



DEPARTMENT OF STATISTICS
367 EVANS HALL
BERKELEY, CALIFORNIA 94720

+1 (510) 394-5077

July 21, 2026

Nicholas M. Andersen, Acting Director
Cybersecurity and Infrastructure Security Agency
Department of Homeland Security

Dear Mr. Andersen:

I write in reference to the auditing recommendations in the 10 July 2026 press release, at the URL <https://www.dhs.gov/news/2026/07/10/dhs-requires-states-adopt-common-sense-election-security-measures-receiving-federal>

[A]fter each federal election, states must conduct a manual audit of at least 5% of all ballots cast. This common-sense process involves physically reviewing a random selection of ballots to confirm that the machine reported results match results on the paper ballots. This will ensure accuracy and detect any machine errors or manipulation.

I have two serious concerns about this. First, examining a random sample of 5% of ballots in general will *not* suffice to ensure accuracy and detect machine errors or manipulation. The sample size required to detect errors or manipulation depends on many things, including the amount of error or manipulation, the number of votes cast, how the random sample is drawn, and how the random sample is used. Second, auditing some contests in an election gives no assurance about any other contests in that election: every contest needs some scrutiny.

How the random sample is drawn is key. For instance, the quotation above suggests that the sample involves selecting entire precincts or scan batches of ballots—groups for which the voting system reports vote subtotals that can be compared to the manual tally. That is inefficient and ineffective.

For a simple illustration, suppose a two-candidate contest (Alice v. Bob) comprises 20 precincts, with 1,000 ballots cast in each precinct. Then each precinct contains 5% of the ballots, so a random sample of 5% of ballots amounts to one precinct selected at random.

Suppose that in reality, in every precinct, Alice received 505 votes and Bob received 495: the true contest totals are $505 \times 20 = 10,100$ for Alice and $495 \times 20 = 9,900$ for Bob, so Alice really won.

But the tabulator was hacked in one precinct. In that precinct, Bob was reported to have received 700 votes and Alice 300, so the reported totals are

$$505 \times 19 + 300 = 9,895 \text{ votes for Alice, and}$$

$$495 \times 19 + 700 = 10,105 \text{ votes for Bob.}$$

Bob is the *reported* winner, even though he didn't really win.

If 5% of the ballots are selected at random by selecting a precinct at random and the precinct is manually recounted and the vote totals are compared to the reported totals, there is only a 5% chance that the sample

will be the hacked precinct, and so only a 5% chance that audit will find anything amiss. This approach to auditing has little chance of noticing a problem, despite the fact that hacking changed the outcome. Even if the sample were 10% of the ballots (2 of the 20 precincts), the chance the sample would include the hacked precinct is only 10%.

If the sample were drawn by selecting individual ballots (rather than entire precincts) and the votes on those ballots were compared to how they were tabulated by the equipment, the chance of discovering at least one discrepancy would be much larger. However, most currently fielded voting systems cannot report how they interpreted any particular ballot, at least for precinct-count systems, where matching the electronic records to physical ballots has a high risk of compromising voter privacy. (There is another method of auditing using individual ballots that does not require comparing the votes on the selected ballots to their machine interpretation, called “ballot polling,” which is more efficient than sampling entire precincts but less efficient than comparing individual ballots to their machine interpretations.)

To continue the toy example, suppose that the auditors are lucky and happen to select the hacked precinct. What happens next? What additional action does finding a discrepancy trigger? Does it automatically require a full hand count of the entire election? In this case, a full hand count would be an appropriate action, because it would correct the reported outcome and show that Alice really won.

Consider the following change to the hypothetical: in the hacked precinct, the true totals are 690 for Bob and 310 for Alice, so even though Bob didn’t receive as many votes as reported, he still really beat Alice, 10,095 to 9,905. If the hacked precinct was selected by the audit, what action does finding the discrepancy trigger? Automatically requiring a full hand count would be a waste of time and effort, because the full hand count would just confirm that Bob won.

More generally, finding a discrepancy between the voting system’s tally and a hand tally does not necessarily cast doubt on the reported result, for instance, if correcting the error would widen the margin of victory. Nor does a discrepancy necessarily signal misconfiguration, malfunction, or hacking: some discrepancies are to be expected even under normal operation. For instance, voters sometimes mark ballots in a way that is an unambiguous vote to a human auditor but likely to be misinterpreted by the voting system. Examples include marks that are too light, “hesitation marks” in a vote target, and marks the system is not programmed to look for (e.g., the voter circled an instead of filling it). Provided the net error in the machine tabulation across all ballots (not just those in the sample) does not exceed the margin, the reported outcome is still correct.

Auditing for the sake of auditing is security theater. An audit should accomplish something. The single most important role for an election audit is to check whether the announced winners really won, i.e., whether the reported winner(s) really got more votes than the other candidates. A properly designed audit that relies on a suitably drawn random sample from a trustworthy and essentially complete record of the validly cast votes can do that, but the sample size required is not the same for every contest.

In general, the sample size required to confirm that an outcome is correct depends on many things, including:

- The number of ballots cast in the contest
- The margin of victory in the contest in votes
- How the sample is drawn. (For instance, by sampling entire precincts is less efficient than sampling

individual ballots. If precincts are the sampling unit, sampling using probability proportional to the number of ballots cast in the precinct generally leads to smaller sample sizes than sampling precincts with equal probability. Stratified sampling, which some state laws would require, generally increases sample sizes.)

- What the votes in the sample are compared to, if anything.

In general, an audit cannot provide assurance that outcomes are correct unless there are provisions for expanding to a full manual recount (to set the record straight) if the sample does not provide sufficiently high confidence that the outcome is right. Incorrect outcomes can only be corrected by a full manual recount: it would be immoral if not illegal to overturn a result on the basis of a sample of ballots alone.

Any audit method that has a known minimum chance of catching and correcting wrong outcomes, regardless of why the outcomes are wrong, is by definition a *risk-limiting audit*. The *risk limit* is the maximum chance that the audit will fail to correct a reported outcome that is wrong.

It is impossible to conduct a genuine risk-limiting audit without a trustworthy record of the validly cast votes: ensuring that only eligible voters cast ballots, that the votes on the ballots reflect the voters' expressed intentions (hand-marked paper ballots, not BMD printout, which can print something other than what the voter selected), and that the ballots are kept secure and organized throughout the canvass and audit are key. There are risk-limiting audit methods for most types of elections in the U.S., including plurality, multi-winner plurality, supermajority, and instant-runoff voting (ranked-choice voting). There are risk-limiting audit methods that can work with every voting system in the U.S., although the efficiency—the number of ballots the audit has to inspect manually before it can stop—depends heavily on what data the voting system can export.

Some jurisdictions (notably, Georgia) have conducted sham risk-limiting audits to cover up large problems and shoddy election administration. Such audits are “garbage in, garbage out,” because the underlying paper trail isn't trustworthy. But the fact that a tool can be misused doesn't mean the tool is faulty: tools need to be used properly. The same is true of hammers and scalpels.

Finally, auditing some contests says nothing about any contests that are not audited: misconfiguration, procedural errors, and hacking can affect some contests but not others. *Every* contest in *every* election should be audited. This is the moral equivalent of “measure twice, cut once,” advice that applies to every cut, not just to a few cuts in a given project.

If I may be of assistance in developing a more rigorous, statistically sound auditing plan, I would be delighted to help.

Sincerely,



Philip B. Stark

Distinguished Professor of the Graduate School