

# An Introduction to Risk-Limiting Audits and Evidence-Based Elections

Prepared for the Little Hoover Commission

Philip B. Stark

Department of Statistics, University of California, Berkeley

stark@stat.berkeley.edu | [www.stat.berkeley.edu/~stark](http://www.stat.berkeley.edu/~stark)

2 July 2018

---

## Revision history

---

2 July 2018	Original sent to Little Hoover Commission
3 July 2018	Added executive summary & table of contents

---

## Contents

- Executive Summary
- Evidence-Based Elections
- Risk-Limiting Audits (RLAs)
- Compliance Audits
- Methods for conducting efficient RLAs
- Resources required for RLAs
- Principles for audit legislation
- Direction for California and the nation

## Executive Summary

- Elections should be conducted in a way that gives the public convincing evidence that reported election outcomes are correct.
- Such evidence can be provided by publicly verifiable *risk-limiting audits* (RLAs) of voter-verifiable paper ballots that a *compliance audit* has demonstrated to be trustworthy.
- There are many ways to conduct risk-limiting audits, involving different ways of drawing samples of ballots and different demands on the voting system and on auditors.
- All RLAs require manually inspecting voter-verifiable paper ballots.
- Every jurisdiction that uses paper ballots could conduct *ballot-polling* RLAs immediately, without changing their voting equipment; the software requirements are minimal and free software is available.

- There are more efficient methods for RLAs, in particular, *ballot-level comparison audits*, which have higher demands on the voting system and require more software, but which involve manually inspecting fewer paper ballots.
- Audit laws should embody a number of principles, including requiring serious checks of the integrity of the paper trail, truly limiting the risk of certifying electoral outcomes that are incorrect, specifying risk limits, specifying how contests subject to RLAs are to be selected, ensuring that the audit cannot be subverted, and providing the public enough information to verify that the audit did not stop prematurely.
- Voter-verifiable paper ballots are necessary for evidence-based elections, but “paper” alone is not sufficient: the mode of marking paper should not make it harder for voters to confirm that the paper record accurately reflects their preferences. For voters without disabilities, hand-marked paper ballots may be substantially more usable than ballot-marking devices (BMDs) for this purpose, especially if the BMD prints only a summary ballot, rather than a full-face ballot. More usability research should be conducted before such BMD systems are deployed as the primary or only method of casting a vote.
- A publicly owned (open-source) voting system built from the ground up to be reliable, secure, and auditable would save taxpayers large sums, increase the trustworthiness of elections, make RLAs more efficient, and encourage innovation.
- California and the US could transition to evidence-based elections quickly and economically by:
  - Requiring the legal ballot of record to be voter-verifiable paper in every jurisdiction.
  - Requiring local election officials to produce *ballot manifests* describing how and where paper ballots (specifically, ballot cards) are stored.
  - Promulgating regulations that ensure and demonstrate that the voter-verifiable paper ballots were preserved complete and intact.
  - Developing and deploying open-source, easily auditable systems based to the extent possible on commodity hardware, and cultivating a robust, competitive market for support.
  - Immediately requiring ballot-polling RLAs of countywide and statewide contests.
  - Phasing in RLAs of smaller contests as voting equipment is replaced.

## Evidence-Based Elections<sup>1</sup>

There is no perfect, infallible way to count votes. All methods—including optical scan, touchscreen, and hand-counting—are subject to errors, procedural lapses,

---

<sup>1</sup>Evidence-based elections were introduced in Stark, P.B., and D.A. Wagner, 2012. Evidence-Based Elections, *IEEE Trans. Security Privacy*, 10, 33–41, doi.org/10.1109/MSP.2012.62 Preprint.

and deliberate manipulation.

If there is a trustworthy, voter-verifiable paper trail of voter intent, that paper trail can be used to check (and if necessary, to correct) the *electoral outcomes* of the contests in an election. *Electoral outcome* means the winning candidates or positions,<sup>2</sup> not the exact numerical tally.

The principle of *evidence-based elections* is that local election officials should not only find the true winner(s) of an election, they should also provide the electorate convincing evidence that they did. Generally, that means that the election must be conducted using voter-verifiable paper ballots; all validly cast ballots must be protected to ensure that they continue to represent the correct outcome; there must be convincing evidence that those ballots were kept inviolate through the audit; and the reported outcomes need to be checked against the paper trail by suitable audits or hand counts.

**Evidence = auditability + auditing.**

### Risk-Limiting Audits<sup>3</sup>

In most states that audit elections, including California, statutory audits provide no assurance that, if a reported outcome is wrong, the error will be detected, much less corrected. Indeed, California's 1% Post-Election Manual Tally (PEMT) does not require local election officials to give greater scrutiny to closer contests, nor to take any particular action if the audit discovers tabulation errors. In 2017, AB 840 further hampered the ability of the 1% PEMT to discover outcome-changing errors by allowing local election officials to exclude provisionally cast ballots and VBM ballots that arrive after election day. The 1% PEMT wastes resources, and does not accomplish anything in particular.<sup>4</sup>

---

<sup>2</sup>Or, for instance, whether there is a runoff.

<sup>3</sup>Risk-limiting audits have been endorsed by the Presidential Commission on Election Administration, the American Statistical Association, the League of Women Voters, Common Cause, Verified Voting Foundation, and many other organizations concerned with election integrity. They are required by law in Colorado and Rhode Island, and have been tested in California, Ohio, and Denmark. They were developed in 2007; the first publication is Stark, P.B., 2008. Conservative Statistical Post-Election Audits, *Ann. Appl. Statistics*, 2, 550–581. Reprint. Since then, there have been extensions for other social choice functions (e.g., proportional representation, see Stark, P.B., and V. Teague, 2014. Verifiable European Elections: Risk-limiting Audits for D'Hondt and Its Relatives, *JETS: USENIX Journal of Election Technology and Systems*, 3, 18–39. [https://www.usenix.org/system/files/jets/issues/0301/overview/jets\\_0301\\_stark\\_update\\_9-10-15.pdf](https://www.usenix.org/system/files/jets/issues/0301/overview/jets_0301_stark_update_9-10-15.pdf)), for auditing any number of contests simultaneously, for different types of voting equipment, etc. For a general but still somewhat technical introduction, see Stark, P.B., and M. Lindeman, A Gentle Introduction to Risk-Limiting Audits, *IEEE Security and Privacy*, 10, 42–49, doi:10.1109/MSP.2012.56 Preprint

<sup>4</sup>I have heard many people claim that it checks whether the machines are working correctly. That assertion lacks appropriate nuance. Machines never work perfectly. The question is whether they worked well enough, in this election, to find the true winner(s). That is the question a risk-limiting audit answers. Moreover, since votes cast in person and votes cast by mail are not necessarily tallied by the same equipment, the 1% PEMT (after AB 840) can

In contrast, a *risk-limiting audit* (RLA) is any post-election procedure that offers the following statistical guarantee:

*If a full manual tally of the complete voter-verifiable paper trail would show a different electoral outcome, there is a known, pre-determined minimum chance that the procedure will lead to a full manual tally.*

If the procedure does lead to a full manual tally, the result of that tally replaces the reported result, thereby correcting it.

The maximum chance that the procedure will *not* lead to a full manual tally if that tally would show a different outcome is called the *risk limit*. Equivalently, the risk limit is the largest chance that the audit will fail to correct an outcome that is incorrect, where *incorrect* means that a full manual tally of the voter-verifiable paper trail would find different winner(s).

For instance, a RLA with a risk limit of 5% has at least a 95% chance of requiring a full manual tally, if that tally would show a different outcome.

There are many methods for conducting risk-limiting audits. For instance, a full handcount is a risk-limiting audit, with a risk limit of zero. But by inspecting randomly selected ballots and using appropriate statistical methods, it is possible to conduct risk-limiting audits much more efficiently—when the electoral outcome is correct.<sup>5</sup>

## Compliance Audits

A risk-limiting audit of an untrustworthy paper trail, or any audit that purports to ascertain voter intent from an electronic record or from an artifact that the voter did not have the opportunity to check, is “security theater.” There is little reason to believe that a full manual tally of such records would reveal the true winner(s). It is therefore crucial to base audits on voter-verifiable paper records; to ensure that those records include every validly cast vote exactly once, and no others (checking the determination of eligibility, in particular); to ensure that those records remain complete and intact from the election through the audit; and to assess the evidence that they are trustworthy. Absent affirmative evidence that the paper trail is a trustworthy record of voter intent—that it accurately reflects the intent of every voter who legitimately cast a ballot in the contests under audit, and no others—the audit might simply confirm the incorrect outcome.

The process of assessing the trustworthiness of the paper trail is called a *compliance audit*. Compliance audits should include the following steps, among others:

---

entirely omit any number of machines from the audit, if those machines are used only to tally ballots cast by mail.

<sup>5</sup>When the outcome is incorrect, the audit is intended to have a large probability of requiring a full manual tally, so it generally will not save labor then.

- Ballot accounting.
  - Check that the number of ballots sent to polling places equals the number returned voted plus the number returned spoiled, plus the number returned unvoted.<sup>6</sup>
  - Check that the number of ballots returned from each polling place does not exceed the number of voters registered at that polling place or the number of pollbook signatures at the polling place
  - Check that the number of ballots of each style corresponds to the number of ballots of each style reported by the voting system. Ballot counts for this purpose should be based on the physical paper, not on the voting system: the audit needs external touchstones to check the voting system.
- Eligibility.
  - Check signature verification on vote-by-mail ballots, especially if signature verification was automated.
  - Check the disposition of provisional ballots to ensure that all that were validly cast (and no others) were included in the results.
  - Check that each voter received the correct ballot style based on her eligibility. For vote-by-mail ballots, there should be a record of the ballot style mailed to the voter; for in-person voting, this might require recording (e.g., in pollbooks) the ballot style given to the voter. For provisionally cast ballots, this might be more complicated.
- Physical chain of custody.
  - Record seal numbers whenever a batch of ballots is sealed.
  - Check physical seals for signs of tampering whenever a batch is unsealed.<sup>7</sup>
  - Use numbered seals that are hard to forge or bypass; check seal numbers against the numbers recorded when the boxes or bags were sealed; and log the result.
  - Review custody logs. Check that at least two staff accompanied the ballots whenever ballots were not locked securely and under surveillance.
  - Review surveillance video of secure ballot storage facility to ensure there was no unauthorized access to ballots.
- Due diligence regarding processes, equipment, etc.
  - Review voting equipment event logs.
  - Review any complaints made by voters or anomalies or problems noted by pollworkers.

Some of these steps are formally or informally part of the canvass procedure in some jurisdictions. Ideally, the Secretary of State would require these steps

---

<sup>6</sup>For systems that print ballots on demand, check that the paper stock (sheets cast, spoiled, and still blank) adds up to the number of sheets sent to the polling place or vote center.

<sup>7</sup>Good procedure is to photograph each seal after it is applied, check each seal against its photograph before breaking the seal (perhaps taking a second photograph to document the state of the seal), and record any discrepancies.

(and others) to be conducted in a way that is publicly verifiable, and require jurisdictions to report the results. Protocols around physical seals and physical chain of custody are uneven at best. Before the election, there should also be scrutiny of voter registration databases and changelogs for those databases. Pre-election “logic and accuracy testing” should also include checking the usability each ballot style by ordinary voters. I am not aware of any jurisdiction that currently checks the usability of ballots, despite high-profile, high-consequence ballot usability problems, e.g., “butterfly ballots” in Florida.

Compliance audits should be part of any recount, not just a precursor to risk-limiting audits. Absent a compliance audit, there might not be much reason for the public to trust that a recount will find the true winner(s).

### Methods for conducting efficient risk-limiting audits

The basic strategy behind current methods for risk-limiting audits is to start by acknowledging that the reported electoral outcome might be incorrect, then to examine more and more randomly selected ballots until either (a) the evidence is convincing that a full manual tally would confirm the reported outcome, or (b) there has been a full manual tally.

There is more than one way to do this. Two basic approaches are *ballot-polling audits* and *comparison audits*. Both can be conducted by randomly selecting either groups of ballots (*batch-level audits*) or individual ballots (*ballot-level audits*).<sup>8</sup>

*Ballot-polling audits* are like exit polls, but instead of asking voters how they voted, they manually examine randomly selected voter-verifiable ballots.<sup>9</sup> If a sufficiently large random sample of ballots shows a sufficiently large margin in favor of the reported winner, that is evidence that the reported winner really won.<sup>10</sup> Ballot-polling audits have the advantage that they require very little of the voting system: just the reported winners, and access to the ballots. They do require local election officials to organize the paper trail well enough to draw a random sample of ballots.

*Comparison audits* compare how the voting system tallied groups of ballots to how humans tally the same physical group of ballots. A group might be, for instance, all ballots tallied in a given precinct, or by a given machine. That yields

---

<sup>8</sup>Ballot-level audits tend to require examining fewer ballots in all than audits based on larger batches. Roughly speaking, the number of batches one needs to examine to confirm a contest with a given margin of victory at a given risk limit is about the same, regardless of the batch size. Hence, to attain a given risk limit, an audit that uses batches the size of precincts (say, 500 ballots per batch on average) requires examining about 500 times as many ballots as an audit that uses batches consisting of a single ballot (i.e., a ballot-level audit).

<sup>9</sup>Unlike voters, ballots have to reply, and have to reply truthfully, so ballot-polling audits give strong statistical evidence while exit polls generally suffer from large biases.

<sup>10</sup>How to quantify the strength of the evidence depends on how the sample is drawn, among other things.

a *batch-level comparison audit*. The most efficient comparison audits use groups consisting of individual ballots: *ballot-level comparison audits*. To conduct a ballot-level comparison audit, the voting system must report how it interpreted individual ballots. Such interpretations are called *cast-vote records*. The cast-vote record for a ballot lists the voting system’s interpretation of voter intent for each contest on the ballot. California AB 44 requires new voting systems to be amenable to ballot-level comparison audits, and several vendors now make systems that report cast-vote records that can be linked to the corresponding physical ballot.

One method for conducting a ballot-level comparison audit with a 5% risk limit requires manually inspecting approximately  $7/(\text{diluted margin})$  ballots, unless the audit finds errors in the cast-vote records. The *diluted margin* is the margin of victory in votes, divided by the total number of ballot cards<sup>11</sup> in the population from which the sample is drawn (which must include all ballots cast in the contest, and may include others). For instance, in the 2018 gubernatorial primary in California, Newsom and Cox advanced to the general election. The margin of Cox over Villaraigosa, the runner-up, was 618,215 votes out of 7,060,646 ballots cast, including undervotes. The diluted margin is thus  $618,215/7,060,646 = 8.75\%$ . A ballot-level comparison audit with a risk limit of 5% would have required inspecting approximately  $7/0.0875 = 80$  ballots selected at random from the entire state (assuming the audit did not find any errors). That is a trivial amount of work. Unfortunately, the voting systems in most California counties do not currently support ballot-level comparison audits.

California’s voting systems *do* support ballot-polling risk-limiting audits. The expected number of ballots required to confirm the outcome (namely, that Newsom and Cox won) at 5% risk limit using ballot-polling is 443 (again assuming that the reported results are correct). That is still a trivial amount of work to justify public confidence in the outcome.

Most ways of conducting RLAs require a *ballot manifest* describing how ballots are stored, for instance, “There are 913 boxes of ballots, numbered 1 through 913. Box 1 contains 301 ballots. Box 2 contains 199 ballots. . . .” It is reasonable to require local election officials to construct ballot manifests routinely: if an official cannot keep track of how much paper there is and where it is, she is not doing her job. Some counties might not currently organize their paper flow in a way that makes constructing ballot manifests possible; they would probably need to revise their procedures to conduct RLAs efficiently.

Ballot manifests should be constructed without relying on the voting system to count the paper; otherwise, we are relying on the voting system to check itself. (Moreover, common sources of error in elections are to scan the same box of

---

<sup>11</sup>In many elections in California, a “ballot” consists of two or more “ballot cards” that contain different contests. Sorting the physical ballot cards into homogeneous groups can greatly reduce the number of cards that must be inspected at random to yield a given number of cards that contain a particular contest—and increase the diluted margin.

ballots twice, or to omit a box. Relying on the voting system to construct a manifest would miss such errors.)<sup>12</sup>

## Resources Required for Risk-Limiting Audits

Ballot-polling RLAs require a ballot manifest and the reported results; the hardware and software requirements are minimal; open-source code already exists for all the computations.<sup>13</sup> Batch-level comparison audits using precincts as batches do not save effort compared to ballot-polling RLAs for typical margins and precinct sizes; therefore, I do not recommend them. Ballot-level comparison audits require voting systems that can report cast-vote records for individual ballots in a way that the corresponding physical ballot can be retrieved, and vice versa. Most voting systems currently deployed in California cannot. Ballot-level comparison audits also require exporting those CVRs and “committing” to them in a verifiable way. This involves more software, but Colorado commissioned open-source software that contains much of the needed functionality. Modifying it for California’s needs would not be expensive.<sup>14</sup>

Methods for combining ballot-polling in some counties with ballot-level comparisons in others to produce a RLA of cross-jurisdictional contests have been developed (and software has been written to implement them)<sup>15</sup>, but they have not been used in a real election. RLA methods exist for plurality contests, majority contests, super-majority contests, vote-for-n contests, and proportional representation.<sup>16</sup> There is currently no efficient RLA method for instant-runoff voting.

Probably the most difficult aspect of auditing is to coordinate the actions of

---

<sup>12</sup>However, ballot manifests can be augmented by data from the voting system to facilitate audits, provided the audit is designed to take into account the possibility that the voting system data is incorrect. For instance, there are ways to combine cast-vote records with ballot manifests to make it easier to sample ballots that contain specific contests and still ensure that the procedure is a RLA.

<sup>13</sup>See, e.g., <https://www.stat.berkeley.edu/~stark/Vote/ballotPollTools.htm>

<sup>14</sup>Colorado is likely to make some of the needed modifications itself, to facilitate auditing cross-jurisdictional contests. I would expect that any additional modifications for California’s purposes would cost on the order of \$100k-\$200k. Other states might be willing to share the cost of generally useful enhancements. The underlying calculations are well worked out (see, e.g., <https://github.com/pbstark/S157F17/blob/master/audit.ipynb> or <https://www.stat.berkeley.edu/~stark/Vote/auditTools.htm>); the main issues concern user interfaces for state and local election officials, and some changes to the data structures that the Colorado software (RLATool) currently uses.

<sup>15</sup>Lindeman, M., N. McBurnett, K. Ottoboni, and P.B. Stark, 2018. Next Steps for the Colorado Risk-Limiting Audit (CORLA) Program, <https://github.com/pbstark/CORLA18/blob/old-version1/CO18.pdf>

<sup>16</sup>For examples, see, e.g., Stark, P.B., 2009. Efficient post-election audits of multiple contests: 2009 California tests. 2009 Conference on Empirical Legal Studies. (Preprint), and Stark, P.B., and V. Teague, 2014. Verifiable European Elections: Risk-limiting Audits for D’Hondt and Its Relatives, *JETS: USENIX Journal of Election Technology and Systems*, 3, 18–39. [https://www.usenix.org/system/files/jets/issues/0301/overview/jets\\_0301\\_stark\\_update\\_9-10-15.pdf](https://www.usenix.org/system/files/jets/issues/0301/overview/jets_0301_stark_update_9-10-15.pdf).



different counties, for contests that cross county lines. A pilot risk-limiting audit of a statewide contest in California would be very instructive.

In my experience, it takes about 2 minutes to retrieve a particular randomly selected ballot and transcribe the votes for two or three contests.<sup>17</sup> Additional contests take on the order of 10 seconds each per ballot. The cost of conducting RLAs seems to be very small compared to the overall cost of holding an election. In Colorado, some local election officials report that RLAs are easier than the statutory audits that RLAs replaced, even though the previous audits were limited to tallying at most 500 ballots—far fewer than California’s 1% PEMT generally requires. There are many ways California can reduce the effort required to conduct RLAs, including sorting ballot cards into homogeneous physical batches that contain the same contests, and printing serial numbers on ballots after they have been anonymized (so that voters’ preferences cannot be linked to their identities).

Some vendors are promoting systems that make digital images of ballots, claiming that the images make performing RLAs easier, because fewer (or no) paper ballots need to be inspected. That is simply not true, as a matter of statistics: if a risk-limiting audit relies on images of ballots, it must check that the error in making the images from the voter-verified paper ballots plus the error the system made interpreting those images to make cast-vote records is not large enough to cause the electoral outcome to be wrong. It is a mathematical fact that this requires examining at least as many physical ballots as an audit that just compares cast-vote records to a human reading of the voter-verified ballots, without relying on the digital images.<sup>18</sup>

## Principles for Audit Legislation

A risk-limiting audit law should satisfy a number of principles.

1. It should require rigorous physical custody of ballots, and compliance audits, as discussed above. A RLA that relies on an untrustworthy paper record accomplishes little.
2. It should **require genuine RLAs**: the procedures and calculations should ensure that whenever an outcome is incorrect, the audit has the requisite

---

<sup>17</sup>The process is much faster if serial numbers are printed on the ballots (after the voted ballot has been dissociated from the voter’s identity).

<sup>18</sup>Moreover, (i) there are demonstrations of ways that scanners inadvertently alter images in ways that would change the appearance of voter intent, including erasing votes; and (ii) one common source of large errors in election results is failing to scan a batch of ballots, or scanning the same batch twice. “Picking errors” (where a scanner picks up more than one card) and paper jams can also lead to ballots being omitted or scanned more than once. Expecting digital images to accurately reflect voter intent from every validly cast ballot, exactly once, is wishful thinking, even in the absence of hacking. Of course, hacking the scanners or the image processing software is within the technical ability of many undergraduate computer science students.

chance of leading to a full hand count.<sup>19</sup> That in turn entails a number of things:

- **The audit must ascertain voter intent manually, directly from the human-readable marks on the paper ballots** the voters had the opportunity to verify. It is not adequate to rely on digital images of ballots, paper printed from an electronic record, barcodes, or other artifacts that are not verifiable by the voter and/or are not tamper evident; nor is it adequate to re-tabulate the votes electronically, either from images of the ballots or from the original paper. Digital images, re-printed ballots, and other computer data are not reliable records of voter intent: they can be incomplete, fabricated, or altered accidentally or maliciously, by software bugs, procedural lapses, or hacking. The bill should prohibit relying on such things for the determination of voter intent. Making this prohibition explicit is important because, as mentioned above, voting system vendors are marketing technology that purports to facilitate RLAs by allowing auditors to examine digital images of ballot instead of paper ballots. But relying on those images as an accurate representation of voter intent would in fact undermine RLAs. Relying on an electronic record created by the voting system to accurately reflect voter intent amounts to asking the same doctor for a second opinion (or asking a defendant whether s/he is guilty).
- **The audit must take all validly cast ballots into account.** If ballots are omitted from consideration, for instance, vote-by-mail ballots that did not arrive by election night or provisionally cast ballots, the audit cannot be a genuine risk-limiting audit. The simplest way to do this is to start the audit when the canvass is complete but before the results are certified. Local election officials may want to start the audit sooner, for instance, before all provisionally cast ballots have been resolved. There are a number of ways this can be done and still yield a true RLA. For example, the audit can treat ballots that have not been tallied when the audit starts as if they had the votes that cast the most doubt on the outcome based on the votes already tallied. (In a plurality contest, that would mean treating ballots that had not been adjudicated and tallied as if each ballot showed a valid vote for every losing candidate or position.) Another approach involves stratified sampling, drawing independent random samples from different collections of ballots after those collections are tabulated, and combining the results using statistically appropriate formulae. While these approaches enable election officials to start the audit as soon as election night, the risk calculations are more complicated, as is the “escalation” process of examining more ballots if

---

<sup>19</sup>I do not believe the statute should dictate methods or calculations, only principles. That makes it possible to use improved methods as they are developed and/or as voting systems are replaced.

the initial sample turns out to be insufficient to confirm the outcome of every contest. More generally, there are tradeoffs involved in choosing among RLA methods, both in logistical complexity for election officials and in the ability for the public to verify that the audit was performed correctly. But basing the audit on vote totals that omit provisionally cast ballots or vote-by-mail ballots that were not tabulated as of some date—as AB 840 does—cannot yield a true RLA.

- **The audit must have the ability to correct incorrect outcomes.** This might mean that the audit must take place before results are certified, or that the audit can revise already-certified results.
3. **The risk limit(s) should be in statute.** Allowing the Secretary of State or local election official to choose the risk limits may create a real or apparent conflict of interest.
  4. **The statute should specify how the contests to be audited are selected.**
    - If not every contest will be audited in every election, the selection of contests to audit should involve a random element to ensure that every contest has some chance of being selected, to ensure that a malicious opponent would not be able to predict that any particular race will not be audited.
    - Every contest not audited with a RLA should be audited using a *risk-measuring audit* instead.<sup>20</sup>
    - **The statute must apply to cross-jurisdictional contests, including statewide contests.** Because the point of an RLA is to ensure that reported contest outcomes are correct, every county involved in a particular contest must examine ballots in such a way that the overall cross-jurisdictional procedure is an RLA of that contest. Operationally, auditing cross-jurisdictional contests—whether by an actual RLA or a partial RLA—will entail a number of things:
      - Contest-level results (not merely county-level results) must be known before the audit can conclude.
      - Audits of cross-jurisdictional contests need to be coordinated in some way, so that each county knows when its portion of the audit can stop.
      - The risk calculations need to match the way the sample is drawn. The approach that involves examining the fewest ballots (when the contest outcome is correct) is to sample directly from all

---

<sup>20</sup>Risk-measuring audits are related to risk-limiting audits, but they do not have a pre-specified minimum chance of requiring a full manual tabulation when that tabulation would show a different result. In statistical terminology, a risk-measuring audit reports a p-value for the compound hypothesis that a full count would yield a different electoral outcome, based on the audit data. Equivalently, they report the smallest value for which a risk-limiting audit conducted using that value as its risk limit would have stopped without examining more ballots.

ballots cast in the contest. That requires the Secretary of State or another entity to tell each jurisdiction how many ballots it needs to draw from each cross-jurisdictional contest, in light of the margin and what the audit reveals as it progresses. A rigorous audit can also be based on a stratified sample, which would de-couple the sampling in different jurisdictions. However, stratification generally requires inspecting more ballots to attain the same risk limit, and the risk calculations are more complicated.

**5. The audit sample must not be predictable before the audit starts.**

- The “seed” for selecting the sample must be sufficiently random, e.g., involve rolling 10-sided dice 20 times, with public participation.<sup>21</sup>
- The sample from any collection of ballots should not be selected before election officials have “committed” to the tally of those ballots. For example, nobody should be able to know whether precinct 207 will be audited until the election official has published the tally for precinct 207.<sup>22</sup>

**6. The public must be able to verify that the RLA did not stop prematurely,** not merely “observe” the RLA. Among other things, this requires election officials to:

- Disclose the algorithms used to select the sample, to calculate the risk, and to determine when the audit can stop.
- Provide public opportunity to observe the selection of the “seed” for the random sample.
- Provide adequate evidence that the paper trail of cast ballots is complete and intact (evidence generated in part by the *compliance audit*).
- Provide public opportunity to verify that the correct ballots were inspected during the audit.
- Provide public opportunity to observe the voters’ marks on the ballots that were inspected by the audit<sup>23</sup> For some ways of conducting RLAs, the public needs additional information. For instance, with “ballot-level comparison audits,” the public also needs to see the cast-vote

---

<sup>21</sup>Colorado’s public ceremony is a good model. See <https://youtu.be/ysG4pFFmQ-E>

<sup>22</sup>There are examples (notably, in Cuyahoga County, OH, <https://www.wired.com/2008/03/the-mysterious/>) where election officials altered the tallies in precincts selected for recount after the sample was selected, to ensure that the inspection would not find any discrepancies.

<sup>23</sup>It is important to have published rules governing how human marks on ballots are to be interpreted in audits and recounts. For instance, if a voter makes a write-in vote for a candidate who is also listed on the ballot, is that a valid vote? If a voter marks a vote for a listed candidate and also writes in that candidate’s name, is that a valid vote? If a voter marks a vote for a candidate, crosses through the mark, and marks a vote for a second candidate, is that a valid vote for the second candidate? If a voter makes a stray mark on the ballot that is distinctive enough to identify the ballot, is the ballot valid?

record for each audited ballot and proof that the full set of cast-vote records yields the reported contest results.

### **Direction for California and the nation**

California already requires paper ballots; all we need to conduct evidence-based elections is better procedures for safeguarding ballots, compliance audits, and risk-limiting audits.

**Recommendation 1.** Pass legislation requiring the Secretary of State to develop regulations for the secure, verifiable custody of voted ballots (including protocols for using seals, transporting ballots, and storing ballots), and regulations for “compliance audits” to ensure that eligibility was determined correctly and that the security of the physical ballots was maintained.

RLAs of countywide and statewide contests are reasonably economical using ballot-polling, which requires little software or other infrastructure and no change to voting equipment. To conduct such RLAs would require counties to report final, but uncertified vote totals to the Secretary of State earlier, and would require coordinating the audits of different counties. It would not be expensive, but would make the results of those contests trustworthy, and would be a stepping stone towards auditing all contests to a risk limit.

**Recommendation 2.** Pass legislation immediately requiring risk-limiting audits for countywide and statewide contests and risk-measuring audits for all other contests, based on the ballots selected for the risk-limiting audits. I would recommend using a risk limit no larger than 5 percent for those contests. I would expect that auditing most such contests will require examining far fewer ballots than the current 1% PEMT. The legislation should allow jurisdictions to substitute RLAs for the 1% PEMT if they audit all contests that intersect their jurisdiction with either a RLA (if the contest is entirely contained in jurisdictions willing to conduct RLAs) or a “partial risk-limiting audit” (as defined in SB 360) if the contest is not entirely contained in counties that are willing to conduct RLAs. At some point, RLAs of all contests should become mandatory, but the risk limit for smaller contests might not be the same as that for larger contests.

As California counties replace their voting systems, AB 44 ensures that the new systems will support ballot-level comparison audits. As counties acquire such systems, they can move from ballot-polling audits to ballot-level comparison audits. But the cost of such systems is unnecessarily high, largely as a result of the current approach to certification by the US Election Assistance Commission (EAC) and state laws.

I believe that our nation could (and should) develop a publicly owned, open-source<sup>24</sup> voting system designed from the start to be accessible, secure, reliable, and auditable, and to run on commodity “off-the-shelf” hardware to the extent

---

<sup>24</sup>E.g., licensed under the MIT software license or GPL 3.0.

possible.<sup>25</sup> I estimate that this would cost roughly \$50 million, of which \$30-\$40 million would be to design, prototype, and test the system, and \$10-\$20 million would be to train local election officials to use the system and to seed a competitive market for commercial support.<sup>26</sup> Existing voting system vendors could compete to package, sell, lease, and support the systems. But taxpayers would benefit immediately: current annual maintenance costs greatly exceed the cost of developing, acquiring, and maintaining a system of this sort.

**Recommendation 3.** Pass legislation requiring California, alone or in collaboration with other states and the federal government, to develop an accessible, auditable, secure, open-source voting system based, to the extent possible, on commodity hardware.<sup>27</sup>

I expect that the systems Los Angeles County and San Francisco County are developing will provide good starting points. However, this raises another issue, accessibility.

Every voter should have the opportunity to vote using a method that is usable, given his or her abilities. More than one kind of usability matters: voters must be able to record their preferences accurately and easily, and to verify accurately and easily that the paper record—which law should specify is the official ballot<sup>28</sup>—reflects those preferences. Voters with dexterity or visual disabilities may need assistive technology to cast a vote. There is currently a push for using ballot-marking devices (BMDs) for all voters, and, in particular, devices that do not print a full-face ballot but instead print “selections-only” or “summary” ballots, possibly with a barcode or QR code to assist the voting system in transcribing voter intent.

Local election officials evidently like BMDs for a variety of reasons:

- No ambiguous voter marks, which occur from time to time with hand-marked paper ballots
- Less paper to store
- No need to pre-print ballots in general, nor anticipate how many ballots of a given style, or in a given language, will be needed
- Risk mitigation with respect to lawsuits on behalf of voters with disabilities

---

<sup>25</sup>Most of the hardware a voting system needs can be general purpose office equipment (high-speed scanners and commodity CPUs). Being able to buy or replace components at Costco or Amazon rather than paying premium prices for vendor-branded, special-purpose “voting” hardware would immediately save jurisdictions large sums.

<sup>26</sup>I imagine using tax incentives, small business loans, and other programs to encourage entrepreneurship. These businesses would perform the role that Red Hat, Oracle, IBM, VMWare, and others do in the Linux world—although I think that voting systems should run on local, dedicated, air-gapped hardware, not in the cloud.

<sup>27</sup>This would be facilitated if the US Election Assistance Commission adopted standards for interoperability of voting systems, and guidelines for certifying modules of systems rather than only complete voting systems. There is already movement in this direction, but the EAC currently lacks a quorum.

<sup>28</sup>The Board of Advisors recently passed a resolution to that effect: <https://www.eac.gov/documents/2018/04/27/resolution-2018-03-auditability-of-voter-intent-passed-10-8-4-advisors-resolution-page/>

The Minnesota recounts showed that the number of truly ambiguous voter marks is negligible, so I discount the first concern.

However, my personal experience is that hand-marked paper is more usable by voters without disabilities, not only for recording intent but also for verifying that the paper ballot correctly portrays that intent. We cannot force voters to check that the paper ballot accurately reflects their preferences, but we should not field voting systems that make it unnecessarily hard for voters to check. For a “California-style” ballot with many contests, including some that have similar names, I am not able to check whether a summary ballot accurately reflects my selections unless I make notes about my intended votes, for instance, using a sample ballot. My memory simply isn’t good enough. The extra verification step of checking whether the paper accurately reflects one’s choices is especially important for ballot-marking devices (compared to hand-marked ballots) because the technology might alter the vote as a result of miscalibration, bugs, or hacking. To my knowledge, there has been no testing of BMDs to check their usability for voters to verify that the printed paper (whether a “selections-only” summary ballot or a full-face ballot) captured their intent accurately. There has been usability work on VVPATs (voter-verifiable paper audit trails), which shows that voters are not able to use them well.<sup>29</sup>

**Recommendation 4.** Do not certify voting equipment that requires all voters to use BMDs until and unless there has been adequate testing to confirm that BMDs are as usable (by voters without disabilities) as hand-marked paper ballots for recording voter intent and for verifying that the printed ballot correctly captures voter intent, for California ballots. Every voter should have the opportunity to cast a vote by hand marking a full-face paper ballot.

## Biographical information

I am professor of Statistics and Associate Dean of Mathematical and Physical Sciences at the University of California, Berkeley. I am on the Board of Advisors of the US Election Assistance Commission (appointed by Minority Leader Nancy Pelosi). I am on the Board of Directors of Verified Voting Foundation. I served on Secretary of State Debra Bowen’s Post-Election Audit Standards Working Group in 2007. I have worked closely with state and local election officials in California and Colorado to pilot and deploy risk-limiting audits, which I invented in 2007. The software Colorado uses for RLAs is based on software I wrote. I testified in Jill Stein’s recount suit in Wisconsin, and filed a report in her suit in Michigan. I have testified in two election-related lawsuits in California. I was involved in drafting several pieces of California legislation. I have testified to both houses of the California legislature regarding election integrity.

---

<sup>29</sup>See, e.g., Everett, S.P., 2007. The Usability of Electronic Voting Machines and How Votes Can Be Changed Without Detection, Ph.D. Dissertation, Rice University. <http://chil.rice.edu/alumni/petersos/EverettDissertation.pdf>

The opinions expressed above are my own, not necessarily those of my employer nor of any other organization with which I am affiliated.

### **Acknowledgments**

I am grateful to Robert J. Geller, Joseph Lorenzo Hall, Mark Lindeman, and Ronald L. Rivest for helpful comments.