Risk-Limiting Audits

Making Every Vote Count

A Practical Guide to Risk-Limiting Audits

Washington, DC

Philip B. Stark

31 January 2019

Department of Statistics, University of California, Berkeley

Why audit?

- *Any* way of counting votes can make mistakes
- *Every* electronic system is vulnerable to bugs, configuration errors, & hacking
- **Did error cause losing candidate(s) to appear to win?**

Did reported winner(s) really win?

- If there's a reliable, voter-verified paper trail, can check
- Usually don't need to look at many ballots
- To tell whether soup is too salty, don't have to drink the whole pot—or even 10%: stir, then taste a tablespoon
- Too much salt in soup is like too much error tabulating votes
- Stirring, then tasting a tablespoon is like checking a random sample of ballots

Checking equipment v. checking outcomes

- Sterile scalpel v. patient's condition
- Checking outcomes should be routine
  - No matter how big the margin, need *some* checking
  - May avoid contentious recounts
  - Vote once, count a bit more than once, certify once

What's special about RLAs?

An RLA has a big, known chance of correcting the reported outcome if the reported outcome is wrong.

Risk limit is the largest possible chance that, the audit *won't* correct the reported outcome if the reported outcome is wrong.

- If risk limit is 5%, then if the outcome is wrong, there's a 95% chance the RLA will correct it
- Accuracy standard: did reported winner(s) really win?
- Only checks *tabulation*: assumes that the paper trail is trustworthy

How much auditing do we need?

At least enough to be confident who won. (Maybe don't need accuracy to the last vote.)

- Starting sample size doesn't matter
- What matters is when you *stop* auditing
- **Don't stop until there's convincing evidence the reported outcome is right**

Example: Check the claim that a coin is biased in favor of heads

- Toss 5 times, get HHHHH
- Only 3% chance that would happen if the coin is fair; even lower if coin favored tails
- Therefore, strong evidence coin favors heads
- Just like checking whether reported winner really won

Requirements

- Voter-verified paper trail
  - Any jurisdiction with paper can do an RLA
  - Some voting equipment makes it *easier*, but replacing equipment isn't necessary
- "Ballot manifest": description of how ballots are stored
  - Should be routine
  - "It's the day after the election. Do you know where your ballots are?"
- Manually inspect random sample of paper ballots
  - individual ballots, batches, unstratified, stratified, w/ or w/o replacement
  - polling audits: just need ballots
  - comparison audits: also need to export data from voting system & check totals

Pilots (since 2008)

California: Alameda, Humboldt, Inyo, Madera, Marin (2), Merced, Monterey, Napa, Orange (2), San Luis Obispo, Santa Cruz, Stanislaus, Ventura, Yolo (2)

Colorado: Arapahoe, Boulder, others; now routine statewide

Indiana: Marion

Michigan: Kalamazoo, Lansing, Rochester Hills

New Jersey: Essex, Gloucester, Union

Ohio: Cuyahoga

Rhode Island: Bristol, Cranston, Portsmouth

Virginia: Fairfax

Denmark

Evidence and Trustworthiness

An audit can't be better than the paper trail it uses.

- No paper trail, no audit
- If paper trail not voter-verifiable (e.g., some BMDs), can't verify winner
- If paper trail untrustworthy, audited outcome untrustworthy

5 Cs

- *Create* durable, trustworthy record of voter intent
  - ideally, hand-marked paper ballots with BMDs for voters who benefit from them
- *Care* for the paper record
  - verifiable chain of custody, 2-person custody rules, ballot accounting, good seal protocols, etc.
- *Compliance* audit: establish whether paper trail is trustworthy
  - ballot accounting, including VRDB, pollbooks, etc.; check chain of custody logs, video, etc.; eligibility
- *Check* reported outcome against the paper by auditing
- *Correct* the reported outcome if it is wrong