³ The 1% tally gives essentially no useful information about the accuracy of contests that use instant-runoff or ranked-choice voting methods. It gives a limited amount of information about the accuracy of plurality, majority, and super-majority contests.

- 2) elections officials must be able to match each CVR to the corresponding physical ballot, and vice versa. Generally, this involves either marking ballots as they are tallied or keeping ballots in the same order in which they are tallied, although those are not the only possible means.

For further discussion, see Section V: Recommendations for Modifications to Existing Voting Systems

Conducting Risk-Limiting “Ballot-Polling” Audits of Election Results

During the pilot program, Stark developed a new risk-limiting audit method, called a “Ballot-Polling” audit. The advantage of Ballot-Polling audits is that elections officials need only the overall election results, not precinct subtotals or individual ballot results (CVRs). A ballot-polling audit makes no demands on the vote tabulation system, but when the contest outcome is correct, it does require examining more ballots than a ballot-level comparison audit—especially when margins are small.⁴ As discussed above, a ballot-level risk-limiting *comparison audit* compares the voting system’s interpretation of a given ballot to a manual interpretation of the same ballot.

Ballot-Polling is economical for contests with wide margins, but can become inordinately time consuming if the margin of victory is very small. The hand count workload for ballot-polling audits grows rapidly as the margin shrinks. Ballot-level comparison audits are more efficient for small margins, but such audits require knowing how the voting system interpreted every ballot.

To perform a ballot-polling audit, physical ballots are selected at random and interpreted by the human eye. This selection continues until the sample gives sufficiently high confidence that the winners according to a full hand count would be the same as the winners according to the original machine count.

Ballot-polling audits make their own statistical assessment of who won directly from a random sample of ballots. This kind of audit 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 the outcome of large contests, such as county-wide or statewide contests, especially if the margin of victory is not too small. They do not require any change to current voting technology, and their set-up costs are minor.

⁴ It may require examining fewer ballots than a precinct-level comparison audit. And compared to ballot-level comparison audits, the workload of a ballot-polling audit may be smaller, because its set-up costs are much lower.

“Transitive auditing” Using a Parallel Scan and Independent Tally

Early in 2011, 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 conducted most of the audits for this pilot program by means of a parallel scan and an independent tally of the votes. A parallel scan and independent tally is essentially a second machine count of all ballots using commercial off-the-shelf (COTS) scanners and (in this case) open-source tally software. Fortunately, during the spring and summer of 2011 a team of University of California (UC) researchers led by UC Berkeley Professor David Wagner (Wagner) were developing open source tally software (OpenCount⁵) as part of a separately funded project. Wagner offered early versions of the software to the project team for use during the pilot program. Unlike commercial, federally certified vote-tabulation systems currently in use in California, the OpenCount is capable of reporting a CVR for every ballot in a way that the CVR can be associated to the physical ballot, and vice versa.

For the audits, 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 and the OpenCount CVR constructed from that image. Kai Wang, a PhD student working to develop OpenCount with Wagner, helped operate the OpenCount software to produce a CVR for each ballot. This allowed auditing the interpretation of individual ballots rather than subtotals for entire precincts. Making individual ballots auditable – i.e., creating auditable “batches” of one ballot each – brings very significant 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. If the parallel tally for each audit shows the same winner(s) as the official voting system, the audit can confirm the official results *transitively* (i.e., if A and B agree, and B is correct, then A must also be correct).

Auditing Using Small Batches with Corresponding Vote Totals

Auditing smaller batches is more efficient and cost-effective than auditing entire precincts of ballots. The team tested ways to reduce batch sizes for risk-limiting audits. For example, in Orange County, small batches were created by using the ballots cast on each direct recording electronic (DRE) voting machine as one batch. This helped, but still required auditing hundreds of times more ballots than would have been required if the system had been capable of reporting CVRs for individual ballots.

The project team also experimented with creating stacks of 50 or 100 ballots, feeding those ballots into the voting system, and then re-setting the voting system in order to create an interim

⁵See <https://code.google.com/p/opencount/>

vote tally for each batch. This process was very cumbersome and error prone; moreover, when the electoral outcome is correct, it still requires 50–100 times more hand counting than a ballot-level comparison audit.

The project team concluded that the gains from reducing batch sizes are not worth pursuing unless the batch size can be reduced to an individual ballot. Since current voting systems cannot report a cast vote record (CVR) for each ballot and counties would need to develop procedures to link each CVR to the corresponding physical ballot for the audit, the project team used a parallel scan and independent tally to conduct “transitive” ballot-level comparison audits.⁶

Development of Web-Based Tools and Instructions for Elections Officials

The project team developed a set of web-based tools (<http://statistics.berkeley.edu/~stark/Vote/auditTools.htm>) and step-by-step instructions for elections officials to conduct risk-limiting audits. The website page can be expanded to see explanations of how the audit works, the math behind the tools, and the computer code that implements the tools, so elections officials and the public can better understand risk-limiting audits and can check for themselves that the software implementation is correct. These resources were continually refined and improved during the pilot program.

The 2011 Audits

Following is a summary of each audit conducted during 2011:

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 DRE for polling place voting.

This audit was conducted as follows:

- 1) Ballots cast on eSlate direct recording electronic (DRE) voting machines were audited by creating small batches. Each batch consisted of the votes cast on a single DRE. The project team hand counted the voter-verified paper audit trail (VVPAT) from randomly selected DREs and compared the totals to 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 by selecting individual ballots and comparing the marks

⁶ One ballot-polling audit was conducted during the project, and several comparison audits were conducted using precinct-level batches, prior to receipt of EAC funding.

on the physical ballot with the voting system results. This was possible because the Hart system generated a CVR for each paper ballot, and ballots within precincts were identified by a numerical code and barcode that was effectively unique, linking the CVR to the ballot.

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: 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 the audit was conducted, in part, using batches of ballots (the ballots cast on each DRE machine served as one batch). The Hart eSlate DREs could not produce individual CVRs that could be associated with each voter's selections, so the entire VVPAT on the selected DRE machines had to be tallied by hand, an unwieldy and time-consuming process.

If all of the votes had been cast on paper ballots, the entire audit could have been conducted at the ballot level, which would have required manually examining 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 team had originally planned to conduct a ballot-level comparison audit. To prepare, Monterey County staff Bates-stamped every ballot cast in the contest with a unique identifying number, then scanned the ballot to make digital images to be processed to create CVRs. Unfortunately, the set of images was not ready until the day of the audit, and the software that

was to be used to interpret the scans (TEVS⁷) did not perform as well on the actual images as it had on the test images. There was no time on the day of the audit to tune the software to perform better.

Therefore, instead of conducting a ballot-level comparison audit, the team conducted a **“ballot-polling” audit** (see description above), which relies only on manually interpreting the votes on a random sample of ballots. It requires knowing how many ballots were cast in all and how to find each ballot, but it does not require knowing how the voting system interpreted any subset of the ballots.

Confirming the outcome with 90% confidence required examining 89 ballots selected at random. 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 (pointless) full hand count. 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 they agreed with the audit team’s manual interpretation of each audited ballot.

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. The county uses Premier AccuVote-OS v. 2.0.12 with AutoMARKs for accessibility.

This audit relied on the OpenCount system to construct a CVR for every ballot. This was the first simultaneous risk-limiting audit of two contests. The method Stark used was the “super-simple” method (described above under “Background”), because it uses a relatively simple, easy-to-understand mathematical formula⁸ 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 winners of both measures. There were 10,689 ballots cast in the election, and the narrower of the two margins of victory in the contests was 45%.

Stark, Wang, and San Luis Obispo County elections officials performed the audit in front of 10 public observers, including a media representative. The public was able to see, hear, and

⁷ <https://code.google.com/p/tevs/>

⁸ The formula only requires division, and allows the ballots to be selected at random with equal probabilities. There are methods that are more efficient in that they require smaller sample sizes, but they involve more complicated formulae and can require drawing ballots with different probabilities, depending on the CVR for each ballot. The project team concluded that operational simplicity and transparency was worth some statistical inefficiency, especially since the resulting sample sizes are still quite manageable.

compare the manual interpretation of each ballot audited against the CVR for the ballot. All 16 ballots matched the CVRs, so no escalation was required.

The audit was extremely efficient because it was conducted at the ballot-level – i.e., it compared randomly selected individual ballots to the CVRs for those ballots – rather than selecting entire precincts and comparing a manual tally of the votes in those precincts to machine subtotals for those precincts. The audit confirmed the winner with 90% confidence. Public notice and observation were built into the process.

The initial sample size was set to be large enough that if the audit found no errors in the initial sample, the audit could stop. If the audit had found errors that had inflated the margin, it would have been necessary to audit more ballots selected at random. As discussed above, risk-limiting audits can lead to a full hand count unless the sample provides strong evidence that a full count will merely confirm that the original outcome is correct.

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 contest to be audited using the super-simple method. The audit was successful: The election outcome was confirmed by looking at 90 individual ballots.

Stark developed a set of web-based auditing tools and tested those 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 ballots⁹ cast in the election to produce digital images. The digital images were processed by Kai Wang using the

⁹ Three ballots were not scanned for the audit. Due to time constraints, these three ballots were not added to the scanned ballots prior to providing the ballot images to the project team. Although the three ballots alone could not change the election outcome, since the ballots could have contained votes for the loser, project team treated the ballots as such, slightly narrowing the margin of victory calculation for the audit, which affected the initial sample size calculation. This ensured that the initial sample size took into account the fact that the missing ballots might have been cast for the loser.

OpenCount ballot tally software. The software created a CVR for each ballot and tallied the votes reflected by those CVRs. Ventura County staff kept the ballots in the physical order in which they were scanned so that 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.

Prior to the Ventura audit, Ventura County staff had technical difficulties uploading ballot images for the parallel tally, because of the high bandwidth required. For this reason, Ventura county staff took the extra step of personally delivering a hard drive containing the ballot images by car to UC San Diego for processing using OpenCount. The time spent on this delivery is included in Ventura County's audit costs.

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).

5. Stanislaus County: December 2, 2011

Stanislaus County conducted a risk-limiting audit of City of Oakdale Measure O, in which 3,152 ballots were cast. To prepare for the audit, Stanislaus County staff rented a scanner for a day and scanned all of the paper ballots cast in the election to produce digital images. (One ballot could not be located for scanning; it was treated as a "no" vote by the audit, to ensure that the audit was conservative.) Stanislaus 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 for scanning to make it easier to find individual ballots. The digital images were processed using OpenCount software. The software created a CVR for each ballot and tallied the votes on those CVRs. According to the software, there were 1,728 "yes" votes and 1,391 "no" votes, a margin of 336 with the missing ballot treated as a "no." This corresponds to a diluted margin of $336/3152 = 10.6\%$.

The web tools at <http://statistics.berkeley.edu/~stark/Vote/auditTools.htm> were used to determine an initial sample size for an audit at 10% risk limit, which turned out to be 49 ballots. A seed for the random number generator was selected by drawing film canisters containing numbered slips of paper at random from an opaque bag. The web tools were then used to select the ballots to audit. The human eye interpretation of all 49 ballots matched the CVRs for those ballots, so the audit stopped. It took approximately 1 hour and 5 minutes to conduct the audit.

The statutory 1% audit required a hand tally of all the ballots cast in one of the five precincts that contained the contest. The precincts ranged in size from 452 ballots cast to 792 ballots cast. The average number of ballots – the expected number of ballots the 1% audit would require tallying in this contest – was 630 ballots. Even though the 1% audit examined far more than the 49 ballots the risk-limiting audit examined, the statutory 1% manual tally could have had a chance as large as 80% of not finding a single error even if the machine-count winner had been wrong. In contrast, the risk-limiting audit had a 90% chance of requiring a full hand count if the

machine-count winner had been wrong. Again, this shows the power and efficiency of risk-limiting audits compared to the current statutory 1% audit.

6. Alameda County: December 5, 2011

To prepare for the audit, Alameda County staff used a small county scanner to scan all of the paper ballots cast in the election to produce digital images. Before scanning the ballots, county staff stamped each ballot with an identification number to make it easier to associate CVRs with the physical ballots. The digital images were processed using OpenCount ballot tally software.

The software created a CVR for each ballot and tallied the votes on those CVRs. Even though the numbering would have sufficed, Alameda County staff also kept the ballots in the physical order in which they were scanned to make it easier to associate CVRs with the paper ballots they represented. The ballots were organized into batches for scanning to make it easier to find individual ballots.

Four City of Alameda contests were audited simultaneously: City Council (vote for 3 of 5) and three measures. All votes were cast on paper ballots; 1,374 ballots were cast in all. The OpenCount software found one extra vote for Bukowski for City Council (409 versus 408) and one extra “no” vote for measure F (841 versus 840) compared with the official tally. The web-based audit tools developed for the project were used to determine an initial sample size for an audit at 10% risk limit: 17 individual ballots to be selected at random from the 1,374. Numbered ping-pong balls were drawn at random from a bingo-like tumbler by county staff to generate a seed for the random number generator in the web-based tool. The 17 ballots were retrieved and compared to the CVRs. All 17 ballots matched their CVRs, so the audit stopped without escalation. Two members of the public observed the audit, which took approximately 25 minutes.

7. Merced County: December 12, 2011

Two City of Merced contests were audited simultaneously, Mayor and City Councilmember (vote for 3 of 8). A total of 7,321 ballots were cast in these contests. The reported winner in the mayoral contest was Stan Thurston, with 2,231 votes; the runner-up was Bill Blake with 2,037 votes. The three reported winners of the City Council contest were Noah Lor (3,736 votes), Mark “Tony” Dossetti (3,669 votes) and Mike Murphy (3,375 votes); runner-up was Richard L. Cervantes (2,416 votes). The diluted margin for the two contests was $(2231 - 2037)/7321 = 2.6\%$, the smallest diluted margin among contests that had been audited under the pilot.

To prepare for the audit, Merced County staff used an office scanner they owned to scan all of the paper ballots cast in the election to produce digital images. The digital images were processed by Kai Wang using OpenCount software, which created a CVR for each ballot and tallied the votes on those CVRs. Merced 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 for scanning to make it easier to find individual ballots.

The tools at <http://statistics.berkeley.edu/~stark/Vote/auditTools.htm> were used to determine an initial sample size for an audit at 10% risk limit (which turned out to be 198 ballots), to draw the random sample, and to locate the selected ballots within bundles of stored ballots. The human eye interpretation of all 198 ballots matched the CVRs for those ballots, so the audit stopped without escalation. It took about 3 hours and 15 minutes to conduct the audit.

8. Humboldt County: December 16, 2011

Humboldt County was the first in the program to conduct risk-limiting audits of election results without on-site help from the project team. The county used the draft instructions and web tools developed for the pilot program and conducted a risk-limiting audit of three contests.

Humboldt County works with the Humboldt Transparency Project after each election to confirm election results by scanning ballots and creating a parallel tally using Transparency Project software (TEVS). Humboldt County conducted risk-limiting audits of three contests using the Transparency Project CVRs and results for the contests:

Resort Improvement District #1 (elect 3):

(6 candidates)

Total Ballots Cast	193
Ballots Examined for 1% Manual Tally	72
Ballots Examined for Risk-Limiting Audit	52

Eureka City Schools Trustee Area 4 (elect 1):

(2 candidates)

Total Ballots Cast	5,455
Ballots Examined for 1% Manual Tally	15
Ballots Examined for Audit	34

Ferndale Unified School District (elect 2):

(3 candidates)

Total Ballots Cast	640
Ballots Examined for 1% Manual Tally	89
Ballots Examined for Audit	57

Even though Humboldt County examined 176 ballots for the 1% manual tally of the three contests above, the statutory 1% manual tally left as much as a 50% chance of not finding a single error, even if the machine-count found a wrong winner.

In contrast, the risk-limiting audit involved reviewing fewer ballots – 143 ballots – and guaranteed at least a 90% chance of catching and correcting a wrong outcome. Since no errors were found in the initial sample for each contest, no escalation was needed to confirm results.

As with the prior audits in other counties, the Humboldt County audit showed the power and efficiency of risk-limiting audits compared to the flat 1% manual tally currently required by law.

Section III: The 2012 Pilot Audits

During this phase of the project, the project team worked with six counties to conduct audits as follows:

- 1) Following the June 2012 Presidential Primary Election, audits of small election contests were successfully completed in Madera and Napa counties. See Appendix G for county audit reports and cost comparisons;
- 2) Following the June 2012 Presidential Primary Election, large multi-contest audits were prepared, but not completed, in Marin, Orange, Santa Cruz and Yolo counties. The counties stopped their process before the random sample was selected, because of difficulties with the independent tally software;
- 3) During the remainder of 2012, the audit project team collaborated with the University of California research team that developed the OpenCount ballot tally software. The OpenCount team made several improvements to the OpenCount software (not funded under this grant), and the audit project team made improvements to the auditing procedures used in the pilot program to reduce the time needed to complete a parallel scan and independent tally of the ballots.

Unfortunately, improvements to the software could not be made in time for the next round of audits, which were scheduled to follow the November 2012 election. Therefore, four additional counties that had originally volunteered to participate in the program following the November 2012 Presidential General Election were unable to conduct pilot audits.

In November 2012, the SOS requested a 24-month extension of the project period in order to conduct post-election risk-limiting audits using the remainder of the grant funds following the 2014 state and federal elections. However, the EAC was able to provide only a twelve-month extension, to the end of 2013, for the pilot program.

Section IV: 2013 Project Wrap Up

Ballot Accounting Best Practices

The EAC granted a one-year no-cost extension to the project in order to permit the project team to continue its work in 2013. Because 2013 was a non-election year and no further audits could be conducted, the project team collaborated with the Marin County elections officials and Verified Voting to develop a working document called, “Ballot Accounting Best Practices” for county elections officials.

Establishing the integrity of the audit trail is a prerequisite to conducting a post-election risk-limiting audit — or a recount. The Ballot Accounting Best Practices document provides step-by-step procedures for each ballot reconciliation, tally, and audit process required in California on election day and during the 28-day canvass period to ensure all ballots are accounted for, properly tallied, and ultimately stored in a manner that permits a post-election audit of election results. The document provides generic instructions, which are not voting-system specific, but includes examples of how to complete each step, based on how the process is completed in Marin County using its Accuvote voting system. See **Appendix F** for a copy of the Ballot Accounting Best Practices.

Marin County Multi-Contest Audit Completion

Also during this phase, one of the large multi-contest audits, which was attempted but not completed in 2012, was successfully completed in Marin County. Marin made preparations for its audit in July 2012, but had to postpone due to delays in ballot processing caused by the audit software and the county’s need to prepare for the November 2012 Statewide General Election. In February 2013, Marin County completed its audit autonomously – without onsite help from the project team – using the web-based audit tools and instructions. Marin County’s report and cost comparison sheet can be found in **Appendix H**.

Section V: Recommendations for Modifications to Existing Voting Systems

Voting systems used today make post-election risk-limiting audits unnecessarily difficult for several reasons.¹⁰ First, the systems do not export data in formats intended to be parsed by software. Converting those formats into usable data for audits is labor intensive and error prone. Systems should output vote data in a common data format, such as EML, so that the data can be “digested” by other software. This would also facilitate uniform reporting of election results to the public, to the Secretary of State, and to the news media.

Second, most voting systems cannot report results that correspond to the physical batches of ballots that pass through the system; instead, they report only for precincts or subsets of precincts such as “Precinct 1: ballots cast at the polls” and “Precinct 1: vote-by-mail ballots.” This makes it necessary to manually sort (and possibly re-scan and re-process) large numbers of ballots to conduct risk-limiting post-election audits efficiently.

Third, voting systems should track every page of multi-page ballots. For each precinct or batch, the voting system should report how many pages of each type for each ballot style were processed. In some counties it was difficult to ascertain how many of Ballot Card A and how many of Ballot Card B were in a given precinct or batch of ballots, which made conducting audits difficult.

Fourth, the systems currently report vote tally data at relatively coarse levels of aggregation, such as precincts. As discussed above, auditing at the level of individual ballots is vastly more efficient, faster, less error prone, and more transparent. Voting systems should be capable of recording and reporting a CVR for each ballot. It should be possible to retrieve the CVR for any particular physical ballot, and to retrieve the ballot that corresponds to any particular CVR. In short, voting systems should be capable of exporting a complete list of CVRs together with the information that links each CVR to the corresponding physical ballot.

The above recommendations for voting system modifications would facilitate conducting rapid, accurate, transparent, cost-effective risk-limiting audits.

¹⁰ See Stark, P.B., and D.A. Wagner, 2012. Evidence-Based Elections. *IEEE Security and Privacy*, 10, pp 33–41.

Section VI: Cost-Effectiveness: Risk-Limiting Audits Versus the 1% Manual Tally

Effectiveness

The pilot program proved without a doubt that risk-limiting audits are more effective at discovering and correcting vote tally errors and confirming the winners and losers of an election than the current 1% manual tally law. 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. Increasing the confidence level to 99% would have required roughly twice the counting effort, still only a modest burden compared with the 1% manual tally.

By contrast, California's 1% manual tally law requires elections officials to hand tally 100% of the ballots from 1% of all precincts after each election and leads to the hand counting of tens of thousands of ballots across the state after each election. 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 at confirming election results and incapable of correcting erroneous election results.

Cost

The time it took to conduct the risk-limiting pilot audits was minimal – a few minutes to a few hours – compared to the time it takes to conduct the 1% manual tally. However, because of the limitations of fielded vote tabulation systems, counties that participated in the pilot audits spent a significant amount of time preparing for the audits by re-scanning the ballots cast in the contests to be audited using COTS scanners. Some counties used standard office scanners or scanner/copiers that they had on hand, rather than renting high-speed scanners.

Overall during the pilot program, owing to the limitations of current voting systems, the cost of the post-election risk-limiting audits was more than the cost of California's 1% manual tally requirement. Participating counties submitted spreadsheets, detailing the risk-limiting audit costs compared to the cost of the 1% manual tally for the election. See **Appendix H**.

Improving Cost-Effectiveness

If the next generation of voting systems can produce CVRs for individual ballots, then audits can become very inexpensive. Some voting systems in use today may be able to be reconfigured, with prior review and approval of the Secretary of State of course, to allow the capture and reporting of ballot-level results. Then risk-limiting audits to confirm election results could be conducted in a more cost-effective manner: it would eliminate the need for a second scan and re-processing.

Section VII: Products of the Pilot Program

The post-election risk-limiting audit pilot program led to a number of useful products that should help continue the effort to move toward implementing risk-limiting audits of election results in California and beyond.¹¹

Audit Tools

The project team, led by Stark, developed a set of web-based tools (statistics.berkeley.edu/~stark/Vote/auditTools.htm) and step-by-step instructions for elections officials. The tools explain how the audits work and show the math behind the tools, so that elections officials and the public can understand risk-limiting audits. These tools were refined and improved during the pilot program, and the project team developed a reporting form as an adjunct to the AuditTools website for counties to report the results of risk-limiting audits of election results.

See **Appendix B** for screenshots and printouts of the current AuditTools website content.

Step-by-Step Audit Instructions

The project team developed step-by-step instructions for conducting risk-limiting post-election audits. The instructions were continually refined and improved during the pilot program.

See **Appendix C** for the current instructions.

Ballot Accounting Best Practices

The project team collaborated with Marin County elections officials and Verified Voting to develop a working document, called “Ballot Accounting Best Practices” for county elections officials.

See **Appendix D** for the current version of the Ballot Accounting Best Practices.

OpenCount Software

In order to facilitate conducting risk-limiting audits, given the limits of current voting system ballot tracking and reporting capabilities, the project team conducted a separate scan and tally of the ballots in contests to be audited using the OpenCount software developed under a separate project led by University of California Professor David Wagner.

See **Appendix E** for screenshots and printouts from the OpenCount software repository website.

¹¹ Indeed, some of the tools and findings have been used in Colorado already.

Legislation and Regulations to Require Risk-Limiting Audits in Pilots of New Voting Systems

As a result of the success of this audit program, a requirement was added to SB 360 (Padilla), Chapter 602, Statutes of 2013, effective January 1, 2014, which requires (in part) that any new voting system piloted in live California elections be subject to a risk-limiting post-election audit.

See **Appendix F** for a copy of the relevant California statute and legislation.

The Secretary of State is currently promulgating regulations for SB 360. The regulations detail the risk-limiting audit requirements for voting system pilots.

See **Appendix G** for a copy of the proposed regulations.

Section VIII Next Steps

The Legislature Should Revise California's Election Audit Laws

While there is still work to do to streamline risk-limiting audits so that they can be conducted efficiently for all contests with the voting systems in use today, the following should be considered for California:

- 1) **County Option to Conduct Risk Limiting Audits in Lieu of the 1% Manual Tally.** California counties should be permitted to conduct a post-election risk-limiting audit of election results in lieu of conducting the 1% manual tally currently required by law, since the 1% manual tally was proved ineffective in the audit pilot program. Allowing counties to conduct risk-limiting audits instead of the 1% manual tally will allow counties to partially offset the current time and cost associated with conducting risk-limiting audits. Risk-limiting audit procedures can be further refined – and made more cost-efficient – if more counties conduct the audits and establish ballot accounting and audit procedure routines, much the same way the 1% manual tally is currently part of the election canvass routine. The legislature should specify the confidence level or risk limit counties should use for such risk-limiting audits. While the pilot audits were conducted using a risk limit of 10%, a lower limit, such as 1% (99% confidence), may be preferable in practice.
- 2) **Ballot Polling Audits for Large Contests.** For large (e.g., county-wide or state-wide) contests with modest (but not microscopic) margins of victory ballot-polling risk-limiting audits could be performed efficiently using current voting systems, once procedures are in place to ensure the integrity of the audit trail, in part through sound ballot accounting as described in **Appendix D**. For a copy of Stark's brief paper describing the advantages, disadvantages, requirements and formula for ballot polling audits: "Ballot-polling Risk-limiting Audits in Two Pages (+/-1)" dated August, 24, 2012 see **Appendix I**.¹²
- 3) **Automatic Recounts for Contests with Extremely Narrow Margins?** Risk-limiting audits can confirm the outcome of large contests with far less work than a full recount, even with relatively narrow margins. For instance, for a margin of 0.1%, a ballot-level risk-limiting comparison audit at 90% confidence requires inspecting roughly 5,000 ballots, and may stop there if no errors are found. However, for a contest with an extremely narrow margin of victory, elections officials should consider whether conducting a full hand count is more efficient than inspecting ballots at random. Based on our experience in the pilot program, we think the break-even point is approximately when

¹² A fuller explanation and more detailed work estimates are given in Lindeman, M., P.B. Stark, and V.S. Yates, 2012. BRAVO: Ballot-polling Risk-Limiting Audits to Verify Outcomes. *2012 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE '12)*.

the audit can be expected to examine on the order of 10% of the ballots cast.¹³ This is a ballpark estimate; details of how jurisdictions store their ballots have a large effect on efficiency.

In statewide contests, a full manual recount can take several weeks, even when conducted in all counties simultaneously. Therefore, a risk-limiting audit may save considerable time and effort, even for contests with very small margins, provided that errors did not systematically favor one candidate. If there is reason to believe that errors were systematic and occurred at a rate comparable to the margin—or if the initial stages of a risk-limiting audit show that there are such systematic errors—it may be more efficient to proceed immediately to a full hand count, since the risk-limiting audit would likely progress eventually to a full hand count.

The Secretary of State Should Approve Random Selection Methods

Currently, counties use a variety of methods to randomly select precincts for the 1% manual tally required under current law. Some of the methods, such as drawing slips of paper from a box, generally do not result in random selection despite good-faith attempts to stir the paper well. Other methods, such as relying on the pseudo-random number generator in spreadsheet software, are neither transparent nor reproducible. Moreover, some such generators have been shown to have implementation errors. The Secretary of State should review and approve random selection methods to be used for all election audits, as permitted under Elections Code section 15360(c).¹⁴

¹³ Here is a hypothetical based on the ballot-level comparison audit method enunciated in Lindeman and Stark, 2012. A Gentle Introduction to Risk-Limiting Audits. *IEEE Computing Now*, 10, 42–49. Suppose 1-vote discrepancies (such as erroneously recording an overvote as a vote for a candidate) that overstate the margin occur at the same rate as 1-vote discrepancies that understate the margin, and suppose there are no 2-vote discrepancies (finding that a vote for one candidate was misinterpreted as a vote for a different candidate). Then, for auditing with a risk limit of 10% (i.e., 90% confidence level), each discrepancy the audit finds requires the sample size to increase by $(0.56/\text{margin})$ ballots, on average. For instance, if the margin is 0.1%, on average finding a discrepancy will cause the audit to look at an additional $0.56/0.001 = 560$ ballots. If the intrinsic rate of error in the vote tabulation system is high enough that examining 560 ballots is sufficiently likely to uncover another error, the audit is likely to continue to a full hand count. That makes sense, since an error rate of 1 in 560 ballots is about 0.18%, nearly double the margin, 0.1%. On the other hand, if the intrinsic error rate is small, say one in 10,000 ballots, the audit is unlikely to escalate to a full hand count unless there are systematic errors that favor the apparent winner, thereby making the margin appear larger than it truly was. In the project team’s experience, the rate of error in a well calibrated central-count optical scan (CCOS) system is generally lower than 1 in 10,000 unless something has gone wrong, e.g., a software configuration error or voters using gel pens.

¹⁴ One method is to roll 10-sided dice to generate a 20-digit “seed” to be used in a cryptographically secure pseudo-random number generator. This method was used in several of

Counties Should Test Cross-Jurisdictional Audits

A risk-limiting audit of a contest that crosses jurisdictional boundaries was not conducted as part of the audit program. The project team identified two sets of adjacent counties: Yolo and Sacramento, and Santa Cruz and Monterey, which were willing to conduct joint audits of cross-jurisdictional contests. Those counties may wish to voluntarily conduct cross-jurisdictional risk-limiting audits of one or more contests following the November 2014 election in order to test risk-limiting audits on cross-jurisdictional contests.

States Should Standardize Ballot Accounting Methods

The project team recommends California standardize the ballot accounting process used by counties throughout the state. The project team discovered a number of areas where counties could benefit from seeing the best practices of their counterparts in other counties to ensure all ballots printed, cast, destroyed, counted, unused, etc. are accounted for and stored securely during transit, use, and storage. Protocols for using security seals, ensuring secure chain of custody for ballots and equipment, video monitoring, etc., vary throughout the state and could be improved.

the audits conducted under this grant; such a generator is implemented on the AuditTools web page, together with tools to facilitate finding individual randomly selected ballots.

Section IX Conclusion

The Secretary of State Post-Election Risk-Limiting Audit Pilot Program was designed to test the effectiveness and efficiency of risk-limiting audits. Audits of 17 contests were completed in eleven counties following elections held in 2011 and 2012. All audits successfully confirmed the official election results by reviewing a relatively small number of individual ballots (i.e., a few dozen to a few hundred ballots) cast in each contest that was audited.

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. Audits at 99% confidence would have required roughly twice as much hand counting, still a small burden. By contrast, the 1% manual tally provided little if any evidence that the election outcomes were correctly calculated by the voting system.

The project team met its goals to develop and successfully test standards, procedures, and tools for conducting post-election risk-limiting audits. The success of the program has informed new legislation to reform California election audit and recount laws as well as legislation and regulation establishing requirements for future voting system testing and deployment. The audit tools and procedures developed during the program can be used by elections officials in California and other jurisdictions across the United States to conduct risk-limiting audits of election results. Future adoption of laws requiring risk-limiting post-election audits will build public confidence in election results and reduce the need for voter-requested manual recounts and election contest actions in court.

Table of Appendices

Appendix A

AB 2023 (Saldaña), Chapter 122, Statutes of 2010

Appendix B

Audit Tools Website
statistics.berkeley.edu/~stark/Vote/auditTools.htm

Appendix C

Audit Step-by-Step Instructions

Appendix D

Ballot Accounting Best Practices

Appendix E

OpenCount Ballot Tally Software
<https://code.google.com/p/opencount/>

Appendix F

Elections Code section 19209, added by Chapter 602, Statutes of 2013
SB 360 (Padilla), Chapter 602, Statutes of 2013

Appendix G

Text of Proposed Regulations: Procedures for Approving, and Rejecting Voting Systems for a Pilot Program (2014)

Appendix H

County Pilot Audit Reports

Appendix I

Recent Papers and Other Resources on Risk-Limiting Post-Election Audits