Betting against the losers: supermartingale tests of election outcomes

H.B. Keller Colloquium
Department of Computing and Mathematical Sciences
Caltech

Philip B. Stark

14 April 2025

University of California, Berkeley

Many collaborators including (most recently):

Andrew Appel, Josh Benaloh, Michelle Blom, Andrew Conway, Rich DeMillo, Alexander Ek, Floyd Everest, Amanda Glazer, Alex Halderman, Zhuoqun Huang, Harri Hursti, Wojciech Jamroga, Dan King, Mark Lindeman, Kellie Ottoboni, Aaditya Ramdas, Ron Rivest, Peter Rœnne, Peter Ryan, Laurent Sandrolini, Steve Schneider, Carsten Schürmann, Brian Sheehan, Jake Spertus, Peter Stuckey, Vanessa Teague, Poorvi Vora, Damjan Vukcevic, Ian Waudby-Smith, Ran Xie

Evidence-Based Elections

P.B. Stark and D.A. Wagner

***Abstract*—We propose an alternative to current requirements for certifying voting equipment and conducting elections. We argue that elections should be structured to provide convincing affirmative evidence that the reported outcomes actually reflect how people voted. This can be accomplished with a combination of software-independent voting systems, compliance audits, and risk-limiting audits. Together, these yield a resilient canvass framework: a fault-tolerant approach to conducting elections that gives strong evidence that the reported outcome is correct or reports that the evidence is not convincing. We argue that, if evidence-based elections are adopted, certification and testing of voting equipment can be relaxed, saving money and time and reducing barriers to innovation in voting systems—and election integrity will benefit. We conclude that there should be more regulation of the evidence trail and less regulation of equipment, and that compliance audits and risk-limiting audits should be required.**

***Keywords*-elections, software-independent voting system, risk-limiting audit, resilient canvass framework EDICS SEC-INTE, APP-CRIM, APP-INTE, APP-OTHE.**

I. Introduction

IDEALLY, what should an election do? Certainly, an election should find out who won, but we believe it also should produce convincing evidence that it found the real winners—or report that it cannot. This is not automatic; it requires thoughtful design of voting equipment, carefully planned and implemented voting and vote counting processes, and rigorous post-election auditing.

While approximately 75% of US voters currently vote on equipment that produces a voter-verifiable paper record of the vote, about 25% vote on paperless electronic voting machines that do not produce such a record [1].

Because paperless electronic voting machines rely upon complex software and hardware, and because there is no feasible way to ensure that the voting software is free of bugs or that the hardware is executing the proper software, there is no guarantee that electronic voting machines record the voter's votes accurately. And, because paperless voting machines preserve only an electronic record of the vote that cannot be directly observed by voters, there is no way to produce convincing evidence that the electronic record accurately reflects the voters' intent. Internet voting shares the shortcomings of paperless electronic voting machines, and has additional vulnerabilities.

Numerous failures of electronic voting equipment have been documented. Paperless voting machines in Carteret County, North Carolina irretrievably lost 4,400 votes; other machines in Mecklenburg, North Carolina recorded 3,955 more votes than the number of people who voted; in Bernalillo County, New Mexico, machines recorded 2,700 more votes than voters; in Mahoning County, Ohio, some machines reported a negative total vote count; and in Fairfax, Virginia, county officials found that for every hundred or so votes cast for one candidate, the electronic voting machines subtracted one vote for her [2]. In short, when elections are conducted on paperless voting

EVIDENCE-BASED ELECTIONS: CREATE A MEANINGFUL PAPER TRAIL, THEN AUDIT

Andrew W. Appel* & Philip B. Stark**

CITE AS: 4 GEO. L. TECH. REV. 523 (2020)

TABLE OF CONTENTS

Desirable properties for voting systems

- (strong) software independence (Rivest & Wack)
- contestibility & defensibility (Appel, DeMillo, Stark)

- Fine to use computers in elections where:
  - failures won't stop people from casting votes
  - other processes will (w/ high prob) catch and correct material errors

- Fine to use computers in elections where:
  - failures won't stop people from casting votes
  - other processes will (w/ high prob) catch and correct material errors
- Sensible threshold for materiality: altered an electoral outcome.

- Fine to use computers in elections where:
  - failures won't stop people from casting votes
  - other processes will (w/ high prob) catch and correct material errors
- Sensible threshold for materiality: altered an electoral outcome.
- "*Can* be checked and corrected" isn't enough. Need "*will* be checked and corrected."

Risk-Limiting Audit (RLA)

Limit *risk* that an incorrect outcome will be certified.

Corrects wrong reported outcomes w/ high probability.

Never alters a correct reported outcome.

Risk: maximum chance of certifying the outcome if the outcome is in fact wrong.

Risk-Limiting Audit (RLA)

Limit *risk* that an incorrect outcome will be certified.

Corrects wrong reported outcomes w/ high probability.

Never alters a correct reported outcome.

Risk: maximum chance of certifying the outcome if the outcome is in fact wrong.

RLA cannot *restore* trustworthiness to a poorly run election.

Leverages trustworthiness of the vote record in a well-run election to provide affirmative evidence that the reported winners really won–or correct the results if not.

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group
- Can audits help protect elections that rely on vulnerable technology?

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group
- Can audits help protect elections that rely on vulnerable technology?
- Typical statutory audits: fixed percentage of batches, no consequences
    - Are the machines working properly?

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group
- Can audits help protect elections that rely on vulnerable technology?
- Typical statutory audits: fixed percentage of batches, no consequences
  - Are the machines working properly?
- Detection (state-of-the-art 2007):
  - the audit should have big chance of finding at least one error if outcome is wrong
  - but errors often occur even when the outcome is right!

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group
- Can audits help protect elections that rely on vulnerable technology?
- Typical statutory audits: fixed percentage of batches, no consequences
    - Are the machines working properly?
- Detection (state-of-the-art 2007):
    - the audit should have big chance of finding at least one error if outcome is wrong
    - but errors often occur even when the outcome is right!
- Affirmative evidence:
    - has the audit given strong evidence that the reported winners really won?
    - if not, collect more evidence or hand-count all votes in trustworthy record to see who won

- Origin of RLAs: 2007 California Post Election Audit Standards Working Group
- Can audits help protect elections that rely on vulnerable technology?
- Typical statutory audits: fixed percentage of batches, no consequences
  - Are the machines working properly?
- Detection (state-of-the-art 2007):
  - the audit should have big chance of finding at least one error if outcome is wrong
  - but errors often occur even when the outcome is right!
- Affirmative evidence:
  - has the audit given strong evidence that the reported winners really won?
  - if not, collect more evidence or hand-count all votes in trustworthy record to see who won

"Trustworthy" means a complete, accurate count would show who really won.

Some records born untrustworthy: malleable or vulnerable tech btw voter & record, such as BMDs or client/Internet/server.

Some records born untrustworthy: malleable or vulnerable tech btw voter & record, such as BMDs or client/Internet/server.

To stay trustworthy, need:

- physical inventories of ballots & other materials
- demonstrably secure chain of custody
- appropriate physical security
- eligibility audits
- ballot accounting
- pollbook and participation reconciliation
- comparisons with registration
- trustworthy upper bound on # validly cast cards containing each contest

Some records born untrustworthy: malleable or vulnerable tech btw voter & record, such as BMDs or client/Internet/server.

To stay trustworthy, need:

- physical inventories of ballots & other materials
- demonstrably secure chain of custody
- appropriate physical security
- eligibility audits
- ballot accounting
- pollbook and participation reconciliation
- comparisons with registration
- trustworthy upper bound on # validly cast cards containing each contest

Can't have cyber-resilience w/o some physical security

RLA pseudo-algorithm

Input: upper bound on cards cast, card identifiers, trustworthy collection of cards

Output: strong evidence that reported outcome is correct, the correct outcome, or statement that available cards don't suffice to determine the outcome

```
while (!(full handcount) && !(strong evidence outcome is correct)) {
    examine more ballots
}
```

RLA pseudo-algorithm

Input: upper bound on cards cast, card identifiers, trustworthy collection of cards

Output: strong evidence that reported outcome is correct, the correct outcome, or statement that available cards don't suffice to determine the outcome

```
while (!(full handcount) && !(strong evidence outcome is correct)) {
    examine more ballots
}

if (full handcount) {
    if (missing cards couldn't change the outcome) {
        handcount result replaces reported result
    }
    else {
        declare outcome indeterminate
    }
```

The National Academies of SCIENCES ENGINEERING MEDICINE

Securing the Vote: Protecting American Democracy, 2018

Home

Elections should be conducted with human-readable paper ballots. Paper ballots form a body of evidence that is not subject to manipulation by faulty software or hardware and that can be used to audit and verify the results of an election. Human-readable paper ballots may be marked by hand or by machine (using a ballot-marking device), and they may be counted by hand or by machine (using an optical scanner), the report says. Voters should have an opportunity to review and confirm their selections before depositing the ballot for tabulation. Voting machines that do not provide the capacity for independent auditing – i.e., machines that do not produce a printout of a voter's selections that can be verified by the voter and used in audits – should be removed from service as soon as possible.

States should mandate a specific type of audit known as a "risk-limiting" audit prior to the certification of election results. By examining a statistically appropriate random sample of paper ballots, risk-limiting audits can determine with a high level of confidence whether a reported election outcome reflects a correct tabulation

- Endorsed by NASEM, PCEA, ASA, LWV, CC, VV, . . .
- $\sim$60 pilot audits in about 17 states and DK
- Laws in $\sim$15 states—not all good
- Methods for plurality, multi-winner plurality, supermajority, proportional representation, IRV/RCV, Borda count, STAR-voting, all ‘scoring rules’

The Annals of Applied Statistics
2008, Vol. 2, No. 2, 550–581
DOI: 10.1214/08-AOAS161

CONSERVATIVE STATISTICAL POST-ELECTION AUDITS

BY PHILIP B. STARK

University of California, Berkeley

There are many sources of error in counting votes: the apparent winner might not be the rightful winner. Hand tallies of the votes in a random sample of precincts can be used to test the hypothesis that a full manual recount would find a different outcome. This paper develops a conservative sequential test based on the vote-counting errors found in a hand tally of a simple or stratified random sample of precincts. The procedure includes a natural escalation: If the hypothesis that the apparent outcome is incorrect is not rejected at stage s, more precincts are audited. Eventually, either the hypothesis is rejected—and the apparent outcome is confirmed—or all precincts have been audited and the true outcome is known. The test uses a priori bounds on the overstatement of the margin that could result from error in each precinct. Such bounds can be derived from the reported counts in each precinct and upper bounds on the number of votes cast in each precinct. The test allows errors in different precincts to be treated differently to reflect voting technology or precinct sizes. It is not optimal, but it is conservative: the chance of erroneously confirming the outcome of a contest if a full manual recount would show a different outcome is no larger than the nominal significance level. The approach also gives a conservative P-value for the hypothesis that a full manual recount would find a different outcome, given the errors found in a fixed size sample. This is illustrated with two contests from November, 2006: the U.S. Senate race in Minnesota and a school board race for the Sausalito Marin City School District in California, a small contest in which voters could vote for up to three candidates.

1. Introduction. Votes can be miscounted because of human error (by voters or election workers), hardware or software "bugs" or deliberate fraud. Post-election audits—manual tallies of votes in individual precincts—are intended to detect miscount, especially miscount large enough to alter the outcome of the election.[1] To the best of my knowledge, eighteen states require or allow post-election audits [National Association of Secretaries of State (2007) and Verified Voting

Received October 2007; revised March 2008.

EVT-WOTE 2009

Implementing Risk-Limiting Post-Election Audits in California

Joseph Lorenzo Hall[1,2,*], Luke W. Miratrix[3], Philip B. Stark[3], Melvin Briones[4], Elaine Ginnold[4], Freddie Oakley[5], Martin Peaden[6], Gail Pellerin[6], Tom Stanionis[5], and Tricia Webber[6]

[1]University of California, Berkeley; School of Information
[2]Princeton University; Center for Information Technology Policy
[3]University of California, Berkeley; Department of Statistics
[4]Marin County, California; Registrar of Voters
[5]Yolo County, California; County Clerk/Recorder
[6]Santa Cruz County, California; County Clerk

Abstract

Risk-limiting post-election audits limit the chance of certifying an electoral outcome if the outcome is not what a full hand count would show. Building on previous work [18, 17, 20, 21, 11], we report pilot risk-limiting audits in four elections during 2008 in three California counties: one during the February 2008 Primary Election in Marin County and three during the November 2008 General Elections in Marin, Santa Cruz and Yolo Counties. We explain what makes an audit *risk-limiting* and how existing and proposed laws fall short. We discuss the differences among our four pilot audits. We identify challenges to practical, efficient risk-limiting audits and conclude that current approaches are too complex to be used routinely on a large scale. One important logistical bottleneck is the difficulty of exporting data from commercial election management systems in a format amenable to audit calculations. Finally, we propose a bare-bones risk-limiting audit that is less efficient than these pilot audits, but avoids many practical problems.

Sets of Half-Average Nulls Generate Risk-Limiting Audits: SHANGRLA

ESORICS Voting 20, LNCS

Philip B. Stark

University of California, Berkeley

31 January 2020

Abstract. Risk-limiting audits (RLAs) for many social choice functions can be reduced to testing sets of null hypotheses of the form "the average of this list is not greater than 1/2" for a collection of finite lists of non-negative numbers. Such social choice functions include majority, super-majority, plurality, multi-winner plurality, Instant Runoff Voting (IRV), Borda count, approval voting, and STAR-Voting, among others. The audit stops without a full hand count iff all the null hypotheses are rejected. The nulls can be tested in many ways. Ballot polling is particularly simple; two new ballot-polling risk-measuring functions for sampling without replacement are given. Ballot-level comparison audits transform each null into an equivalent assertion that the mean of re-scaled tabulation errors is not greater than 1/2. In turn, that null can then be tested using the same statistical methods used for ballot polling—applied to different finite lists of non-negative numbers: the SHANGRLA approach reduces auditing different social choice functions and different audit methods to the same simple statistical problem. Moreover, SHANGRLA comparison audits are more efficient than previous comparison audits for two reasons: (i) for most social choice functions, the conditions tested are both necessary and sufficient for the reported outcome to be correct, while previous methods tested conditions that were sufficient but not necessary, and (ii) the tests avoid a conservative approximation. The SHANGRLA abstraction simplifies stratified audits, including audits that combine ballot polling with ballot-level comparisons, producing sharper audits than the "SUITE" approach. SHANGRLA works with the "phantoms to evil zombies" strategy to treat missing ballot cards and missing or redacted cast vote records. That also facilitates sampling from "ballot-style manifests," which can dramatically improve efficiency when the audited contests do not appear on every ballot card. Open-source software implementing SHANGRLA ballot-level comparison audits is available. SHANGRLA was tested in a pilot audit of an instant-runoff contest in San Francisco, CA, in November, 2019.

Keywords: sequential tests, martingales, Kolmogorov's inequality

Acknowledgments: I am grateful to Andrew Conway, Steven N. Evans, Kellie Ottoboni, Ronald L. Rivest, Vanessa Teague, and Poorvi Vora for helpful conversations and comments on earlier drafts. The SHANGRLA software was a collaborative effort that included Michelle Blom, Andrew Conway, Dan King, Laurent Sandrolini, Peter Stuckey, and Vanessa Teague.

The Annals of Applied Statistics
2023, Vol. 17, No. 1, 641–679
https://doi.org/10.1214/22-AOAS1646

ALPHA: AUDIT THAT LEARNS FROM PREVIOUSLY HAND-AUDITED BALLOTS

BY PHILIP B. STARK[a]

Department of Statistics, University of California, Berkeley, [a]*stark@stat.berkeley.edu*

A risk-limiting election audit (RLA) offers a statistical guarantee: if the reported electoral outcome is incorrect, the audit has a known maximum chance (the risk limit) of not correcting it before it becomes final. BRAVO (Lindeman, Stark and Yates (In *Proceedings of the 2011 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE'11)* (2012) USENIX)), based on Wald's sequential probability ratio test for the Bernoulli parameter, is the simplest and most widely tried method for RLAs, but it has limitations. It cannot accommodate sampling without replacement or stratified sampling which can improve efficiency and are sometimes required by law. It applies only to ballot-polling audits which are less efficient than comparison audits. It applies to plurality, majority, supermajority, proportional representation, and instant-runoff voting (IRV, using RAIRE (Blom, Stuckey and Teague (In *Electronic Voting* (2018) 17–34 Springer))) but not to other social choice functions for which there are RLA methods. And while BRAVO has the smallest expected sample size among sequentially valid ballot-polling-with-replacement methods when the reported vote shares are exactly correct, it can require arbitrarily large samples when the reported reported winner(s) really won but the reported vote shares are incorrect. ALPHA is a simple generalization of BRAVO that: (i) works for sampling with and without replacement, with and without weights, with and without stratification, and for Bernoulli sampling; (ii) works not only for ballot polling but also for ballot-level comparison, batch polling, and batch-level comparison audits; (iii) works for all social choice functions covered by SHANGRLA (Stark (In *Financial Cryptography and Data Security* (2020) Springer)), including approval voting, STAR-Voting, proportional representation schemes, such as D'Hondt and Hamilton, IRV, Borda count, and all scoring rules, and (iv) in situations where both ALPHA and BRAVO apply, requires smaller samples than BRAVO when the reported vote shares are wrong but the outcome is correct—five orders of magnitude in some examples. ALPHA includes the family of betting martingale tests in RiLACS (Waudby-Smith, Stark and Ramdas (In *Electronic Voting. E-Vote-ID 2021* (2021) Springer)) with a different betting strategy parametrized as an estimator of the population mean and explicit flexibility to accommodate sampling weights and population bounds that change with each draw. A Python implementation is provided.

Non(c)esuch Ballot-Level Comparison Risk-Limiting Audits

ESORICS, Voting 2022, LNCS

Philip B. Stark(✉)

University of California, Berkeley, CA, USA
stark@stat.berkeley.edu

Abstract. Risk-limiting audits (RLAs) guarantee a high probability of correcting incorrect reported electoral outcomes before the outcomes are certified. The most efficient are *ballot-level comparison audits* (BLCAs), which compare the voting system's interpretation of randomly selected individual ballot cards (*cast-vote records*, CVRs) from a trustworthy paper trail to a human interpretation of the same cards. BLCAs have logistical and privacy hurdles: Individual randomly selected cards must be retrieved for manual inspection; the voting system must export CVRs; and the CVRs must be linked to the corresponding physical cards, to compare the two. In practice, such links have been made by keeping cards in the order in which they are scanned or by printing serial numbers on cards as they are scanned. Both methods may compromise voter privacy. Cards selected for audit have been retrieved by manually counting into stacks or by looking for cards with particular serial numbers. The methods are time-consuming; the first is also error-prone. Connecting CVRs to cards using a unique pseudo-random number ("cryptographic nonce") printed on each card after the voter last sees it could reduce privacy risks, but retrieving the card imprinted with a particular random number may be harder than counting into a stack or finding the card with a given serial number. And what if the system does not in fact print a unique number on each ballot or does not accurately report the numbers it printed? This paper presents a method for conducting BLCAs that maintains the risk limit even if the system does not print a genuine nonce on each ballot or misreports the identifiers it used. The method also allows untrusted technology to be used to retrieve the cards selected for audit—automation that may reduce audit workload even if cards are imprinted with serial numbers rather than putative nonces. The method limits the risk rigorously, even if the imprinting or retrieval technology misbehaves. If the imprinting and retrieval systems behave properly, this protection does not increase the number of cards the RLA has to inspect to confirm or correct the outcome.

Keywords: Risk-limiting audit · Voter privacy

Overstatement-Net-Equivalent Risk-Limiting Audit: ONEAudit

ESORICS Voting 23: LNCS

Philip B. Stark[1] [0000−0002−3771−9604]

University of California, Berkeley, CA USA stark@stat.berkeley.edu

Abstract. A procedure is a risk-limiting audit (RLA) with risk limit α if it has probability at least $1-\alpha$ of correcting each wrong reported outcome and never alters correct outcomes. One efficient RLA method, card-level comparison (CLCA), compares human interpretation of individual ballot cards randomly selected from a trustworthy paper trail to the voting system's interpretation of the same cards (cast vote records, CVRs). CLCAs heretofore required a CVR for each cast card and a "link" identifying which CVR is for which card—which many voting systems cannot provide. This paper shows that every set of CVRs that produces the same aggregate results overstates contest margins by the same amount: they are *overstatement-net-equivalent* (ONE). CLCA can therefore use CVRs from the voting system for any number of cards and ONE CVRs created *ad lib* for the rest. In particular:

– Ballot-polling RLA is algebraically equivalent to CLCA using ONE CVRs derived from the overall contest results.
– CLCA can be based on batch-level results (e.g., precinct subtotals) by constructing ONE CVRs for each batch. In contrast to batch-level comparison auditing (BLCA), this avoids manually tabulating entire batches and works even when reporting batches do not correspond to physically identifiable batches of cards, when BLCA is impractical.
– If the voting system can export linked CVRs for only some ballot cards, auditors can still use CLCA by constructing ONE CVRs for the rest of the cards from contest results or batch subtotals.

This works for every social choice function for which there is a known RLA method, including IRV. Sample sizes for BPA and CLCA using ONE CVRs based on contest totals are comparable. With ONE CVRs from batch subtotals, sample sizes are smaller than than for BPA when batches are homogeneous, approaching those of CLCA using CVRs from the voting system, and much smaller than for BLCA: A CLCA of the 2022 presidential election in California at risk limit 5% using ONE CVRs for precinct-level results would sample approximately 70 ballots statewide, if the reported results are accurate, compared to about 26,700 for BLCA. The 2022 Georgia audit tabulated more than 231,000 cards (the expected BLCA sample size was ≈103,000 cards); ONEAudit would have audited ≈1,300 cards. For data from a pilot hybrid RLA in Kalamazoo, MI, in 2018, ONEAudit gives a risk of 2%, substantially lower than the 3.7% measured risk for SUITE, the "hybrid" method the pilot used.

Keywords: Risk-limiting audit, BPA, card-level comparison audit, batch-level comparison audit

You can do RLAs for IRV

The Process Pilot of Risk-Limiting Audits for the San Francisco District Attorney 2019 Instant Runoff Vote

Michelle Blom,* Andrew Conway,† Dan King,‡ Laurent Sandrolini,§
Philip B. Stark,¶ Peter J. Stuckey‖ and Vanessa Teague**

April 2, 2020

The City and County of San Francisco, CA, has used Instant Runoff Voting (IRV) for some elections since 2004. This report describes the first ever process pilot of Risk Limiting Audits for IRV, for the San Francisco District Attorney's race in November, 2019. We found that the vote-by-mail outcome could be efficiently audited to well under the 0.05 risk limit given a sample of only 200 ballots. All the software we developed for the pilot is open source.

1. Introduction

Post-election audits test a reported election result by randomly sampling paper ballots.[1] A *Risk Limiting Audit (RLA)* of a trustworthy paper trail of votes either finds strong statistical evidence that the reported outcome is correct, or reverts to a full manual tabulation to set the record straight.[2] (The outcome is the political result—i.e., who won—not the exact vote counts.) The maximum chance that a RLA will fail to correct the reported outcome if the reported outcome is wrong is the *risk limit.* RLAs are

*School of Computing and Information Systems, University of Melbourne. michelle.blom@unimelb.edu.au
†Silicon Econometrics Pty. Ltd., andrewelections@greatcactus.org

Adaptively Weighted Audits of Instant-Runoff Voting Elections: AWAIRE

2023 E-Vote-ID

Alexander Ek[1], Philip B. Stark[2], Peter J. Stuckey[3], and Damjan Vukcevic[1](✉)

[1] Department of Econometrics and Business Statistics, Monash University, Clayton, Australia
damjan.vukcevic@monash.edu
[2] Department of Statistics, University of California, Berkeley, CA, USA
[3] Department of Data Science and AI, Monash University, Clayton, Australia

Abstract. An election audit is *risk-limiting* if the audit limits (to a pre-specified threshold) the chance that an erroneous electoral outcome will be certified. Extant methods for auditing instant-runoff voting (IRV) elections are either not risk-limiting or require cast vote records (CVRs), the voting system's electronic record of the votes on each ballot. CVRs are not always available, for instance, in jurisdictions that tabulate IRV contests manually.

We develop an RLA method (AWAIRE) that uses adaptively weighted averages of test supermartingales to efficiently audit IRV elections when CVRs are not available. The adaptive weighting 'learns' an efficient set of hypotheses to test to confirm the election outcome. When accurate CVRs are available, AWAIRE can use them to increase the efficiency to match the performance of existing methods that require CVRs.

We provide an open-source prototype implementation that can handle elections with up to six candidates. Simulations using data from real elections show that AWAIRE is likely to be efficient in practice. We discuss how to extend the computational approach to handle elections with more candidates.

Adaptively weighted averages of test supermartingales are a general tool, useful beyond election audits to test collections of hypotheses sequentially while rigorously controlling the familywise error rate.

Stylish Risk-Limiting Audits in Practice

Amanda K. Glazer[1,2], Jacob V. Spertus[3], Philip B. Stark[4]

Abstract: Risk-limiting audits (RLAs) can use information about which ballot cards contain which contests (*card-style data*, CSD) to ensure that each contest receives adequate scrutiny, without examining more cards than necessary. RLAs using CSD in this way can be substantially more efficient than RLAs that sample indiscriminately from all cast cards. We describe an open-source Python implementation of RLAs using CSD for the Hart InterCivic Verity voting system and the Dominion Democracy Suite® voting system. The software is demonstrated using all 181 contests in the 2020 general election and all 214 contests in the 2022 general election in Orange County, CA, USA, the fifth-largest election jurisdiction in the U.S., with over 1.8 million active voters. (Orange County uses the Hart Verity system.) To audit the 181 contests in 2020 to a risk limit of 5% without using CSD would have required a complete hand tally of all 3,094,308 cast ballot cards. With CSD, the estimated sample size is about 20,100 cards, 0.65% of the cards cast—including one tied contest that required a complete hand count. To audit the 214 contests in 2022 to a risk limit of 5% without using CSD would have required a complete hand tally of all 1,989,416 cast cards. With CSD, the estimated sample size is about 62,250 ballots, 3.1% of cards cast—including three contests with margins below 0.1% and 9 with margins below 0.5%.

Assorters: connect outcome correctness to means of lists

Plurality contest. Alice is the reported winner; Bob, Candy lost.

A $\rightarrow 1$

B $\rightarrow 0$

C, invalid $\rightarrow 1/2$

Alice beat Bob if mean $> 1/2$

A $\rightarrow 1$

C $\rightarrow 0$

B, invalid $\rightarrow 1/2$

Alice beat Candy if mean $> 1/2$

Scoring rule: assign at most u points to each candidate. Highest sum wins.

Alice is the reported winner; Bob, Candy reportedly lost.

List for Alice v Bob:

$$\text{value for } i\text{th card} = \frac{\text{score(Alice)} - \text{score(Bob)} + u}{2u}$$

Alice beat Bob if mean $> 1/2$.

List for Alice v Candy:

$$\text{value for } i\text{th card} = \frac{\text{score(Alice)} - \text{score(Candy)} + u}{2u}$$

Alice beat Candy if mean $> 1/2$.

Lists are nonnegative.

Supermajority: the issue must receive at least a fraction f of the valid votes to pass.

List to confirm that the issue won, if it reportedly won:

$$\text{value for } i\text{th card} = \begin{cases} \frac{1}{2f}, & \text{card shows a vote for issue} \\ 0, & \text{card shows a vote against issue} \\ 1/2, & \text{card does not have a valid vote.} \end{cases}$$

Issue received more than f of the valid votes if mean $> 1/2$.

List is nonnegative and bounded by $1/(2f)$.

Similar way to characterize winner of D'Hondt, Hamiltonian, STAR-Voting, IRV/RCV, ...

Only special cases of STV so far.

Heuristic: in a series of fair or unfavorable games, it is unlikely that your fortune will ever reach a large multiple of your initial stake.

Example: betting on a fair coin.

- Start with $1 stake.
- Bet on the outcome of the next coin toss:
    - If coin lands heads, you get back your bet, doubled.
    - If coin lands tails, you lose your bet.
    - Can bet up to your current fortune, but can't borrow.
    - Can use any betting scheme whatsoever
- If you go broke, you're out.

Chance fortune ever reaches $10 is at most 10% (1/10).
Chance it ever reaches $20 is at most 5% (1/20).
Chance it ever reaches k is at most $1/k$.

Heuristic: in a series of fair or unfavorable games, it is unlikely that your fortune will ever reach a large multiple of your initial stake.

Example: betting on a fair coin.

- Start with \$1 stake.
- Bet on the outcome of the next coin toss:
  - If coin lands heads, you get back your bet, doubled.
  - If coin lands tails, you lose your bet.
  - Can bet up to your current fortune, but can't borrow.
  - Can use any betting scheme whatsoever
- If you go broke, you're out.

Chance fortune ever reaches \$10 is at most 10% (1/10).
Chance it ever reaches \$20 is at most 5% (1/20).
Chance it ever reaches k is at most $1/k$.

If your fortune reaches \$100, strong evidence that the coin isn't fair.

Betting on ballots

Create set of repeated casino games, at least one of which is fair or unfavorable if any reported winner didn't really win. (SHANGRLA + ALPHA)

- Each game involves betting on the next number sampled at random from a list.
- Each game involves a different list.
- If the outcome is wrong, the mean of at least one of the lists is $\leq 1/2$.
- Start with a stake of $1 in each game.
- Bet using any "non-anticipating" strategy (can't peek into the future).
- If your fortune gets to $20 in every game, audit stops.
- If you go broke in any game, do a full hand count.
- If you don't get to $20 in every game (or get bored), full hand count.

Betting on the mean

μ: hypothesized upper bound on the mean.

In tth game, wager λ_t on the value X_t that will be drawn next. (Constraints on λ_t ensure wealth ≥ 0)

μ_t is (upper bound on) the mean of the list after $t-1$st draw, if the null is true.

Bet pays $\lambda_t(X_t - \mu_t)$: expected value ≤ 0

Wealth after tth wager is

$$M_t := \prod_{i=1}^{t} (1 + \lambda_i(X_i - \mu_i)), \tag{1}$$

If null is true, expected to break even or lose money on each bet.

For sampling with replacement, μ_t is the same in each draw.

For sampling without replacement,

$$\mu_t = \frac{1}{N - t + 1}\left(N\mu - \sum_{j=1}^{t-1} X_j\right).$$

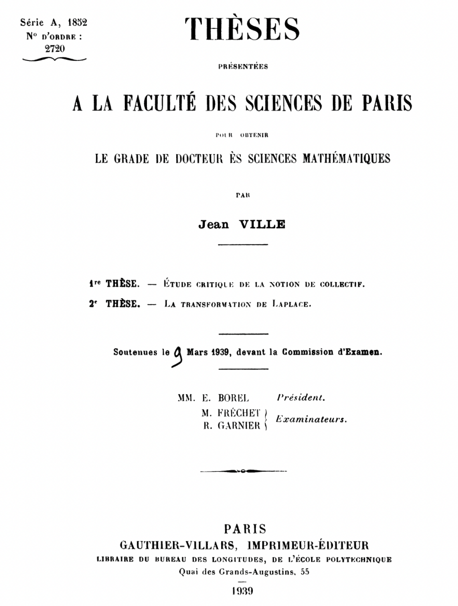

Série A, 1852
N° D'ORDRE :
2720

THÈSES

PRÉSENTÉES

A LA FACULTÉ DES SCIENCES DE PARIS

POUR OBTENIR

LE GRADE DE DOCTEUR ÈS SCIENCES MATHÉMATIQUES

PAR

Jean VILLE

1re THÈSE. — Étude critique de la notion de collectif.

2e THÈSE. — La transformation de Laplace.

Soutenues le 9 Mars 1939, devant la Commission d'Examen.

MM. E. BOREL *Président.*
M. FRÉCHET, R. GARNIER *Examinateurs.*

PARIS
GAUTHIER-VILLARS, IMPRIMEUR-ÉDITEUR
LIBRAIRE DU BUREAU DES LONGITUDES, DE L'ÉCOLE POLYTECHNIQUE
Quai des Grands-Augustins, 55

1939

Ville (1939): In a sequence of fair or sub-fair wagers in which you aren't allowed to borrow money, the chance you ever multiply your bankroll by k is at most $1/k$

At most 5% chance you get to \$20 in every game if any reported winner didn't really win.

Thus, RLA with risk limit 5%.

Better betting strategies → more efficient audits: current research.

Unlikely to make much money unless the game is favorable.

If game is favorable, what betting strategy will grow wealth fastest?

Picking the bet λ_t

BSTJ 1956

A New Interpretation of Information Rate

By J. L. KELLY, JR.

(Manuscript received March 21, 1956)

If the input symbols to a communication channel represent the outcomes of a chance event on which bets are available at odds consistent with their probabilities (i.e., "fair" odds), a gambler can use the knowledge given him by the received symbols to cause his money to grow exponentially. The maximum exponential rate of growth of the gambler's capital is equal to the rate of transmission of information over the channel. This result is generalized to include the case of arbitrary odds.

Thus we find a situation in which the transmission rate is significant even though no coding is contemplated. Previously this quantity was given significance only by a theorem of Shannon's which asserted that, with suitable encoding, binary digits could be transmitted over the channel at this rate with an arbitrarily small probability of error.

INTRODUCTION

Shannon defines the rate of transmission over a noisy communication channel in terms of various probabilities.[1] This definition is given sig-

Journal of the Royal Statistical Society Series B: Statistical Methodology, 2024, **86**, 1–27
https://doi.org/10.1093/jrsssb/qkad009
Advance access publication 16 February 2023
Discussion Paper

ROYAL STATISTICAL SOCIETY DATA | EVIDENCE | DECISIONS — Journal of the Statistics Society Series B Statistical Methodology B

Estimating means of bounded random variables by betting

Ian Waudby-Smith[1] and Aaditya Ramdas[1,2]

[1]Department of Statistics & Data Science, Carnegie Mellon University, Pittsburgh, PA, USA
[2]Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA, USA

Address for correspondence: Aaditya Ramdas, Machine Learning, Carnegie Mellon University, 5000 Forbes Ave, 132 Baker Hall, Pittsburgh, PA 15213, USA. Email: aramdas@cmu.edu

Read before The Royal Statistical Society at the Discussion Meeting organized by the Research Section on Tuesday, 23 May 2023, Dr Robin Evans in the Chair.

Abstract

We derive confidence intervals (CIs) and confidence sequences (CSs) for the classical problem of estimating a bounded mean. Our approach generalizes and improves on the celebrated Chernoff method, yielding the best closed-form "empirical-Bernstein" CSs and CIs (converging exactly to the oracle Bernstein width) as well as non-closed-form "betting" CSs and CIs. Our method combines new composite nonnegative (super) martingales with Ville's maximal inequality, with strong connections to testing by betting and the method of mixtures. We also show how these ideas can be extended to sampling without replacement. In all cases, our bounds are adaptive to the unknown variance, and empirically vastly outperform prior approaches, establishing a new state-of-the-art for four fundamental problems: CSs and CIs for bounded means, when sampling with and without replacement.

Mathematical Finance, Vol. 1, No. 1 (January 1991), 1–29

UNIVERSAL PORTFOLIOS

THOMAS M. COVER[1]

Departments of Statistics and Electrical Engineering, Stanford University, Stanford, CA

We exhibit an algorithm for portfolio selection that asymptotically outperforms the best stock in the market. Let $\mathbf{x}_i = (x_{i1}, x_{i2}, \ldots, x_{im})^t$ denote the performance of the stock market on day i, where x_{ij} is the factor by which the jth stock increases on day i. Let $\mathbf{b}_i = (b_{i1}, b_{i2}, \ldots, b_{im})^t$, $b_{ij} \geq 0$, $\Sigma_j b_{ij} = 1$, denote the proportion b_{ij} of wealth invested in the jth stock on day i. Then $S_n = \Pi_{i=1}^n \mathbf{b}_i^t \mathbf{x}_i$ is the factor by which wealth is increased in n trading days. Consider as a goal the wealth $S_n^* = \max_{\mathbf{b}} \Pi_{i=1}^n \mathbf{b}^t \mathbf{x}_i$ that can be achieved by the best constant rebalanced portfolio chosen *after* the stock outcomes are revealed. It can be shown that S_n^* exceeds the best stock, the Dow Jones average, and the value line index at time n. In fact, S_n^* usually exceeds these quantities by an exponential factor. Let $\mathbf{x}_1, \mathbf{x}_2, \ldots,$ be an arbitrary sequence of market vectors. It will be shown that the nonanticipating sequence of portfolios $\hat{\mathbf{b}}_k = \int \mathbf{b} \Pi_{i=1}^{k-1} \mathbf{b}^t \mathbf{x}_i \, d\mathbf{b} / \int \Pi_{i=1}^{k-1} \mathbf{b}^t \mathbf{x}_i \, d\mathbf{b}$ yields wealth $\hat{S}_n = \Pi_{k=1}^n \hat{\mathbf{b}}_k^t \mathbf{x}_k$ such that $(1/n) \ln(S_n^* / \hat{S}_n) \to 0$, for every bounded sequence $\mathbf{x}_1, \mathbf{x}_2, \ldots,$ and, under mild conditions, achieves

$$\hat{S}_n \sim \frac{S_n^* (m-1)! (2\pi/n)^{(m-1)/2}}{|J_n|^{1/2}},$$

where J_n is an $(m-1) \times (m-1)$ sensitivity matrix. Thus this portfolio strategy has the same exponential rate of growth as the apparently unachievable S_n^*.

KEYWORDS: portfolio selection, robust trading strategies, performance weighting, rebalancing

436 IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. 70, NO. 1, JANUARY 2024

Tight Concentrations and Confidence Sequences From the Regret of Universal Portfolio

Francesco Orabona, *Member, IEEE*, and Kwang-Sung Jun

***Abstract*—A classic problem in statistics is the estimation of the expectation of random variables from samples. This gives rise to the tightly connected problems of deriving concentration inequalities and confidence sequences, i.e., confidence intervals that hold uniformly over time. Previous studies have shown that it is possible to convert the regret guarantee of an online learning algorithm into concentration inequalities, but these concentration results were not tight. In this paper, we show regret guarantees of universal portfolio algorithms applied to the online learning problem of betting give rise to new implicit time-uniform concentration inequalities for bounded random variables. The key feature of our concentration results is that they are centered around the maximum log wealth of the best fixed betting strategy in hindsight. We propose numerical methods to solve these implicit inequalities, which results in confidence sequences that enjoy the empirical Bernstein rate with the optimal asymptotic behavior while being never worse than Bernoulli-KL confidence bounds. We further show that our confidence sequences are never vacuous with even one sample, for any given target failure rate $\delta \in (0, 1)$. Our empirical study shows that our confidence bounds achieve the state-of-the-art performance, especially in the small sample regime.**

***Index Terms*—Confidence sequence, regret, universal portfolio.**

I. INTRODUCTION

CONSIDER the problem of constructing valid confidence intervals for bounded random variables. We are interested in estimating the conditional mean μ of a sequence of random variables $X_1, X_2, \ldots$, such that for any i we have $0 \leq X_i \leq 1$ almost surely and

$$\mathbb{E}[X_t | X_1, \ldots, X_{t-1}] = \mu. \quad (1)$$

the empirical mean deviates from the expectation μ. A related concept is the one of producing $(1-\delta)$-level *confidence sequence* [1], recently revived by [2], which is a sequence of confidence sets $I_t = I_t(X_1, \ldots, X_t)$ such that

$$\mathbb{P}\{\mu \in I_t, \forall t \geq 1\} \geq 1 - \delta.$$

We are interested in I_t being intervals, i.e., $I_t = [\ell_t, u_t]$ for some $\ell_t < u_t$. The time-uniform property makes these estimates more useful in the real world because they are "safer" to be used, allowing one to monitor the average of samples being collected and then stop gathering new samples at any time point in an adaptive way (e.g., when the required precision has been reached). This safety aspect was recently argued in much more general terms by [3].

In addition to conventional applications of testing in clinical trials and social sciences, confidence sequences have recently found numerous applications in statistics and machine learning. For example, A/B testing, or A/B/N testing more generally, is a form of randomized experimentation where two or more versions of a variable (e.g., layouts, contents, recommended products in web pages or mobile applications) are shown to different segments of website visitors at the same time to determine which version leaves the maximum impact and drives business metrics [4]. Time-uniform confidence sequences play an important role here as they allow one to monitor the data being collected (e.g., clicks or business metrics) to determine when to stop and conclude which alternative is more attractive. Given the limited amount of A/B testings that can be run in reality, it is crucial that the

SEQUENTIAL TESTS OF STATISTICAL HYPOTHESES

By A. Wald

Columbia University

Table of Contents

Ann. Math. Statist. 16 117-186, 1945. 10.1214/aoms/1177731118

Wald's sequential probability ratio test (SPRT)

- Developed during WWII, but seen as so useful it was kept secret until 1945
- Error control is special case of Ville's inequality
- Also proved (essential) optimality

Admissible anytime-valid sequential inference must rely on nonnegative martingales

Aaditya Ramdas[1], Johannes Ruf[2], Martin Larsson[3], Wouter M. Koolen[4]

[1] Departments of Statistics and Machine Learning, Carnegie Mellon University
[2] Department of Mathematics, London School of Economics
[3] Department of Mathematics, Carnegie Mellon University
[4] Machine Learning Group, CWI Amsterdam

aramdas@cmu.edu, j.ruf@lse.ac.uk
martinl@andrew.cmu.edu, wmkoolen@cwi.nl

November 8, 2022

Abstract

Confidence sequences, anytime p-values (called p-processes in this paper), and e-processes all enable sequential inference for composite and nonparametric classes of distributions at arbitrary stopping times. Examining the literature, one finds that at the heart of all these (quite different) approaches has been the identification of nonnegative (super)martingales. Thus, informally, nonnegative (super)martingales are known to be sufficient for *anytime-valid* sequential inference, even in composite and nonparametric settings. Our central contribution is to show that nonnegative martingales are also universal—after appropriately defining *admissibility*, we show that all admissible constructions of confidence sequences, p-processes, or e-processes must necessarily utilize nonnegative martingales. Our proofs utilize several modern mathematical tools for composite testing and estimation problems: max-martingales, Snell envelopes, transfinite induction, and new Doob-Lévy martingales make appearances in previously unencountered ways. Informally, if one wishes to perform anytime-valid sequential inference, then any existing approach can be recovered or dominated using nonnegative martingales. We provide several nontrivial examples, with special focus on testing symmetry, where our new constructions render past methods inadmissible. We also prove the subGaussian supermartingale to be admissible.

Keywords: Admissibility; composite nonnegative supermartingale; p-process; confidence sequence; Doob-Lévy martingale; e-process; max-martingale; optional stopping; Snell envelope; Ville's inequality.

E-value: nonnegative random variable ϵ that has expected value 1 under the null.

E-process: stochastic process $(\epsilon_t)_{t\in\mathbb{N}}$ such that for an arbitrary stopping time τ, ϵ_τ is an E-value

- If ϵ is an E-value, $1 \wedge 1/\epsilon$ is a P-value
- If (ϵ_t) is an E-process, $(1 \wedge 1/\epsilon_t)$ is a sequence of anytime-valid P-values

Blending Bayesian and Frequentist perspectives

Rigorous frequentist error control, but can use Bayesian reasoning, priors, etc., to set the bets.

- If the null is false, good prior helps grow wealth quickly and reject with a small sample size.
- If the null is true, no matter how you bet, unlikely to reject the null.

Common misconceptions

- RLAs are “tabulation audits”
- RLAs work by checking whether the vote shares in the sample “match” the reported vote shares.
- RLAs assume that errors/malfunctions/problems are distributed randomly.
- RLAs are worthwhile even if the paper trail isn’t trustworthy.

RLA isn't "tabulation audit." Doesn't check tabulation: checks who won.

RLA isn't "tabulation audit." Doesn't check tabulation: checks who won.

	ballot 1	ballot 2	ballot 3	ballot 4	ballot 5	ballot 6	ballot 7	ballot 8
machine	Alice	Alice	Alice	Alice	Bob	Bob	Bob	none

RLA isn't "tabulation audit." Doesn't check tabulation: checks whether accurate tabulation would find the same winners.

	ballot 1	ballot 2	ballot 3	ballot 4	ballot 5	ballot 6	ballot 7	ballot 8
machine	Alice	Alice	Alice	Alice	Bob	Bob	Bob	none
reality	Bob	none	none	none	none	none	Alice	Alice

RLA isn't "tabulation audit." Doesn't check tabulation: checks whether accurate tabulation would find the same winners.

	ballot 1	ballot 2	ballot 3	ballot 4	ballot 5	ballot 6	ballot 7	ballot 8
machine	Alice	Alice	Alice	Alice	Bob	Bob	Bob	none
reality	Bob	none	none	none	none	none	Alice	Alice

Every ballot was mistabulated, and the totals are wrong:

tally	Alice	Bob	non-vote
machine	4	3	1
reality	2	1	5

RLA isn't "tabulation audit." Doesn't check tabulation: checks whether accurate tabulation would find the same winners.

	ballot 1	ballot 2	ballot 3	ballot 4	ballot 5	ballot 6	ballot 7	ballot 8
machine	Alice	Alice	Alice	Alice	Bob	Bob	Bob	none
reality	Bob	none	none	none	none	none	Alice	Alice

Every ballot was mistabulated, and the totals are wrong:

tally	Alice	Bob	non-vote
machine	4	3	1
reality	2	1	5

But Alice really won, so it's appropriate for an RLA to stop without a full handcount.

Addressing other misconceptions

- RLAs do *not* work by comparing the vote shares in the sample to the reported vote shares: they assess the evidence that the reported winners won.
- RLAs *do* work whether errors/misbehavior is "random" or adversarial: the risk calculations assume "worst-case"
- Applying RLA procedures to untrustworthy vote records *cannot* provide evidence that reported outcomes are correct: security theater.

Overstatement-net-equivalence

- Any change to the CVRs that keeps vote totals the same (more generally, that keeps assorter totals the same) does not affect correctness
- Any set of CVRs that yields in the same reported winner(s) can be used for auditing
- “Transitive” audits and ONEAudit

Phantoms-to-zombies

- Replace any missing data with worst-case data
- Can only increase the measured risk: can't compromise the risk Limit
- Useful for dealing with unaccounted-for ballots, unretrievable ballots, missing CVRs, ...

Formalizing RLAs as an interactive proof

Trust assumptions:

- upper bound on the number of validly cast cards in each contest
- no cards altered or added (losing cards is OK)
- identifiers on cards can't be added or changed once audit starts

Don't need to trust that prover reported identifiers accurately, labeled every card, didn't re-use labels, etc.

Abuse

RLAs increasingly used to distract from a "game-over" problem: no trustworthy, organized, complete record of expressed preferences of eligible voters who validly cast ballots.

E.g., GA SoS Raffensperger claimed that a (deeply flawed) audit of one contest in 2020 based on untrustworthy paper trail "reaffirmed that the state's new secure paper ballot voting system accurately counted and reported results." https://sos.ga.gov/news/historic-first-statewide-audit-paper-ballots-upholds-result-presidential-race

When Audits and Recounts Distract from Election Integrity: The 2020 U.S. Presidential Election in Georgia

Philip B. Stark[1][0000−0002−3771−9604]

University of California, Berkeley, CA USA pbstark@berkeley.edu

Abstract. The U.S. state of Georgia was central to efforts to overturn the results of the 2020 Presidential election, including a phone call from then-president Donald Trump to Georgia Secretary of State Brad Raffensperger asking Raffensperger to 'find' 11,780 votes. Raffensperger has maintained that a '100% full-count risk-limiting audit' and a machine recount agreed with the initial machine-count results, which proved that the reported election results were accurate and that 'no votes were flipped.' While there is no indication of widespread fraud, there is reason to distrust the election outcome: the two machine counts and the manual 'audit' tallies disagree substantially, even about the number of ballots cast. Some ballots in Fulton County, Georgia, were included in the original count at least twice; some were included in the machine recount at least thrice. Audit handcount results for some tally batches were omitted from the reported audit totals: reported audit results do not include every vote the auditors counted. In short, the two machine counts and the audit were not probative of who won because of poor processes and controls: a lack of secure physical chain of custody, ballot accounting, pollbook reconciliation, and accounting for other election materials such as memory cards. Moreover, most voters voted with demonstrably untrustworthy ballot-marking devices, so even a perfect handcount or audit would not necessarily reveal who really won. True risk-limiting audits (RLAs) and rigorous recounts can limit the risk that an incorrect electoral outcome will be certified rather than being corrected. But no procedure can limit that risk without a trustworthy record of the vote. And even a properly conducted RLA of some contests in an election does not show that any other contests in that election were decided correctly. The 2020 U.S. Presidential election in Georgia illustrates unrecoverable errors that can render recounts and audits 'security theater' that distract from the more serious problems rather than justifying trust.

2 The 2020 audit

Secretary of State Brad Raffensperger claimed, "Georgia's historic first statewide audit reaffirmed that the state's new secure paper ballot voting system accurately counted and reported results."[2] Moreover, "[] we did a 100 percent risk-limiting audit with a hand recount which proved the accuracy of the count and also proved that the machines were accurately counting it, and that no votes were flipped."[3] VotingWorks Executive Director Ben Adida claimed "Georgia's first statewide audit successfully confirmed the winner of the chosen contest and should give voters increased confidence in the results."[4] Per the official report of the audit, "[t]he audit confirmed the original result of the election, namely that Joe Biden won the Presidential Contest in the State of Georgia. The audit [] provides sufficient evidence that the correct winner was reported."[5]

Secretary Raffensperger has also used the recount and audit in his defense against a lawsuit that seeks to provide all Georgia voters the option to hand-mark paper ballots in person, rather than being compelled to use BMDs (Curling et al. v. Raffensperger et al., Civil Action No. 1:17-CV-2989-AT, U.S. District Court for the Northern District of Georgia, Atlanta Division). Raffensperger


411412
411413
411418
421224
421422
411221
421222
421225
421226
511221
521225
561341
511211
421426