

Estimating the Margin of Victory for Instant-Runoff Voting

David Cary

Abstract

A general definition is proposed for the margin of victory of an election contest. That definition is applied to Instant Runoff Voting (IRV)¹ and several estimates for the IRV margin of victory are described: two upper bounds and two lower bounds. Given round-by-round vote totals, the time complexity for calculating these bounds does not exceed $O(C^2 \log C)$, where C is the number of candidates. It is also shown that calculating the larger and more useful of the two lower bounds can be viewed, in part, as solving a longest path problem on a weighted, directed, acyclic graph.

Worst-case analysis shows that neither these estimates, nor any estimates based only on tabulation round-by-round vote totals, are guaranteed to be within a constant factor of the margin of victory. These estimates are calculated for IRV elections in Australia and California. Pseudo code for calculating these estimates is provided.

1 Introduction

New procedures are being developed for post-election audits involving the manual tally of random samples of ballots. Such audit procedures often target plurality elections and base the number of random samples in part on the contest margin of victory. Risk-limiting audits [13] [5] are examples of such audits. A challenge to extending these audit procedures to IRV is being able to feasibly calculate the margin of victory.

Calculating the margin of victory exactly for an IRV election in the general case may quickly become infeasible as the number of candidates and the number of voters increases. However, for some auditing procedures it is sufficient to calculate a lower bound for the margin of victory.

2 Margin of Victory Definition

There are several ways the notion of a margin of victory for single-winner plurality elections can be extended to other election methods. Generally, a margin of victory describes in some sense how much the set of ballots that was counted would have to change in order to produce a different winner. The concept of a margin of victory can be extended in different ways depending on what kinds of changes are considered and how those changes are measured. Definition 1 considers changes that occur as the result of adding ballots and removing ballots from the set being counted. A change to the set of ballots is measured by the sum of the number of ballots added and the number of ballots removed.

Definition 1. The *margin of victory* for a single election contest is the minimum total number of ballots that must in some combination be added and removed in order for the set of contest winner(s) to change with some positive probability.

A single-winner plurality election is a vote-for-one contest with one winner. A multi-winner plurality election is a vote-for-N contest with N winners. In a multi-winner plurality election, a voter can vote for multiple candidates, but a voter can give at most one vote to any one candidate. So for both single-winner and multi-winner plurality elections, adding or removing a ballot can change the vote total difference between two candidates by at most one vote. Definition 1 is consistent with typical usage of the term for plurality elections:

- For a single-winner plurality contest, the margin of victory is the difference of the vote totals of two candidates with the most votes, the winner and a runner-up.
- In a multi-winner plurality election, the margin of victory is the difference in the vote totals between a

winner with the lowest vote total and a runner-up, a loser with the greatest vote total.

For plurality elections, the election winners can be changed by adding ballots that are counted for a loser, removing ballots that were counted for a winner, or any combination of adding and removing such ballots.

Any combination of added ballots and removed ballots will result in a tie between a winner with the fewest votes and a runner-up, provided that the number of added ballots plus the number of removed ballots is equal to the margin of victory described in the previous bullet points, the added ballots all vote for the runner-up and not the winner, and the removed ballots all vote for the winner and not the runner-up. However, any lesser number of additions and removals of ballots in any combination and counting for any candidates can reduce the difference in vote totals between a winner and a loser, but it can never give a loser as many votes as a winner. Since any lesser number of ballots can not change who the winners are, the amounts described in the previous bullet points satisfy the Definition 1.

An election method using vote-for-1 balloting with N winners is called the limited vote. Definition 1 is also consistent with the usual notion for limited vote contests. Since a voter can give a candidate at most one vote, adding or removing a ballot can change the vote total difference between two candidates by at most one vote.

In contrast, Definition 1 may differ from some notions of a margin of victory for cumulative voting, which uses vote-for- N balloting with N winners, but allows a voter to assign multiple votes to a given candidate. This highlights that Definition 1 measures the margin of victory in units of ballots, rather than votes or points.

The following items provide additional explanation about this definition:

- For simplicity of exposition, it is assumed that any election method discussed will resolve any ties by a random selection from among the tied candidates, where each of the tied candidates has a positive probability of being selected. With this assumption, the margin of victory may be the minimum number of ballots to add and remove that will create a tie, the resolution of which will result in different winners. Various adjustments to what follows can typically accommodate other methods of handling ties.
- This definition allows for scenarios that change a ballot. Changing an existing ballot is treated as removing the ballot and adding a new ballot, contributing a count of two ballots towards the margin of victory.

- This definition is applicable when all ballots are equally weighted. The definition and related procedures would need to be generalized for elections where ballots can have different weights, for example in a corporate election where the ballots are weighted by the number of shares owned. However, being equally weighted still allows different ballots to count for different numbers of votes, as in a multi-winner plurality election.
- For some election methods, the margin of victory can be determined by only considering various ways of adding ballots. Plurality elections are an example of this. For other methods, including IRV, determining the margin of victory can require considering a combination of both added and removed ballots.
- This definition is also applicable to some forms of Single Transferable Vote (STV) elections, also known as choice voting. STV can help ensure greater proportional representation when used for multi-winner elections. IRV can be considered a single-winner special case of STV.

For many types of elections, including IRV elections, special consideration may be needed when applying this margin of victory definition to contests with fewer than two candidates. For example, the margin of victory might be considered to be infinity or undefined, if there is just one candidate and the list of eligible candidates is finalized before ballots are examined. In that case, it might be impossible for the one candidate to lose the election. On the other hand, if any ballot can introduce a new candidate, the margin of victory might be the number of votes for the one tallied candidate, because adding that many ballots for a new candidate would create a tie. As a practical matter, this paper will assume there are at least two candidates. Similarly, it is assumed that there is at least one ballot that counts as a vote for at least one candidate. For exceptional cases that do not satisfy these assumptions, some of the assertions, formulas, and algorithms may not be valid as presented and should be appropriately augmented if those cases need to be handled.

3 Estimating the Margin of Victory with Lower and Upper Bounds

In situations where it may be infeasible to calculate the margin of victory exactly, it may be satisfactory to estimate the margin of victory with an upper bound or a lower bound. When determining the sample size for a post-election audit, including deciding whether to expand the sample size, instead of using the exact margin of victory, a lower bound of the margin of victory might

be used conservatively instead. The lower bound would be conservative to the extent that its use will never result in a smaller random sample being audited compared to using the exact margin of victory, if it were known.

An upper bound for the margin of victory is of interest for auditing purposes because it also provides an upper bound on how far away a lower bound might be from the actual, but unknown margin of victory. If an upper bound and lower bound are relatively far apart, it might be worthwhile to first invest in some additional computational effort in order to try to find a larger lower bound, which might allow significant reductions in audit sample sizes. It might also be worthwhile to first invest in some additional computational effort in order to try to find a smaller upper bound, which might confirm that the lower bound is a relatively good estimate. If an upper and lower bound are relatively close to each other, that indicates that both are relatively close to the actual margin of victory and any improvement in the lower bound might be expected to have negligible impact on the audit sample size.

An upper bound for the margin of victory can be found by demonstrating that adding and removing a specific combination of ballots actually changes the winner with some positive probability. The total number of additions and removals is an upper bound for the margin of victory because the margin of victory is by definition the minimum such number.

A number L can be shown to be a lower bound for the margin of victory if it can be shown that every combination of added ballots and removed ballots is incapable of changing the winner whenever the number of added ballots plus the number of removed ballots is less than L . Since the margin of victory is the minimum number of such added and removed ballots that *can* change the winner, the margin of victory can not be less than L , and so L is a lower bound.

4 IRV Elections and Notation

The tabulation of an IRV contest for a single winner is modeled as follows:

- A voter ranks candidates from highest, most preferred to lowest, least-preferred.
- A voter can rank any number of candidates or may be constrained in how many candidates are allowed to be ranked.
- A voter is not allowed to give two candidates the same ranking.
- The tabulation is conducted in rounds. One candidate is eliminated each round.
- Any candidate that has not been eliminated is designated a continuing candidate.
- At the beginning of each round, votes are tallied. Each ballot counts as one vote for the highest-ranked, most preferred, continuing candidate. If all ranked candidates have been eliminated, the ballot does not count for any candidate. As an efficiency measure, for rounds other than the first, this vote tally can be calculated by transferring each ballot which counted in the previous round for the candidate that was just eliminated. Each such ballot is transferred to count for the candidate that is now the highest-ranked, most preferred, continuing candidate on the ballot.
- In each round, one candidate with the fewest number of votes is eliminated. If two or more candidates are tied for having the fewest votes, one of those candidates is randomly chosen for elimination, provided each tied candidate has a positive probability for being chosen.
- Rounds continue until all candidates have been eliminated.
- The contest winner is the last candidate to be eliminated. This is the one continuing candidate at the beginning of the last round and the one candidate that is not eliminated in the second-to-last round.

This model of an IRV election offers some simplified terminology for describing the margin of victory. Actual implementations [10] [7] [9] of IRV may differ from this model in several ways:

- A winner is declared and the tabulation is halted before all candidates have been eliminated, or even before all but two candidates have been eliminated.
 - Eliminating more than one candidate simultaneously in a single round under certain circumstances may be allowed or required.
 - The boundaries of rounds may vary. For example, a candidate might be eliminated at the beginning of the next round, based on the vote totals of the previous round.
- Despite these differences, the model used here is consistent with such implementations, provided that:
- The order of elimination is consistent. The order of elimination might be truncated due to an early declaration of a winner, or collapsed due to multiple eliminations, but the order of elimination is never reversed.

Notation	Description
C	the number of candidates
c	one of the candidates
r	one of the rounds
$VT(r, c)$	the vote total in round r for candidate c .
$E(r)$	the candidate that is eliminated at the end of round r . E denotes the elimination order for the IRV tabulation. $E(C)$ is the contest winner.
$CO(r, k)$	the continuing candidate in round r , just before the elimination, with the k th lowest vote total. CO represents the candidate order by increasing vote total, and in case of ties, by increasing elimination order. $CO(r, 1)$ is always the candidate eliminated in round r .
$VTO(r, k)$	the k th lowest vote total in round r among continuing candidates. VTO represents the vote totals in order , based on CO :
$MoSE(r)$	the margin of single elimination for round r , defined as the difference in votes for the two candidates with the fewest votes, if $r < C$. When $r = C$, there is just one candidate, and the margin is the votes for that candidate.
MoV	the margin of victory for the contest

Table 1: Notation for IRV tabulation

- The vote tallies for each candidate are the same for the corresponding round in this model.
- The winner is the same.

If these conditions hold, the calculations described for this model can still be applicable to those implementations. Beginning in Section 7, this model is extended to consider simultaneous elimination of multiple candidates in a round.

This model is unusual in that the last round is a round that starts with just one continuing candidate, the winner. This extends the tabulation to more rounds than are typically considered for an actual IRV contest.

Some IRV implementations will stop as soon as one candidate has a majority of all votes counting for continuing candidates, even though there are three or more continuing candidates at the beginning of the round. While this may be sufficient to establish the winner, it can limit the ability to estimate the margin of victory. The number of votes that the winner has above 50% is not the margin of victory. The number of votes above 50% reflects a margin for terminating versus continuing the tabulation, but it does not reflect a margin for the leading candidate to win versus lose. The latter is what matters for calculating the margin of victory using Definition 1.

Some IRV specifications [4] will continue until the round in which just two continuing candidates remain. This is sufficient for calculating the estimates presented

in this paper, if only single eliminations are performed. However it is sometimes convenient for the purposes of terminology and algorithm to extend the tabulation model to an additional round with only one continuing candidate. This is particularly true when considering the tabulation options as a directed, acyclic graph, which is discussed beginning in Section 8.

This model of the IRV tabulations does not generally accommodate IRV elections that limit the number of rounds to a predetermined number. The extent to which the results given here can be adapted to those elections will depend on the specific circumstances.

During the tabulation of an IRV contest, each round starting with M continuing candidates is decided as if it were a limited vote, vote-for-one contest with $M - 1$ winners, with the winners advancing to the next round. The one loser for the round is the runner-up and is eliminated in that round. The margin of victory for the $M - 1$ winners of the round is the difference in votes for the two candidates with the fewest votes for the round. That margin of victory for the round will be called the margin of single elimination for the round, to help distinguish it from the margin of victory for the entire contest. Because adding or removing a ballot can change the difference of vote totals for two candidates by at most one vote, the winners of the round and the elimination for the round can only be changed by adding and removing some combination of ballots that is greater than or equal to that

margin of single elimination.

By adding ballots that are equal in number to the margin of single elimination and that rank the loser of the round as the most preferred candidate, the loser can tie with a winner of the round, changing the elimination for that round with some positive probability. Because those added ballots would always count for the loser in any previous round, they could not change which candidates were eliminated in previous rounds.

Some notation describing an IRV contest and its tabulation is shown in Table 1. If there are C candidates, there are C rounds which are numbered $1, 2, \dots, C$. Round r begins with $C + 1 - r$ continuing candidates and has a vote total for each of them. At the end of round r , after the elimination for the round, there are $C - r$ continuing candidates.

An alternate way of describing $VTO(r, k)$ is:

$$VTO(r, k) \equiv VT(r, CO(r, k)) \quad (1)$$

The definition for $MoSE(r)$ can be expressed as:

$$MoSE(r) \equiv \begin{cases} VTO(r, 2) - VTO(r, 1) & \text{if } r < C, \\ VTO(C, 1) & \text{if } r = C. \end{cases} \quad (2)$$

Any larger value for $MoSE(C)$ will also support the estimates presented here, so as a practical matter, the number of ballots cast may be used in lieu of actually performing a tally for round C .

If $VT(r, c)$ and $E(r)$ are known for all values of r and c , the values of $CO(r, k)$ and $VTO(r, k)$ can be determined in $O(C^2 \log C)$ time, since that involves sorting at most C candidates for each of the C tabulation rounds.

5 Challenges of Calculating IRV MoV

Bartholdi and Orlin demonstrated that strategic voting under IRV is NP-complete [3]. The strategic voting they considered removed potential sources of complexity by considering just a single voter trying to vote strategically while having perfect knowledge of how everyone else would vote. This version of strategic voting is related to the problem of calculating the margin of victory, but with the restriction that ballots can not be removed or changed, only added to the set being tabulated.

Many election methods exhibit both unstable and unresponsive behaviors. Unstable behavior is highlighted in close contests, where relatively small changes in voting can result in large changes in outcome. Unresponsive behavior is highlighted in other contests where any changes to voting, up to a certain threshold, do not change the outcome at all. The margin of victory is one measure of the unresponsiveness of a contest. The difficulty of calculating an IRV margin of victory is related

to IRV being composed of a series of elimination rounds. Instability from a round, viewed as a limited vote contest, is diffused into later rounds at the ballot level using the rule that votes are transferred to the most preferred continuing candidate. Whether the instability is amplified or dampened in later rounds can vary depending on the details of ballot rankings.

IRV, like simple two-round runoff elections are non-monotonic, meaning changing the voted ranking of a candidate to a more preferred ranking can cause that candidate to lose. However non-monotonicity is not by itself the reason why an IRV margin of victory is so difficult to calculate. Calculating the margin of victory for some monotonic election methods, such as the Condorcet-compliant Schulze method [12], is more difficult than calculating the margin of victory for some non-monotonic election methods, such as top-two instant runoffs (all but two candidates are eliminated in the first round).

The margin of victory estimates presented in the following sections can be efficiently calculated because they avoid these complications. This is done in part by using only round-by-round vote totals and by exploiting some of the structure that IRV imposes on those vote totals.

6 Easy Estimates for the IRV MoV

6.1 Last-Two-Candidates Upper Bound

Definition 2. The *Last-Two-Candidates upper bound for the margin of victory*, $MoVUBLTC$, is the margin of single elimination between the last two continuing candidates in round $C - 1$:

$$MoVUBLTC \equiv MoSE(C - 1) \quad (3)$$

The Last-Two-Candidates upper bound is an upper bound for the margin of victory because adding that many ballots counting for the non-winning candidate in that round creates a tie and the random selection to resolve that tie determines the winner. If the added ballots list the non-winning candidate as the first choice, those ballots will count for that candidate in all previous rounds as well, and so will not affect the elimination order of any previous rounds. So the additional ballots will not change which candidates are the two continuing candidates in round $C - 1$, but they will change the winner with some positive probability.

$$MoV \leq MoVUBLTC \quad (4)$$

$MoVUBLTC$ can be calculated in constant time and constant space, if the vote totals $VT(r, c)$ and the elimination order $E(r)$ are given in an appropriate form, for example in arrays.

6.2 Winner-Survival Upper Bound

The Last-Two-Candidates upper bound is the margin by which the winner avoids being eliminated when there is just one other continuing candidate. A similar margin of winner survival can be calculated for every round. The minimum such margin is an upper bound for the margin of victory.

Definition 3. The *margin of survival* in round r for continuing candidate c , $MoS(r, c)$, is the difference in vote totals in round r between candidate c and the candidate eliminated in round r .

$$MoS(r, c) \equiv VT(r, c) - VT(r, E(r)) \quad (5)$$

Definition 4. The *Winner-Survival upper bound for the margin of victory*, $MoVUBWS$, is the smallest margin of any round by which the winner avoids being singly eliminated.

$$MoVUBWS \equiv \min_{1 \leq r \leq C-1} MoS(r, E(C)) \quad (6)$$

To demonstrate that $MoVUBWS$ is an upper bound for the margin of victory, let round r be the first round in which the winner has a margin of survival equal to $MoVUBWS$. There are at least $MoVUBWS$ ballots counting for the winner in the first round because the winner's margin of survival in the first round is at least that large. Removing $MoVUBWS$ of those ballots reduces the winner's vote total by $MoVUBWS$ votes in every round up to and including round r . Other candidate vote totals are not changed. With the ballots removed, the winner can not be eliminated before round r , because the winner's margin of survival in those rounds was greater than $MoVUBWS$. Since all other vote totals stay the same in those rounds, the elimination order before round r does not change. In round r , the winner is now tied to be eliminated, changing the winner with a positive probability. So $MoVUBWS$ is an upper bound. Also, $MoVUBWS$ is the minimum of some values which include $MoVUBLTC$, so $MoVUBWS$ is at least as good of an estimate as $MoVUBLTC$ is.

$$MoV \leq MoVUBWS \leq MoVUBLTC \quad (7)$$

$MoVUBWS$ can be calculated in $O(C)$ time and constant space, if the vote totals $VT(r, c)$ and the elimination order $E(r)$ are given in an appropriate form, for example in arrays.

6.3 Single-Elimination-Path Lower Bound

Definition 5. The *Single-Elimination-Path lower bound for the margin of victory*, $MoVLBSEP$, is the minimum margin of single elimination for any round:

$$MoVLBSEP \equiv \min_{1 \leq r \leq C} MoSE(r) \quad (8)$$

Equation 8 gives the same result whether the minimum ranges over $1 \leq r < C$ or $1 \leq r \leq C$. That is because for every candidate c , $VT(r, c)$ is a non-decreasing function of r as long as candidate c is a continuing candidate. So every candidate has that candidate's highest vote total in the round that candidate is eliminated. Also, the winning candidate has a vote total in each round that is at least as large as the vote total of the candidate that is eliminated in that round. As a result, the winner has a vote total in round C that is at least as large as every vote total for every candidate in every round, so $MoSE(C)$ is at least as large as every other margin of single elimination.

The reason $MoVLBSEP$ is a lower bound for MoV is because in order to change the winner of an IRV election, the winner has to be eliminated in a round before the last round. If fewer than $MoVLBSEP$ ballots are added or removed, then none of the margins of single elimination can be reduced to zero, because all of the margins of single elimination are at least as big as $MoVLBSEP$. If none of the margins can be reduced to zero, none of the eliminations change, so the winner can not change either.

$$MoVLBSEP \leq MoV \quad (9)$$

$MoVLBSEP$ can be calculated in $O(C)$ time and constant space, if the vote totals in order, $VTO(r, k)$, are given in an appropriate form, for example in arrays. Pseudo code in Figure 3 of Section 11 provides an example implementation that uses $VT(r, c)$ and $CO(k)$ in lieu of $VTO(r, k)$ per Equation 1. $MoVLBSEP$ can be calculated in $O(C^2)$ time and constant space, if only the vote totals $VT(r, c)$ are so given. That calculation can be performed by searching the vote totals of each round for the two lowest vote totals, which are then used to calculate $MoSE(r)$.

7 Multiple Eliminations

Some implementations of IRV eliminate multiple candidates in a single round when it can be shown, based on the vote totals of that round, that those candidates will necessarily be eliminated before any other continuing candidates using just single eliminations. In such a situation, the order of elimination of those candidates among themselves can not change the outcome of the election.

For example, consider a round where the lowest vote totals are 20, 30, 35, and 200. The margin of single elimination for this round is $10 = 30 - 20$. So $MoVLBSEP$ for the election as a whole has to be less than or equal to 10. However, the lowest three candidates could all be eliminated in a single round as a multiple elimination because their combined vote total, $20 + 30 + 35 = 85$, is

less than the next highest vote total, 200. The difference, $200 - 85 = 115$, is the margin of multiple elimination.

In the extreme case that the votes for two of the eliminated candidates are all transferred to the third candidate, that candidate will have 85 votes, which is still 115 votes less than the next higher vote total of 200 votes. At least 115 ballots would have to be added or removed in some combination in order to prevent those lowest three candidates from being the next three candidates to be eliminated. When determining a lower bound for the margin of victory in this case, the small margin of single elimination, 10, and the margins of single elimination for the other two candidates can be ignored in favor of the larger margin of multiple elimination, 115.

Definition 6. The *margin of multiple elimination* in round r to simultaneously eliminate the k candidates with the lowest vote totals, $MoME(r, k)$, is the difference of the next higher vote total and the sum of the vote totals for the k candidates:

$$MoME(r, k) \equiv VTO(r, k + 1) - \sum_{i=1}^k VTO(r, i) \quad (10)$$

for $1 \leq k \leq C - r$

This definition is extended to round C with $MoME(C, 1) = MoSE(C)$.

When considering multiple eliminations, the single elimination notation for designating rounds will still be used. With this notation, eliminating k candidates in round r has the effect of notationally advancing to round $r + k$ for the next tabulation round.

Note that $MoME(r, 1) = MoSE(r)$, reflecting that a multiple elimination of a single candidate is the same as the usual single elimination of that candidate. Unless specifically noted otherwise, multiple eliminations will be considered to include single eliminations. However there is one notable distinction to remember. A multiple elimination of k candidates is usable when $MoME(r, k) > 0$. However, single eliminations are usable in the additional case when $MoME(r, 1) = 0$. Since the margin of a single elimination is never negative, the following definition applies.

Definition 7. A *multiple elimination of k candidates is usable* if and only if $k = 1$ or $MoME(r, k) > 0$.

This sense of being usable is separate and independent of whether the multiple elimination is legally permitted for a primary tabulation. Usable multiple eliminations can help estimate a margin of victory regardless of whether those multiple eliminations are legally required, optional, or prohibited for a primary tabulation. The validity of a margin of victory estimate is independent of which equivalent method is used to determine the winner.

A margin of multiple elimination is a different kind of margin than a margin of victory. A margin of multiple elimination is a bounding estimate of worst-case behavior that can be used to prospectively justify use of a multiple elimination, based only on a limited set of summary vote totals. In contrast, a margin of victory is an optimum value, evaluated retrospectively, which can depend on additional information about detail rankings beyond what is exposed by summary vote totals.

All of the k candidates of a multiple elimination are guaranteed to be the next eliminated candidates using single elimination. Even in the extreme case that any $k - 1$ of those candidates are eliminated and all of their ballots are transferred to the remaining candidate, that remaining candidate will still have a deficit of $MoME(r, k)$ votes to avoid being the next eliminated candidate. Adding and removing fewer than $MoME(r, k)$ ballots might change the order in which those k candidates are next eliminated, but that will not be enough to prevent them from being eliminated next in some order. Once those k candidates are next eliminated, the rest of the tabulation, including determination of the winner, will proceed independently of the order in which any candidates were previously eliminated.

The reason for that independence is that when transitioning from one round to the next, it is sufficient for a tabulation to maintain a state consisting only of the set of continuing candidates and the set of ballots being counted. Subsequent rounds are otherwise independent of the vote totals of previous rounds and the order and manner in which any candidates were previously eliminated. As noted in Section 4, it is common practice for IRV tabulations to carry additional state to the next round, such as the candidate for which each ballot last counted or which candidate was most recently eliminated. However that additional state is used only as an efficiency measure and does not change the vote totals or eliminations in subsequent rounds.

Since subsequent rounds are dependent on which candidates were previously eliminated, but not the order or manner in which they were previously eliminated, multiple eliminations are consistent with an optimal substructure that separates what happens in previous rounds from what happens in subsequent rounds. Such an optimal substructure enables the dynamic programming algorithms discussed in Sections 8 and 9.

8 Longest Path Leads to Lower Bound

When a round has several usable multiple eliminations, an IRV tabulation algorithm might choose the one that eliminates the most candidates. The winner will not change, regardless of which usable multiple elimination is chosen. So eliminating the most candidates possible

is simply a heuristic greedy strategy aimed at reducing the number of rounds or the tabulation effort. However, when estimating the margin of victory, it is advantageous to select the multiple elimination that imposes the least constraint on the lower bound. To determine which multiple eliminations to use in a round, it is necessary to consider not only the margin of elimination in that round, but the possible margins of elimination in subsequent rounds as well.

Overall, finding a lower bound for the margin of victory can become a question of what is a best sequence of multiple eliminations. A lower-bound estimate for the margin of a sequence of eliminations is the smallest of its component margins of eliminations, similar to the approach that used *MoVLBSEP* to estimate the margin for the sequence of eliminations consisting only of single eliminations. A sequence with the largest such lower-bound estimate for its margin could provide a better lower-bound estimate for the contest margin of victory.

Once the values for $MoME(r, k)$ are known, the calculation of an improved lower-bound estimate of MoV can be seen as solving a longest path problem on a weighted, directed, acyclic graph. The graph vertices correspond to the rounds of the IRV tabulation. The edges of the graph represent eliminating k candidates in round r and advancing effectively to round $r + k$. An edge is in the graph if and only if the elimination is a usable multiple elimination: $k = 1$ or $MoME(r, k)$ is positive. The weight of an edge is the value $MoME(r, k)$. The graph is acyclic because eliminations always advance to a later round. The length of a path is not the sum of the weights of the edges, but the minimum of the weights of the edges. Any path can be identified by its starting vertex and the sequence of elimination counts, the values of k , used to traverse that path. The optimization goal is to find a path from the round 1 vertex to the round C vertex that maximizes the length of the path.

This longest path problem is specialized in several ways:

- The topological order of the vertices is explicitly given by the problem: round sequence order.
- The path that connects all of the vertices is part of the graph: the sequence of single eliminations.
- Part of the problem is discovering which of the $C \cdot (C - 1)/2$ possible edges are actually part of the graph.
- The length of a path is the minimum of its edge weights rather than their sum.

This problem can be solved with a longest path algorithm that is adapted to the special requirements listed

above. The algorithm can use either of two possible approaches: starting at round 1 and working forward to the last round, or starting at the last round and working backward to the first round [6]. Both approaches are similar in nature and have the same computational complexity. The method presented in detail in this paper is the second approach, starting at the last round and working in reverse round order.

This problem can also be understood as a bottleneck problem rather than a length problem. Each multiple elimination imposes a downward constraint or bottleneck on the lower-bound estimate of the margin of victory. Each path of eliminations imposes a downward constraint, a bottleneck rating, which is measured by the smallest bottleneck of its component parts. The optimization goal is to find a path with the largest bottleneck rating. That will be a path of eliminations that avoids very small bottlenecks as much as possible.

Another, optional perspective for understanding and validating solutions to this problem is an algebraic framework for best-path or optimum-distance algorithms using semirings [8]. Semirings offer a mathematical basis for generalizing path length optimization problems to other types of problems. In this case, semirings can be used to express that the problem seeks the longest path, but path length is the minimum, rather than the sum, of a path's component segments. A semiring that accomplishes this is $(\mathbb{Z}_{\geq 0} \cup \{+\infty\}, \max, \min, 0, +\infty)$, where $\mathbb{Z}_{\geq 0}$ is the set of non-negative integers, the semiring addition of two numbers is the larger of the two, and their semiring product is the smaller of the two. The identity element for semiring addition is 0 and the identity element for semiring multiplication is $+\infty$. In the general framework, semiring addition expresses how optimization comparisons are made, while semiring multiplication expresses how paths are measured in relation to their components. Semirings with \max as addition and \min as multiplication are known as bottleneck semirings because they are used to represent bottleneck optimization problems.

9 Best-Path Lower Bound

To find an estimated margin of best elimination overall, first define the estimated margin of best elimination restricted to partial paths that start in round r and go to the last round. This is defined recursively in reverse round order by considering the best paths that start in round r , begin by eliminating k candidates, and go to the last round.

Definition 8. The *estimated margin of best elimination starting in round r and initially eliminating k candidates*, $EMoBE(r, k)$, and the *estimated margin of best elimina-*

tion starting in round r , $EMoBE(r)$, are defined recursively as follows:

For $r = C$:

$$EMoBE(C) \equiv MoME(C, 1) \quad (11)$$

For any round r with $1 \leq r \leq C - 1$, and for any k with $1 \leq k \leq C - r$:

$$EMoBE(r, k) \equiv \min(MoME(r, k), EMoBE(r + k)) \quad (12)$$

$$EMoBE(r) \equiv \max_{k=1}^{C-r} EMoBE(r, k) \quad (13)$$

The estimated margin of best elimination starting in round r , is the largest margin that can be so found, among all possible values of k .

Since $EMoBE(C) \geq 0$ and $MoME(r, 1) \geq 0$, by induction it is true for all rounds that:

$$EMoBE(r, 1) \geq 0 \quad (14)$$

$$EMoBE(r) \geq 0 \quad (15)$$

Note that the definition of $EMoBE(r)$ involves all values of k , regardless of whether the multiple elimination of k candidates is usable. However, the value of $EMoBE(r)$ only depends values of k for usable multiple eliminations.

In each round r , there is always at least one value of k for which an elimination of k candidates is usable and:

$$EMoBE(r) = EMoBE(r, k) \quad (16)$$

If $EMoBE(r) = 0$, let $k = 1$. Single eliminations are always usable.

If $EMoBE(r) > 0$, let k be one of the values for which $EMoBE(r) = \min(MoME(r, k), EMoBE(r + k))$. Then $MoME(r, k) \geq EMoBE(r) > 0$, so the elimination of k candidates is usable.

As a result, an election tabulation starting from round r can always be completed with usable multiple eliminations, each with a margin of elimination greater than or equal to $EMoBE(r)$.

Definition 9. The *Best-Path lower bound for the margin of victory*, $MoVLBBP$, is the final result of this calculation in reverse round order:

$$MoVLBBP = EMoBE(1) \quad (17)$$

Because the election can be tallied from the beginning by only doing eliminations with margins greater than or equal to $MoVLBBP = EMoBE(1)$, any combination of adding and removing less than $MoVLBBP$ ballots can not change those eliminations in a way that could change

the winner, so $MoVLBBP$ really is a lower bound for MoV :

$$MoVLBBP \leq MoV \quad (18)$$

Since $MoSE(C) = MoME(C, 1) = EMoBE(C)$ and $MoSE(r) = MoME(r, 1)$, by induction on the reverse round order, the following holds:

$$\min_{j=r}^C MoSE(j) \leq EMoBE(r) \quad (19)$$

because:

$$\begin{aligned} \min_{j=r}^C MoSE(j) &= \min(MoSE(r), \min_{j=r+1}^C MoSE(j)) \\ &\leq \min(MoME(r, 1), EMoBE(r + 1)) \\ &= EMoBE(r, 1) \\ &\leq EMoBE(r) \end{aligned} \quad (20)$$

Applying Equation 19 with $r = 1$ gives:

$$MoVLBBP \leq MoVLBBP \quad (21)$$

If there are no opportunities for simultaneous elimination of two or more candidates, then $MoME(r, k) \leq 0$ for all r and for all $k \geq 2$, causing the inequalities of Equations 19 and 21 to become equalities.

In terms of the longest path problem of Section 7, $EMoBE(r)$ is the length of the longest path from the round r vertex to the round C vertex, $MoVLBBP$ is the length of the longest path, and $MoVLBBP$ is the length of the path consisting only of single eliminations.

$MoVLBBP$ can be calculated in $O(C^2)$ time and $O(C)$ space, if the vote totals in order, $VTO(r, c)$, are given in an appropriate form, for example in arrays. Pseudo code in Figure 4 of Section 11 provides an example implementation that uses $VT(r, c)$ and $CO(k)$ in lieu of $VTO(r, k)$ per Equation 1. $MoVLBBP$ can be calculated in $O(C^2 \log C)$ time and $O(C)$ space, if only the vote totals $VT(r, c)$ are so given. The space is still $O(C)$ because the results of sorting the candidates within a round are only needed one round at a time.

10 Estimates for Early-Terminated Tabulations

When a candidate has a majority of all of the votes for continuing candidates in a particular round, that candidate is assured of being the winner, even if the tabulation continued for all C or $C - 1$ rounds. Sometimes an IRV tabulation is stopped at the first round in which a candidate has such a majority of votes. Even without the vote totals for later rounds, some estimates for the margin of victory are possible.

Definition 10. The *majority winner in round r* , is the candidate, if one exists, that has more than 50% of all votes for continuing candidates in that round.

A majority winner in round r is the IRV winner because the vote total for that candidate can not decrease in later rounds and the number of votes for all continuing candidates in later rounds can not increase. So a majority winner in round r will be a majority winner in all subsequent rounds as well.

The Last-Two-Candidates upper bound can be adapted to rounds with a majority winner by allowing for the possibility that all votes for any candidates other than the two leading candidates might transfer to the majority winner in later rounds.

Definition 11. The *round r Majority-Winner upper bound for the margin of victory*, $MoVUBMW(r)$, is the difference in votes between the two leading candidates, plus the sum of votes for all other candidates, provided that round r has a majority winner.

$$MoVUBMW(r) \equiv VTO(r, C + 1 - r) \quad (22)$$

$$- VTO(r, C - r)$$

$$+ \sum_{k=1}^{C-1-r} VTO(r, k)$$

The Majority-Winner upper bound is also equal to twice the difference between 50% of the votes for all continuing candidates in round r and the votes for the second leading candidate in round r . The Majority-Winner upper bound is an upper bound for the margin of victory because if that many ballots were added and those ballots ranked the second leading candidate of round r as the most preferred choice, even if votes for all other candidates transferred to the majority winner, the second leading candidate would not be eliminated before round $C - 1$ and would at least tie with the majority winner in round $C - 1$.

If the tabulation were continued for all rounds, the candidate eliminated in round $C - 1$ would have at least $VTO(r, C - r)$ votes, because that candidate either would be the second leading candidate of round r or would have survived when the second leading candidate of round r was eliminated. Likewise, in round $C - 1$, the majority winner would have at most $VTO(r, C + 1 - r) + \sum_{k=1}^{C-1-r} VTO(r, k)$ votes. Which means that a Majority-Winner upper bound is never a better estimate than the Last-Two-Candidates upper bound:

$$MoVUBLTC \leq MoVUBMW(r) \quad (23)$$

The Winner-Survival upper bound can be adapted to tabulations that stop in round r with a majority winner.

Definition 12. The *round r Majority-Winner-Survival upper bound for the margin of victory*, $MoVUBMWS(r)$, is the minimum of the margins of the majority winner's survival in the first $r - 1$ rounds and the Majority-Winner upper bound for round r , provided that round r has a majority winner, w .

$$MoVUBMWS(r) \equiv \min\left(\min_{1 \leq s \leq r-1} MoS(s, w), MoVUBMW(r)\right) \quad (24)$$

The Majority-Winner-Survival upper bound is an upper bound because it is the minimum of some upper bounds. Because of Equations 6 and 23, the Majority-Winner-Survival upper bound for round r is never a better estimate than the Winner-Survival upper bound:

$$MoVUBWS \leq MoVUBMWS(r) \quad (25)$$

There is not a good way to adapt the Single-Elimination-Path lower bound to tabulations that stop as soon as a majority winner is detected. However, if round r has a majority winner, then eliminating all $C - r$ candidates except that winner is a usable multiple elimination. The Best-Path lower bound can be adapted by restricting the best path algorithm to only consider elimination paths that end with a multiple elimination of at least $C - r$ candidates. Such elimination paths do not pass through the vertices for any rounds after round r but before round C .

Definition 13. The *round r Majority-Winner-Best-Path lower bound for the margin of victory*, $MoVLBMWBP(r)$, is the largest margin of any multiple elimination path that does not pass through the vertices for any rounds greater than r and less than C , provided that round r has a majority winner, w .

Because the Majority-Winner-Best-Path lower bound finds the best path of a subset of those paths considered by the Best-Path lower bound, the Majority-Winner-Best-Path lower bound is never a better estimate, and hence is a lower bound.

$$MoVLBMWBP(r) \leq MoVLBBP \quad (26)$$

A special case of the Majority-Winner-Best-Path lower bound is the case for $r = 1$. In that case, the only elimination path considered is the path consisting of just one multiple elimination of $C - 1$ candidates. The Majority-Winner-Best-Path lower bound is equal to $MoME(1, C - 1)$. That is also equal to twice the difference between the first-round votes for the majority winner and 50% of the first-round votes for all continuing candidates. An equation for $MoVLBMWBP(1)$ can be written with parallelism to Equation 22 showing that it is what the margin of elimination would be in round $C - 1$

```

calculate_MoVUBLTC() {
    r = C-1;
    return VT[r, CO[r,2]] - VT[r, CO[r,1]];
}

```

Figure 1: Pseudo code to calculate the Last-Two-Candidates upper bound, *MoVUBLTC*

```

calculate_MoVUBWS() {
    movubws = VT[C, E[C]];
    for (r = 1; r <= C-1; r++) {
        mos = VT[r, E[C]] - VT[r, E[r]];
        if (mos < movubws) {
            movubws = mos;
        }
    }
    return movubws;
}

```

Figure 2: Pseudo code to calculate the Winner-Survival upper bound, *MoVUBWS*

```

calculate_MoVLBSEP() {
    movlbsep = VT[C, E[C]];
    for (r = 1; r <= C-1; r++) {
        mose = VT[r, CO[r,2]] - VT[r, CO[r,1]];
        if (mose < movlbsep) {
            movlbsep = mose;
        }
    }
    return movlbsep;
}

```

Figure 3: Pseudo code to calculate the Single-Elimination lower bound, *MoVLBSEP*

```

calculate_MoVLBBP() {
    EMoBE[C] = VT[C, E[C]];
    for (r = C-1; r >= 1; r--) {
        /* sort candidates for round r here
           if CO is not provided as input */
        elim_votes = 0;
        EMoBE[r] = -1;
        for (k = 1; k <= C-r ; k++) {
            elim_votes += VT[r, CO[r, k]];
            mome = VT[r, CO[r, k+1]] - elim_votes;
            if (k == 1 || mome > 0) {
                if (mome <= EMoBE[r + k]) {
                    emoberk = mome;
                }else{
                    emoberk = EMoBE[r + k];
                }
            }
            if (emoberk > EMoBE[r]) {
                EMoBE[r] = emoberk;
            }
        }
    }
}
return EMoBE[1];
}

```

Figure 4: Pseudo code to calculate the Best-Path lower bound, *MoVLBBP*

in the extreme case that all first round votes for candidates other than the two leading first round candidates were transferred to a candidate other than the winner.

$$\begin{aligned}
 MoVLBMWBP(1) = VTO(1, C) & \quad (27) \\
 & - VTO(1, C - 1) \\
 & - \sum_{k=1}^{C-2} VTO(1, k)
 \end{aligned}$$

11 Pseudo Code

The pseudo code in Figures 1 - 4 illustrates the algorithms for calculating the four principle estimates of *MoV*. There are four routines, each returning one of the estimates.

The pseudo code generally uses C-like syntax, but does not explicitly define variables. Details about the allocation and persistence of data are ignored.

The routines assume the following data is available: *E*, *VT*, and *CO* as arrays, and *C* as a scalar, corresponding to the values in Table 1. These routines assume that the number of candidates is at least 2.

12 Worst-Case Analysis of Estimates

The following examples show that any estimates based only on tabulation round-by-round vote totals do a poor job of estimating the IRV margin of victory in the worst cases:

- The ratios of the best estimates and the margin of victory, $MoVUBWS/MoV$ and $MoV/MoVLBBP$, are unbounded and can grow exponentially with the number of candidates.
- There are no margin of victory estimates, based only on tabulation round-by-round vote totals, that have better worst-case performance.

This is shown with a sequence of pairs of IRV contests, parametrized by the number of candidates, $C \geq 5$. The rankings for ballots of each pair of contests are described in Table 2. Ballots for contest 1 of each pair do not include the optional rankings for candidate a_1 , denoted as $[a_1 >]$ in Table 2, while ballots for contest 2 do include those optional rankings. Because the optional a_1 rankings are never the most preferred rankings and a_1 is eliminated first, the tabulation of contest 2 never counts votes for the optional a_1 rankings. So both contests have the same:

Number of Ballots	Ballot Rankings
2	$a_0 > [a_1 >] w$
1	$a_1 > a_0 > w$
2	$a_2 > [a_1 >] a_0 > w$
4	$a_3 > [a_1 >] a_0 > w$
...	...
2^{C-4}	$a_{C-3} > [a_1 >] a_0 > w$
$2^{C-3} + 2$	$b > [a_1 >] w$
$2^{C-2} + 1$	w

Table 2: Worst Case Ballot Rankings

- candidates: $a_0, a_1, a_2, \dots, a_{C-3}, b, w$
- number of voters, $2^{C-1} + 4$
- winner, w
- elimination sequence: $a_1, a_2, \dots, a_{C-3}, a_0, b, w$
- tabulation round-by-round vote totals
- Best-Path lower bound, $MoVLBBP = 1$
- Winner-Survival upper bound, $MoVUBWS = 2^{C-3}$

However, $MoV = 2^{C-3}$ for contest 1, and $MoV = 1$ for contest 2.

Because these two contests have the same tabulation round-by-round vote totals, no margin of victory estimate that relies only on those vote totals can distinguish between the two contests. No such lower bound can be higher than the margin of victory for contest 2, $MoV = 1 = MoVLBBP$, and no such upper bound can be lower than the margin of victory for contest 1, $MoV = 2^{C-3} = MoVUBWS$.

The rest of this section confirms these values for $MoVLBBP$, MoV , and $MoVUBWS$.

With three exceptions, each candidate a_k is the most preferred candidate on 2^{k-1} ballots with rankings $a_k > [a_1 >] a_0 > w$. The three exceptions are: there are two ballots with a_0 ranked most preferred, ballots ranking a_0 as most preferred omit the subsequent redundant ranking for a_0 , and ballots ranking a_1 as most preferred omit the redundant optional ranking for a_1 .

As the tabulation proceeds during the first $C - 3$ rounds, candidates a_k are eliminated, starting with a_1 , and all votes are transferred to a_0 . In each of those

rounds, a_0 avoids elimination by a margin of 1 vote, the margin of single elimination. In round $C - 2$, a_0 has accumulated all $2^{C-3} + 1$ votes that started with any of the a_k candidates, while candidates b and w still have their original vote totals, $2^{C-3} + 2$ and $2^{C-2} + 1$ respectively. Candidate a_0 is then eliminated by a margin of one vote and all $2^{C-3} + 1$ votes are transferred from a_0 to w . In round $C - 1$, w defeats b by a margin of 2^{C-2} votes, $3(2^{C-3}) + 2$ to $2^{C-3} + 2$.

All of the first $C - 2$ rounds have a margin of single elimination = 1. The only other usable multiple eliminations occur in the first $C - 3$ rounds and eliminate all continuing a_k candidates, leaving only candidates b and w . Those multiple eliminations have a margin = 1, since all continuing a_k candidates have a total of $2^{C-3} + 1$ votes. As a result, there are no paths of multiple elimination that provide a better lower bound than provided by the path of single elimination, and the Best-Path lower bound, $MoVLBBP = 1$.

Since candidate w does not receive any additional votes from the first round to round $C - 2$, the margin of winner survival decreases during those rounds. Candidate w survives by 2^{C-3} votes in round $C - 2$ and by 2^{C-2} votes in round $C - 1$. So the Winner-Survival upper bound, $MoVUBWS = 2^{C-3}$.

In order to change the winner in contest 1, without the optional a_1 rankings, w has to be eliminated while there is at least one other continuing candidate. Looking at the detail ballot rankings, none of the a_k candidates is ranked at a higher preference than w on more than $2^{C-3} + 1$ ballots. In round 1, w has $2^{C-2} + 1$ votes, and so w has at least that many votes in any round up until w would be eliminated. In order for w to be eliminated while one of the a_k candidates is still a continuing candidate requires adding or removing ballots to overcome a margin of at least $2^{C-3} = (2^{C-2} + 1) - (2^{C-3} + 1)$. If none of the a_k candidates is a continuing candidate, b is the only other candidate that could be a continuing candidate when w is eliminated. But the pair-wise margin between w and b , as seen in round $C - 1$, is 2^{C-2} votes. There are no other ways to defeat w in contest 1, so a lower bound for the margin of victory for contest 1 is the least of those alternatives, 2^{C-3} . Since that is also the Winner-Survival upper bound, the margin of victory for contest 1, $MoV = 2^{C-3}$.

In contest 2, with the optional a_1 rankings, if a single ballot is added that ranks a_1 as the most preferred candidate, then in the first round there is a three-way tie for elimination between a_0 , a_1 , and a_2 , each with two votes. If a_0 or a_2 is chosen for elimination in the first round, then all of the a_k votes are accumulated by a_1 , until round $C - 2$ when a_1 and b are tied for elimination, each with $2^{C-3} + 2$ votes. If b is chosen for elimination, votes for b are transferred to a_1 , and a_1 defeats w

Margin of Victory, Range of Best Estimates Australia 2010 House of Representatives




Figure 5: Margin of victory range bounded by best estimates for Australia 2010 House of Representatives, using the Best-Path lower bound and the Winner-Survival upper bound

Margin of Victory, Range of Best Estimates California Local IRV Elections, 2004-2010

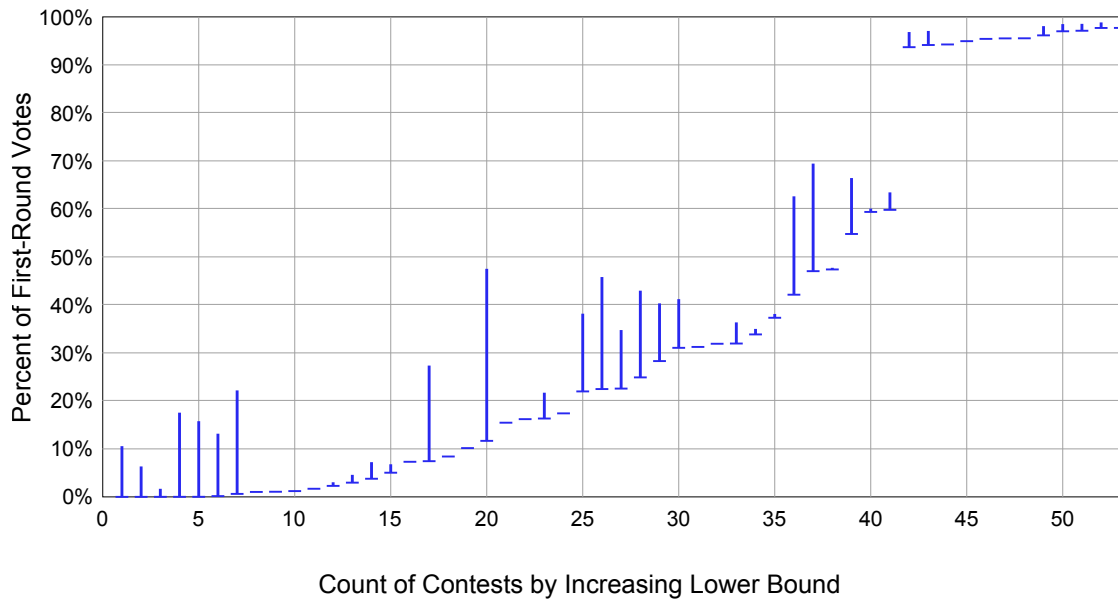


Figure 6: Margin of victory range bounded by best estimates for California local IRV elections, 2004-2010, using the Best-Path lower bound and Winner-Survival upper bound or using the first-round Majority-Winner-Best-Path lower bound and the first-round Majority-Winner-Survival upper bound

in round $C - 1$, $2^{C-2} + 4$ to $2^{C-2} + 1$. This shows that adding a single ballot can change the outcome of the election with a positive probability, so an upper bound for the margin of victory for contest 2, is 1. Since that is also the Best-Path lower bound, the margin of victory for contest 2, $MoV = 1$.

The ballots in these examples do not rank all candidates. However, no ballots are exhausted in any of the analysis, so the rankings can be completed in any manner without changing the analysis. So these examples can be made applicable to IRV elections where all candidates must be ranked. These particular examples do not directly apply to IRV contests where at most three candidates can be ranked.

13 Estimates Applied to Elections

While these estimates may perform poorly against specially constructed worst-case examples, for some uses the effectiveness of these estimates with real-world elections is more relevant. The behavior of elections can be influenced by many factors that are context dependent. To gain some initial insight into how these estimates perform, estimates for the margin of victory were calculated for two sets of IRV elections. The first set consists of the 150 contests for the 2010 election of Australia's House of Representatives [2]. The second set consists of the 53 IRV contests in California since 2004. That set consists of 36 contests in San Francisco [11], some occurring each year from 2004 to 2010, and another 17 contests in the cities of Oakland, Berkeley, and San Leandro [1].

All estimates are expressed as a percentage of total first-round votes for candidates. For each contest, the best available lower bound and upper bound were calculated. For most contests these estimates were the Best-Path lower bound and the Winner-Survival upper bound. For twelve San Francisco elections only first-round vote totals² were published. Those elections were all won by significant margins in the first round. For these contests, the first-round Majority-Winner-Best-Path lower bound and the first-round Majority-Winner-Survival upper-bound were used as the best estimates.³ The lower bounds for those elections were at least 40% of the total votes. In all other cases for California contests, the published cast vote records were used to perform a full tabulation using only single elimination and to then calculate the estimates.

The estimates are presented graphically in Figure 1 for the Australia elections and in Figure 2 for the California elections. The contests are ordered by increasing lower bound estimate and display the interval between the lower bound and the upper bound.

These two sets of election data represent some diverse

uses of IRV. The Australia contests were all concurrent, while the California contests cover seven years of elections. The Australia contests were for a national legislature, while the California contests cover a variety of local positions, including mayors, city council members, a school board members, and various specialty positions such as auditor and city attorney. The Australia contests were partisan elections with one candidate per party, while the California contests are nominally non-partisan. Australia requires voters to rank all candidates. California contests limited voters to three rankings and those rankings were optional. The reporting of results was also different. Australia provided reports that only used single elimination and reported round-by-round vote totals for $C - 1$ rounds. Some reporting of California contests used multiple eliminations and stopped when a majority winner was identified.

There are many ways to analyze this data, depending in part on the intended use of the estimates. Full evaluation of these estimates for risk-limiting audits will depend on the specific risk-limiting audit method being used as well as the patterns of discrepancies identified by an audit.

Contests where the lower bound and upper bound are both equal to the margin of victory do not display a vertical bar. This happened in 71 (47%) of the Australia contests and 17 (32%) of the California contests. The lower bound was greater than 10% for 85 (57%) of the Australia contests and 35 (66%) of the California contests. This suggests that these estimates may be sufficient to feasibly conduct risk-limiting audits for a significant portion of IRV contests.

On the other hand, the ratio of the upper bound to the lower bound was greater than 10 for two (1.3%) of the Australia contests, with lower bounds of 0.14% and 2.83%. That ratio was greater than 10 for seven (13%) of the California contests, all with lower bounds less than 1% of the total first-round vote, five with lower bounds less than 0.02%. This suggests that these estimates will not be sufficient to feasibly conduct risk-limiting audits for all IRV contests.

14 Conclusions

Four principle estimates for the margin of victory of an IRV contest have been defined: the Single-Elimination-Path lower bound, the Best-Path lower bound, the Winner-Survival upper bound, and the Last-Two-Candidates upper bound. These estimates satisfy the relationship:

$$\begin{aligned} MoVLBSEP &\leq MoVLBBP \leq & (28) \\ & MoV \\ &\leq MoVUBWS \leq MoVUBLTC \end{aligned}$$

These estimates can be efficiently calculated from election tabulation summary totals, without looking at individual ballots. When adapting post-election auditing procedures from plurality elections to IRV elections, the Best-Path lower bound is recommended as a starting point for use as an estimate of the margin of victory.

There are two kinds of approximations made during the calculation of the Best-Path lower bound that can cause it to differ from the actual margin of victory. The first approximation is the allowance for extreme cases when a multiple elimination of two or more candidates is performed. The second approximation occurs because the Best-Path lower bound guarantees not just that the winner does not change, but that none of the multiple eliminations in a best path of multiple eliminations will change either. This approximation is reflected in the use of the minimum function to calculate the bottleneck imposed by an elimination path.

Because of these approximations, there can on occasion be IRV elections where the Best-Path lower bound does not estimate the margin of victory adequately and the use of the Best-Path lower bound would result in unnecessarily large audit sample sizes. In those situations, additional techniques can be used to possibly improve the lower-bound estimate of the margin of victory. By considering the actual rankings of ballots cast there may be significant opportunities to raise the lower-bound estimate, resulting in reduced audit sample sizes.

15 Acknowledgments

I thank members and associates of Californians for Electoral Reform, especially Jonathan Lundell, and members of the San Francisco Voting System Task Force, especially Jim Soper, for their helpful advice and feedback about this subject and this paper. I also thank the anonymous reviewers for their helpful feedback, particularly the idea for the Winner-Survival upper bound.

References

- [1] ALAMEDA COUNTY REGISTRAR OF VOTERS. Previous elections and results - November 2, 2010 - General Election. http://www.co.alameda.ca.us/rov/past_elections/election_info2010-11-02.htm, November 2010.
- [2] AUSTRALIAN ELECTORAL COMMISSION. Distribution of preferences by candidate by division. <http://results.aec.gov.au/15508/Website/Downloads/HouseDopByDivisionDownload-15508.csv>, September 2010.
- [3] BARTHOLDI, J. J., AND ORLIN, J. B. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8 (1991), 341–354.
- [4] California Assembly Bill 1121. http://leginfo.ca.gov/pub/09-10/bill/asm/ab_1101-1150/ab_1121_bill_20090825_amended_sen_v95.pdf, August 2009.
- [5] HALL, J. L., STARK, P. B., MIRATRIX, L. W., BRIONES, M., GINNOLD, E., OAKLEY, F., PEADEN, M., PELLERIN, G., STANIONIS, T., AND WEBBER, T. 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).
- [6] HUANG, L. Advanced dynamic programming in semiring and hypergraph frameworks. COLING - Tutorials, <http://www.cis.upenn.edu/~lhuang3/coling.pdf>, 2008.
- [7] Minneapolis, Minnesota, Code of Ordinances, Chapter 167 - Municipal Elections: Rules of Conduct. http://library.municode.com/HTML/11490/level3/COOR_TIT8.5EL_CH167MUELRUCO.html, April 2008.
- [8] MOHRI, M. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata, Languages and Combinatorics* 7, 3 (2002), 321–340.
- [9] The Charter of the City of Oakland, California, Article XI - Elections, Section 1105 - Ranked Choice Voting. http://library.municode.com/HTML/16308/level2/THCHOA_ARTXIEL.html, February 2007.
- [10] San Francisco, California Charter, Article XIII - Elections, Section 102 - Instant Runoff Elections. <http://library.municode.com/HTML/14130/level1/ARTXIIEEL.html>, March 2002.
- [11] SAN FRANCISCO DEPARTMENT OF ELECTIONS. City and County of San Francisco : Election archives by year. <http://www.sfgov2.org/index.aspx?page=1671>, November 2010.
- [12] SCHULZE, M. A new monotonic and clone-independent single-winner election method. *Voting Matters* 17 (2003), 9–19.
- [13] STARK, P. B. Conservative statistical post-election audits. *The Annals of Applied Statistics* 2, 2 (2008), 550–581.

Notes

¹In some jurisdictions, Instant Runoff Voting is also known as Ranked-Choice Voting (RCV), single-seat Ranked-Choice Voting, the alternative vote, or preferential voting.

²The published vote totals were sometimes just first-choice vote totals. These could differ from the true first-round totals to the extent that if a voter did not vote for a first choice but did vote for a second or third choice, the second-choice vote would count as a first-round vote. The possibility of the likely slight adjustments to the published vote totals are ignored by the data presented here.

³When votes were reported only for the winning candidate and a single block of one or more unspecified write-in candidates, it was assumed that no write-in candidate received more than one vote, which tended to expand the interval between lower bound and upper bound.