

Election Manipulation 100

Michelle Blom¹, Peter J. Stuckey², and Vanessa J. Teague¹

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

`michelle.blom,vjteague@unimelb.edu.au`

² Faculty of Information Technology,
Monash University
Clayton, Australia

`peter.stuckey@monash.edu.au`

Abstract. The true election margin for an Instant Runoff Voting (IRV) election can be hard to compute, because a small modification early in the elimination sequence can alter the outcome and result in a candidate winning the last round by a large margin. It is often assumed that the true margin is the last-round margin, that is half the difference between the two candidates who remain when everyone else is eliminated, though it is well known that this need not be the case. Perceptions of confidence in the outcome, and even formal policies about recounts, often depend on the last-round margin. There is already some prior work on how to compute the true election margin efficiently for IRV, and hence how to find the minimal manipulation. In this work we show how to manipulate an election efficiently *while also producing a large last-round margin*. This would allow a successful manipulation to evade detection against naive methods of assessing the confidence of the election result. This serves as further evidence for accurate computations of the exact margin, or for rigorous Risk Limiting Audits which would detect a close or wrong election result (respectively) regardless of the last-round margin.

1 Introduction

Instant Runoff Voting (IRV), also known as Alternative Vote (AV), is a system of preferential voting in which voters rank candidates in order of preference. Given candidates a , b , and c , each vote cast in an IRV election is a (possibly partial) ranking over the candidates. A vote with the ranking $[a, c, b]$ expresses a first preference for candidate a , a second for c , and a third for b . The tallying of votes proceeds by distributing each vote to its first ranked candidate. The candidate with the smallest number of votes is eliminated, with their votes redistributed to subsequent, less preferred candidates. Elimination proceeds in this fashion, until a single candidate w remains, who is declared the winner. IRV is used for all lower house parliamentary elections across Australia, parliamentary elections in Fiji and Papua New Guinea, presidential elections in Ireland and Bosnia/Herzegovina, and local elections in numerous locations world-wide, including the UK and United States [9].

The last round margin (LRM) of an IRV election – the difference in tallies of the final two remaining candidates, divided by two and rounded up – is commonly used as an indicator of how close the election was. Blom *et al.* [4] have shown that the true margin of victory (MOV) of the election – the smallest number of votes one would have to alter to change who won the election – is generally equal to the last round margin, but not always. In some cases, the MOV can be much smaller than the last round margin. The Australian Electoral Commission (AEC) use the “margin between the two leading candidates” after all remaining candidates have been eliminated, and their preferences distributed, to determine whether an automatic recount of cast votes should be performed.³ The AEC definition of a “margin”, in this context, is the difference in tallies of two candidates (not divided by two). When this margin is less than 100 votes, an automatic recount is triggered. In traditional paper-based elections, where counting proceeds by hand and scrutineers are present to oversee the counting of ballots, these margins play a major role in determining whether further scrutiny of the outcome is warranted.

In this paper, we put ourselves in the shoes of a potential adversary seeking to change the outcome of an IRV election – an election in which voters have cast paper ballots, and these ballots have been consequently scanned and counted using software. We assume that this adversary has sufficient access to the systems used to execute the IRV counting algorithm, complete knowledge of the election profile (the rankings present on each vote), and the ability to change each ballot’s electronic record. The adversary wants to alter the smallest number of these electronic records so that their desire of changing the outcome is realised, while at the same time ensuring that the last round margin of the manipulated election is larger than a given threshold. Our adversary does not want to change too many votes, as the more votes that are modified, the greater the likelihood that the manipulation will be discovered. To realise this adversary, we adapt the margin computation algorithm of Blom *et al.* [4] to compute the smallest number of votes that it must change to both alter an election outcome, and create a manipulated election with desirable properties (such as a large last round margin). Note that throughout this paper we use the terms vote and ballot interchangeably – each ballot is equivalent to a single vote in the context of IRV.

Using the Australian New South Wales (NSW) 2015 Legislative Assembly election as a case study, we report the number of votes that this adversary would have needed to modify in each seat to change the candidate who won, while controlling the margin by which they won. Here, as a running example, we consider the smallest manipulation in order to achieve a last round margin of victory of at least 100 votes, in order to prevent an automatic recount, *election manipulation 100*.

It is obvious that it is always possible to achieve a last-round margin of at least x by finding the Margin of Victory MOV , making those minimal changes, and then altering an extra x ballots in favour of the desired candidate. The interesting cases are those in which a last-round margin of x can be achieved by manipulating fewer than $MOV + x$ votes. We find that this is possible in a small

³ <https://www.aec.gov.au/Elections/candidates/files/hor-recount-policy.pdf>

number of examples—these seats would make natural targets for manipulation. We find empirically that in natural elections it’s very often the case that $\text{LRM} = \text{MOV}$. After manipulation, however, we find that the MOV of the manipulated election is often much smaller than its LRM. The effect of this declines when the manipulation is small rather than large, but could be considered a genuine indication that someone is manipulating the results.

The assumption of complete knowledge of the election profile, and the ability to change any vote, is a strong one. We consider, in our concluding discussion, how we can analyse the likelihood of election-changing manipulations in a context where our adversary *does not* have knowledge of the complete election profile. In this setting, the adversary may have seen a portion of cast votes, after scanning, and is able to modify the rankings of future scanned votes as they are scanned and their electronic record created. How likely is it that such an adversary can choose appropriate manipulations, in this context, and achieve the election of a desired candidate with an appropriate margin? We consider how we might design a series of experiments to answer this question.

The remainder of this paper is structured as follows. We discuss related work on margin computation for IRV, manipulation, and auditing in Section 2. Preliminaries and definitions are provided in Section 3. Section 4 summarises the margin computation algorithm of Blom *et al.* [4] and how it can be adapted to add side constraints – such as ensuring a last round margin of at least a given threshold – on the nature of any acceptable manipulation of an election. We demonstrate this adapted algorithm on case study – the 2015 NSW Legislative Assembly election – in Section 5. We conclude in Sections 6 and 7 with a discussion of how to model and analyse a weaker adversary, without complete knowledge of the rankings on every ballot.

2 Related Work

Blom *et al.* [4] present a branch-and-bound algorithm (denoted *margin-irv*) for efficiently computing the margin of victory in an IRV election, improving upon an existing method by Magrino *et al.* [8]. Blom *et al.* [3] extend this work to compute the margin of victory over candidates (MOVC) for an IRV election. That work computes the smallest number of votes that must be changed in order to change the winner of the election to one of a given subset of candidates. In this paper, we extend the *margin-irv* algorithm of Blom *et al.* [4] to compute the smallest number of vote changes required to yield an election with desired properties, such as a last round margin of at least a certain size. We then demonstrate this extended algorithm on the New South Wales 2015 Legislative Assembly Parliamentary Election.

Since the MOV is the minimum number of vote changes necessary to successfully manipulate the election result, the election result can be shown to be correct if there are fewer than MOV manipulations.

A number of methods have been developed for auditing various kinds of elections [1], and for first past the post (FPTP) elections in particular. Risk Limiting

Initially, all candidates remain standing (are not eliminated)
While there is *more than one* candidate standing
 For every candidate c standing
 Tally (count) the votes in which c is the highest-ranked
 candidate of those standing
 Eliminate the candidate with the smallest tally
The winner is the one candidate not eliminated

Fig. 1: The IRV counting algorithm: the candidate with the smallest tally is repeatedly eliminated, with the ballots in their tally redistributed to remaining candidates according to their next preference.

Audits (RLAs) [7, 10] have been applied to a number of such elections, including four 2008 elections in California [6] and elections in over 50 Colorado counties in 2017. RLAs provide strong statistical evidence that the reported outcome of an election is correct, or revert to a manual recount if it is wrong. The probability that the audit fails to detect a wrong outcome is bounded by a *risk limit*. Lindeman *et al.* [7] present a ballot-polling RLA for FPTP elections, which has consequently been adapted by Blom *et al.* [2] for IRV. Several approaches for designing a risk-limiting comparison audit of an IRV election have also been proposed [10]. A genuine RLA would defeat the attack described in this paper, because it would detect a wrong election result with high probability. Our proposed variety of manipulation works only against naive recount triggers based on last-round margins.

3 Preliminaries

Votes are tallied in an IRV election in a series of rounds (see Figure 1). In each round, the candidate with the smallest number of votes (their tally) is eliminated, with the last remaining candidate declared the winner of the election. All votes in an eliminated candidate’s tally are distributed to the next most-preferred (remaining) candidate in their ranking.

Let \mathcal{C} be the set of candidates in an IRV election \mathcal{B} . We refer to sequences of candidates π in list notation (e.g., $\pi = [c_1, c_2, c_3, c_4]$), and use such sequences to represent both votes and the order in which candidates are eliminated. An election \mathcal{B} is defined as a multiset⁴ of votes, each vote $b \in \mathcal{B}$ a sequence of candidates in \mathcal{C} , with no duplicates, listed in order of preference (most preferred to least preferred). Let $first(\pi)$ denote the first candidate appearing in sequence π (e.g., $first([c_2, c_3]) = c_2$). In each round of vote counting, there are a current set of eliminated candidates \mathcal{E} and a current set of candidates still standing $\mathcal{S} = \mathcal{C} \setminus \mathcal{E}$. The winner c_w of the election is the last standing candidate.

⁴ A multiset allows for the inclusion of duplicate items.

Ranking	Count
[a]	55
[c, a]	30
[b, c]	36
[c]	15

(a)

Candidate	Round 1	Round 2
a	55	55
b	36	—
c	45	81

(b)

Ranking	Count
[a]	55
[c, a]	25
[b, c]	41
[c]	15

(c)

Candidate	Round 1	Round 2
a	55	80
b	41	41
c	40	—

(d)

Table 1: IRV example, with (a) the number of votes cast with each listed ranking over candidates a, b, c , and (b) tallies after each round of vote counting (c) the number of votes recorded after manipulation, and (d) the tallies after each round of vote counting in the manipulated election

Each candidate $c \in \mathcal{C}$ has a *tally* of votes. Votes are added to this tally upon the elimination of a candidate $c' \in \mathcal{C} \setminus \{c\}$, and are redistributed from this tally upon the elimination of c .

Definition 1. Tally $t_{\mathcal{S}}(c)$ Given candidates $\mathcal{S} \subseteq \mathcal{C}$ are still standing in an election \mathcal{B} , the tally for candidate $c \in \mathcal{C}$, denoted $t_{\mathcal{S}}(c)$, is defined as the number of votes $b \in \mathcal{B}$ for which c is the most-preferred candidate of those remaining. Let $p_{\mathcal{S}}(b)$ denote the sequence of candidates mentioned in b that are also in \mathcal{S} .

$$t_{\mathcal{S}}(c) = |\{b \mid b \in \mathcal{B}, c = \text{first}(p_{\mathcal{S}}(b))\}| \quad (1)$$

Definition 2. Margin of Victory (MOV) The MOV in an election with candidates \mathcal{C} and winner $c_w \in \mathcal{C}$, is the smallest number of votes whose ranking must be modified (by an adversary) so that a candidate $c' \in \mathcal{C} \setminus \{c_w\}$ is elected.

Definition 3. Last Round Margin (LRM) The LRM of an election, in which two candidates $\mathcal{S} = \{c, c'\}$ remain with $t_{\mathcal{S}}(c)$ and $t_{\mathcal{S}}(c')$ votes in their tallies, is equal to half the difference between the tallies of c and c' rounded up.

$$LRM = \left\lceil \frac{|t_{\mathcal{S}}(c) - t_{\mathcal{S}}(c')|}{2} \right\rceil \quad (2)$$

Example 1. Consider the example election shown in Table 1 between candidates a, b and c . Their initial tallies are 55, 36, and 45 votes, respectively, and b

is eliminated first. Candidates a and c subsequently have tallies of 55 and 81 votes, giving c the victory with a last round margin of 13 votes. A seemingly comfortable victory.

But let's examine what occurs if we change 5 of the $[c, a]$ votes to $[b, c]$ votes. Now the initial tallies are 55, 41, and 40 votes, respectively, and c is eliminated first. Candidates a and b subsequently have tallies of 80 and 41 votes, giving a the victory with a last round margin of 20 votes.

Note that the apparent comfortable victory of c originally is an illusion, the actual MOV for this election is 5, as the manipulation illustrates. Interestingly even though we only manipulate 5 votes, now a wins the election with a last round margin of 20 votes! The actual margin of victory in the manipulated election is 1, demonstrated by the fact that if we change 1 of the $[b, c]$ votes to a $[c]$ vote, the first round tallies of each candidate are $\{a : 55, b : 40, c : 41\}$, and c is eliminated.

4 Computing the MOV for an IRV election

A description of both the *margin-irv* algorithm, and the original branch-and-bound method of Magrino *et al.* [8], can be found in Blom *et al.* [4, 3]. We summarise this algorithm in this section, and describe how it can be modified to compute the smallest number of vote changes required to both (i) bring about a change in the outcome of the election, and (ii) produce a manipulated election profile with certain properties, modelled as side constraints. We consider the following two side constraints in this paper:

- The LRM of the manipulated election must be at least T_{LRM} votes;
- The eliminated candidate e in each round must have Δ fewer votes in their tally than the candidate with the next smallest tally.

Consider an IRV election \mathcal{B} with candidates \mathcal{C} and winner $w \in \mathcal{C}$. The *margin-irv* algorithm starts by adding $|\mathcal{C}| - 1$ partial elimination sequences to a search tree, one for each of alternate winner $c \in \mathcal{C} \setminus \{w\}$. These partial sequences form a frontier F , with each sequence containing a single candidate – an alternate winner. Note that a partial sequence $[a, b, c]$ represents an election outcome in which a and b are the last two candidates eliminated, and c the winner. All other candidates are assumed to have been eliminated in some prior round.

For each partial sequence $\pi \in F$, we compute a lower bound on the number of vote changes required to realise an elimination sequence that *ends* in π . These lower bounds are used to guide construction of the search tree, and are computed by both solving an Integer Linear Program (ILP), and applying several rules for lower bound computation. These rules are described in Blom *et al.* [4]. The ILP, denoted DISTANCETO, computes a lower bound on the smallest number of vote changes required to transform the election \mathcal{B} , with an elimination sequence π' , to one with an elimination sequence that ends in π . When applied to a complete order π , containing all candidates, DISTANCETO exactly computes the smallest number of votes changes required to realise the outcome π . The largest of the

lower bounds computed by the rules of Blom *et al.* [4] and the DISTANCETO ILP is assigned to each partial sequence π as it is added to F . The DISTANCETO ILP is defined in Section 4.1. To enforce additional constraints on the nature of any manipulated election, we add these constraints to each ILP solved.

The partial sequence $\pi \in F$ with the smallest assigned lower bound is selected and *expanded*. For each candidate $c \in \mathcal{C}$ that is not already present in π , we create a new sequence with c appended to the front. For example, given a set of candidates e , f , and g , with winning candidate g , the partial sequence $\pi = [f]$ will be expanded to create two new sequences $[e, f]$ and $[g, f]$. We evaluate each new sequence π' created by assigning it a lower bound on the number of votes required to realise any elimination order ending in π' .

While exploring and building elimination sequences, *margin-irv* maintains a running *upper bound* on the value of the true margin. Without any side constraints designed to inject desirable properties into a manipulated election, this upper bound is initialised to the last round margin of the original election. To enforce additional constraints on the properties of any manipulated election, we need to manipulate at least as many, and often more, votes than required to simply change the original outcome. Consequently, we must set the upper bound maintained by *margin-irv* to a higher value. In this context, we set the initial upper bound to the total number of votes cast in the election. This is clearly always a correct upper bound on any manipulation.

When a sequence π containing all candidates is constructed, the DISTANCETO ILP computes the exact number of vote manipulations required to realise it, while satisfying all desired side constraints. If this number is lower than our current upper bound, the upper bound is revised, and all orders in F with a lower bound greater than or equal to it are pruned from consideration (removed from F). This process continues until F is empty (we have considered or pruned all possible alternate elimination sequences). The value of the running upper bound is the true margin of victory (with side constraints) of the election.

4.1 DistanceTo with Side Constraints

We now present the DISTANCETO Integer Linear Program (ILP) used to compute lower bounds on the degree of manipulation required to realise an election outcome ending in a given candidate sequence, and the (exact) smallest number of vote changes required to realise a given (complete) alternate elimination sequence. This ILP, without added side constraints, was originally presented by Magrino *et al.* [8].

Let \mathbf{R} denote the set of possible (partial and total) rankings R of candidates \mathcal{C} that could appear on a vote, N_R the number of votes cast with ranking $R \in \mathbf{R}$, and N the total number of votes cast. Let $\mathcal{R}_{j,i}$ denote the subset of rankings in \mathbf{R} ($\mathcal{R}_{j,i} \subset \mathbf{R}$) in which c_j is the most preferred candidate still standing (i.e., that will count toward c_j 's tally) at the start of round i (in which candidate c_i is eliminated). For each $R \in \mathbf{R}$, we define variables:

q_R integer number of votes to be changed into R ;

m_R integer number of votes with ranking R in the unmodified election to be changed into something other than R ; and
 y_R number of votes in the modified election with ranking R .

Given a partial or complete order π , the DISTANCETO ILP is:

$$\min \sum_{R \in \mathbf{R}} q_R$$

$$N_R + q_R - m_R = y_R \quad \forall R \in \mathbf{R} \quad (3)$$

$$\sum_{R \in \mathbf{R}} q_R = \sum_{R \in \mathbf{R}} m_R \quad (4)$$

$$\sum_{R \in \mathcal{R}_{i,i}} y_R \leq \sum_{R \in \mathcal{R}_{j,i}} y_R \quad \forall c_i, c_j \in \pi . i < j \quad (5)$$

$$n \geq y_R \geq 0, \quad N_R \geq m_R \geq 0, \quad q_R \geq 0 \quad \forall R \in \mathbf{R} \quad (6)$$

Constraint (3) states that the number of votes with ranking $R \in \mathbf{R}$ in the new election is equal to the sum of those with this ranking in the unmodified election and those whose ranking has *changed to* R , minus the number of votes whose ranking has been *changed from* R . Constraint (5) defines a set of *special elimination constraints* which force the candidates in π to be eliminated in the stated order. Constraint (4) ensures that the total number of votes cast in the election does not change as a result of the manipulation.

The above ILP does not include any additional side constraints – properties that we want the manipulated election to satisfy besides resulting in a different winner to that of the original election. We show, in Section 5, that manipulated elections found by *margin-irv* in this setting are almost always evidently close, with a last round margin of 0 or 1 vote. This makes sense as the algorithm is trying to manipulate as few votes as possible, breaking any ties in favour of an alternate outcome. An adversary with the ability to modify electronic records of cast votes, however, will want to create a manipulated election that is not evidently close. An election with a tie in the final round of counting, or a difference of several votes in the tallies of the final two remaining candidates, is likely to be closely scrutinised. Australian IRV elections with a last round margin of less than 100 votes, for example, trigger an automatic recount.

Given the widespread use of the last round margin as the indicator of how close an IRV election is, rather than the true MOV of the election, our adversary can use this to their advantage. Consider a candidate elimination sequence π , containing at least two candidates from a set \mathcal{C} . Let the last two candidates in the sequence π be denoted by c_k and c_{k+1} , with $|\mathcal{C}| = k + 1$. Adding the following side constraint to DISTANCETO ensures that the last round margin of any manipulated election is greater than or equal to T_{LRM} votes.

$$\sum_{R \in \mathcal{R}_{k,k}} y_R \leq \sum_{R \in \mathcal{R}_{k+1,k}} y_R + 2T_{LRM} \quad (7)$$

We can add any number of desired side constraints to this ILP to inject desirable properties into any manipulated election. In this paper we consider two side constraints: requiring the last round margin of the manipulated election to be equal to or greater than a given threshold T_{LRM} ; and ensuring no ties arise in the manipulated election when determining which candidate to eliminate in each round. The latter constraint can be modelled by requiring the tally of the eliminated candidate e in each round i to contain Δ fewer votes than that of the candidate with the next smallest tally in round i .

$$\Delta + \sum_{R \in \mathcal{R}_{i,i}} y_R \leq \sum_{R \in \mathcal{R}_{j,i}} y_R \quad \forall c_i, c_j \in \pi . i < j \quad (8)$$

Constraint (8) modifies the set of special elimination constraints (Constraint 5) with the addition of the Δ constant on the left hand side.

4.2 Selecting a Desired Winner

An adversary is likely to have a goal of electing a specific candidate, or one of a set of specific candidates, in place of the original winner. Blom *et al.* [3] show that we can compute the smallest number of vote changes necessary to elect a *specific alternate* winner – a candidate from a given set \mathcal{C}' – by adjusting the way we construct our initial frontier F in the branch-and-bound algorithm described above. Consider an election with candidates \mathcal{C} and winner $w \in \mathcal{C}$. If we are interested in simply changing the candidate who wins to any candidate that is not w , we add $|\mathcal{C}| - 1$ partial sequences to our frontier, one for each alternate winner. As described above, each of these sequences contains just one candidate – the alternate winner in question. In the setting where we want to elect a candidate from the set \mathcal{C}' , we create, and add to our frontier, a partial candidate sequence for each of the candidates in \mathcal{C}' . The remainder of the algorithm remains unchanged. The use of this restricted frontier, in conjunction with a DISTANCETO ILP containing side constraints, allows us to compute a minimal manipulation of votes required to elect a specific candidate with, for example, a large last round margin. In the case study below, we consider an adversary that simply wants to change the election winner to any alternate candidate.

5 Case Study: The NSW 2015 State Election

In the 2015 NSW State Election, 4.56 million votes were cast across 93 IRV elections, one in each of 93 different electorates. Table 2 considers these 93 elections, recording the number of votes cast, the last round margin, the true margin of victory, and the last round margin of the manipulated election found by *margin-irv* without the addition of side constraints. In all but five elections (Ballina, Heffron, Lismore, Maitland, and Willoughby) the MOV is the LRM, showing that they are almost always equal. In all but one election (Ballina), minimal manipulation results in an evidently close election. Ballina shows that, while uncommon, an adversary can achieve a large last round margin without performing any more manipulation than necessary to alter the winner of the seat.

Table 2: LRM, MOV, and LRM of the manipulated election (denoted LRM*), for each seat in the 2015 NSW lower house election (no added side constraints).

Seat	C	B	MOV	LRM	LRM*	Seat	C	B	MOV	LRM	LRM*
Albury	5	46335	5840	5840	0	M-Fields	7	47183	3519	3519	0
Auburn	6	43781	2265	2265	0	Maitland	6	47826	4012	5446	0
Ballina	7	47454	1130	1267	1248	Manly	5	47287	10806	10806	0
Balmain	7	46952	1731	1731	0	Maroubra	5	46492	4717	4717	1
Bankstown	6	42899	5542	5542	0	Miranda	6	49454	5881	5881	0
Barwon	6	47707	5229	5229	0	Monaro	5	46202	1122	1122	0
Bathurst	5	48632	7267	7267	1	M-Druitt	5	44948	6343	6343	0
B-Hills	5	49266	10023	10023	0	Mulgoa	5	48257	4336	4336	0
Bega	5	47658	3663	3663	1	Murray	8	46387	8574	8574	0
Blacktown	5	46262	5565	5565	0	M-Lakes	6	48252	3627	3627	0
B-Mntns.	6	47608	3614	3614	1	Newcastle	7	48136	3132	3132	0
Cabramatta	5	47691	7613	7613	0	Newtown	7	45392	3536	3536	0
Camden	5	48152	8217	8217	0	N-Shore	7	46247	8517	8517	0
C-belltown	5	45124	3096	3096	0	N-lands	6	48340	11969	11969	0
Canterbury	5	47631	6610	6610	0	Oatley	5	48119	3006	3006	0
Castle Hill	5	48092	13160	13160	0	Orange	5	48784	10048	10048	0
Cessnock	5	45822	9187	9187	0	Oxley	5	46514	4591	4591	0
Charlestown	7	48919	5532	5532	0	Parramatta	7	47447	5509	5509	0
Clarence	8	47181	4069	4069	0	Penrith	8	47577	2576	2576	0
C-Harbour	5	45162	5824	5824	1	Pittwater	5	48345	11430	11430	1
Coogee	5	46322	1243	1243	0	P-M.quarie	5	49231	8715	8715	0
C-mundra	5	47160	9247	9247	0	P-Stephens	5	47037	2088	2088	0
Cronulla	5	50333	9674	9674	0	Prospect	5	47195	1458	1458	0
Davidson	5	49147	12960	12960	0	Riverstone	5	46945	5324	5324	0
Drummoyne	6	46818	8099	8099	0	Rockdale	6	46240	2004	2004	0
Dubbo	7	46582	8680	8680	0	Ryde	5	48286	5153	5153	0
East Hills	5	47449	189	189	0	S-Hills	7	47874	3774	3774	0
Epping	6	49532	7156	7156	0	S-harbour	7	50995	7519	7519	0
Fairfield	5	45921	6998	6998	0	S-Coast	5	45788	4054	4054	1
Gosford	6	48259	102	102	0	Strathfield	5	46559	770	770	0
Goulburn	6	48663	2945	2945	0	S-Hill	7	47073	3854	3854	0
Granville	6	45212	837	837	0	Swansea	8	48200	4974	4974	0
Hawkesbury	8	46856	7311	7311	1	Sydney	8	42747	2864	2864	1
Heathcote	6	51128	3560	3560	0	Tamworth	7	49004	4643	4643	0
Heffron	5	46367	5824	5835	0	Terrigal	5	48871	4053	4053	0
Holsworthy	6	47126	2902	2902	1	T-Entrance	5	47953	171	171	0
Hornsby	6	49834	8577	8577	1	Tweed	5	44185	1291	1291	0
Keira	5	50599	8164	8164	0	U-Hunter	6	47296	866	866	0
Kiama	5	47686	3856	3856	0	Vaocluse	5	46145	9783	9783	0

Continued

Seat	$ C $	$ \mathcal{B} $	MOV	LRM	LRM*	Seat	$ C $	$ \mathcal{B} $	MOV	LRM	LRM*
Kogarah	6	46421	2782	2782	0	W-Wagga	6	46610	5475	5475	0
Ku-ring-gai	5	48436	10061	10061	0	Wakehurst	6	47894	10770	10770	0
L-M.quarie	7	47698	4253	4253	0	Wallsend	5	49631	9418	9418	0
Lakemba	5	44728	8235	8235	0	Willoughby	6	47302	10160	10247	0
Lane Cove	6	48622	7740	7740	1	Wollondilly	6	47182	7401	7401	1
Lismore	6	47046	209	1173	1	Wollongong	7	49702	3367	3367	0
Liverpool	5	45291	8495	8495	1	Wyong	7	46070	3720	3720	0
L-derry	5	45928	3736	3736	0						

We now add our side constraints (7,8) to each DISTANCETO ILP solved by *margin-irv*, with T_{LRM} chosen such that the difference between the tallies of the last two remaining candidates, in each election, is at least 100 votes ($T_{LRM} = 50$), and $\Delta = 1$. This would avoid an automatic recount, and ties when determining which candidate to eliminate. Table 3 reports, for all 93 seats, the minimal manipulation MAN (i.e. number of ballots changed) required to change the winner of each election while ensuring these constraints hold, the original last round margin of the election (LRM), the last round margin of the manipulated election (LRM*), and the margin of victory of the manipulated election (MOV*). MOV* represents the smallest number of vote changes required to change the winner of the manipulated election (to any alternate winner, not necessarily back to its original winner).

In many cases, we can create a manipulated election where the LRM* is not only at least 50 votes (leading to a difference in the tallies of the last two remaining candidates of 100 votes), but has a MOV equal to it. This is reflective of most IRV elections – the MOV is generally equal to the LRM. In others, the manipulated elections are much closer than the LRM suggests.

To ensure a LRM* of at least T_{LRM} votes, we often have to manipulate a further T_{LRM} votes on top of those we must change to simply change the winning candidate. For $T_{LRM} = 50$, this is the case for all of the 93 seats with the exception of Ballina and Lismore – we can find a minimal manipulation, just enough to change the winning candidate, that also has an LRM of at least 50.

Table 3: Minimal MANipulation compared to LRM of the original election and MOV and LRM of the manipulated (*) election, for each seat in the 2015 NSW lower house election (side constraints requiring LRM* to be at least 50 votes, and tie breaking with $\Delta = 1$, added).

Seat	MAN	LRM	LRM*	MOV*	Seat	MAN	LRM	LRM*	MOV*
Albury	5890	5840	50	50	M-Fields	3569	3519	50	50
Auburn	2315	2265	50	50	Maitland	4062	5446	50	1
Ballina	1130	1267	1229	1	Manly	10856	10806	50	1
Balmain	1781	1731	50	50	Maroubra	4767	4717	51	51
Bankstown	5592	5542	50	50	Miranda	5931	5881	50	50

Continued

Seat	MAN	LRM	LRM*	MOV*	Seat	MAN	LRM	LRM*	MOV*
Barwon	5279	5229	50	50	Monaro	1172	1122	51	51
Bathurst	7317	7267	51	51	MDruitt	6393	6343	50	50
B-Hills	10073	10023	50	50	Mulgoa	4386	4336	51	51
Bega	3713	3663	51	51	Murray	8624	8574	50	1
Blacktown	5615	5565	50	50	M-Lakes	3677	3627	50	50
B-Mtns.	3664	3614	51	51	Newcastle	3182	3132	50	50
Cabramatta	7663	7613	50	50	Newtown	3586	3536	50	50
Camden	8267	8217	50	50	N-Shore	8567	8517	50	1
C-belltown	3146	3096	50	50	N-Tablelands	12019	11969	50	1
Canterbury	6660	6610	50	50	Oatley	3056	3006	51	51
Castle Hill	13210	13160	50	1	Orange	10098	10048	50	1
Cessnock	9237	9187	50	50	Oxley	4641	4591	50	50
Charlestown	5582	5532	50	51	Parramatta	5559	5509	50	50
Clarence	4119	4069	50	50	Penrith	2626	2576	50	50
C-Harbour	5874	5824	51	51	Pittwater	11480	11430	50	50
Coogee	1293	1243	50	50	P-Macquarie	8765	8715	50	50
C-mundra	9297	9247	50	50	P-Stephens	2138	2088	50	50
Cronulla	9724	9674	50	1	Prospect	1508	1458	50	50
Davidson	13010	12960	50	50	Riverstone	5374	5324	50	50
Drummoyne	8149	8099	50	50	Rockdale	2054	2004	50	50
Dubbo	8730	8680	50	50	Ryde	5203	5153	51	51
East Hills	239	189	50	50	S-Hills	3824	3774	50	50
Epping	7206	7156	50	1	S-harbour	7569	7519	50	50
Fairfield	7048	6998	50	50	S-Coast	4104	4054	51	51
Gosford	152	102	50	50	Strathfield	820	770	50	50
Goulburn	2995	2945	50	50	S-Hill	3904	3854	50	1
Granville	887	837	50	50	Swansea	5024	4974	50	52
Hawkesbury	7361	7311	50	50	Sydney	2914	2864	51	51
Heathcote	3610	3560	50	50	Tamworth	4693	4643	51	51
Heffron	5874	5835	50	1	Terrigal	4103	4053	50	50
Holsworthy	2952	2902	51	51	T-Entrance	221	171	51	51
Hornsby	8627	8577	50	1	Tweed	1341	1291	51	51
Keira	8214	8164	50	50	U-Hunter	916	866	50	50
Kiama	3906	3856	50	50	Vaocluse	9833	9783	50	1
Kogarah	2832	2782	50	50	W-Wagga	5525	5475	50	50
Ku-ring-gai	10111	10061	50	1	Wakehurst	10820	10770	50	50
L-M.quarie	4303	4253	50	50	Wallsend	9468	9418	50	50
Lakemba	8285	8235	50	1	Willoughby	10210	10247	51	2
Lane Cove	7790	7740	50	50	Wollondilly	7451	7401	50	51
Lismore	209	1173	50	1	Wollongong	3417	3367	50	50
Liverpool	8545	8495	50	1	Wyong	3770	3720	50	50
L-derry	3786	3736	50	50					

In some cases we can perform a small amount of additional manipulation, beyond that required to simply change the election outcome, and receive a much larger increase in the LRM*. Imagine that our adversary desired an even larger LRM* – say, 5% of the total votes cast. Table 4 reports the new number of votes changes (MAN) required to manipulate the 93 NSW elections to ensure that

both the LRM* is at least 5% of the total number of votes cast, and eliminated candidates do not appear in any ties (with $\Delta = 1$).

We have boldened the elections in which the apparent change in the election outcome is much greater than the degree of manipulation performed. In Heffron, for example, just changing the winner requires 5825 vote changes (12.6% of the total votes cast). The LRM for Heffron is 5835. If we only change 5825 votes, our manipulated election will have a LRM* of 1 vote. By changing a further 117 votes, we can create an election with both a different winner *and* a LRM* of 2319 votes. In Ballina, performing an additional 133 vote manipulations yields an increase in the LRM* of 1145 votes (from 1229 to 2374).

When performing just enough manipulation to ensure a LRM* of 50 votes (and a change in winner), the MOV* and LRM* of the manipulated elections substantially differ in 19 of the 93 elections. When performing substantially more manipulation, to ensure a LRM* of 5% of the total votes cast, the MOV* and LRM* of the manipulated elections substantially differ in 50 of the 93 elections.

Table 4: Minimal MANipulation and LRM the original election and LRM and MOV of the manipulated (*) election, for each seat in the 2015 NSW lower house election (side constraints requiring LRM* to be at least 5% of the total cast votes, and tie breaking with $\Delta = 1$, added).

Seat	MAN	LRM	LRM*	MOV*	Seat	MAN	LRM	LRM*	MOV*
Albury	8157	5840	2317	2317	M-Fields	5878	3519	2360	2360
Auburn	4454	2265	2190	2190	Maitland	6278	5446	2462	1
Ballina	1263	1267	2374	1	Manly	13171	10806	2365	1
Balmain	3075	1731	3196	1	Maroubra	7042	4717	2326	2326
Bankstown	7687	5542	2145	1989	Miranda	8354	5881	2473	1449
Barwon	7615	5229	2386	633	Monaro	3432	1122	2311	2311
Bathurst	9699	7267	2433	2235	M-Druitt	8591	6343	2248	1679
B-Hills	12487	10023	2464	2080	Mulgoa	6749	4336	2414	2414
Bega	6046	3663	2384	2384	Murray	10893	8574	2320	1
Blacktown	7879	5565	2314	2314	M-Lakes	6040	3627	2413	675
B-Mntns.	5947	3614	2381	1	Newcastle	5278	3132	2407	2
Cabramatta	9998	7613	2385	669	Newtown	5806	3536	2270	131
Camden	10625	8217	2408	2408	N-Shore	10830	8517	2313	1
C-belltown	5353	3096	2257	2257	N-Tablelands	14386	11969	2417	1
Canterbury	8992	6610	2382	1	Oatley	5412	3006	2407	2407
Castle Hill	15565	13160	2405	1	Orange	12487	10048	2440	1
Cessnock	11479	9187	2292	1401	Oxley	6917	4591	2326	2326
Charlestown	7978	5532	2446	1990	Parramatta	7881	5509	2373	1747
Clarence	6428	4069	2360	2360	Penrith	4955	2576	2379	2379
C-Harbour	8082	5824	2259	2259	Pittwater	13847	11430	2418	1353
Coogee	3560	1243	2317	1461	P-Macquarie	11177	8715	2462	2462
C-mundra	11605	9247	2358	9	P-Stephens	4440	2088	2352	2352

Continued

Seat	MAN	LRM	LRM*	MOV*	Seat	MAN	LRM	LRM*	MOV*
Cronulla	12191	9674	2517	1	Prospect	3818	1458	2360	2360
Davidson	15418	12960	2458	290	Riverstone	7672	5324	2348	2348
Drummoyne	10440	8099	2341	680	Rockdale	4316	2004	2312	2312
Dubbo	11009	8680	2330	1	Ryde	7567	5153	2415	2415
East Hills	2562	189	2373	2373	S-Hills	6168	3774	2394	2394
Epping	9633	7156	2477	1	S-harbour	10069	7519	2550	1
Fairfield	9295	6998	2297	1	S-Coast	6343	4054	2290	860
Gosford	2515	102	2413	2413	Strathfield	3098	770	2328	2328
Goulburn	5379	2945	2434	2434	S-Hill	5487	3854	2354	1
Granville	3098	837	2261	2261	Swansea	7384	4974	2410	2370
Hawkesbury	9654	7311	2343	2343	Sydney	5001	2864	2138	1934
Heathcote	6116	3560	2557	2557	Tamworth	7093	4643	2451	2451
Heffron	5942	5835	2319	1	Terrigal	6497	4053	2444	2444
Holsworthy	5258	2902	2357	2357	T-Entrance	2569	171	2399	2399
Hornsby	11069	8577	2492	1	Tweed	3500	1291	2210	2210
Keira	10694	8164	2530	970	U-Hunter	3231	866	2365	1748
Kiama	6241	3856	2385	2385	Vaucluse	12091	9783	2308	1
Kogarah	5104	2782	2322	2322	W-Wagga	7806	5475	2331	2331
Ku-ring-gai	12483	10061	2422	621	Wakehurst	13165	10770	2395	1329
L-M.quarie	6638	4253	2385	2385	Wallsend	11900	9418	2482	1939
Lakemba	10472	8235	2237	1	Willoughby	12525	10247	2366	1
Lane Cove	10171	7740	2432	42	Wollondilly	9760	7401	2360	1062
Lismore	2449	1173	2353	1	Wollongong	5852	3367	2486	330
Liverpool	10760	8495	2265	2265	Wyong	6024	3720	2304	2304
L-derry	6033	3736	2297	2297					

These results demonstrate that, in the presence of manipulation, the LRM of an election is generally not a good indicator of how close the election was or whether its result should be audited or not. Then again, neither is its MOV. A clever adversary with sufficient access to change electronic records of cast votes will be able to design a manipulation that results in both a sizable LRM and MOV. To ensure that both the LRM and MOV of an election is sufficiently large, however, requires more manipulation than just desiring a large LRM, or just desiring a change in winner.

6 Modelling a Weaker Adversary

A likely practical scenario for election manipulation is one in which the adversary has partial knowledge of the ballot profiles and the opportunity to manipulate (some of) the rest. This would be the case, for example, if a corrupt scanner were able to modify ballot images or interpretations without the paper record being subsequently audited.

There are various models for an adversary with the power to manipulate a restricted number of votes, which is particularly relevant in contexts in which a small manipulation can change the outcome [5].

An interesting question to address in this context is whether a manipulation computed for, say, the first half of the ballots, could then be simply doubled and applied successfully to the second half. Obviously this is not true in general, if

there is some systematic difference between earlier and later votes (for example, if later votes come from a geographically distinct area from the earlier ones). It is an interesting practical question to understand how to extrapolate successful manipulations from a subset of ballots to the whole election, given reasonable assumptions about the information contained in the initial sample. Of course, other data, such as from past elections, could also be available to an attacker.

7 Concluding Remarks

We show how to compute successful manipulations that are designed specifically to avoid triggering a recount based on last-round margin, an inaccurate but commonly used assessment of the closeness of an IRV election.

The attack shown in this paper would be detected (with high probability) by a genuine Risk Limiting Audit, or by a recount triggered from the properly-computed true Margin of Victory rather than the last-round margin.

References

1. T. Antonyan, S. Davtyan, S. Kentros, A. Kiayias, L. Michel, N. Nicolaou, A. Russell, and A. A. Shvartsman. State-wide elections, optical scan voting systems, and the pursuit of integrity. *IEEE Transactions on Information Forensics and Security*, 4(4):597–610, 2009.
2. M. Blom, P. J. Stuckey, and V. Teague. Ballot-Polling Risk Limiting Audits for IRV Elections. In *International Conference on Electronic Voting (EVOTE-ID)*, pages 17–34, 2018.
3. M. Blom, P. J. Stuckey, and V. Teague. Computing the Margin of Victory in Preferential Parliamentary Elections. In *International Conference on Electronic Voting (EVOTE-ID)*, pages 1–16, 2018.
4. M. Blom, P. J. Stuckey, V. Teague, and R. Tidhar. Efficient Computation of Exact IRV Margins. In *European Conference on AI (ECAI)*, pages 480–487, 2016.
5. Anthony Di Franco, Andrew Petro, Emmett Shear, and Vladimir Vladimirov. Small vote manipulations can swing elections. *Communications of the ACM*, 47(10):43–45, 2004.
6. J.L. Hall, L.W. Miratrix, P.B. Stark, M. Briones, E. Ginnold, F. Oakley, M. Peaden, G. Pellerin, T. Stanionis, and T. Webber. Implementing risk-limiting post-election audits in California. In *Proc. 2009 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE '09)*, Montreal, Canada, August 2009. USENIX.
7. 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.
8. T. R. Magrino, R. L. Rivest, E. Shen, and D. A. Wagner. Computing the margin of victory in IRV elections. In *USENIX Accurate Electronic Voting Technology Workshop*, USENIX Association Berkeley, CA, USA, 2011.
9. R. Richie. Instant Runoff Voting: What Mexico (and Others) Could Learn. *Election Law Journal*, 3:501–512, 2004.
10. A.D. Sarwate, S. Checkoway, and H. Shacham. Risk-limiting audits and the margin of victory in nonplurality elections. *Politics, and Policy*, 3(3):29–64, 2013.