Estimating the Margin of Victory for an IRV Election

Part 1

by David Cary

November 6, 2010

Summary

New procedures are being developed for post-election audits involving manual recounts of random samples of ballots. These new audit procedures often target plurality elections and base the number of random samples on the margin of victory. One of the challenges of extending these new audit procedures to Instant Runoff Voting (IRV)* is knowing how to calculate the margin of victory.

This paper offers a general definition for an election's 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 increases. However, for many auditing procedures it is sufficient to calculate a lower bound for the margin of victory.

This paper describes how to calculate one upper bound and two lower bounds for the margin of victory for an IRV election. Given summary election vote totals, the time complexity for calculating these bounds does not exceed $O(C^2 \log C)$, where C is the number of candidates.

Definition

The **margin of victory** for an election is the minimum number of ballots that must be added and/or removed in order for the election winner(s) to change with some positive probability.

This definition is consistent with typical usage of the term for plurality elections:

- In a single-winner plurality election, the margin of victory is the difference of the vote totals of the two candidates with the most votes.
- In a multi-winner plurality election producing w winners, the margin of victory is the difference in the vote totals between the winner with the lowest vote total and the non-winner with the greatest vote total.

For plurality elections, the margin of victory can be demonstrated by either adding ballots that are counted for the runner-up, removing ballots that were counted for the (lowest) winner, or any combination of adding and removing such ballots.

---

* In some jurisdictions, Instant Runoff Voting (IRV) is also known as Ranked Choice Voting (RCV).

The following points 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 non-zero probability of being selected. With this assumption, the margin of victory may be the minumum number of ballots to add or remove that will create a tie or to avoid a tie. Minor adjustments to what follows can typically accommodate other methods of handling ties.
- When there is a tie and different winners can result, depending on how the tie is randomly broken or resolved, the margin of victory is zero. This is because a different outcome can result without adding or removing any ballots, just redoing the random tie-breaking.
- This definition allows an existing ballot to be changed. Changing an existing ballot is treated as removing the ballot and adding a new ballot, contributing a count of 2 ballots towards the margin of victory.
- It is assumed that all ballots are equally weighted with one vote in the election tally. The procedures can be readily generalized to elections where ballots can have different weights, for example in a corporate election where the ballots are weighted by the number of shares owned.
- For some election methods, calculating the the margin of victory only requires consideration of scenarios where ballots are added. Plurality is an example of this. For other methods, including IRV, the margin of victory can require a combination of added and removed ballots.
- This definition is also applicable to some forms of Single Transferable Vote (STV) elections, also known as choice voting. STV offers greater proportional representation through multi-winner elections. Most of the contents of this paper can be extended, for example, to the Weighted Inclusive Gregory method, involving fractional transfers. IRV can be considered a single-winner special case of STV.

IRV Elections and Notation

An IRV election tally is modeled as follows:

- In a single-winner plurality election, the margin of victory is the difference of the vote totals between the two candidates with the most votes.
- A voter ranks candidates from highest, most preferred to lowest, least-preferred.
- A voters can rank any number of candidates or may be constrained to rank at most R candidates.
- A voter is not allowed to give two candidates the same ranking.
- The tally 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 continuing candidate. If all ranked candidates have been

eliminated, the ballot does not count for any candidate. For the rounds other than the first, this vote tally can be calculated by taking the votes that were counted for the candidate that was eliminated in the previous round and transferring each ballot to count for the highest-ranked continuing candidate on the ballot.

- A candidate with the fewest number of votes in that round is eliminated. If there is a tie between two or more candidates with the fewest votes, one of those candidates is randomly chosen for elimination.
- Rounds continue until all but one candidate has been eliminated. The last continuing candidate is the winner.

This model of an IRV election offers some simplified terminology for calculating the margin of victory. Actual implementations of IRV may differ from this model in several ways:

- A winner may be declared before all but one candidate has been eliminated.
- It may be allowed or required to eliminate more than one candidate simultaneously in a single round in certain circumstances.
- The boundaries of rounds may vary. For example, a candidate might be eliminated at the beginning of the next round, rather that at the end of the current round.

Despite these differences, the model used here can still be consistent with such implementations, to the extent that:

- The order of elimination is consistent. The order of elimination might be truncated due to a early declaration of a winner or collapsed due to multiple eliminations, but the order of elimination is never reversed.
- The vote tallies for each candidate are the same for corresponding rounds.
- The winner is the same.

If these conditions hold, the calculations described for this model can still be applicable to those implementations.

Using this model, if there are C number of candidates, there will always be C-1 rounds which are numbered 1, 2, … , C-1. Round r will have tallies for C + 1 – r candidates. At the end of round r, after the eliminations for the round, there will be C – r continuing candidates.

Define the following:

VT(r, c) is the **vote total** in round r for candidate c.

E(r) is the candidate that is eliminated at the end of round r. E denotes the **elimination order** for the IRV tally. Extend this to a full permutation of candidates by defining E(C) is the election winner.

CO(r, k) is the continuing candidate in round r, just before the elimination, with the kth lowest vote total. CO represents the **candidate order** by increasing vote total, and in case of ties, by increasing elimimination order. CO(r, 1) is always the candidate eliminated in round r.

VTO(r, k) is the kth lowest vote total in round r among continuing candidates. VTO repesents the **vote total order**, based on CO:

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

MoSE(r) is the **margin of single elimination** for round r, defined as the difference in votes between the two candidates with the fewest votes:

$$MoSE(r) = VTO(r, 2) - VTO(r, 1) \quad (2)$$

MoV is the **margin of victory** for the election.

An Upper Bound for the Margin of Victory:

Upper bounds for MoV can be determined by demonstrating that adding and/or removing a specific combination of ballots actually changes the winner. The total number of additions and removals in that specific case is an upper bound for MoV because the MoV is the minimum such number needed to change the winner.

Define the **first upper bound for the MoV,** as:

$$\mathbf{MoVUB1} = MoSE(C-1) \quad (3)$$

This is the difference in vote totals between the last two continuing candidates in the last round. It is an upper bound for the margin of victory because adding that many ballots for the non-winning candidate in that round would create a tie and hence change the possible outcome. As a result:

$$MoV \leq MoVUB1 \quad (4)$$

MoVUB1 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.

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 MoV. If an upper bound and lower bound are relatively far apart, it might be worthwhile to invest in some additional computational effort to try to find a better, higher lower bound that may allow significant reductions in audit sample sizes. On the other hand, if an upper and lower bound are relatively close to each other, any improvement in the lower bound might be expected to have negligible impact on the audit sample size.

A Lower Bound for the Margin of Victory:

The **first lower bound for the margin of victory, MoVLB1,** is defined as being the smallest margin of single elimination:

$$\mathbf{MoVLB1} = \min \{ MoSE(k) : k = 1, \ldots , C-1\} \quad (5)$$

The reason MoVLB1 is a lower bound for MoV is because in order to change the winner of an IRV election, there has to be a change of which candidate is eliminated in at least one of the rounds. If fewer than MovLB1 ballots are added or removed, then none of the margins of single elimination can be reduced to zero, and none of the eliminations can change.

$$\text{MoVLB1} \leq \text{MoV} \tag{6}$$

MoVLB1 can be calculated in O(C) time and constant space, if either the vote totals VT(r, c) and candidate order CO(r, k) or just the vote total order VTO(r, k) are given in an appropriate form, for example in arrays. MoVLB1 can be calculated in $O(C^2)$ time and constant space, if only the vote totals VT(r, c) are given in an appropriate form, for example in arrays.

A Second Lower Bound for the Margin of Victory

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 have to be eliminated before any other continuing candidates. In such a situation the order of elimination of those candidates among themselves can not change the outcome of the election.

For example, consider the end of a round where the lowest vote totals are 20, 30, 35, and 200. The margin of single elimination for this round is 30 – 20 = 10, so MoVLB1 for the election as a whole has to be less than or equal to 10. However, these 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 this case, the small margin of single elimination, 10, can be ignored in favor of the larger margin of multiple elimination, 115, when estimating a lower bound for the MoV.

Define the **margin of multiple elimination** in round r to simultaneously eliminate k candidates to be:

$$\mathbf{MoME}(\mathbf{r},\mathbf{k}) = \text{VTO}(r, k+1) - \sum_{i=1}^{k} \text{VTO}(r, i) \quad \text{for } k = 1, \ldots, C-r \tag{7}$$

The margin of multiple elimination is the sum of the k lowest vote totals, the votes for the candidates to be eliminated, subtracted from the next higher vote total.

[When considering multiple eliminations, the single elimination notation for designating rounds will still be used. This numbering of rounds also reflects the number of continuing candidates. With this notation, eliminating k candidates in a round has the effect of notationally advancing k single elimination rounds for the next 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 allowed whenever MoME(r, k) > 0. However, single eliminations are allowed in the additional case when MoME(r, 1) = 0. Multiple eliminations are not allowed for k ≥ 2, when MoME(r, k) ≤ 0.

When there are several multiple eliminations that are allowed in a round, an IRV algorithm might choose the one that eliminates the most candidates. However, when estimating the MoV, it is advantageous to select the multiple elimination that imposes the least constraint on the lower bound estimate. To determine which multiple eliminations in a round to use, 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.

To do that, define the most restrictive margin of elimination for the best elimination sequence for the rest of the tally, or more succinctly, the **margin of best elimination, MoBE(r),** recursively in reverse round order, starting with r = C:

$$\text{MoBE}(C) = \text{total number of ballots} + 1 \tag{8}$$

For any round r, $1 \le r \le C-1$, define:

$$\text{MoBE}(r) = \max\{ \min(\text{MoME}(r, k), \text{MoBE}(r+k)) : k = 1, \dots , C-r \} \tag{9}$$

If k candidates were simultaneously eliminated in round r, that would impose a constraint of MoME(r, k) for that round and a constraint of MoBE(r+k) for all subsequent rounds, with a combined constraint being the minimum of the two, min(MoME(r, k), MoBE(r+k)).

Since MoME(r, 1) ≥ 0, by induction it is true for all rounds that:

$$\text{MoBE}(r) \ge 0 \tag{10}$$

In each round r, there is always at least one value of k for which doing a multiple elimination of k candidates is allowed and MoME(r, k) ≥ MoBE(r). To show this, there are two cases to consider:

<u>Case 1:</u> MoBE(r) > 0.

By the definition of MoBE(r), there is at least one value of k for which MoBE(r) = min(MoME(r, k), MoBE(r+k)) and so MoME(r, k) ≥ MoBE(r). Let k be one of those values. Then MoME(r, k) ≥ MoBE(r) > 0, so doing a multiple elimination of k candidates is allowed.

<u>Case 2:</u> MoBE(r) = 0.

Then MoBE(r, 1) = MoSE(r) ≥ 0 = MoBE(r) and doing a single elimination is always allowed.

As a result, an election tally starting from round r can always be completed using allowed multiple eliminations, each with a margin of elimination greater than or equal to MoBE(r). The next section, "Tracking Information for MoVLB2", describes how to identify one such sequence of multiple eliminations.

It is also true for any rounds r and t that:

$$\text{MoBE}(r) \le \text{MoBE}(t) \quad \text{if } r \le t \tag{11}$$

The **second lower bound for the margin of victory, MoVLB2,** is the culmination of this calculation in reverse round order**:**

$$\textbf{MoVLB2} = \text{MoBE}(1) \tag{12}$$

Because the election can be tallied from the beginning by only doing eliminations with margins greater than or equal to MoVLB2 = MoBE(1), any combination of adding and removing less than MoVLB2 ballots can not change those eliminations, so MoVLB2 really is a lower bound for MoV:

$$\text{MoVLB2} \leq \text{MoV} \tag{13}$$

Since MoME(r,1) = MoSE(r), we have inductively on the reverse round order:

$$\min \{ \text{MoSE}(k) : k = r, \ldots, C\text{-}1\} \leq \text{MoBE}(r) \tag{14}$$

and so:

$$\text{MoVLB1} \leq \text{MoVLB2} \tag{15}$$

If there are no opportunities for simultaneous elimination of two or more candidates, then MoME(r, k) ≤ 0 for all r and for all k ≥ 2, and both of the above inequalities become equalities.

MoVLB2 can be calculated in $O(C^2)$ time and O(C) space, if the vote totals VT(r, c) and order of vote totals VTO(r, k) are given in an appropriate form, for example in arrays. MoVLB2 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.

Tracking Information for MoVLB2

There is some tracking information about the calculation of MoVLB2 and MoBE(r) that will record exactly which multiple elimination choices can be used to tally the IRV election while assuring that all margins of elimination are at least as big as MoVLB2. The tracking information is specified by the following definitions.

BestK(r) is the number of candidates to simultaneously eliminate in round r and still satisfy MoBE(r), preferring to eliminate more when there is a choice:

$$\textbf{BestK(r)} = \max \{ k : k = 1, \ldots, C - r \text{ and } (k = 1 \text{ or } \text{MoME}(r, k) > 0) \text{ and } \text{MoBE}(r) = \min(\text{MoME}(r, k), \text{MoBE}(r+k)) \} \tag{16}$$

SourceR(r) is the earliest round that is the source of the value for MOBE(r):

$$\textbf{SourceR}(\textbf{r}) = \begin{cases} \text{r} & \text{if MoME}(\text{r}, \text{BestK}(\text{r})) \leq \text{MoBE}(\text{r}+\text{BestK}(\text{r})) \\ \text{SourceR}(\text{r}+\text{BestK}(\text{r})) & \text{otherwise} \end{cases} \tag{17}$$

The function SourceR is idempotent, meaning for any round r:

$$\text{SourceR(SourceR(r)) = SourceR(r)} \quad (18)$$

It is also true that for any round r:

$$\text{MoBE(r) = MoME(SourceR(r), BestK(SourceR(r)))} \quad (19)$$

This means that the most restrictive margin of elimination in the best elimination sequence starting in round r happens in round SourceR(r) when eliminating BestK(Source(r)) candidates in that round. It also follows then, that:

$$\text{MoVLB2 = MoME(SourceR(1), BestK(SourceR(1)))} \quad (20)$$

With this information a MoVLB2-consistent sequence of eliminations can be reconstructed using an effective-round function, **ER(s)**, that designates what the corresponding single-elimination round is for the actual round s of a tally that may involve multiple eliminations per round (here s does not conform to the single-elimination round notation, but ER(s) does):

Initialize the effective round ER(1) = 1.
Starting with s=1, conduct round s as follows:
- Eliminate BestK(ER(s)) candidates, with a margin of multiple elimination equal to MoME(ER(s), BestK(ER(s)).
- Stop, if ER(s) = C-1.
- Set ER(s+1) = ER(s) + BestK(ER(s)).
- Increment s.
- Repeat for the next round.

Conclusions

Three estimates for the margin of victory for an IRV election have been defined and shown to satisfy the relationship:

$$\text{MovLB1} \leq \text{MoVLB2} \leq \text{MoV} \leq \text{MoVUB1} \quad (21)$$

These estimates can be efficiently calculated from election tally summary totals, without looking at individual ballots.

When adapting post-election auditing procedures from plurality elections to IRV elections, MoVLB2 is recommended as a starting point for use as an estimate of the margin of victory.

There may on occasion be IRV elections where MoVLB2 does not estimate the MoV adequately and its use would result in unnecessarily large audit sample sizes. In those situations, additional techniques can be used to refine the lower bound estimates of the margin of victory. By considering the actual ballots cast and avoiding the extreme-case assumptions associated with MoVLB2, there may be significant opportunities to raise these lower bound estimates, resulting in reduced audit sample sizes. Part 2 of this paper will describe some of those techniques.