The Risk Limit of Bayesian Audits

Kellie Ottoboni

E-VOTE-ID 2018

October 2, 2018

University of California, Berkeley
DEPARTMENT OF STATISTICS

Why audit?

Provide assurance that reported outcomes are correct by examining some or all of a voter-verifiable paper record.

Why audit?

Provide assurance that reported outcomes are correct by examining some or all of a voter-verifiable record.

32 states have some sort of law mandating post-election audits.

- Which contests?
- How are samples of ballots drawn?
- Is vote-by-mail included?
- Are audit results binding?

Risk-limiting audits

Statistical check that tabulation errors would not change the electoral outcome.

Risk limit: chance of failing to correct a wrong outcome

https://www.verifiedvoting.org/resources/voting-equipment/

RLAs are hypothesis tests.

H_0: The reported winner is **wrong**.

Risk = P-value

RLAs are hard.

- Multiple pools of ballots

RLAs are hard.

- Multiple pools of ballots
- Heterogeneous voting systems

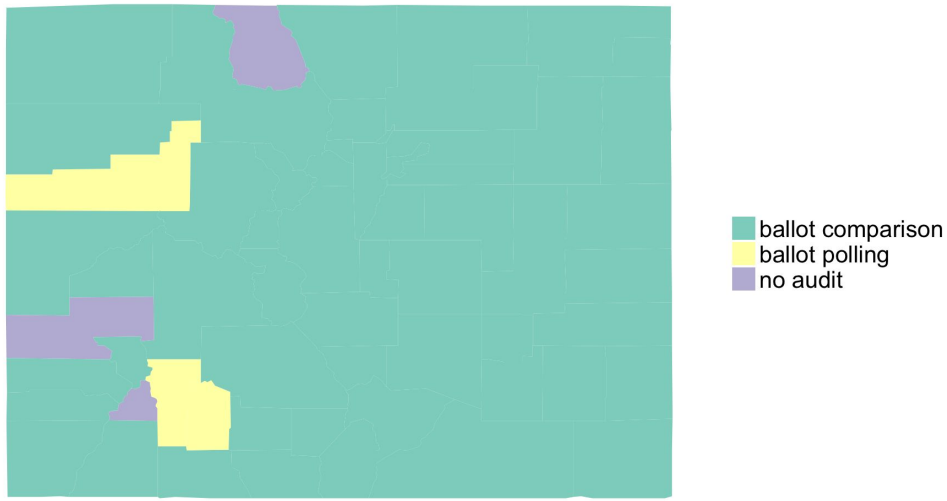

RLAs are hard.

- Multiple pools of ballots
- Heterogeneous voting systems
- Complex social choice functions

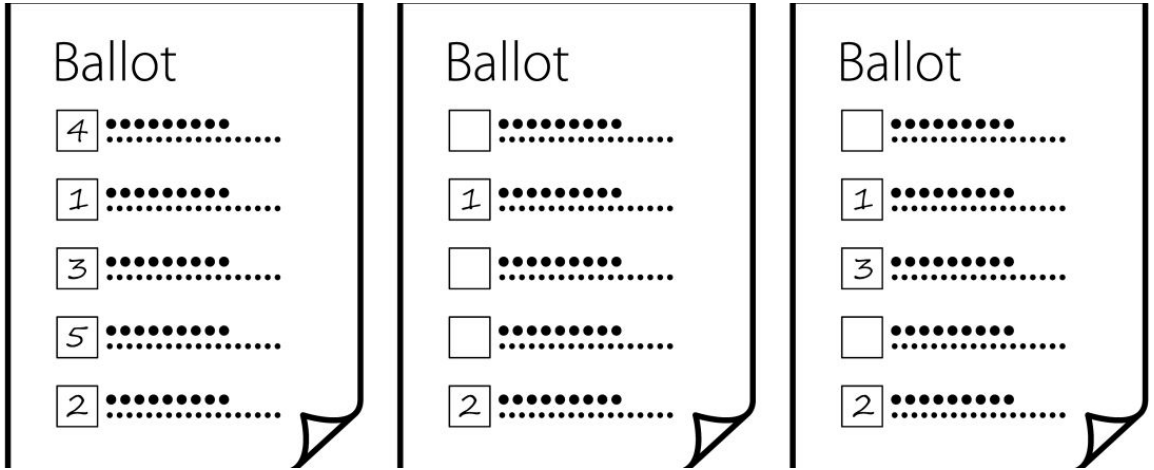

Bayesian Audits

(Rivest and Shen, 2012)

Bayesian audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

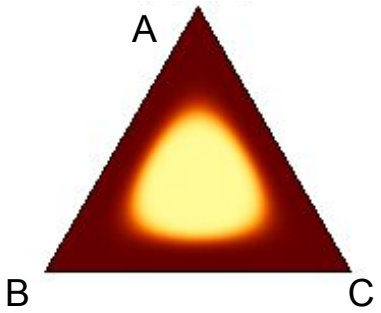

1. Model voter preferences as random with a prior distribution.

Bayesian audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

1. Model voter preferences as random with a prior distribution.
2. Sample infinitely many elections (collections of voter preferences) from the prior.

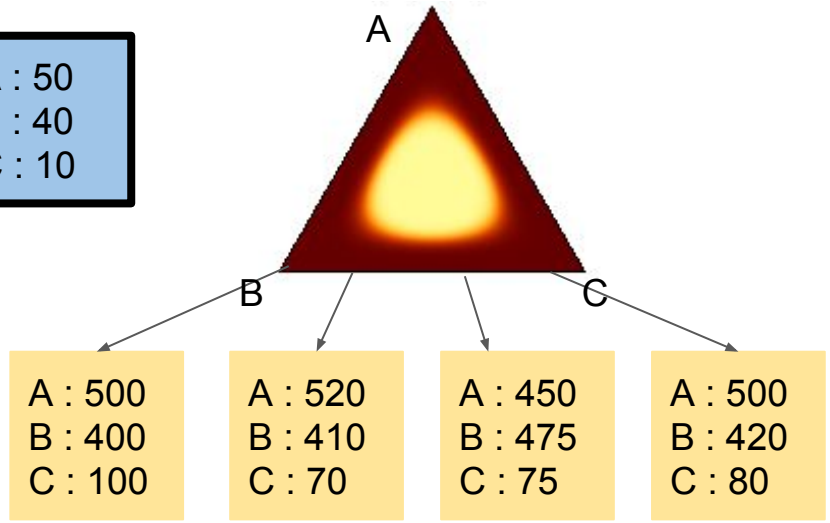

Bayesian audits

A : 500 B : 400 C : 100	A : 50 B : 40 C : 10

1. Model voter preferences as random with a prior distribution.
2. Sample infinitely many elections (collections of voter preferences) from the prior.
3. Audit each hypothetical election and discard those for which the audit sample differs from observed.

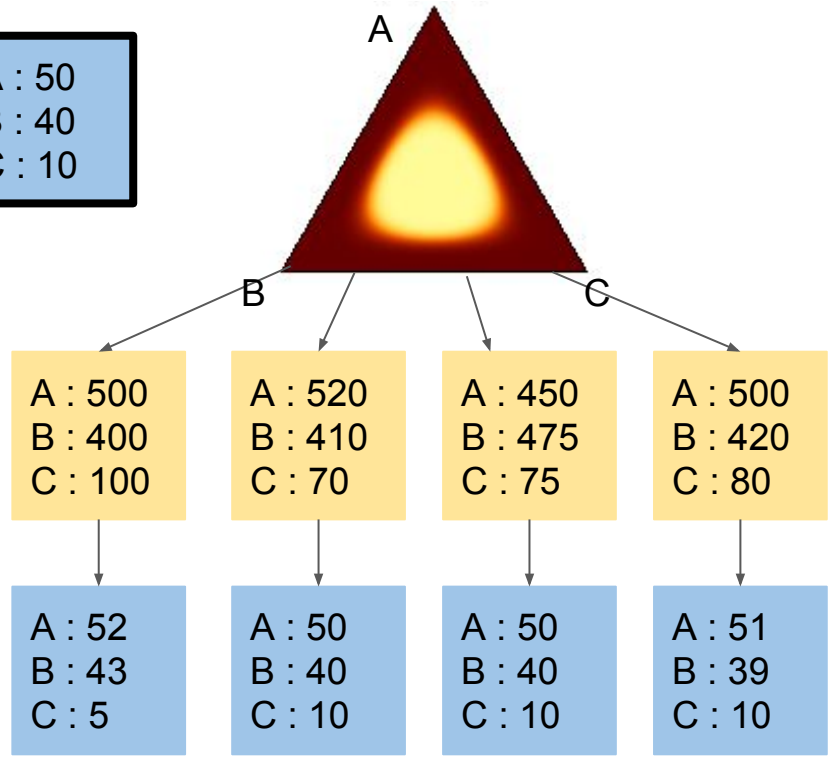

Bayesian audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

1. Model voter preferences as random with a prior distribution.
2. Sample infinitely many elections (collections of voter preferences) from the prior.
3. Audit each hypothetical election and discard those for which the audit sample differs from observed.

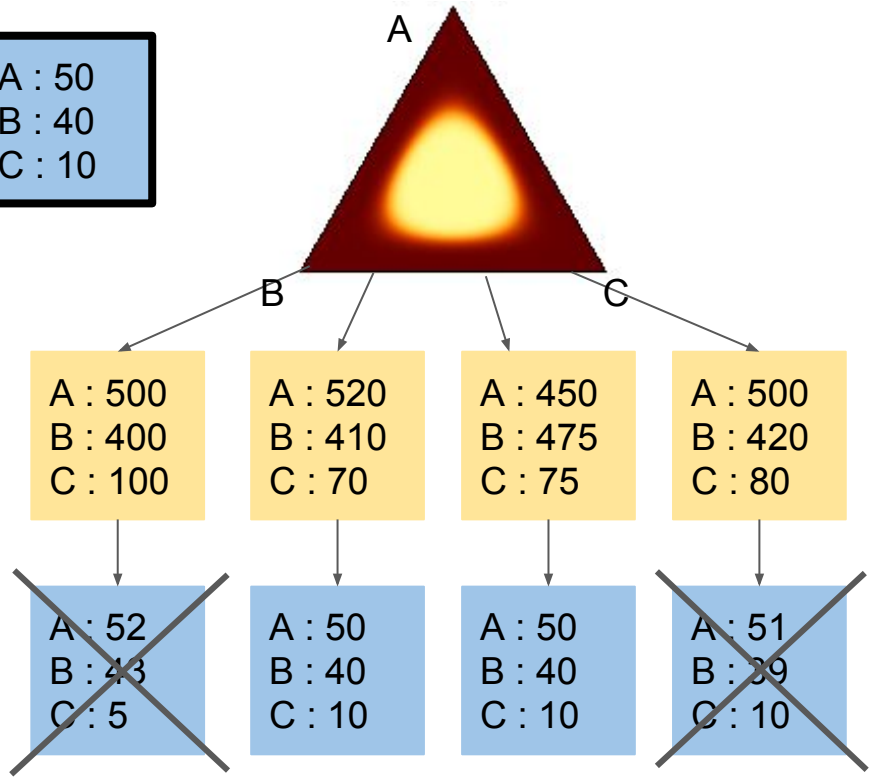

Bayesian audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

1. Model voter preferences as random with a prior distribution.
2. Sample infinitely many elections (collections of voter preferences) from the prior.
3. Audit each hypothetical election and discard those for which the audit sample differs from observed.
4. Calculate the fraction of elections whose winner differs from the reported winner in the actual election.

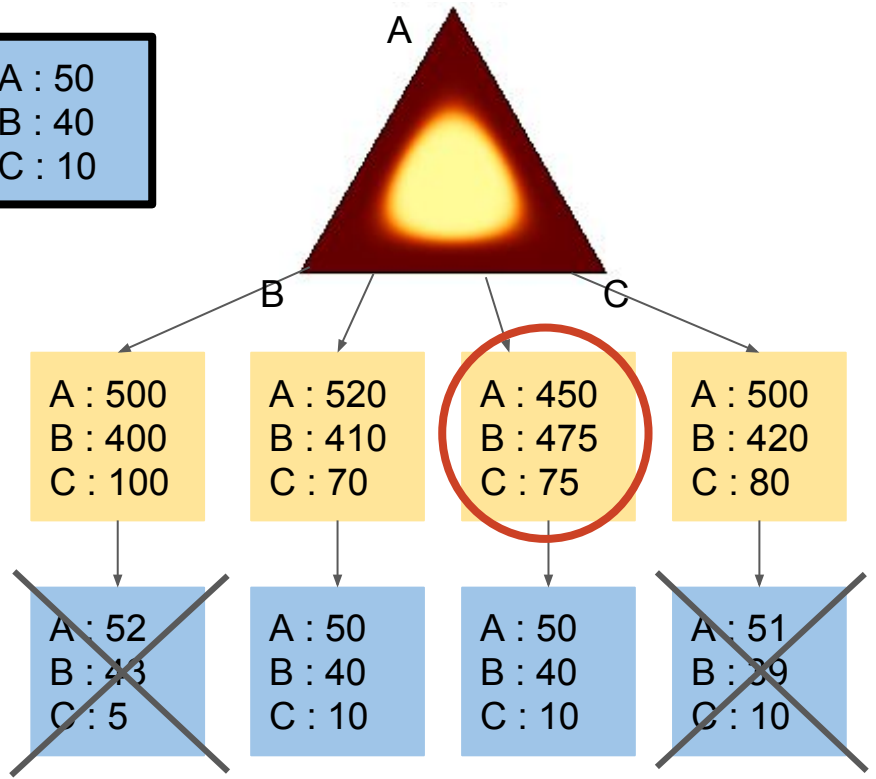

Posterior probability = 1/2

Risk-limiting audits

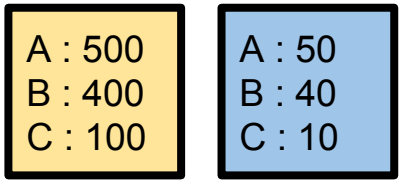

1. Assume the reported outcome is wrong.

Risk-limiting audits

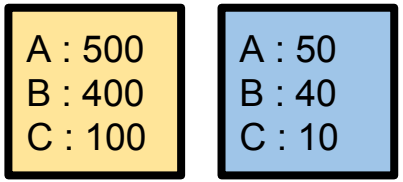

1. Assume the reported outcome is wrong. Consider all possible ways.

A : 499
B : 501
C : 0

A : 400
B : 500
C : 100

A : 450
B : 475
C : 75

A : 475
B : 500
C : 25

Risk-limiting audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

1. Assume the reported outcome is wrong. Consider all possible ways.
2. For each collection of voter preferences, find the probability that the audit will stop without a full hand count.

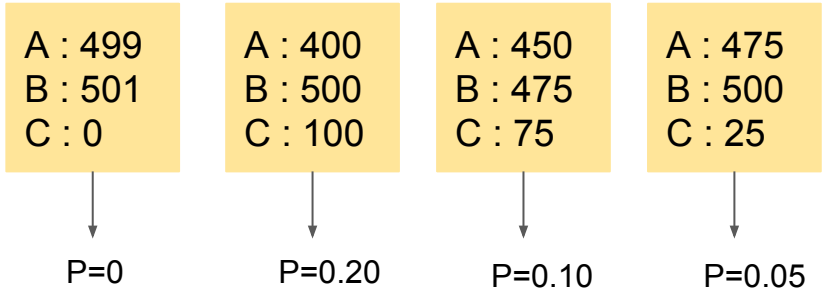

Risk-limiting audits

A : 500
B : 400
C : 100

A : 50
B : 40
C : 10

1. Assume the reported outcome is wrong. Consider all possible ways.
2. For each collection of voter preferences, find the probability that the audit will stop without a full hand count.
3. Risk is the maximum.

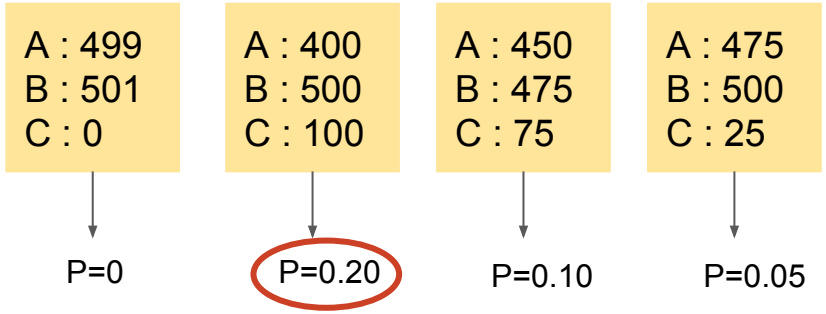

Risk = 0.20

RLAs

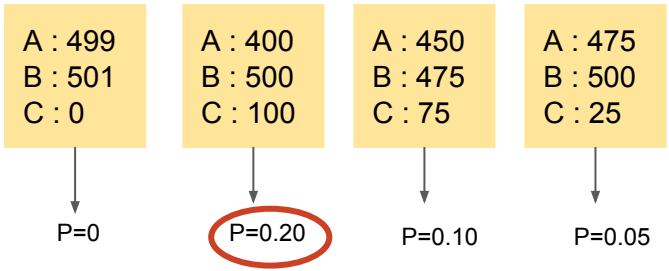

Bayesian audits

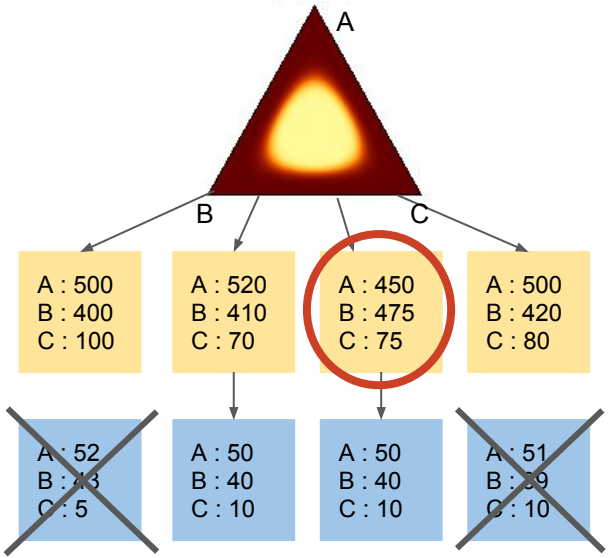

- Fixed, unobserved voter preferences vs a prior distribution
- Worst case chance vs average over hypothetical elections
- Bayesian audits don’t require computing P-values

When is a Bayesian audit risk-limiting?

Not always.

Risk <= Probability that the audit stops, assuming a tie.

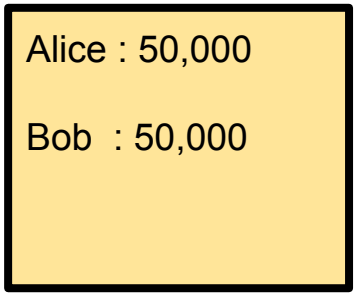

Run 10,000 Bayesian ballot-polling audits and record the posterior probability when the audit stops.

Not always.

Risk <= Probability that the audit stops, assuming a tie.

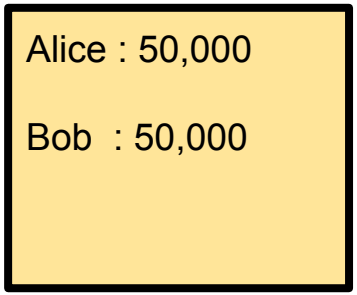

Run 1000 Bayesian ballot-polling audits and record the posterior probability when the audit stops.

Desired risk limit	Empirical risk	Factor
0.001	0.0065	6.5
0.002	0.0141	7
0.005	0.0303	6.1
0.01	0.0575	5.8
0.02	0.1052	5.3
0.05	0.2255	4.5
0.1	0.3837	3.8

Decision theory

Θ : Set of all possible voter preferences

$\mathcal{D}$: Set of all possible audit rules

$R(\theta, \delta)$: Risk of using the audit rule δ when voter preferences are θ

Decision theory

Θ : Set of all possible voter preferences

$\mathcal{D}$: Set of all possible audit rules

$R(\theta, \delta)$: Risk of using the audit rule δ when voter preferences are θ

$$\underbrace{\inf_{\delta \in \mathcal{D}} \underbrace{\sup_{\theta \in \Theta} R(\theta, \delta)}_{\text{Risk from an RLA}}}_{} = \underbrace{\inf_{\delta \in \mathcal{D}} \underbrace{\int_{\theta \in \Theta} R(\theta, \delta) d\pi(\theta)}_{\text{Average risk over least favorable prior}}}_{}$$

A way forward?

- Identify least favorable priors -- usually not an “uninformative” prior

- Restrict to audit rules that give the desired risk limit

- Special case: the BRAVO RLA (Lindeman et al, 2012) is equivalent to a Bayesian audit with a certain prior (Vora, unpublished)

Thanks!


kellieotto@berkeley.edu


www.kellieottoboni.com


@kellieotto