

RLAs for 2-Seat STV Elections: Revisited*

Michelle Blom¹[0000-0002-0459-9917], Peter J. Stuckey²[0000-0003-2186-0459],
Vanessa Teague³[0000-0003-2648-2565], and Damjan
Vukcevic⁴[0000-0001-7780-9586]

¹ School of Computing and Information Systems, University of Melbourne,
Parkville, Australia

`michelle.blom@unimelb.edu.au`

² Department of Data Science and AI, Monash University, Clayton, Australia

³ Thinking Cybersecurity Pty. Ltd., Melbourne, Australia

⁴ Department of Econometrics and Business Statistics, Monash University,
Clayton, Australia

Abstract. Single Transferable Vote (STV) elections are a principled approach to electing multiple candidates in a single election. Each ballot has a starting value of 1, and a candidate is elected if they gather a total vote value more than a defined quota. Votes over the quota have their value reduced by a *transfer value* so as to remove the quota, and are passed to the next candidate on the ballot. Risk-limiting audits (RLAs) are a statistically sound approach to election auditing which guarantees that failure to detect an error in the result is bounded by a limit. A first approach to RLAs for 2-seat STV elections has been defined. In this paper we show how we can improve this approach by reasoning about lower bounds on transfer values, and how we can extend the approach to *partially* audit an election, if the method does not support a full audit.

1 Introduction

Single Transferable Vote (STV) elections are widely used around the world for multiple candidate contests. Risk-limiting audits (RLAs) are very complex for STV elections. Prior work has demonstrated that RLAs for some 2-seat STV elections, where at least one candidate has more than quota’s worth of votes on their first preferences, are possible [1]. This existing work proposed two approaches for undertaking RLAs for 2-seat STV. The first tackles the case where this *first-round winner criterion* is satisfied, while the second presents a general method that applies when it does not. The latter method was generally not successful in forming an audit to verify the correctness of both winners.

This paper presents an improved method for 2-seat STV RLAs where the contest satisfies the first-round winner criterion. This new method is able to form audits for a greater number of contests, and reduces the expected sample sizes needed for these audits by 15 to 19% across the contests considered in

* This work was partially supported by the Australian Research Council: Discovery Project DP220101012, OPTIMA ITTC IC200100009.

our evaluation. The original method introduced assertions for (i) verifying that the first winner achieved a quota on their first preferences, (ii) verifying an upper bound on the transfer value of this first winner, and (iii) using this upper bound, verifying that the second winner could not have possibly lost to any of the reported losers. We improve this method by introducing a new assertion that verifies a non-trivial *lower bound* on this first winner’s transfer value, and using both bounds to fine tune the assertions formed in (iii). We additionally show how contests that perform a preliminary batch-elimination prior to electing any candidates can be audited using this scheme.

A full RLA, verifying the correctness of both winners, may not be possible for a given 2-seat contest. It may be desirable to perform some kind of audit to verify some aspects of the reported outcome. We show how the ‘general’ method can be re-framed as a five-stage process that forms a *partial* RLA. This process aims to establish (i) a subset of candidates as definite losers, and (ii) a subset of candidates as definite winners. The remaining candidates are *possible* winners. The first three stages of this revised general method are drawn from the work of [1]. In this paper we add a fourth stage that looks for opportunities to reduce the expected sample size of the partial audit. The final stage summarises what aspects of the reported outcome are verified by the audit, and which are not.

The remainder of this paper is structured as follows. [Section 2](#) describes the variant of STV that we consider, and assertion-based RLAs. Three sections follow that consider different election scenarios and how to audit them: [Section 3](#) covers auditing of batch elimination, [Section 4](#) covers the improved *first-round winner* method and an evaluation against the existing approach, and [Section 5](#) shows how we can partially audit elections where no candidate has a quota on first preferences. We conclude in [Section 6](#).

2 Preliminaries

We consider a variant of STV, modelled on how STV is typically implemented in the United States. We describe this variant in [Section 2.1](#).

We define an STV election as per [Definition 1](#). We define a ballot b as a sequence of candidates π , listed in order of preference (most popular first), without duplicates but without necessarily including all candidates. We use list notation (e.g., $\pi = [c_1, c_2, c_3, c_4]$). The notation $\text{first}(\pi) = \pi(1)$ denotes the first item (candidate) in sequence π .

Definition 1 (STV Election). *An STV election E is a tuple $E = (\mathcal{C}, \mathcal{B}, \mathcal{Q}, N)$ where \mathcal{C} is a set of candidates, \mathcal{B} the multiset of ballots cast⁵, \mathcal{Q} the election quota (the number of votes a candidate must attain to win a seat—usually the Droop quota—[Equation 1](#)), and N the number of seats to be filled.*

$$\mathcal{Q} = \left\lfloor \frac{|\mathcal{B}|}{N + 1} \right\rfloor + 1 \tag{1}$$

⁵ A multiset allows for the inclusion of duplicate items.

Ranking	Count	Ranking	Count		
$[c_1, c_3]$	8,001	$[c_3, c_4]$	5,000	$N = 2$	
$[c_1]$	1,000	$[c_4, c_1, c_2]$	3,950	$Q = 7,001$	
$[c_2, c_3, c_4]$	3,000	$[c_5, c_2]$	50	$t_{c_1,1} = 9,001$	$t_{c_4,1} = 3,950$
				$t_{c_2,1} = 3,000$	$t_{c_5,1} = 50$
Total			21,001	$t_{c_3,1} = 5,000$	

Table 1: An STV election, stating the number of ballots cast with each listed ranking over candidates c_1 to c_5 . The quota, and first-preference tallies are listed.

Definition 2. Projection $\sigma_{\mathcal{S}}(\pi)$ We define the projection of a sequence π onto a set \mathcal{S} as the largest subsequence of π that contains only elements of \mathcal{S} . (The elements keep their relative order in π .) For example: $\sigma_{\{c_2, c_3\}}([c_1, c_2, c_4, c_3]) = [c_2, c_3]$ and $\sigma_{\{c_2, c_3, c_4, c_5\}}([c_6, c_4, c_7, c_2, c_1]) = [c_4, c_2]$.

The tabulation of STV elections proceeds in rounds (see Section 2.1). Initially, all candidates are awarded the ballots on which they are the first ranked candidate. We call a candidate c 's tally at this stage their *first-preference tally*, denoted $t_{c,1}$. We use $t_{c,r}$ to denote a candidate's tally at the start of round r of tabulation. In the election of Table 1, candidates c_1 to c_5 have first-preference tallies of 9001, 3000, 5000, 3950, and 50 votes, respectively.

2.1 'US' style STV

Each ballot cast in the election starts with a value of 1. In any given round of tabulation, if no candidate's tally is equal to or above the election's quota, the candidate with the smallest tally is eliminated. All the ballots in the eliminated candidate's tally are redistributed to the next most preferred *eligible* candidate on the ballot. These ballots are transferred at their current value. At any stage where ballots are distributed from one candidate to another, the only candidates that are eligible to receive votes are those that have not yet been eliminated or elected to a seat, and who have less than a quota's worth of votes in their tally.

Eliminations proceed as described above until at least one candidate's tally equals or exceeds the election's quota. At this stage, these candidates are elected to a seat. These candidates will be elected to a seat in order of their *surplus*. For each such candidate, we define their *surplus* as the difference between their current tally and the quota. The ballots sitting in the tally pile of this candidate are *reweighted* and distributed to the next most preferred eligible candidate on the ballot. For a candidate c , elected to a seat in round r of tabulation, we define the *transfer value* τ_c used to reweight the ballots in their tally as shown in Equation 2, where $t_{c,r}$ denotes the tally of c at the start of round r .

$$\tau_c = \frac{t_{c,r} - Q}{t_{c,r}} \quad (2)$$

For a ballot $b \in \mathcal{B}$, whose current value is $v_{b,r}$, its new value when it leaves the tally pile of candidate c upon their election to a seat becomes $\tau_c v_{b,r}$.

Tabulation proceeds by seating candidates whose tally reaches or exceeds a quota and distributing their votes to eligible candidates, and eliminating candidates when no candidate has a quota. This process continues until either all seats have been filled, or the number of candidates still standing equals the number of seats left to be filled. These remaining candidates are then elected.

Consider the election in [Table 1](#). The quota is 7001 votes. Candidate c_1 has a quota on first preferences, and is elected to the first seat. Their transfer value is $\tau = (9001 - 7001)/9001 = 0.222$. The 8,001 $[c_1, c_3]$ ballots are given to candidate c_3 , adding 1,776.222 votes to c_3 's tally. The 1,000 $[c_1]$ ballots become *exhausted*. Candidate c_3 now has 6,776.222 votes. As no candidate has a quota's worth of votes, the candidate with the smallest tally is eliminated. Here, this is candidate c_5 on 50 votes. These 50 ballots are given to c_2 at their current value of 1. Candidate c_2 now has 3,050 votes. Still no candidate has a quota's worth of votes. Candidate c_2 is eliminated next. The 50 $[c_5, c_2]$ ballots become exhausted, and the 3,000 $[c_2, c_3, c_4]$ ballots are given to c_3 , who now has 9,776.222 votes. Candidate c_3 has achieved a quota, and is elected to the second seat.

Batch elimination We also consider a variation of the above process in which a batch elimination step is first performed. We first determine if there are any candidates for which there is no mathematical possibility for them to win. For each candidate, we compute the number of ballots on which they are ranked, and compare this tally to the tally of the N candidates with the highest first-preference tally. Consider the election in [Table 1](#). The two candidates with the highest first-preference tallies are c_1 on 8,001 votes, and c_3 on 5,000 votes. Candidate c_5 is ranked on 50 ballots. Candidates c_2 and c_4 are ranked on 7,000 and 11,950 ballots, respectively. There is no possibility for c_5 to win, so they are eliminated in the first round, and their 50 $[c_5, c_2]$ ballots are given to c_2 .

Tabulation then proceeds as described above. Candidate c_1 is elected, giving c_3 1,776.2 votes and a tally of 6,776.2. Candidate c_2 is eliminated, giving 3,000 votes to c_3 . Candidate c_3 is elected to the second seat on a tally of 9,776.2.

2.2 Assertion-based Approaches to Risk-Limiting Audits

SHANGRLA [\[3\]](#) provides a general framework for RLAs, using *assertions* as ‘building blocks’. An *assertion* is a statement about the full set of ballots in an election. These are typically expressed as an inequality about some property that would be consistent with a particular election outcome. An example of an assertion is “Alice received more votes than Bob”. In the SHANGRLA framework, we need to design a set of assertions such that, if they are all true, they imply that the reported winner really won the election. To conduct an audit, we statistically test each assertion using general statistical methods that form part of the framework. Assertions need to have a specific mathematical form to fit into the framework. In general, any linear combination of tallies (counts of different types of ballots) can be converted into a SHANGRLA assertion [\[2\]](#). All of the assertions we develop in this paper are of this form.

Table 2: Assertions verifying the batch elimination of the UWIs and Nikiforakis in the 2021 BoE election in Minneapolis, Minnesota, and their sample sizes.

Assertion	Sample Size
AG(S. Brandt, UWIs)	20
AG(S. Pree-Stinson, UWIs)	35
AG(S. Brandt, K. Nikiforakis)	27
AG(S. Pree-Stinson, K. Nikiforakis)	69
Total cost:	69

3 Context: Batch Elimination First

We first consider how we can verify, in an RLA, that the candidates eliminated as part of an initial batch elimination did indeed have no mathematical possibility of winning. To do so, we use the existing AG assertion of Blom *et al* [1].

AG(w, l) Verifies that candidate w always has a higher tally than l by showing that w 's first-preference tally is higher than the maximum tally l could achieve while w is still standing: $t_{w,1} > |\{b : b \in \mathcal{B}, \text{first}(\sigma_{\{w,l\}}(b)) = l\}|$.

For a US-STV election $E = (\mathcal{C}, \mathcal{B}, \mathcal{Q}, N)$, let $Top = \{c_1, \dots, c_N\} \subset \mathcal{C}$ denote the N candidates with the highest first-preference tallies, and $Batch \subset \mathcal{C}$ the set of batch eliminated candidates. We verify that candidate $c \in Batch$ cannot win by showing that **AG**(c_i, c) for all $c_i \in Top$.

Example 1. We consider the 2021 Board of Estimates and Taxation election in Minneapolis, Minnesota. This two-seat STV election involved four candidates – S. Brandt, S. Pree-Stinson, P. Salica, and K. Nikiforakis – and a number of undeclared write-ins (UWIs). The first-preference tallies were 42672 votes for S. Brandt, 25597 votes for S. Pree-Stinson, 20786 votes for P. Salica, 5815 votes for K. Nikiforakis and 755 votes for UWIs. The quota or election threshold was 31876, and out of 145337 ballots, 49712 of these were invalid. The UWIs and K. Nikiforakis were eliminated in the first round with too few mentions to have a mathematical possibility of winning. S. Brandt was then elected, P. Salica was eliminated, leaving S. Pree-Stinson as the second winner. We verify this batch elimination in an RLA with the assertions shown in Table 2, alongside the expected number of ballots required to audit them.⁶ \square

Example 2. Batch Elimination First can change the result of an election. Consider the two-seat STV with candidates $\{w, a, b, c_1, c_2, c_3, c_4, c_5\}$ and ballots $[w] : 15001, [a] : 6875, [b] : 3125, [c_1, w, b] : 1000, [c_2, w, b] : 1000, [c_3, w, b] : 1000, [c_4, w, b] : 1000, [c_5, w, b] : 1000$. The Droop Quota is 10001. Without batch elimination we give a seat to w , then all the votes for w exhaust; then each of the c_i are eliminated leaving tallies $a : 6825$ and $b : 8125$, and finally b is seated.

⁶ For all sample size estimations, we assume a risk limit of 10%, an error rate of 2 overstatements per 1000 ballots, and the ALPHA risk function of SHANGRLA [3].

With a batch elimination all of c_1, \dots, c_5 are eliminated first, none of them has enough mentions to beat w , a , or b . Then w gets a seat with 20001 votes in its tally and each of its surviving votes are transferred at value 0.5. The tallies for a and b are then $a : 6875$ and $b : 5625$, so finally a is seated. \square

4 Context: First Round Winner

Let $E = (\mathcal{C}, \mathcal{B}, \mathcal{Q}, N = 2)$ denote a US-STV election with winners $w_1, w_2 \in W$ in which candidate w_1 is elected to the first seat, in the first round of tabulation (i.e., $t_{w_1,1} \geq \mathcal{Q}$), after any batch elimination has taken place.

Prior Work: In the approach of [1], an IQ assertion is formed to verify that w_1 has a quota on their first preferences, $\text{IQ}(w_1)$. They then establish an estimated upper bound, $\bar{\tau}_{w_1}$, on the transfer value of ballots from w_1 using an assertion of the form $\text{UT}(w_1, \bar{\tau}_{w_1})$. Using this upper bound they create assertions to show that w_2 will always have a higher tally than all other candidates using NL assertions. The method continues to increase the transfer value upper bound $\bar{\tau}_{w_1}$, until the sample size of the resulting audit increases, or $\bar{\tau}_{w_1}$ reaches $2/3$ (the maximum transfer value in a 2 seat STV election).

New Approach: We vary this approach by introducing additional types of assertions to reduce the expected sample sizes required in an audit. The assertions (beyond AG) we use are (starred assertions are new):

- IQ**(c) Verifies that candidate c 's first-preference tally is equal to or greater than a quota: $t_{c,1} \geq \mathcal{Q}$.
- UT**($c, \bar{\tau}_c$) Assumes that candidate c has been elected on their first preferences, and verifies that the transfer value for c is less than $\bar{\tau}_c$: $t_{c,1} < \mathcal{Q}/(1 - \bar{\tau}_c)$.
- LT***($c, \underline{\tau}_c$) Assumes that c has been elected on their first preferences, and verifies that the transfer value for c is greater than $\underline{\tau}_c$: $t_{c,1} > \mathcal{Q}/(1 - \underline{\tau}_c)$.
- AG***($w, l, W, \underline{\tau}, \bar{\tau}$) An extension of the AG assertion [1]. The assertion shows that candidate w will always have higher tally than candidate l in the context where the candidates in W have already been elected to a seat with lower and upper bounds on their transfer values $\underline{\tau}$ and $\bar{\tau}$.

The assertion compares the minimum tally of candidate w in this context, with the maximum tally of l . In the original AG assertion [1], the minimum tally of w consists only of those ballots on which w is ranked first. The AG* assertion retains this and adds further counts to this minimum tally by including contributions from some ballots where w is not ranked first. Specifically, for all ballots b where $\text{first}(\sigma_{\mathcal{C}-W}(b)) = w$, we reduce them in value by taking a product of transfer value lower bounds for the candidates in W that precede w in its ranking, and add these to w 's minimum tally.⁷

⁷ In the actual election, there is a scenario where candidate w does not get these ballots in their vote count: if w is our next winner after those in W . In such a case, w will be a winner, and thus we can ignore it since the context where we use these assertions is precisely to show that w is a winner.

For each ballot $b \in \mathcal{B}$, we define its contribution to the minimum tally of w , and the maximum tally of l , as follows.

$$C_{min}^{AG^*}(b, w, W, \underline{\tau}, \bar{\tau}) = \begin{cases} 1 & \text{first}(b) = w \\ \prod_{k \in W'} \underline{\tau}_k & \text{first}(\sigma_{\mathcal{C}-W}(b)) = w \\ & \text{and } W' = \{c \in W : c \text{ precedes } w \text{ in } b\} \\ 0 & \text{otherwise} \end{cases}$$

$$C_{max}^{AG^*}(b, l, W, \underline{\tau}, \bar{\tau}) = \begin{cases} 0 & l \text{ does not occur in } b \\ 0 & w \text{ appears before } l \text{ in } b \\ \text{maxt}(b, l, W, \bar{\tau}) & \text{first}(b) \in W \\ 1 & \text{otherwise} \end{cases}$$

where $\text{maxt}(b, l, W, \bar{\tau}) = \max\{\bar{\tau}_c : c \in W \text{ precedes } l \text{ in } b\}$. We define the minimum tally of w , $t1_w^{min}$, and the maximum tally of l , $t1_l^{max}$, as follows:

$$t1_w^{min} = \sum_{b \in \mathcal{B}} C_{min}^{AG^*}(b, w, W, \underline{\tau}, \bar{\tau}) \quad (3)$$

$$t1_l^{max} = \sum_{b \in \mathcal{B}} C_{max}^{AG^*}(b, l, W, \underline{\tau}, \bar{\tau}) \quad (4)$$

We say that $AG^*(w, l, W, \underline{\tau}, \bar{\tau})$ iff $t1_w^{min} > t1_l^{max}$. Note $AG(w, l) \equiv AG^*(w, l, \emptyset, -, -)$. **NL*** $(w, l, W, \underline{\tau}, \bar{\tau}, G^*, O^*)$ An extension of the NL assertion [1]. Establishes that candidate w will always have a higher tally than candidate l under the assumptions that: (i) the candidates in W *have already been seated*, with lower and upper bounds on their transfer values $\underline{\tau}$ and $\bar{\tau}$; (ii) G^* denotes the candidates $g \in \mathcal{C}$ for which $AG^*(g, l, W, \underline{\tau}, \bar{\tau})$ holds; and (iii) O^* the candidates $o \in \mathcal{C}$ for which $AG^*(w, o, W, \underline{\tau}, \bar{\tau})$ holds. This contrasts with the assumptions underlying the original NL assertion, which only assumes that the candidates W are seated *at some point*. The assertion compares the minimum tally of w , in this context, against the maximum tally of l .

We define w 's minimum tally at a point at which they could be eliminated, where it is assumed that O^* have been prior eliminated. This minimum tally includes all ballots b where $\text{first}(\sigma_{\mathcal{C}-O^*}(b)) = w$, at value 1, and all ballots b where $\text{first}(\sigma_{\mathcal{C}-W}(b)) = w$, at a reduced value. For the maximum tally of l , we include all ballots on which l precedes w in their ranking, or l appears and w does not, excluding those on which a candidate $g \in G^*$ precedes l . For each ballot $b \in \mathcal{B}$, we define its contribution to the minimum tally of w , and the maximum tally of l , as follows.

$$C_{min}^{NL^*}(b, w, W, \underline{\tau}, \bar{\tau}, O^*) = \begin{cases} 1 & \text{first}(\sigma_{\mathcal{C}-O^*}(b)) = w \\ \prod_{k \in W'} \underline{\tau}_k & \text{first}(\sigma_{\mathcal{C}-W}(b)) = w \\ & \text{and } W' = \{c \in W : c \text{ precedes } w \text{ in } b\} \\ 0 & \text{otherwise} \end{cases}$$

$$C_{max}^{\text{NL}^*}(b, l, W, \underline{\tau}, \bar{\tau}, G^*) = \begin{cases} 0 & l \text{ does not occur in } b \\ 0 & w \text{ appears before } l \text{ in } b \\ 0 & a g \in G^* \text{ appears before } l \text{ in } b \\ \text{maxt}(b, l, W, \bar{\tau}) & \text{first}(b) \in W \\ 1 & \text{otherwise} \end{cases}$$

We define the minimum tally of w , $t2_w^{\text{min}}$, and the maximum tally of l , $t2_l^{\text{max}}$, as follows:

$$t2_w^{\text{min}} = \sum_{b \in \mathcal{B}} C_{min}^{\text{NL}^*}(b, w, W, \underline{\tau}, \bar{\tau}, O^*) \quad (5)$$

$$t2_l^{\text{max}} = \sum_{b \in \mathcal{B}} C_{max}^{\text{NL}^*}(b, l, W, \underline{\tau}, \bar{\tau}, G^*) \quad (6)$$

We say that $\text{NL}^*(w, l, W, \underline{\tau}, \bar{\tau}, G^*, O^*)$ iff $t2_w^{\text{min}} > t2_l^{\text{max}}$.

Figure 1 outlines the procedure used to generate the assertions \mathcal{A} of our new RLA for two-seat STV elections satisfying the first-round winner criterion. The prior approach [1] involved a single loop in which an upper bound on the first winner transfer value was incremented, and a candidate audit formed for each of these potential values for this upper bound. The original AG assertions were computed prior to this loop, as they did not take into account upper or lower bounds on the first winner’s transfer value. Our new approach involves two loops – the outer loop (steps 5-34) over potential values for the lower bound on the first winner’s transfer value, $\underline{\tau}_{w_1}$, and the inner loop (steps 10-29) over potential values for the upper bound on the first winner’s transfer value, $\bar{\tau}_{w_1}$. For each candidate value of $\underline{\tau}_{w_1}$, the inner loop searches for a value for $\bar{\tau}_{w_1}$ that results in the cheapest audit. The outer loop searches for a value for $\underline{\tau}_{w_1}$ for which the inner loop yields the cheapest overall audit. As per Blom *et al* [1], the first assertion we create is $\text{IQ}(w_1)$ to verify that our first winner, w_1 , does indeed achieve a quota on their first preferences (step 1). Where a group elimination has taken place, and our resulting election satisfies the first-round winner criterion, the $\text{IQ}(w_1)$ assertion verifies that w_1 has a quota on the basis of their first-preference tally *and* any votes distributed to them from the group eliminated candidates.

AG^* assertions, which are used to help us form the NL^* assertions required to show that w_2 beats all of the original losers, are formed inside the inner loop (step 14), allowing us to take advantage of both lower and upper bounds on the first winner’s transfer value. When forming each NL^* , we add an AG^* to our audit only if it allows us to reduce the expected ASN of the NL^* we are trying to form, and where the ASN of the NL^* without the AG^* is higher than that of the AG^* itself. In this way, we do not add AG^* assertions to our audit where their benefit, in terms of making a NL^* easier to audit, is outweighed by their cost.

Note that if our final audit contains an AG^* and NL^* with the same winner and loser, we remove the NL^* from our audit as it is redundant.

Example 3. Consider again the 2021 Board and Estimates and Taxation election (Minneapolis, Minnesota). **Example 1** presents the first stage of an RLA for

this election, identifying the assertions required to check that the two batch eliminated candidates did not have a mathematical possibility of winning. After the distribution of these eliminated candidates' ballots, we have a three candidate election that satisfies the first winner criterion.

S. Brandt is elected at this stage, with two remaining candidates (S. Pree-Stinson and P. Salica) vying for the second seat. Using the algorithm in [Figure 1](#), we form the assertion $\text{IQ}(S.\text{Brandt})$. The ASN for this assertion is 34 ballots.

We then enter the outer loop at step 5 with a lower bound on S. Brandt's transfer value set to 0. (Where $\underline{\tau}_{w_1}$ is 0 we actually do not compute the associated LT^* assertion as it is not necessary). Then, starting with an upper bound on S. Brandt's transfer value set to his actual transfer value plus δ (with $\delta = 0.05$), we enter the inner loop of the algorithm in [Figure 1](#) at step 10. The UT assertion required to show that S. Brandt's transfer value is less than, in this case, 0.3311, has an ASN of 131 ballots. AG^* assertions are computed (step 14), and the NL^* assertion required to show that S. Pree-Stinson never loses to P. Salica in the context where S. Brandt is seated first (steps 16 to 24). The NL^* assertion has an ASN of 402 ballots. (In this case, none of the AG^* assertions were found to be helpful in reducing the margin of this NL^* assertion, and $\mathbf{AG}' \leftarrow \emptyset$ in step 22). At this stage, we have an RLA for the election that costs 402 ballots. The inner loop is repeated, with the upper bound on S. Brandt's transfer value set to 0.3811. The required UT assertion now costs 60 ballots. However, when proceeding to create the NL^* assertion required to show that S. Pree-Stinson never loses to P. Salica, the assertion now costs 628 ballots. The new candidate configuration for our RLA is more costly, at 628 ballots, than the previous one, at 402 ballots. So, we break out of our inner loop at step 29.

We repeat our outer loop, with the lower bound on S. Brandt's transfer value now 0.1406 (or half of his actual transfer value, as per step 33). The LT^* assertion required to show that his transfer value is greater than this lower bound has an ASN of 59 ballots. We enter the inner loop at step 5 with the upper bound on S. Brandt's transfer value again set to 0.3311. The ASN of the UT assertion for this bound is, as before, 131 ballots. After proceeding through steps 14–24, we form an NL^* assertion to show that S. Pree-Stinson never loses to P. Salica that now costs 247 ballots (again, we opt not to make use of any computed AG^* assertions). We now have a configuration for our audit that costs 247 ballots in total. When incrementing $\bar{\tau}$ for S. Brandt to 0.3811, we do not improve upon this ASN (in fact, it will increase to 285). We break out of the inner loop at step 29.

The outer loop will be repeated with $\underline{\tau}$ for S. Brandt increased to 0.1906. The required LT^* assertion for this bound will cost 87 ballots. By working through the inner loop, as before, we are able to find an audit configuration costing 217 ballots. Repeating the outer loop again with $\underline{\tau}$ for S. Brandt increased to 0.2406 gives us an audit costing 194 ballots. The LT^* assertion will cost 184 ballots in this audit, the UT 131 ballots, and the required NL^* 194 ballots. Incrementing $\underline{\tau}$ for S. Brandt again to 0.2906, in step 33, puts us beyond his actual transfer value of 0.2811. The outer loop condition fails, and we finish with an audit costing 194 ballots. This audit contains the assertions listed in [Table 3](#). \square

```

1   $iq \leftarrow \text{IQ}(w_1) \triangleright$  Form assertion to verify that  $w_1$  has a quota on first preferences
2   $\underline{\tau}_{w_1} \leftarrow 0 \triangleright$  Lower bound on transfer value for first winner  $w_1$ 
3   $ASN \leftarrow \infty \triangleright$  ASN of our audit
4   $\mathcal{A} \leftarrow \emptyset \triangleright$  Assertions in our audit
5  while  $\underline{\tau}_{w_1} < \tau_{w_1}$  do  $\triangleright$   $\tau_{w_1}$  is the reported transfer value for  $w_1$ 
6     $lt \leftarrow \text{LT}^*(w_1, \underline{\tau}_{w_1}) \triangleright$  Form  $\text{LT}^*$  assertion, denoted  $lt$ .
7     $\bar{\tau}_{w_1} \leftarrow \tau_{w_1} + \delta \triangleright$  Upper bound on transfer value for first winner  $w_1$ 
8     $\mathcal{A}' \leftarrow \emptyset$ 
9     $ASN' \leftarrow ASN$ 
10   while  $\bar{\tau}_{w_1} < 2/3$  do
11      $ut \leftarrow \text{UT}(w_1, \bar{\tau}_{w_1}) \triangleright$  Form  $\text{UT}$  assertion, denoted  $ut$ .
12      $\mathcal{A}'' \leftarrow \{lt, ut, iq\}$ 
13      $ASN'' \leftarrow \max(lt.ASN, ut.ASN, iq.ASN)$ 

     $\triangleright$  Compute  $\text{AG}^*$  assertions between  $w_2$  and each  $l \in \text{losers}$ , and
    between each  $l, l' \in \text{losers}$  such that  $l \neq l'$ 
14    $\text{AG} \leftarrow [\text{AG}^*(c, l, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}]) \mid \forall c \in \text{losers} \cup \{w_2\}, l \in \text{losers}, c \neq l]$ 
15    $O^* \leftarrow [c \mid \text{AG}^*(w_2, c, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}]) \in \text{AG}]$ 

     $\triangleright$  Find  $\text{NL}^*$  assertions to show that  $w_2$  never loses to each  $l \in \text{losers}$ 
16   for each  $l \in \text{losers}$  do
17      $G^* \leftarrow [c \mid \text{AG}^*(c, l, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}]) \in \text{AG}]$ 

18      $t2_{w_2}^{min} \leftarrow \sum_{b \in \mathcal{B}} C_{min}^{\text{NL}^*}(b, w_2, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}], O^*)$  (Equation 5)
19      $t2_l^{max} \leftarrow \sum_{b \in \mathcal{B}} C_{max}^{\text{NL}^*}(b, l, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}], G^*)$  (Equation 6)

20     if  $t2_{w_2}^{min} > t2_l^{max}$  then
21        $nl \leftarrow \text{NL}^*(w_2, l, [w_1], [\underline{\tau}_{w_1}], [\bar{\tau}_{w_1}], G^*, O^*)$ 
22        $\text{AG}' \leftarrow \text{AG}^*$  assertions between  $w_2$  and  $o \in O^*$ , and between
        $g \in G^*$  and  $l$ , that were used to reduce the ASN of  $nl$ 
23        $\mathcal{A}'' \leftarrow \mathcal{A}'' \cup \{nl\} \cup \text{AG}'$ 
24        $ASN'' \leftarrow \max(ASN'', nl.ASN, a.ASN \forall a \in \text{AG}')$ 

25   if  $ASN'' < ASN'$  then
26      $\mathcal{A}' \leftarrow \mathcal{A}''$ 
27      $ASN' \leftarrow ASN''$ 
28      $\bar{\tau}_{w_1} \leftarrow \bar{\tau}_{w_1} + \delta$ 
29   else break
30   if  $ASN' < ASN$  then
31      $\mathcal{A} \leftarrow \mathcal{A}'$ 
32      $ASN \leftarrow ASN'$ 
33      $\underline{\tau}_{w_1} \leftarrow \underline{\tau}_{w_1} + \delta$  if  $\underline{\tau}_{w_1} > 0$  and  $\frac{\tau_{w_1}}{2}$  otherwise
34   else break

```

Fig. 1: Algorithm for generating assertions \mathcal{A} for the revised RLA of a two-seat STV election satisfying the first-round winner criterion. The two reported winners of the election are w_1 and w_2 , and losers denotes the remaining candidates. Given an assertion, a , we use the notation “ $a.ASN$ ” to denote its ASN.

Table 3: Assertions verifying the election of S. Brandt and S. Pree-Stinson in the 2021 BoE election in Minneapolis, Minnesota, and their sample sizes.

Assertion	ASN	Assertion	ASN
Batch elimination		Election of S. Brandt and S. Pree-Stinson	
AG(S. Brandt, UWIs)	20	IQ(S. Brandt)	34
AG(S. Pree-Stinson, UWIs)	35	LT*(S. Brandt, 0.2406)	184
AG(S. Brandt, K. Nikiforakis)	27	UT(S. Brandt, 0.3311)	131
AG(S. Pree-Stinson, K. Nikiforakis)	69	NL*(S. Brandt, P. Salica, ...)	194
Total cost:		194	

4.1 Evaluation

We contrast the expected cost (ASN) of our new two-seat STV RLA (for elections satisfying the first-round winner criterion) relative to the existing method [1]. For sample size estimations, we use a risk limit of 10%, an expected error rate of 2 overstatements per 1000 ballots, and the ALPHA risk function of SHANGRLA [3]. For the algorithm shown in Figure 1, we use a value of 0.05 for δ . Table 4 contrasts the expected sample sizes required by our 2-seat STV RLAs across a set of real 2-seat STV instances—four BoE elections held in Minneapolis, Minnesota between 2009 and 2021, and four elections held as part of the Australian Senate election in 2016 and 2019—and a series of US and Australian (NSW) IRV elections re-imagined as 2-seat STV contests. All these instances satisfy the first-round winner criterion. Selected instances from the full set of 92 NSW Legislative Assembly (NSW-LA) elections, and 23 US IRV elections, are shown in Table 4.

Across the full set of 92 NSW-LA elections (re-imagined as 2-seat STV), no RLA could be formed using the prior approach [1] for 8 instances. With the new method, four of these instances become auditable—although in one case, Lismore, the cost is still quite high at 2180 ballots. Across the remaining 84 instances, the new approach reduces required sample sizes by 15% on average. Across the full set of 23 US IRV elections (re-imagined as 2-seat STV), no RLA could be formed for six instances using the prior approach [1]. Using the new approach, three of these instances become auditable, with sample sizes of 3925, 350, and 166, as shown in Table 4. Again, the new method reduces sample sizes for the remaining 17 elections by 15% on average. For the Australian Senate and Minneapolis STV elections, the new method reduces required sample sizes by 19%, on average.

5 Scenario: General Method

The *general method* of Blom *et al* [1] describes how we can form an RLA for a 2-seat STV election where no candidate has a quota on their first preferences. We do not, in this paper, present an improvement to this approach—in the sense of enabling audits for instances that we could not previously audit. We do, however, show how we can adapt the method to perform *partial* audits of elections where a full RLA, verifying both winners, is not possible. The experiments of Blom *et*

Table 4: ASNs for our 2 seat STV RLAs, comparing the original method of [1] against the revised method. All instances satisfy the first-round winner criterion. Where the revised method improves on the original, ASNs are in bold.

Instance	$ C $	$ \mathcal{B} $	Q	ASN	
				Prior RLA	New RLA
MN BoE 2009	7	32086	10696	191	100
MN BoE 2013	5	48855	16286	33	31
MN BoE 2017	4	69694	23232	23	23
MN BoE 2021	5	95625	31876	402	194
AU Senate'16 ACT	22	254767	84923	77	58
AU Senate'19 ACT	17	270231	90078	131	98
AU Senate'16 NT	19	102027	34010	60	60
AU Senate'19 NT	18	105027	35010	58	58
US IRV elections re-imagined as 2 seat STV					
MN Mayor 2013	36	79415	26472	73	73
Aspen'09 Mayor	5	2528	843	43	41
Berkeley'10 D1 CC	5	5700	1901	60	29
Oakland'10 D4 CC	8	20994	6999	82	64
Oakland'10 Mayor	11	119607	39870	–	3925
Oakland'10 D6 CC	4	12911	4304	–	350
Pierce'08 CE	5	299132	99711	–	166
NSW'19 Legislative Assembly elections re-imagined as 2 seat STV					
Ballina	6	50127	16710	66	66
Bathurst	6	50833	16945	81	57
Clarence	6	49355	16452	147	84
Coffs Harbour	8	47333	15778	1225	515
Cootamundra	6	47448	15817	–	–
Heffron	5	50010	16671	1778	211
Holsworthy	6	48244	16082	20	20
Ku-ring-gai	6	48730	16244	202	99
Lake Macquarie	6	50082	16695	129	73
Lane Cove	6	50941	16981	132	109
Lismore	7	48145	16049	–	2180
Manly	6	48316	16106	363	150
Newcastle	8	50319	16774	214	173
North Shore	9	47774	15925	470	184
Northern Tablelands	4	48678	16227	–	143
Oxley	5	48540	16181	98	80
Pittwater	8	49119	16374	163	110
Summer Hill	6	48785	16262	–	110
Tamworth	6	50578	16860	–	129
Vaucluse	7	46023	15342	–	325
Wallsend	5	51351	17118	149	83
Willoughby	8	47857	15953	–	–
Wollondilly	8	50989	16997	120	88
Wollongong	7	51435	17146	205	123

al [1] demonstrate that forming full RLAs for this class of 2-seat STV elections is challenging, and generally not possible with existing methods.

A partial RLA can be used to verify *some* aspects of the election outcome. For example, that some reported losers did indeed lose, and that one of the reported winners did indeed win. In this paper, we reframe the general method into five stages. We still use the original assertion types, AG and NL, as we are not assuming that one or more candidates have been *previously* seated. Stages 1, 2, and 3 are present in the general method of Blom *et al* [1]. Stages 4 and 5 are introduced in this paper to (i) describe how we can form a partial RLA for a 2-seat STV election when a full RLA cannot be formed (Stage 5) and (ii) reduce the required sample size of the resulting partial or full RLA (Stage 4).

1. **Form AG Assertions** For each pair of candidates $c, c' \in \mathcal{C}$, we determine whether we can form the assertion $\text{AG}(c, c')$. We keep track of each AG that we can form, and its cost.
2. **Rule out candidates (find Definite Losers)** We use the AG assertions that we formed in Stage 1 to determine whether some candidates definitely lost the election. All candidates $c \in \mathcal{C}$ for which there exists *at least two* other candidates $c', c'' \in \mathcal{C} - \{c\}$ such that $\text{AG}(c', c)$ and $\text{AG}(c'', c)$ *definitely lost* the election. We denote this set of candidates DL , and the set of AG assertions required to show that these candidates definitely lost as \mathcal{A}_{DL} . This set will contain two AG assertions for each definite loser. The maximum sample size required to audit any assertion in this set is denoted the Stage 2 sample size.
3. **Rule out alternate winner pairs** We consider all pairs of candidates from the set $\mathcal{C} - DL$, excluding the pair of reported winners, as potential alternate winner outcomes. We follow the approach of Blom *et al* [1], and attempt to rule out each of these alternate winner pairs with an NL assertion. For a pair (c_1, c_2) , we first assume that c_1 is seated *at some point*, and look for another candidate c' that never loses to c_2 in this context:

$$\text{NL}(c', c_2, [c_1], G, O)$$

where G is the set of candidates g for which $\text{AG}(g, c_2)$ and O the candidates o for which $\text{AG}(c', o)$. We compare the cost of this NL assertion with one formed when we assume that c_2 is seated at some point, and look for a c' that never loses to c_1 :

$$\text{NL}(c', c_1, [c_2], G, O)$$

where G is the set of candidates g for which $\text{AG}(g, c_1)$ and O the candidates o for which $\text{AG}(c', o)$. The cheapest NL, assuming we are able to form at least one, is used to rule out the outcome of c_1 and c_2 winning together. We denote the set of assertions used to rule out candidate pairs in this stage, \mathcal{A}_3 . This set includes the formed NL assertions and any AG assertion used to reduce the margin of those NL assertions. The maximum sample size required to audit an assertion in this set the Stage 3 sample size.

4. **Reduce audit sample size** Ruling out a candidate c in Stage 2, by looking for two other candidates who tallies are always greater than c , may be unnecessarily costly. We may have been able to rule out all alternate winner pairs involving c with cheaper NL assertions in Stage 3.

While the Stage 2 sample size is higher than that of Stage 3, we take the current ‘most difficult to rule out’ candidate in DL , d . Let ASN_d^2 denote the sample size required to rule out d as a potential winner in Stage 2. We form a set of alternate winner pairs by pairing d with all candidates in $\mathcal{C} - DL$. We perform the Stage 3 process over this new pair set. If the sample size required to rule out these pairs, ASN_d^3 , is less than ASN_d^2 , we:

- (a) Remove the assertions formed in Stage 2 to rule out d from \mathcal{A}_{DL} ;
- (b) Add the new assertions formed to rule out all alternate winner pairs involving d to \mathcal{A}_3 ;
- (c) Update the Stage 2 sample size, excluding the cost of ruling out d ;
- (d) Update the Stage 3 sample size to include ASN_d^3 ; and,
- (e) Remove d from DL .

If $ASN_d^3 \geq ASN_d^2$, or we could not rule out the new set of alternate winner pairs, we do not change our audit and move to Stage 5. Otherwise, we take the next most difficult to rule out candidate in DL , and repeat Stage 4.

5. **Summarise what can (and cannot) be audited** If we have been able to rule out each alternate winner pair with an NL, we have a full RLA. This RLA contains the assertions in \mathcal{A}_{DL} and \mathcal{A}_3 . If we were not able to rule out every alternate winner pair, we consider whether the ones we could rule out imply that some additional candidates definitely lost or definitely won. Let Rem denote the set of alternate winner pairs that we could not rule out.

- **Definite Winners** If there is a candidate c present in *every* remaining pair in Rem , we add this candidate to a set DW .
- **Definite Losers** Includes all candidates in DL (Stage 2) and any $c \in \mathcal{C} - DL$ (excluding the reported winners) that is not present in any of the alternate winner pairs in Rem . Each such c is added to DL .
- **Potential Winners** All candidates in the set $\mathcal{C} - DL$ are potential winners. These are reported losers and winners whose elimination or election we could not verify.

Our partial RLA contains the assertions in the set $\mathcal{A}_{DL} \cup \mathcal{A}_3$ and can be used to establish that the candidates in the set DL definitely lost, and that the candidates in the set DW definitely won.

Example 4. Consider the 2022 Australian Senate election for ACT, a 2-seat STV election that does not satisfy the first-round winner criterion. The quota for this election was 95073. None of the 23 candidates have a quota on their first preferences. GALLAGHER won her seat after the elimination of 11 candidates. After a further 9 candidates were eliminated, POCOCK won the second seat. In Stage 1, we can form 44 AG assertions with ASNs ranging from 14 to 1141. In

Stage 2, we use some of these AGs to mark 16 candidates as *definite losers*. The maximum ASN of the assertions we use in this stage is 775. In Stage 3, there are 20 alternate winner pairs that we can form with the remaining 7 candidates. We are able to rule out all but three of these with NL assertions, requiring a sample size of 420 ballots. The most expensive candidate to rule out as a winner in Stage 2 requires a sample of 775 ballots. In Stage 4, we take this ‘expensive to rule out candidate’, d , and rule out all alternate winner pairs involving d with NL assertions instead. These NL assertions have a maximum ASN of 47 ballots. Our overall audit cost reduces to 420 ballots. Our Stage 5 summary indicates that we can show that GALLAGHER correctly won, and that four of the remaining 22 candidates, including POCOCK, are potential winners. \square

We have collected data for 587 3-4 seat STV elections taking place in 2017 and 2022 to elect local councillors in Scotland. These elections involve 3 to 13 candidates. When viewed as 2-seat STV elections, 428 of these instances involve an elimination in the first round. A full RLA can be formed for 68 of these 428 instances. For this set of 68 instances, performing Stage 4 reduces required sample sizes by 58% on average for 20 elections of the 68 (reductions range from 5% to 97%), and makes no difference on required sample size for the remaining 48. For the 360/428 instances for which a full RLA could not be formed, performing Stage 4 reduces the required sample size of the partial RLA in 122/360 of these instances (by 49% on average, from 2% to 99%).

6 Concluding Remarks

Auditing STV elections is a challenging problem, but one of very real interest given the common use of STV throughout the world. The main challenge arises as ballots can change their value across tabulation. In this paper we have shown how reasoning about both lower and upper bounds on transfer values may improve our ability to audit 2-seat STV elections. The revisions substantially reduce the number of ballots expected to be required to audit an election, and in some cases makes it possible to audit an election that the previous method [1] could not. We also show how to effectively audit batch elimination, as well as partially audit elections where no candidate gets a quota initially. While significant advances are still required to get to the point of auditing large Australian Senate elections, STV elections with 6 seats and over 100 candidates and say 4 million ballots, the new techniques we develop here help us on the path to this goal.

References

1. Blom, M., Stuckey, P.J., Teague, V., Vukcevic, D.: A first approach to risk-limiting audits for single transferable vote elections. In: Workshop on Advances in Secure Elections VOTING’22 (2022)
2. Blom, M. *et al*: Assertion-based approaches to auditing complex elections, with application to party-list proportional elections. In: Electronic Voting. E-Vote-ID 2021. LNCS, vol. 12900, pp. 47–62. Springer (2021)
3. Stark, P.B.: Sets of half-average nulls generate risk-limiting audits: SHANGRLA. In: Financial Cryptography 2020. LNCS, vol. 12063, pp. 319–336 (2020)