RAIRE:

A Branch and Bound Approach to Auditing Instant-Runoff Voting Elections

Peter J. Stuckey 20/8/2025 joint work with Michelle Blom and Vanessa Teague

Contents

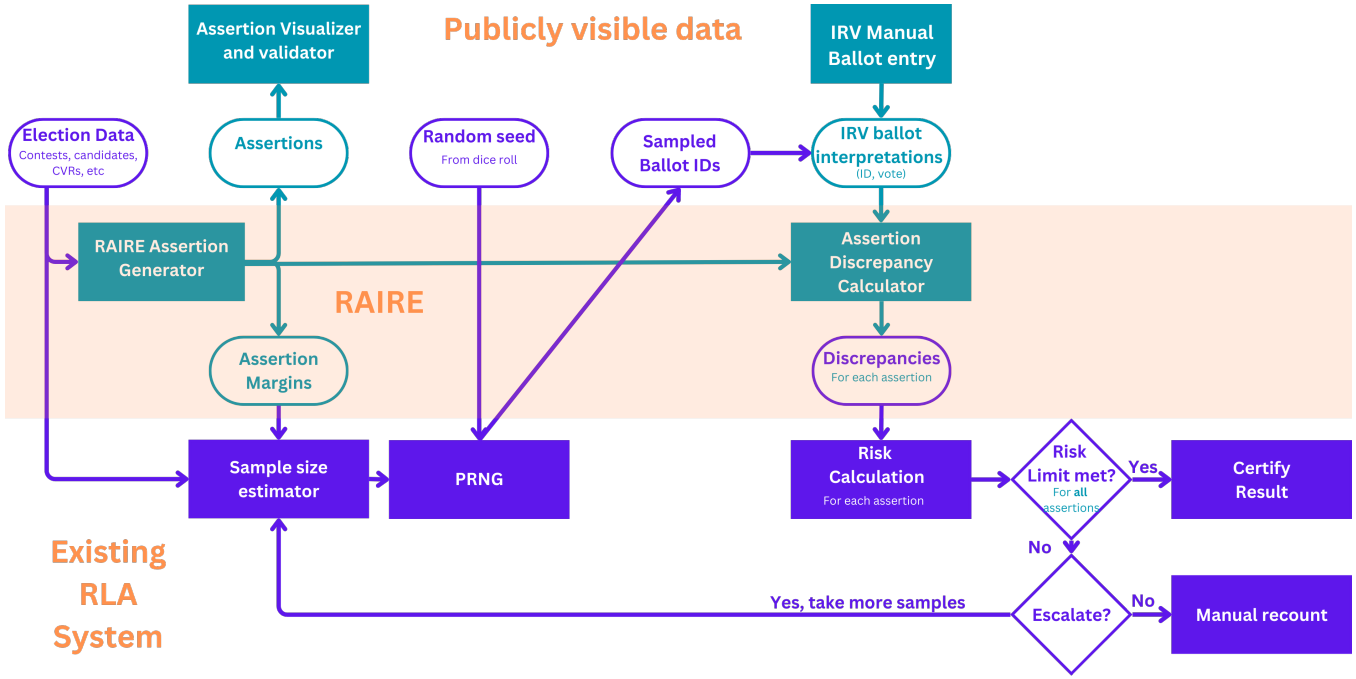

Figure 1.1: A high-level overview of an IRV RLA using RAIRE. Purple elements are existing RLA components. Green indicates new components and data structures for IRV.

Instant Runoff Voting

(Alternatively Ranked Choice Voting)

- Each voter ranks (a subset of) the candidates: most preferred to least preferred
- Every ballot is placed in a pile for its first (most preferred) candidate
- The candidate with the least votes is eliminated
- All the ballots in the pile of the eliminated candidate are
  - Redistributed to the pile of the next (non-eliminated) candidate, or
  - Exhausted if no such candidate exists
- Eventually only one candidate remains: the winner

Instant Runoff Voting

Example

- Consider 4 candidates: Alice, Bob, Chuan, Diego, and ballots
- (A,B,C,D):500 copies, (B,A,C):100 copies, (B,D,A):100, (C):200, (C,D):200, (D,C):400
- Tallies: Alice:500, (A,B,C,D):500 Bob:200, (B,A,C):100 (B,D,A):100 Chuan:400, (C):200 (C,D):200 Diego:400 (D,C):400 Bob eliminated
- Tallies: Alice:600, Chuan: 400, Diego:500 Chuan eliminated
- Tallies: Alice:600, Diego:700. Alice eliminated
- Diego is the winner and elected
- Elimination order: [Bob, Chuan, Alice, Diego]

Why Instant Runoff Voting

- The winner ends up with >50% support among voters
- More efficient (no runoff elections)
- Encourages positive campaigning
- Encourages minor/diverse parties (they can receive votes without changing result)

- Advocate to change Sweden to Single Transferable Vote (multi-seat version)
- Also advocate for mandatory voting!

Contents

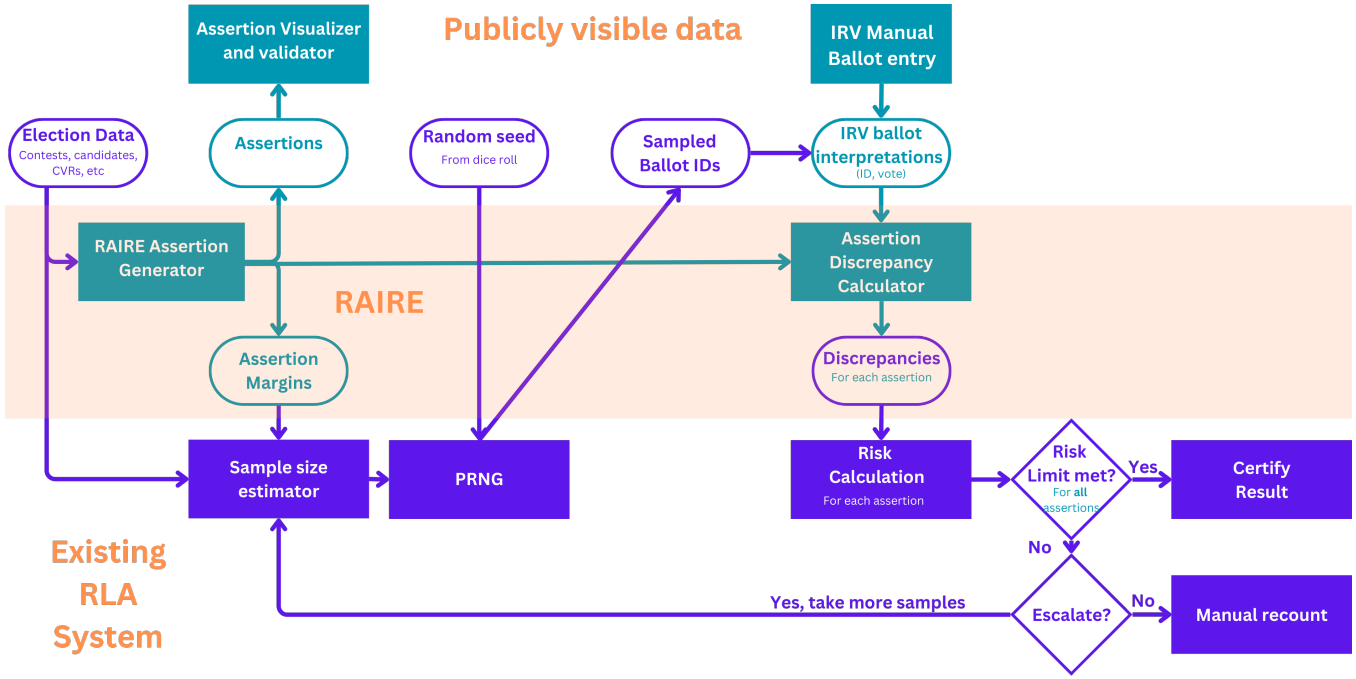

Figure 1.1: A high-level overview of an IRV RLA using RAIRE. Purple elements are existing RLA components. Green indicates new components and data structures for IRV.

Risk Limiting Audits

(Ballot Polling — there are other types)

- Assuming a paper record of each ballot of the election stamped with an ID
- AND an electronic record of each ballot with ID used to calculate the winner W
- A **risk limiting** (post-election) **audit** (**RLA**) with risk limit A is
  - A statistical sampling method for physical ballots that checks them against the electronic records, until there is no more than A chance that W was not the correct winner
- The RLA either stops accepting the result, or demands a full hand recount of the election.

Risk Limiting Audit

Example

- Consider a simple first past the post election with recorded ballots and risk limit 5%
  - Alice: 500, Bob: 200, Chuan: 300, Diego: 400
- We randomly sample IDs and check the physical ballot versus the record until we have enough evidence to be
  - 95% certain Alice had more votes than Bob
  - 95% certain Alice had more votes than Chuan
  - 95% certain Alice had more votes than Diego
- Note that the statistics are not linear. It takes the same no of samples for
  - Alice: 5000, Bob: 2000, Chuan: 3000, Diego: 4000

Risk Limiting Audits

- For most elections are efficiently audited
- RLAs are increasingly in use
  - Legally required in many (non-federal) US elections
  - US election system is quite fraught
- US is increasingly using IRV elections
- But there was no method for RLAs for IRV!
- Indeed conjecture it was infeasible

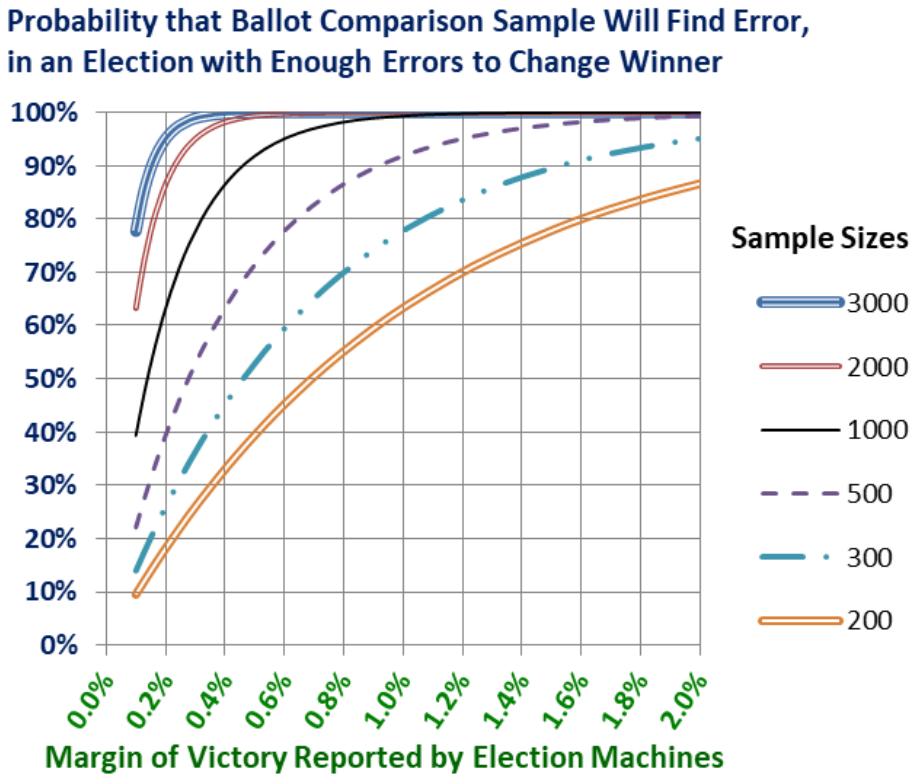

Image from RLA wiki page

Contents

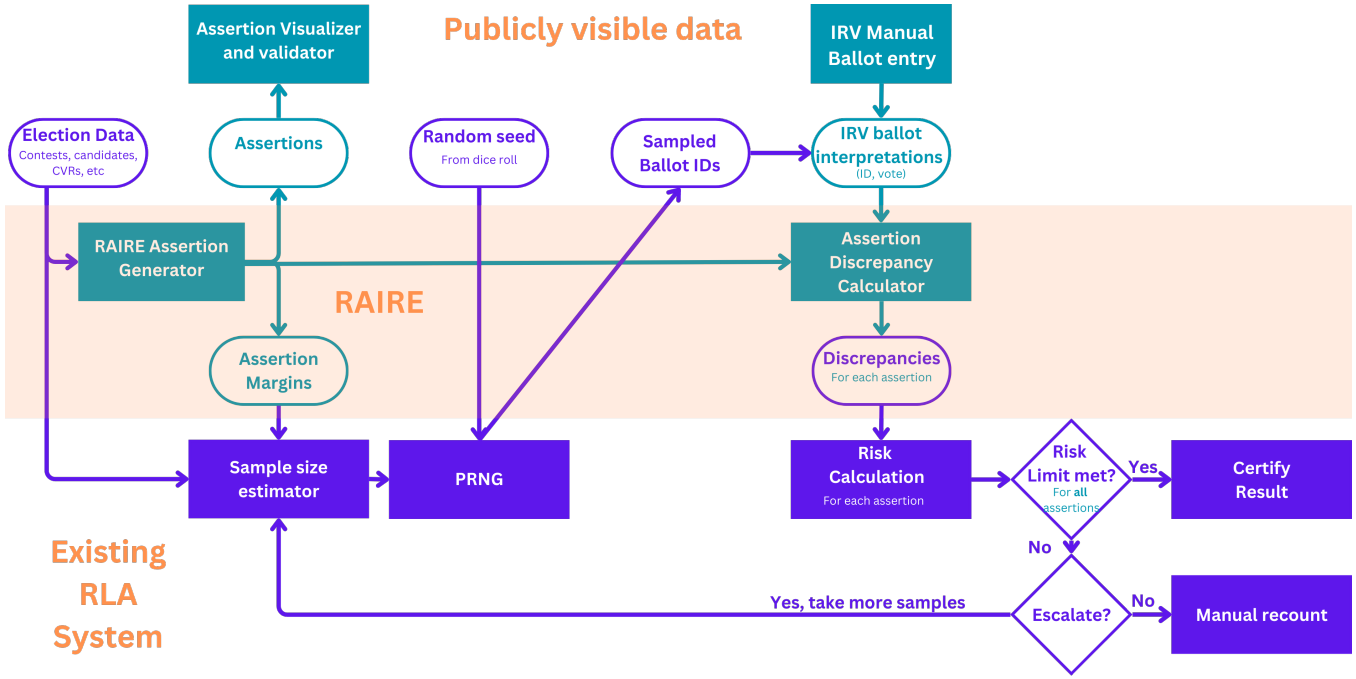

Figure 1.1: A high-level overview of an IRV RLA using RAIRE. Purple elements are existing RLA components. Green indicates new components and data structures for IRV.

RAIRE: Risk-limiting Audits for Instant Runoff Elections

- How do we generate a risk limiting audit for an instant runoff election?
  - First Attempt
  - Elimination Orders and Trees
  - Election Assertions
  - Selecting a set of Assertions
-

First Attempt

- An IRV election is a sequence of first-past-the-post elections [Bob, Chuan, Alice, Diego]
  - Statistically determine that Bob has less votes than each of Chuan, Alice, Diego
  - Determine that Chuan has less votes than each of Alice and Diego
  - Determine that Alice has less votes than Diego
- Problem: **too specific**
  - Audits the election order, not the result
- Consider adding (Elka):5, (Francois):6 to the election: very difficult to audit + irrelevant

Elimination Orders

- To certify the winner of an IRV we need to
  - Eliminate any elections where they dont win
    - In other words, discount elections orders where they are not last
- Note that IRV elections have a nice property
  - Given remaining candidates C, we can determine their piles of ballots
    - **Independent** of the order of elimination of the other candidates
- Hence we can reason about suffixes of election orders

Elimination order tree

- Tree shows suffixes or elimination orders in **reverse**
- We can reason about a node
  - e.g. 5 where only Alice and Chuan remain
- We aim find assertions that remove all possible leaves of alternate winner trees (eliminate a frontier)

Elimination Tree with Winner Alice

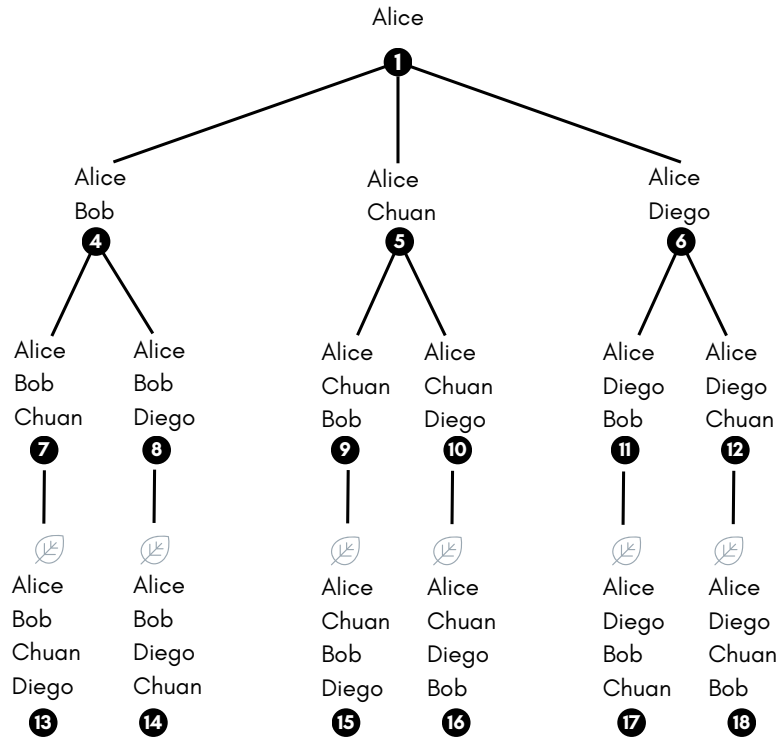

Figure 2.3: All possible elimination orders where Alice is the ultimate winner. The leaves of this elimination tree represent *complete* elimination orders.

Election assertions

- Statements about the IRV election that will prohibit alternate elimination orders
- We will statistically verify that they are each 1 - A likely during the audit
- Hence we can certify the result with no more than A risk
- Our aim is to choose a set of assertions
  - Which prohibits all alternate winner elimination orders
  - For the least expected audit cost

Not eliminated before (NEB) assertions

- Candidate A **NEB** Candidate B if
  - Candidate A's minimum tally (first round preferences) is greater than
  - Candidate B's maximum tally without A eliminated
- Example: (A,B,C,D):100, (B,D,C):40, (C,B,D):40, (C,D):45
  - Alices minimum tally = 100, Bob's maximum tally (while A around) = 80
  - Alice NEB Bob!
  - Not Alice NEB Chuan (Chuan's maximum tally is 125)
- A NEB B **rejects** any elimination order with A eliminated before B

NEB example: Alice NEB Chuan

- Alice NEB Chuan
  - Means Chuan cannot win
  - Many elimination order with Bob winning are also eliminated but not all

Implications of NEB Assertions

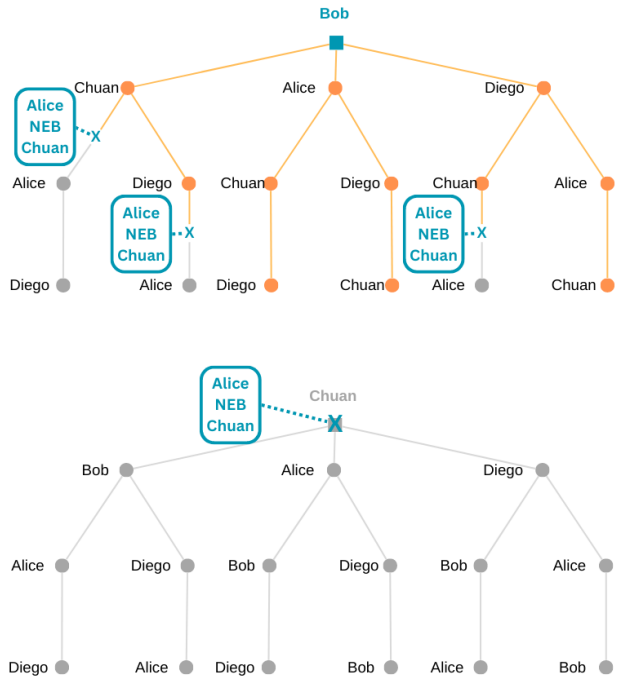

Figure 3.1: Implications of *Alice NEB Chuan* in the trees where Bob and Chuan win. Chuan cannot win, but Bob might.

Not eliminated next (NEN) assertions

- Candidate NEN A > B while remaining candidates C
  - If candidate A always has higher tally than B while candidates C remain
- Example: (A,B,C,D):100, (B,D,C):40, (C,B,D):40, (C,D):45
  - NEN Alice > Chuan when {Alice, Bob, Chuan, Diego} remain
    - Alice min 100, Chuan max 80
  - NEN Bob > Diego when {Bob, Diego} remain
    - Bob min 180, Diego max 45
- NEN A > B while C, rejects any partial order where A is eliminated before B while all of C remain

Combined Example

- Alice NEB Chuan
- Alice NEB Diego
- Alice > Bob if only {Alice, Bob} remain

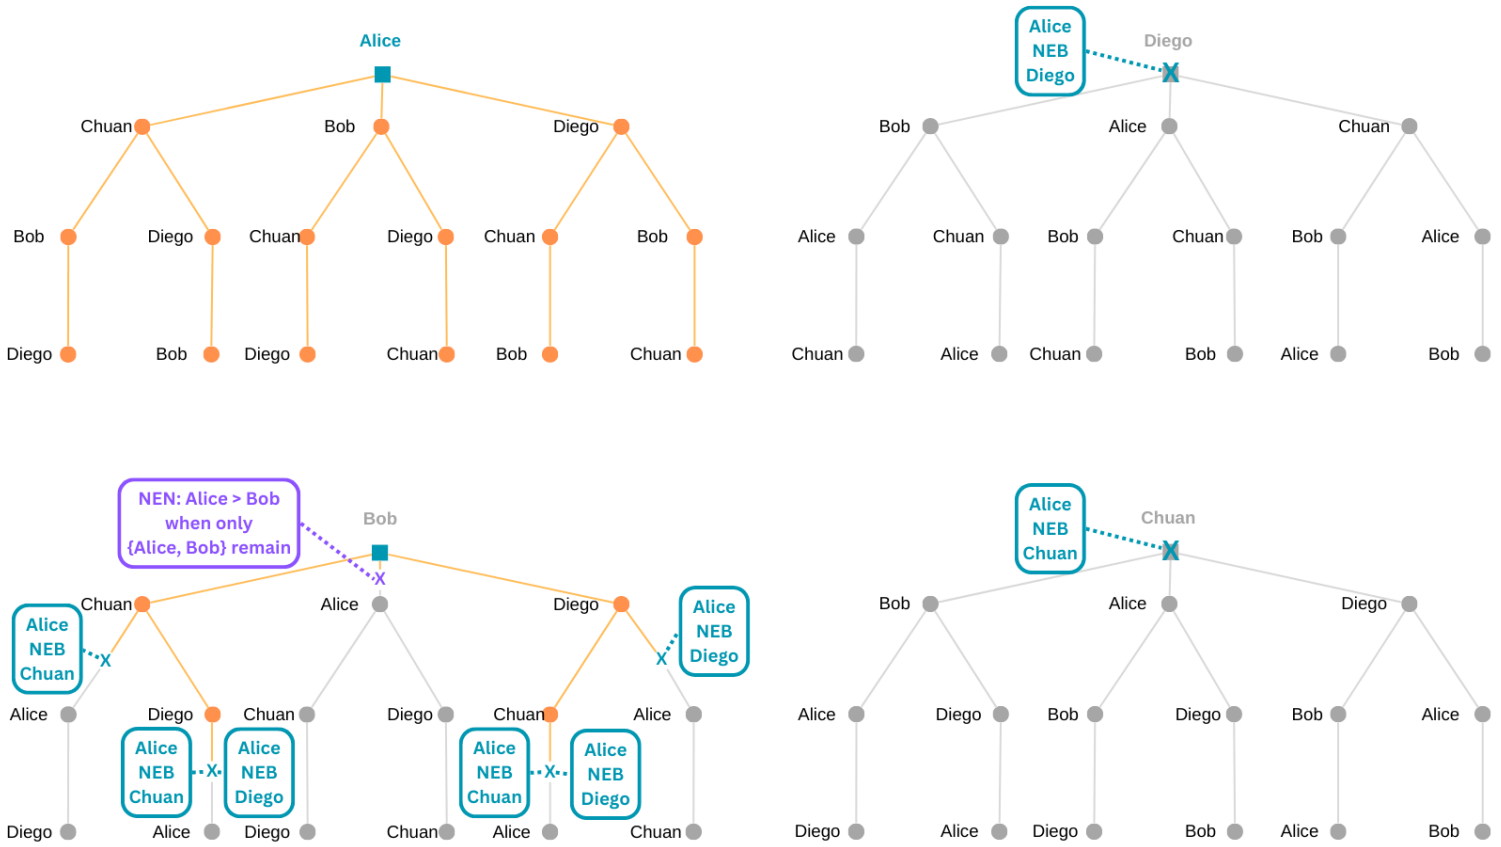

Figure 3.2: Implications of *Alice NEB Chuan* combined with *Alice NEB Diego* and *NEN: Alice > Bob if only {Alice,Bob} remain*. Alice is the only possible winner.

Choosing a set of assertions

- ASN(Assertion): estimated ballots required to reject the negation of Assertion!
- F: Frontier of elimination order nodes to be eliminated: e.g. [Bob], [Chuan], [Diego]
- BestAssertion(Order) returns the assertion eliminating suffix Order with least ASN
  - ∞ if none, attached to each order as expanded
- LB: lower bound on ballots required for entire audit. Initially 0
- S: set of assertions selected, initially empty

Choosing a set of assertions

- Choose the frontier node O where ASN(BestAssertion(O)) is highest
  - If its lower than LB, add BestAssertion(O) to S, delete O
  - Otherwise if O is a leaf node (complete order)
    - Find the assertion A of its parents (and itself) with least ASN,
      - add it to S, update LB, delete all children of this parent node from F
  - Else replace O in F with all of its children elimination order (one more to suffix)
- Return S when frontier F is empty

Does it work?

- MOV = margin of victory
- Polls % = % of ballots examined to certify
- ASN % = estimated % of ballots required to certify
- 36! =3.72×10^{41}

Election	$\|\mathcal{C}\|$	$\|\mathcal{B}\|$	MOV	EO				SE				WO			
				α 0.01		α 0.05		α 0.01		α 0.05		α 0.01		α 0.05	
				Polls %	ASN %	Polls %	ASN %	Polls %	ASN %	Polls %	ASN %	Polls %	ASN %	Polls %	ASN %
Berkeley 2010 D7 CC	4	4,682	364 (7%)	6.7	7.2	3.9	4.7	7.5	7.2	4	4.7	8.7	22.4	4.9	14.7
Berkeley 2010 D8 CC	4	5,333	878 (16%)	∞	∞	∞	∞	2.9	4.2	2	2.8	1.3	1.8	0.8	1.2
Oakland 2010 D6 CC	4	14,040	2,603 (19%)	4.0	4.4	3	2.9	0.7	0.9	0.5	0.6	0.4	0.5	0.3	0.3
Pierce 2008 CC	4	43,661	2,007 (5%)	3.1	2.2	1.8	1.4	3.1	2.2	1.8	1.4	3.2	4.1	1.8	2.7
Pierce 2008 CAD	4	159,987	8,396 (5%)	0.3	0.5	0.2	0.3	0.3	0.5	0.2	0.3	0.5	1.2	0.3	0.8
Aspen 2009 Mayor	5	2,544	89 (4%)	62.4	71.8	52.7	46.9	62.4	71.8	54.8	46.9	∞	∞	∞	∞
Berkeley 2010 D1 CC	5	6,426	1,174 (18%)	2.4	1.7	1.6	1.1	2.4	1.7	1.6	1.1	1.1	1.1	0.8	0.7
Berkeley 2010 D4 CC	5	5,708	517 (9%)	7.5	7	6	4.7	28.7	40.7	17.8	26.6	4.9	7.3	3.8	4.8
Oakland 2012 D5 CC	5	13,482	486 (4%)	11.2	10.3	7.3	6.7	15.1	10.3	11.8	6.7	∞	∞	∞	∞
Pierce 2008 CE	5	312,771	2,027 (1%)	11.6	15.1	7.6	9.8	11.6	15.1	7.6	9.8	∞	∞	∞	∞
San Leandro 2012 D4 CC	5	28,703	2,332 (8%)	9.3	9.7	6.3	6.3	9.3	9.7	6.3	6.3	1.1	4.4	0.8	2.9
Oakland 2012 D3 CC	7	26,761	386 (1%)	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
Pierce 2008 CAS	7	312,771	1,111 (0.4%)	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
San Leandro 2010 Mayor	7	23,494	116 (0.5%)	∞	∞	92.9	∞	∞	∞	92.9	∞	∞	∞	∞	∞
Berkeley 2012 Mayor	8	57,492	8,522 (15%)	94.6	∞	77	∞	2.3	2.6	1.6	1.7	0.2	0.2	0.1	0.2
Oakland 2010 D4 CC	8	23,884	2,329 (10%)	∞	∞	76.4	∞	∞	∞	∞	∞	0.9	3.1	0.6	2
Aspen 2009 CC	11	2,544	35 (1%)	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
Oakland 2010 Mayor	11	122,268	1,013 (1%)	∞	∞	∞	∞	21.5	23.8	15	15.5	∞	∞	∞	∞
Oakland 2014 Mayor	11	101,431	10,201 (10%)	∞	∞	∞	∞	∞	∞	∞	∞	0.8	19.8	0.5	12.9
San Francisco 2007 Mayor	18	149,465	50,837 (34%)	∞	∞	∞	∞	0.03	0.03	0.02	0.02	0.01	0.01	0.01	0.01
Minneapolis 2013 Mayor	36	79,415	6,949 (9%)	∞	∞	∞	∞	∞	∞	∞	∞	0.5	3.1	0.3	2.1
Balmain NSW 2015	7	46,952	1,731 (3.7%)	∞	∞	∞	∞	83.8	∞	65.4	82	5.2	31.6	2.7	20.6
Campbelltown NSW 2015	5	45,124	3,096 (6.9%)	13.6	12.2	8.4	8	∞	∞	∞	∞	1.3	1.7	0.9	1.1
Gosford NSW 2015	6	48,259	102 (0.2%)	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
Lake Macquarie NSW 2015	7	47,698	4,253 (8.9%)	27.7	22.8	14.5	15	6.9	7.8	3.2	5.1	0.7	1.6	0.5	1
Sydney NSW 2015	8	42,747	2,864 (6.7%)	∞	∞	∞	∞	3.3	4.6	2.2	3	1.6	6.9	1	4.5

Contents

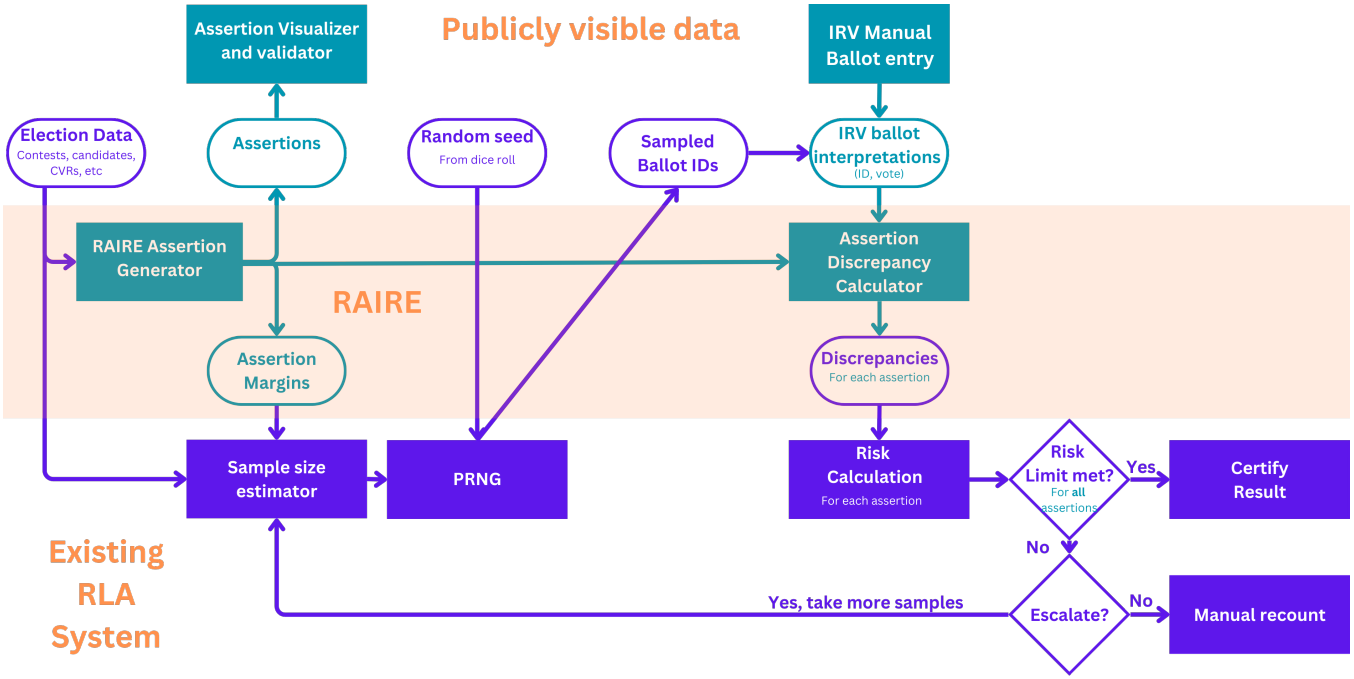

Figure 1.1: A high-level overview of an IRV RLA using RAIRE. Purple elements are existing RLA components. Green indicates new components and data structures for IRV.

Trial

- RAIRE was trialed in an election in
  - San Francisco District Attorney 2019 Instant Runoff Vote
- Just tested the approach on the mailed ballots (which are the only ones with an ID)

Real Life

- Colarado electoral commission
  - Integrated a new version of RAIRE into their election auditing suite
- Since 2023 IRV elections in Colorado have used RAIRE for auditing election

- We expect more states to take up the approach!

Contents

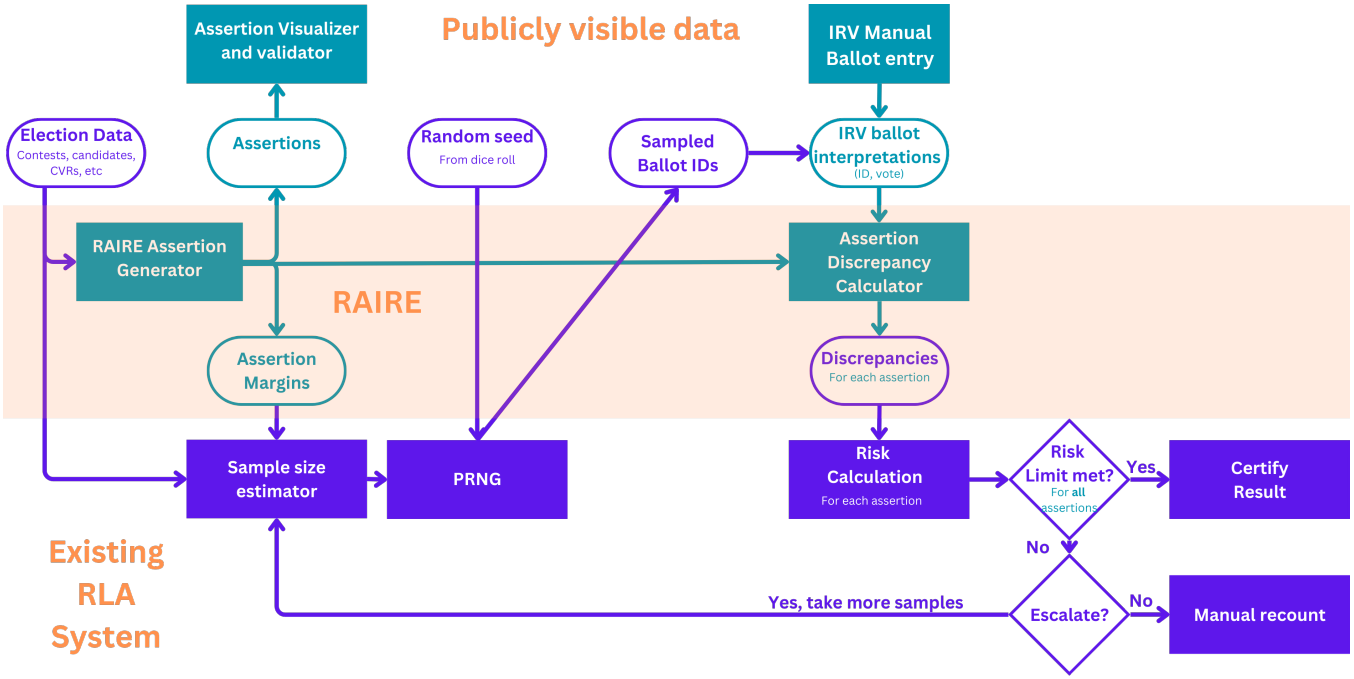

Figure 1.1: A high-level overview of an IRV RLA using RAIRE. Purple elements are existing RLA components. Green indicates new components and data structures for IRV.

Conclusion

- RLAs for IRV require
  - Combinatorics and Statistics working together
- An interesting example of real-world need directly driving research
- One of the quickest example of real-world impact of any of my research
-

Future Work

- We have a replacement (AWAIRE) for RAIRE which have some advantages
  - We can also do comparison auditing (dont need a record of each ballot)
  - The new audits adapt as more ballots are pulled
  - Based on adaptively weighted super-martingales for disjunctions of assertions
  - Combinatorics is now more buried inside the statistics
- We still have no way to RLA Single Transferable Vote elections (for more than 2 seats)

Questions