How Election Officials Were Conned by Risk-Limiting Audits — And the Easy Solution

July 9, 2026



Executive Summary

Risk-limiting audits (RLAs), developed by statistician Philip B. Stark, offer a statistical method to check election results—but only when applied to a complete, intact, and trustworthy voter-verifiable paper trail. All post-election audits, including traditional fixed-percentage checks, share a critical limitation: they cannot create trustworthiness where none exists.

When election officials claim that any audit proves accurate results, they produce a false sense of security. This illusion misleads citizens, sustains insecure voting systems, and creates opportunities for undetected manipulation. **Even after the inventor of RLAs publicly exposed this misuse as creating a false sense of security for unverifiable systems, election**

officials continue to rely on and promote these audits, passing the misleading assurances to the public and dismissing citizen concerns.

The only reliable path to verifiable election accuracy is direct hand-counting of paper ballots by citizens at the precinct level on election night, with full transparency and public oversight. Audits alone are insufficient and can serve as a pathway for bad actors to introduce or hide errors while evading detection.

Philip B. Stark and the Origin of Risk-Limiting Audits



Professor

Philip B. Stark

University of California, Berkeley

AREA: Mathematical and Physical Sciences

SPECIALTY: Mathematics, Applied Mathematics, and Statistics

Philip B. Stark is a Distinguished Professor of Statistics at the University of California, Berkeley, specializing in uncertainty quantification and election auditing. In 2007–2008, as part of California's Post-Election Audit Standards Working Group, he developed risk-limiting audits (RLAs). Key publications include "Conservative Statistical Post-Election Audits" (2008) and "A Gentle Introduction to Risk-Limiting Audits" (with Mark Lindeman, 2012). Stark also produced open-source tools for sampling and risk calculation and advised on implementations such as Colorado's first statewide RLA in 2017.

Precise Mechanics of Risk-Limiting Audits

An RLA is a statistical procedure that manually examines randomly sampled ballots from a trustworthy paper trail. It limits the probability of certifying an incorrect outcome to a pre-specified risk limit (for example, 5%). If discrepancies suggest the reported result may be wrong, the audit expands until the risk limit is met or a full hand count corrects the outcome.

RLAs require strict prerequisites: a complete voter-verifiable paper trail (preferably hand-marked ballots), a compliance audit confirming chain of custody and ballot accounting, a ballot manifest, and verifiable random sampling. Two approaches exist—ballot-polling (hand review without machine records) and ballot-comparison (matching hand counts to machine records). Sampling increases as needed. These features make RLAs statistically stronger than traditional audits, but they provide **zero assurance** without trustworthy underlying records.

Stark's Resignation from Verified Voting

On November 21, 2019, Philip B. Stark resigned from the Verified Voting board. The full text of his resignation letter follows:

> *21 November 2019*
> *Dear colleagues—*
>
> *With sadness and disappointment, I am resigning from the board of Verified Voting.*
>
> *I believe that Verified Voting has lost its way.*

It has been providing cover for inherently untrustworthy voting systems—and the officials who bought them, the companies that make them, and any officials who might contemplate buying them in the future—by conducting "risk-limiting audits" of untrustworthy paper records, creating the false and misleading impression that relying on untrustworthy paper for a RLA can confirm election outcomes (and debasing the meaning of "RLA" in the process).

This contradicts the most basic principle of Evidence-Based Elections: the need to establish that the paper trail is trustworthy.

Several months ago, I asked VV to revise its published policy on BMDs to clarify that parallel testing cannot show that BMD printout is trustworthy. The current policy on parallel testing of BMDs is counter to the advice of David Jefferson, Ron Rivest, Andrew Appel, Rich DeMillo, and me (and perhaps other board and AB members). I have repeated the request. Yet, months later, the original document stands.

VV is promoting the "shiny" part of auditing—the RLA procedure—at the expense of a far more fundamental requirement for trustworthy elections: a trustworthy paper trail. Whitewashing inherently untrustworthy elections by overclaiming what applying RLA procedures to an untrustworthy paper trail can accomplish sets back election integrity. This is security theater, not election integrity.

Indeed, VV publicly claimed that the pilot RLA in Georgia confirms the election outcomes, despite the fact that it was conducted primarily on universal-use BMDs with serious usability and security defects. I understand that the press quotation from Marian about the audit was not approved in advance by her, but she has clarified by email that she

believes it is true.

I do not.

That statement has done damage to a case trying to hold Georgia SoS accountable for historical neglect of election integrity and its ill-advised decision to buy universal-use BMDs, a case in which I have submitted several reports, and which involves Rich DeMillo more centrally. I predict it will also become part of the vendor's advertising for universal-use BMDs. It amounts to a product endorsement from VV that their systems can be proved to get the correct outcome by the "gold standard," a RLA.

I asked VV to issue a public clarification (press release or policy statement) that RLAs cannot show that BMD output is trustworthy—nor can any other audit establish that BMD printout is trustworthy—and that basing an audit on BMD output cannot confirm election outcomes. I asked a week ago, to allow time for the statement to be used in court filings to counter claims that Georgia election officials have made in court documents, quoting VV's participation in the "audit." No such statement has been made.

I also tried to prevent VV from making a similarly misleading claim about auditing in Philadelphia—but it did: Today, VV is conducting a "RLA" relying on the Philadelphia ExpressVote XL printout, and claimed on Twitter that it will "confirm election outcomes." That is patently false. But I'm sure it will be in ES&S press releases and promotional materials within a few days.

VV should be making it clear that trustworthy paper is prerequisite to conducting a RLA, and that auditing the tabulation provides no evidence that outcomes are correct if the paper trail is not trustworthy. Applying RLA procedures to

BMD output might show that tabulation errors are not large enough to produce the apparent margin, but it can't confirm election outcomes, because there is no trustworthy touchstone. Evidence is what matters, not mindlessly applying a procedure. RLAs are not magic. They do not cure all the ills of any system that produces some kind of paper "backup."

Indeed, the need for evidence that the paper trail is trustworthy applies to hand-marked paper ballots as well as to BMD printout: VV should be demanding convincing evidence that the paper trail is trustworthy, and helping jurisdictions develop laws, regulations, and procedures to make it possible to generate that evidence. That requires a compliance audit that checks the chain of custody of the paper, ballot accounting, eligibility determinations, signature verification, among many other things. A compliance audit of a universal-use BMD system cannot provide affirmative evidence that the paper trail is trustworthy. No procedure can. ("Trustworthy" means that a full hand count of the paper would show who really won. It does not mean that every last vote was captured and retained accurately.)

Verified Voting is providing cover for bad actors (election officials and vendors) and inherently untrustworthy voting systems (poorly designed, inaccessible universal-use BMDs) by conducting "risk-limiting audits" of paper trails that cannot in principle be established to be trustworthy even if there were good laws and regulations for compliance audits. VV is helping election officials who purchased poorly designed, unnecessarily expensive, insecure, universal-use BMD systems justify their purchases—despite the fact that virtually every expert on our board and advisory board recommended against purchasing universal-use BMDs. The election officials are now touting the fact that VV helped them conduct (i.e., pilot) RLAs to brag that their systems are trustworthy after all—and to

claim in court that the audits prove their systems are trustworthy. VV is giving vendors quotes, talking points, and—in effect—product endorsements for insecure, inaccessible systems. VV is on the wrong side.

Our message to jurisdictions that buy poorly designed, insecure, universal-use BMD systems should be, "We tried to warn you. You need a better voting system." Instead, we're saying, "Don't worry: VV will teach you to sprinkle magic RLA dust and fantasies about parallel testing on your untrustworthy election. All will be fine; you can use our authority and reputation to silence your critics."

For well over a year, VV has been doing things that contradict and undermine my research, my publications, my expert reports, and my public testimony. Whatever good VV might be doing on other fronts, I cannot continue to support the organization.

I hope you will consider resigning, too.

Regards,
Philip

Misuse of RLAs and Traditional Post-Election Audits

As of 2026, seven states require RLAs by law, with about fifteen authorizing or piloting them. All fifty states perform some post-election audit. Traditional fixed-percentage audits review a set portion of ballots or precincts (often 1–5%) without statistical guarantees or required escalation.

Both RLAs and traditional audits fail when the paper trail is untrustworthy—such as records from universal ballot-marking

devices. Real-world examples include pilots in Georgia (2020) and Philadelphia, where claims of “confirmed” outcomes were made despite fundamental weaknesses. Traditional audits are even less effective due to weaker sampling and no mandated full recount. In all cases, audits check only a subset of records. This limited scope allows bad actors to manipulate results in ways that avoid detection in the sampled portion.

The Core Vulnerability — Untrustworthy Audit Trails

A trustworthy paper trail directly captures voter intent and maintains unbroken chain of custody. Audits examine only samples and cannot retroactively establish integrity. Claiming that an audit proves accuracy without this foundation is security theater. It creates the dangerous belief that “checking some of the garbage” suffices. Throwing ballots into complex systems and auditing a fraction of them is an absurd method for verifying the people’s will. Bad actors can exploit this by concentrating manipulation outside the audited sample, knowing the process will likely pass.

National Implications of False Assurance

Election officials’ routine claims that audits demonstrate accurate results represent a critical and dangerous error. These statements foster public complacency while sustaining systems vulnerable to undetected error or fraud. Elections are the highest-value targets in our Republic. False assurance from audits breeds opportunities for manipulation at scale. Bad actors can alter outcomes without triggering full recounts, undermining governance, policy, and the rule of law.

The critical mistake is not merely technical. Officials, having been given false assurances about the power of audits, sincerely believe and transmit those assurances. Citizens receive and accept them. The result is a closed loop of deception that shields unverifiable systems from scrutiny and chastises those who question them.”

The only legitimate way for American citizens to trust election results is direct hand-counting of paper ballots by citizens themselves at the precinct level on election night. This approach eliminates reliance on machines, limits absentee ballots to verified necessities, uses small precincts (≤1,500 voters), requires bipartisan citizen counting under video recording, and makes all images, tallies, and records publicly available for immediate verification. Results are posted locally the same night. Public audits by citizens and students provide ongoing checks. This method replaces blind trust in partial audits with verified accuracy through transparent, observable counting where cast. Audits cannot substitute for this foundational process. Continued dependence on them invites exploitation and erodes our Republic.

Conclusion

Risk-limiting audits improve statistical rigor over traditional audits only when applied to trustworthy hand-marked paper trails. However, all audits share the same fatal flaw: they provide **no affirmative evidence** without verified records and can **mask manipulation**.

The continued reliance on audits—despite the inventor’s public condemnation—demonstrates that the false claims persist. **This is the whistle that must be blown: audits cannot and do not provide the verified accuracy citizens require.**

The solution is clear: return to simple, transparent, precinct-based hand counting by the people themselves, with full public access to records. Only verified accuracy—through direct observation and counting—restores legitimate trust in election results. Without this shift, the illusion of assurance will sever the final thread of our Constitutional Republic.