Risk-Limiting Audits

Ronald L. Rivest
MIT

NASEM Future of Voting
December 7, 2017

Risk-Limiting Audits (RLAs)

- Assumptions
- What do they do? What do they not do?
- How do RLAs work?
- Extensions
- References

(Assumption) Voter-Verified Paper Ballots

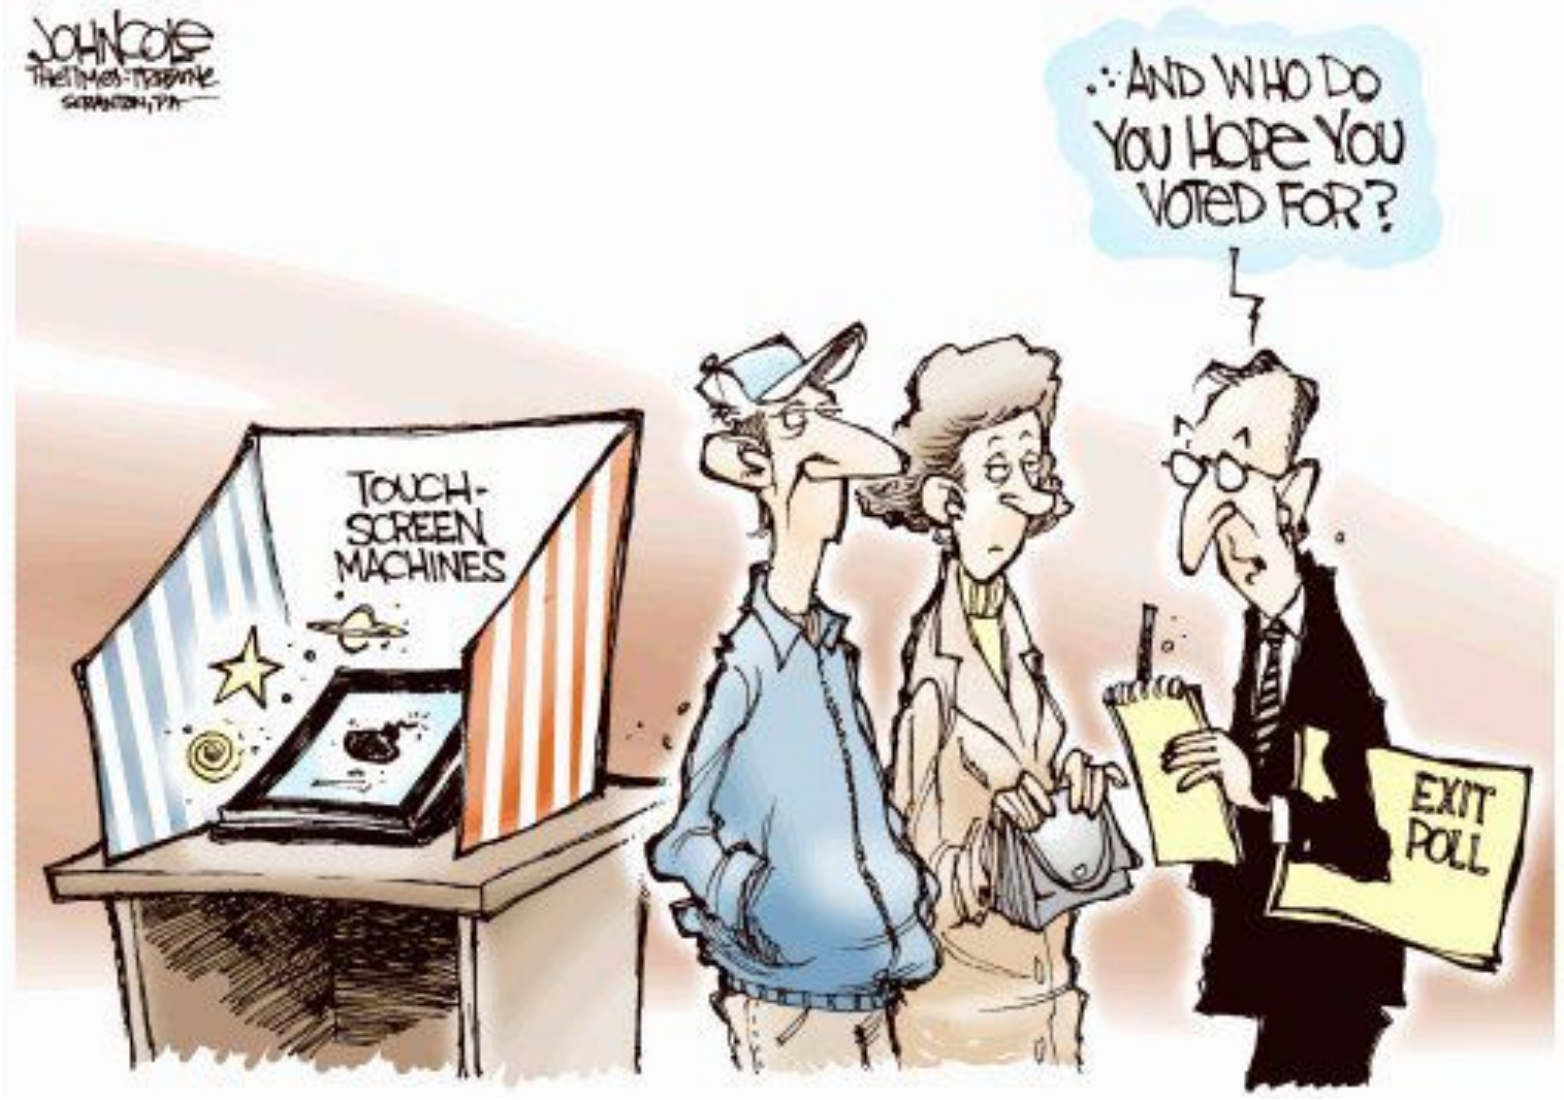

And Who Do You **Hope** You Voted For?

(Assumption) Optical scanners used, for efficient tabulation


DOMINION
VOTING
Security Key
OPEN
System Ready
BALLOTS CAST: 72
Cast
Return

Scanners

- (Assumption) Scanner produces one electronic **CVR** (cast vote record) for each ballot scanned.
- CVRs are tabulated to produce **reported contest outcome** (aka contest winner).

(Concern) Scanners may introduce systematic errors

Causes of scanner errors

- Differences in **interpretation** between machine interpretation and hand interpretation. Voter intent rules.
- **Stray marks** (e.g. caused by folds)
- **Configuration errors**
- **Programming errors**
- Hacking (**adversarial attack**)

(Response to concern)
Statistical “Risk-Limiting” Audit

What does a RLA do?

- A **risk-limiting audit** provides **statistical assurance** that a reported outcome is "correct."
- Here "correct" means "the result that would be obtained by examining all ballots by hand."
- A RLA does so **efficiently** using a **hand examination** of a **random sample** of the cast paper ballots.
- Good for in-person voting and vote-by mail.
- Really, a "risk-limiting tabulation audit" (RLTA).

What a RLA does **not** do

- A RLA does not address:
  - **correctness of the *tally* (as opposed to the outcome)**
  - **voter eligibility**
  - **voter authentication**
  - **usability**
  - **privacy**
  - **chain of custody**

Who is a RLA for?

- **Losing candidates** – to convince them that "they lost fair and square"
- **The winner –** to provide a mandate
- **The public** – to assuage doubts about "rigged elections"

- **All evidence produced by the election and the audit should be published (as much as can be published without violating voter privacy).**

(Ballot-level) Sampling

Ballot manifest

- A **ballot manifest** describes the set of cast paper ballots, and how they are organized in storage (e.g. in 100-count batches, one envelope per batch, 15 batches/container).
- The ballot manifest defines the set of ballots to be sampled from for an audit.

Random Seed Generation

- Because of adversaries, random sample should be determined **after** ballot manifest and CVRs are committed to.
- A good process may start by rolling 20 decimal dice.

- Then seeding a PRNG (pseudo-random number generator) to pick ballots at random.

Overall RLA structure

RLA structure

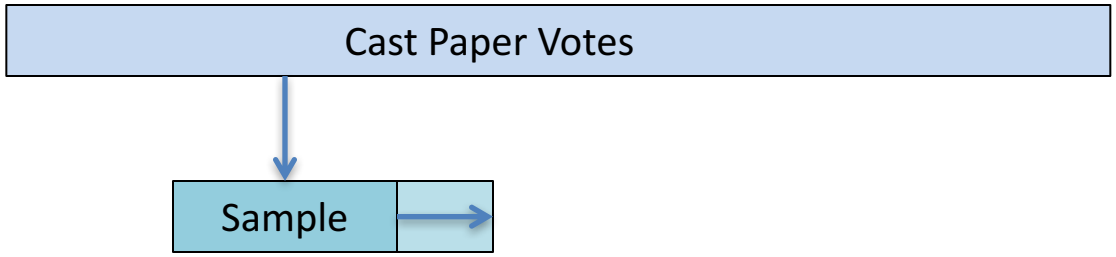

1. Draw initial random sample of paper votes.
2. Interpret them by hand.
3. Stop if reported outcome is now confirmed to desired confidence level.
4. If all ballots now examined, you are done.
5. Otherwise increase sample size **(escalate);** return to 2.

Two (ballot-level) auditing paradigms

- **Ballot-polling audits**:
  Uses the randomly selected cast paper ballots only.
  Like ``exit poll'' of ballots…
- **Comparison audits**:
  Compares randomly selected paper ballots with corresponding electronic records (CVRs) for all contests under audit.

- Comparison audit more efficient by a factor of roughly (1 / margin-of-victory).

What is ``Risk”?

What is ``Risk''??

- Risk is defined as:
  - **the probability that an incorrect reported outcome will be accepted by audit as correct**

What is “Risk Limit”?

What is ``Risk Limit''?

- **Risk Limit** is upper bound on acceptable risk.
- With 5% risk limit, there is at least a 95% chance that an incorrect reported outcome will be detected and fixed (by escalation to hand interpretation of all cast paper ballots), and at most a 5% chance that an incorrect reported outcome will be accepted as correct.

One RLA stopping rule

- For comparison audit
- n = sample size
- m = margin (difference between winner and loser vote-count, divided by number of ballots in population being sampled)
- O_1, O_2 = number of sampled ballots revealing overstatement of margin by one or two
- U_1, U_2 = number of sampled ballots revealing understatement of margin by one or two

One RLA stopping rule (cont.)

- Stop audit when:

  $n > (4.8 + 1.4(O_1 + 5O_2 - 0.6U_1 - 4.4U_2)) / m$

  This is for a risk limit of 0.10.
- For example, with no discrepancies:

  $n > 4.8 / m$

  (This formula used for CO initial sample sizes.)
- Example: if $m = 0.05$

  $4.8 / m = 96$

Extensions

- Many contests, not just one. Audits interact.
- Many jurisdictions, not just one. Sampling is for a distributed contest is distributed.
- Contests have overlapping domains; may be arbitrary relationship (not necessarily nested).
- Some jurisdictions have equipment supporting comparison audits; some don’t. Need blended method if contest spans both.
- Vote-by-mail mixes ballots of different styles.

Conclusions

- Risk-limiting audits can detect and correct tabulation errors.
- Mathematical foundations exist; extensions to handle all real-world scenarios in best way being worked on.
- RLAs work in practice (Colorado!)

References

- "A Gentle Introduction to Risk-Limiting Audits" by Mark Lindeman and Philip B. Stark
  IEEE Security & Privacy **10,** 5 (Sep-Oct. 2012), 42—49.

The End

Thanks for your attention!