# Risk-limiting audits for IRV elections

Michelle Blom, Peter J. Stuckey and Vanessa Teague

The University of Melbourne

October 5, 2018

# Instant Runoff Voting (IRV)

- Preferential voting scheme
- A set of candidates $\mathcal{C}$, one winner
- Each vote is a ranking over $\mathcal{C}$
- Each vote can be a *partial* ranking
  e.g., [Mary Hill, Joe Smith, John Citizen]

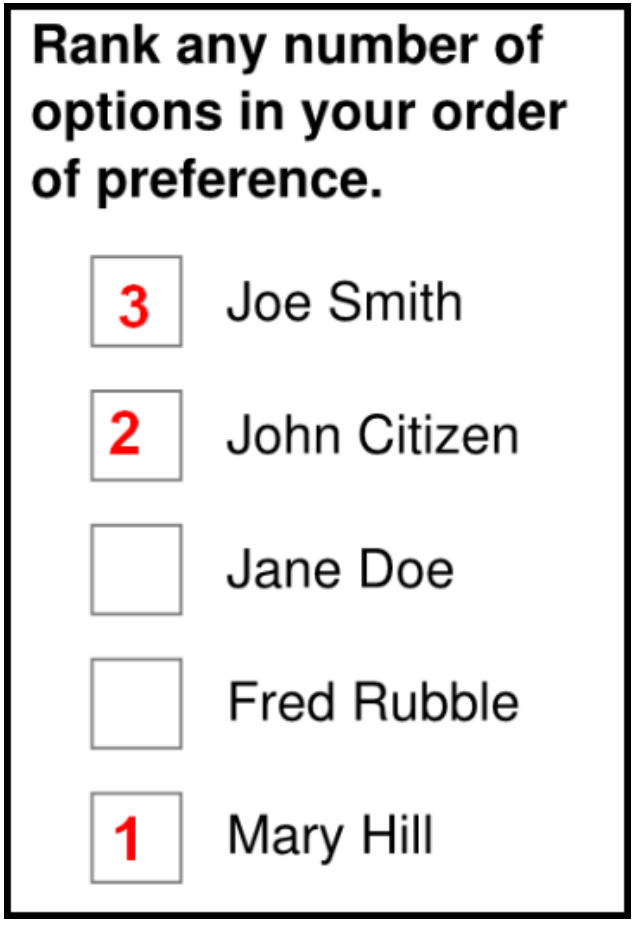

# Instant Runoff Voting (IRV) – An Example

4 candidates, 60 votes

| Ranking | Count |
|---|---|
| $[c_2, c_3]$ | 4 |
| $[c_1]$ | 20 |
| $[c_3, c_4]$ | 9 |
| $[c_2, c_3, c_4]$ | 6 |
| $[c_4, c_1, c_2]$ | 15 |
| $[c_1, c_3]$ | 6 |

(a) Initial tallies

| Candidate | Rnd1 | Rnd2 | Rnd3 |
|---|---|---|---|
| $c_1$ | 26 | 26 | 26 |
| $c_2$ | 10 | 10 | — |
| $c_3$ | 9 | — | — |
| $c_4$ | 15 | 24 | 30 |

(b) Tallies after each round of counting

# BRAVO Ballot-polling Risk Limiting Audits [LSY12], for first-past-the-post

- Given an announced election outcome (*i.e.* a set of winners with tallies), conduct a random audit of the paper ballots until either we're confident the outcome is right, or we revert to a full manual recount.
- Choose a *Risk Limit* $\alpha$.
- Guarantee: If the outcome is wrong, we detect it with probability at least $1 - \alpha$.
- Works great for first-past-the-post, even with multiple winners, but what about IRV?

# BRAVO rough overview

- Maintain a running statistic $T_{wl}$ for each pair of apparent winner $w$ and loser $l$.
- A ballot that shows a valid vote for winner $w$ increases $T_{wl}$ (by an amount dependent on the reported votes for the two candidates).
- A ballot showing a valid vote for the loser $l$ decreases it.
- When each statistic exceeds a threshold, dependent on the risk limit, we know that we have seen enough evidence to reject the hypothesis that $l$ beat $w$.

# Can we apply BRAVO to IRV?

- When you pick out a random IRV ballot, it's not obvious whether it supports or undermines an announced election outcome.

# First idea: audit the whole elimination sequence

- For every IRV elimination, run a BRAVO audit.
- The $k$-th elimination has $n - k - 1$ winners and 1 loser.
- Valid, but hideously inefficient.
- We don't really need to know that the *elimination order* was right, only the final winner.

# Optimisation 1: Batch eliminations

- We can simultaneously eliminate candidates $E$ if the sum of tallies of these candidates is less than the tally of the next lowest candidate.
- Audit using BRAVO by joining all the candidates in $E$ into one losing candidate.
- Sometimes this is a huge improvement, *e.g.* if the differences between candidates in $E$ are small.
- Sometimes it's not.

# Winner-Only audits: a simple idea that actually works

We wish to eliminate the hypothesis that winner $w$ is eliminated before loser $l$.

1. Winner $w$ always has *at least* their first-preference votes.
2. Loser $l$ has *at most* the votes that mention $l$ somewhere, and don't prefer $w$.
3. So run BRAVO with $T_{wl}$ counting case 1 as a vote for $w$ and case 2 as a vote for $l$.

If we can reject the hypothesis that $l$ beat $w$, that proves (except with probability at most $\alpha$) that $l$ was eliminated before $w$.

# In the paper: combining Batch Elimination, Winner-Only, and Full Elimination order auditing

- Combining a series of facts to check that the announced winner is correct.
- Confidence follows directly from BRAVO: if the announced outcome is wrong, one of the BRAVO audits will detect it with probability at least $1 - \alpha$.
- Efficiency is heuristic. Often good, but no guarantees.
  *e.g. (good):* Berkeley Mayor 2012 (57,492 votes):
  Expect 115 samples for $\alpha = 0.05$.
  *e.g. (bad):* Gosford NSW 2015, Pierce County assessor 2008:
  no simulated audit smaller than recounting.

**Questions?**


M. Lindeman, P.B. Stark, and V. Yates.
BRAVO: Ballot-polling risk-limiting audits to verify outcomes.
In *Proceedings of the 2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (EVT/WOTE '11)*. USENIX, 2012.