DiffSum – A Simple Post-Election Risk-Limiting Audit

Ronald L. Rivest
MIT CSAIL, rivest@mit.edu

September 1, 2015

We present DiffSum, a simple risk-limiting post-election ballot-polling audit. See [3, 2, 1] for background.

You wish to check that candidate A really won a plurality election against candidate B. You may sample the n cast paper ballots without replacement.

Procedure DiffSum:

1. [**Choose** c] Let d be the number of decimal digits in n, and choose $c = d + \delta$ where δ controls the error rate (the chance of the audit accepting an incorrect outcome):

δ	0	1	2	3	4
max error rate	22%	15%	10%	6%	4%

2. [**Begin**] Draw an initial sample of 24 ballots.
3. [**Tally**] Determine the number a of votes for A in your sample, and the number b of votes for B.
4. [**Stop?**] Stop the audit (accept A as winner) if $a > b$ and
$$(a-b)^2 > c \cdot (a+b) \ . \tag{1}$$
5. [**Continue?**] If $a + b = n$, stop (you have just completed a full recount). Otherwise, enlarge your random sample and return to step 3.

Remarks: The initial size 24 of the sample in step 2 is arbitrary. In step 5 the increase in sample size is also arbitrary; it could be by a single ballot.

The name "DiffSum" was chosen because (1) says
$$(\text{difference})^2 > c \cdot (\text{sum}) \ . \tag{2}$$

Efficiency: Let m be the true margin (the fraction of votes cast for A minus the fraction cast for B). In a sample of size $s = a+b$, the expected value of $a-b$ is sm. Thus, DiffSum is expected to stop when $(sm)^2 > cs$ or
$$s > c/m^2 \tag{3}$$

DiffSum is approximately as efficient as Bravo—compare (3) with the estimate $2\ln(1/\alpha)/m^2$ for Bravo [2] (here α is the risk limit). Moreover, DiffSum does not need an initial estimate of the vote shares, and Bravo is inefficient when this estimate is inaccurate.

Error rate: The error rate bounds given in Step 1 are based on extensive simulations for $\delta = 0$ to 4, $d = 3$ to 7, $n = 10^d$, and $c = d+\delta$. We measured the error rate over 10,000 simulated elections in each case. Each simulation estimated the error rate when the election was a tie, a worst-case scenario; with more realistic margins the error rate drops dramatically, so that in practice even $c = d$ should give very reliable audits.

Example: An election with $n = 50,000$ votes can be audited using $c = 7$ for a risk limit of $\alpha = 10\%$. For $m = 0.20$, DiffSum examines about 175 ballots (estimated), Bravo (with $\alpha = 0.10$) examines about 115 (estimated). In simulations for this election, DiffSum with $c = 7$ examines about 157 ballots on average, and has an error rate of less than 0.04%; DiffSum with $c = 5$ examines about 112 ballots on average, and has an error rate of less than 0.2%. Bravo examines about 119 ballots on average, and has an error rate of approximately 2.5%.

Extension: In practice, one should cease random sampling once a significant number (say 4%) of the ballots have been sampled, when switching over to a full hand recount becomes more economical.

With more candidates, let DiffSum check that the sample winner beats the sample's strongest loser.

Conclusion: DiffSum is exceptionally simple, and appears quite comparable to Bravo in terms of efficiency and error rate. Further simulations and analysis would be helpful.

Acknowledgment: I thank Philip Stark for helpful comments.

References

[1] M. Lindeman and P. B. Stark. A gentle introduction to risk-limiting audits. *IEEE Security and Privacy*, 10:42–49, 2012.

[2] M. Lindeman, P. B. Stark, and V. S. Yates. BRAVO: Ballot-polling risk-limiting audits to verify outcomes. In *Proc. 2012 EVT/WOTE*, 2012.

[3] L. Norden, A. Burstein, J. L. Hall, and M. Chen. Post-election audits: Restoring trust in elections. Technical report, Brennan Center for Justice and Samuelson Law, Technology & Public Policy Clinic, 2007.