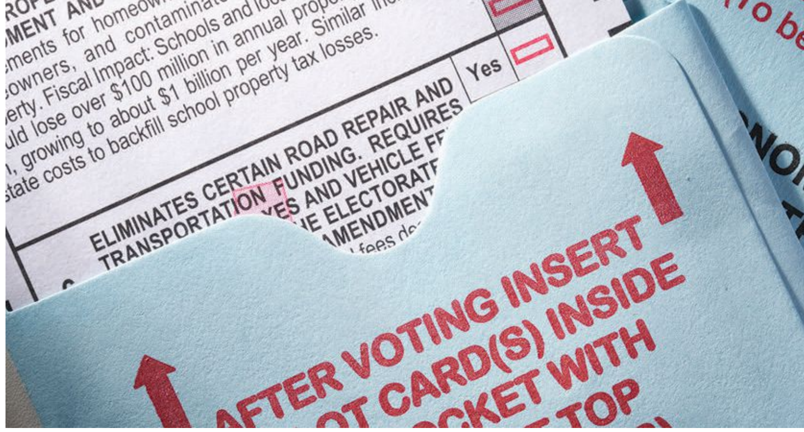

A Bayesian Approach to Election Audits

—AUDREY MALAGON AND RONALD L. RIVEST

Election security is a growing concern as we approach a major election year. Across the country, jurisdictions are turning to voter-marked paper ballots and risk-limiting audits (RLAs) to ensure reported outcomes align with voter intent. During a risk-limiting audit, voter-verified and cast paper ballots are manually sampled and inspected by auditors. The procedure ensures that, if the reported outcome differs from what a full hand tally of the ballots would reveal, then the maximum probability that the RLA accepts the reported outcome is less than a predetermined risk-limit. Philip Stark, Professor of Statistics at the University of California, Berkeley, invented risk-limiting audits in 2008 and continues to promote them as an election security measure. An introduction to his work can be found in the 2012 "A Gentle Introduction to Risk-Limiting Audits" (bit.ly/2MGoSfD).

Rivest and Shen (bit.ly/2N5ufEs) explore a Bayesian approach to post-election auditing with similar goals. A Bayesian audit upper bounds instead the upset probability—the probability that similar collections of paper ballots, taken from a suitably defined Bayesian posterior distribution, give outcomes different than the reported one.

Before any audit, election officials tabulate the collection C of n cast paper ballots to obtain a reported outcome R, which may be wrong (due to error or fraud), but should be correct and thus satisfy

$$R = \text{Outcome}(C).$$

A statistical audit (RLA or Bayesian) aims to give confidence in R's correctness starting with a hand examination and interpretation of a random sample S, of size s, from C:

$$S \leftarrow \text{Sample}(C, s).$$

A Bayesian statistical election audit uses S to estimate the probability p that collections C' similar to C have outcomes different than R. We call p the expected loss of the audit. The audit stops and accepts R as correct if p is less than a given loss limit (e.g., 1%). Otherwise the audit repeats using a larger sample, likely escalating to examine all ballots if R is incorrect. A smaller loss limit may require a larger sample but gives more confidence in the correctness of the election outcome.

A Bayesian audit estimates p by probabilistically "reversing the sampling" ("restoring") to obtain hundreds of ballot collections C' similar to C, and estimating p as the fraction for which

$$R \neq \text{Outcome}(C').$$

Restoring starts with S, then successively adds $n - s$ votes back, in a random manner, to obtain C':

$$C' \leftarrow \text{Restore}(S, n)$$

Which votes does Restore add back? To ensure that C' is similar to C, Restore adds copies of votes randomly selected from the growing sample. To enable restoration of votes even for candidates not in the sample, Restore adds to S one vote for every candidate (even those with votes) when it starts, and removes one vote for every candidate when it ends. These extra votes determine the Bayesian prior.

Since Restore picks votes to copy at random, it may return a somewhat different result C' each time. But each such C' should be similar to the original C, and this similarity improves with the size s of the initial sample S. Variations in C' reflect the uncertainty the auditor has about C, and thus about Outcome(C), knowing only S.

A laptop can restore S to C' quickly, as restoration does not sample or examine by hand any paper ballots. We call constructing C' in this way drawing from the posterior distribution defined by S and the prior. Methods based on Dirichlet-multinomial distributions give even greater efficiency.

Bayesian audits are not necessarily risk-limiting, but nonetheless may serve as a useful alternative when an RLA is not possible. The ability of a Bayesian audit to probabilistically reverse the sampling process is a powerful tool for auditing. Since a Bayesian audit uses only vote copying, it is independent of the tabulation method and works for complex voting methods like IRV. It extends to handle stratified audits, ballot-level comparison audits (stratifying by reported vote), multi-jurisdiction audits, and audits where jurisdictions have different types of equipment. More information about Bayesian audits can be found in (arxiv.org/abs/1801.00528). ■

Audrey Malagon is Batten Associate Professor of Mathematics at Virginia Wesleyan University. She is interested in the mathematics of election security and writing about mathematics for the public. Ronald L. Rivest is an Institute Professor at the Massachusetts Institute of Technology, with interests in algorithms, cryptography, election integrity, and climate change.

GET MORE

Want to learn more about statistical audits of elections and how mathematicians are getting involved? Join us at the Mathematical Association of America Invited Paper Session "Can Mathematics Help Us Trust Our Elections Again?," Friday, January 17, 2020 from 8–10:30 AM at the Joint Mathematics Meetings.