

# Ballot Length in Instant Runoff Voting

Kiran Tomlinson<sup>1</sup>, Johan Ugander<sup>2</sup>, Jon Kleinberg<sup>1</sup>

<sup>1</sup>Cornell University

<sup>2</sup>Stanford University

kt@cs.cornell.edu, jugander@stanford.edu, kleinberg@cornell.edu

## Abstract

Instant runoff voting (IRV) is an increasingly-popular alternative to traditional plurality voting in which voters submit rankings over the candidates rather than single votes. In practice, elections using IRV often restrict the *ballot length*, the number of candidates a voter is allowed to rank on their ballot. We theoretically and empirically analyze how ballot length can influence the outcome of an election, given fixed voter preferences. We show that there exist preference profiles over  $k$  candidates such that up to  $k - 1$  different candidates win at different ballot lengths. We derive exact lower bounds on the number of voters required for such profiles and provide a construction matching the lower bound for unrestricted voter preferences. Additionally, we characterize which sequences of winners are possible over ballot lengths and provide explicit profile constructions achieving any feasible winner sequence. We also examine how classic preference restrictions influence our results—for instance, single-peakedness makes  $k - 1$  different winners impossible but still allows at least  $\Omega(\sqrt{k})$ . Finally, we analyze a collection of 168 real-world elections, where we truncate rankings to simulate shorter ballots. We find that shorter ballots could have changed the outcome in one quarter of these elections. Our results highlight ballot length as a consequential degree of freedom in the design of IRV elections.

## Introduction

Instant runoff voting (IRV) has grown in popularity over the last two decades as an alternative to plurality voting for governmental and organizational elections. Also referred to as ranked choice voting (RCV), single transferrable vote (STV), alternative vote, preferential voting, or the Hare method, IRV allows voters to submit rankings over the candidates rather than voting for a single option. IRV determines a winner from these rankings by repeatedly eliminating the candidate who has the fewest ballots ranking them first; the ballots that listed this eliminated candidate first have their votes reallocated to the next candidate on their list. This process continues, repeatedly eliminating candidates, until only one is left—the winner.

Proponents of IRV argue that it allows voters to report their full preferences, mitigates vote-splitting when similar candidates run, encourages civility in campaigning, and

saves money compared to holding separate runoff elections (FairVote 2022; Lewyn 2012). Many local elections in the United States use IRV, including in Minneapolis, San Francisco, Oakland, Santa Fe, and New York City, as well as statewide elections in Maine and Alaska. IRV is also used in other countries, including Australia and Ireland.

However, IRV has vocal opponents who believe it to be too confusing for voters (Langan 2004; Saltsman and Paxton 2021), leading to outright bans on the use of IRV in Florida (Florida Legislature 2022) and Tennessee (Tennessee Legislature 2022). One particular issue critics point to is the complexity of a ballot that asks voters to rank every candidate, especially when the number of candidates is large. One official tasked with running Utah’s first IRV election raised this as her primary concern after the election:

My concerns with the current RCV law are that we would recommend the number of rankings be limited to three or five instead of an unlimited number based on the number of candidates. So although you can list as many candidates as file on the ballot, I think it is a bit confusing to voters [...] For instance, in Minneapolis they rank three. In St. Paul, they rank five. They don’t usually have them rank as many candidates as there are. (Swensen 2021, Salt Lake County Clerk)

Indeed, many municipalities have different numbers of ranking slots on their IRV ballots, what we call *ballot length*: Oakland uses three, Alaska four, and New York City five. The count goes on: ballot length six would have been mandated by the failed 2019 Ranked Choice Voting Act proposing IRV for US Congressional elections (US Congress 2019). In Maine, voters can rank all of the candidates—even if there are 15 of them. In fact, plurality voting can be viewed as IRV with ballot length one: losing candidates are repeatedly “eliminated” (without redistribution) until the candidate with a plurality is declared the winner.

While making ballots shorter does make them simpler, it also strays from a goal of IRV: allowing voters to express their complete preferences over the candidates. Critics of IRV also raise concerns about *ballot exhaustion* during the IRV algorithm, where all candidates ranked by a voter have been eliminated and that vote no longer contributes to subsequent tallies (Burnett and Kogan 2015).<sup>1</sup> Ballot length is

<sup>1</sup>In plurality, any vote not cast for the winner is “exhausted.”

therefore subject to competing desires: shorter ballots are easier to fill out and simpler to print, but less informative about voter preferences.

Despite the apparent trade-offs involved in ballot length, there has been very little investigation of how these trade-offs might work. As noted above, plurality voting can be seen as IRV with ballot length one, and so the fact that plurality and IRV can produce different outcomes already indicates that ballot length can have important consequences. But aside from early work looking at simulations and a few real-world elections (Kilgour, Grégoire, and Foley 2020; Ayadi et al. 2019) we do not have much insight into the consequences of ballot length more generally. Perhaps, for example, there are underlying structural properties to be discovered that constrain how many winners are possible as we vary the ballot length. Or perhaps “anything goes,” and if we specify which candidate we’d like to see win at each possible ballot length, we can construct a fixed set of rankings that produce each desired winner at the corresponding length.

**Overview of Results.** In this paper, we show that the effect of ballot length essentially behaves like the latter extreme, where almost every sequence of outcomes is possible. In particular, we prove that modulo a simple feasibility constraint, it is possible to pick any sequence of candidates (with repetitions allowed), and to have this be the sequence of winners at ballot lengths 1, 2, 3, .... For example, there are voter preferences such that one candidate wins if the election is run with odd ballot length and another wins with even ballot length. We make a central assumption that voters have fixed ideal rankings and report as long a prefix of their ideal ranking as the ballot allows. Given  $k$  candidates, we show that up to  $k - 1$  of them can win as the ballot length varies from 1, . . . ,  $k - 1$  and voter preferences remain fixed. Moreover, we establish exact matching lower bounds on the number of voters required to produce  $k - 1$  distinct winners.

We also consider how these results are affected if we make standard modeling assumptions about voters. If we model voters abstractly as exhibiting single-peaked or single-crossing preferences, we prove that  $k - 1$  distinct winners across ballot lengths cannot be achieved. We also consider voters who rank candidates according to a shared one-dimensional ideological spectrum; since such voters are both single-peaked and single-crossing, there cannot be  $k - 1$  distinct winners in these cases. We find through simulation that in this one-dimensional case, ballot lengths above  $k/2$  almost always produce the same winner as full IRV ballots.

Finally, we use data from 168 real-world elections from PrefLib (Mattei and Walsh 2013) (most of them originally conducted using IRV), and we find that different winners across ballot lengths is a phenomenon that occurs commonly: in 25% of the PrefLib elections at least two different candidates win as the ballot length is varied by truncation. However, truly pathological cases with  $k - 1$  winners appear to be extremely rare: we observe at most three distinct winners across ballot lengths, and that occurs only once in the 168 PrefLib elections. But even with these real-world voter preferences, more than three winners can occur; by resampling ballots in the PrefLib elections, we ob-

serve cases with four, five, and even six different winners across ballot lengths. We note that one third of the elections initially used ballot length of at most four, where it is impossible to have more than three different winners across ballot lengths. An extended version of the paper including all proofs omitted for space is available online (Tomlinson, Ugander, and Kleinberg 2022). Our code and data are available at <https://github.com/tomlinsonk/irv-ballot-length>.

## Related Work

There has been considerable work on what happens when individual voters choose not to rank all the candidates—a practice sometimes called *voluntary truncation*—in contrast with *forced truncation* (i.e., ballot length restrictions) (Kilgour, Grégoire, and Foley 2020). In many voting systems including IRV, election outcomes can change dramatically as voters independently choose to rank more or fewer candidates (Saari and Van Newenhizen 1988). This matter has been studied from a computational angle as the *possible winners* problem, which asks, given a collection of partial ballots, which candidates could become winners as those ballots are filled out (Konczak and Lang 2005; Chevaleyre et al. 2010; Baumeister et al. 2012; Xia and Conitzer 2011; Ayadi et al. 2019). There is also a wide array of research on how partial ballots can be used for strategic voting and campaigning (Baumeister et al. 2012; Narodytska and Walsh 2014; Menon and Larson 2017; Kamwa 2022; Fishburn and Brams 1984). On the empirical side, voluntary truncation is a concern since it can lead to ballot exhaustion (Burnett and Kogan 2015). In political science, voluntary truncation is also referred to as *under-voting* (Neely and Cook 2008). Several studies have asked whether different demographic groups are more likely to under-vote and how this could have a disenfranchising effect (Neely and Cook 2008; Coll 2021; Hoffman et al. 2021). There has also been research on “over-voting” in IRV, which refers to ranking a single candidate in more than one position (e.g., first and second), especially its correlation with underrepresented voting populations (Neely and Cook 2008; Neely and McDaniel 2015).

In contrast, we investigate what happens when all voter preferences are truncated as a result of ballot length. That is, we focus on a question of election design rather than on voter choice. In this direction, Ayadi et al. (2019) investigated how often IRV with short ballots produces the full-ballot winner in the Mallows model and in five PrefLib elections. However, all five PrefLib elections they studied produced the full-ballot winner at all ballot lengths—in analyzing a larger collection of 168 PrefLib elections, we find multiple winners across ballot lengths in 25% of them. Ayadi et al. also examined several other interesting facets of IRV ballot length, including a low-communication IRV protocol (a form of online, per-voter ballot length customization) and the complexity of the possible winners problem under truncated ballots. The issue of ballot length in IRV was also touched on by Kilgour, Grégoire, and Foley (2020), who examined its effect in simulation for  $k = 4, 5$ , and 6 candidates, where they found up to  $k - 2$  distinct winners across ballot lengths. We prove that in fact  $k - 1$  winners are possible for all  $k \geq 3$ . Ballot length has been considered in con-

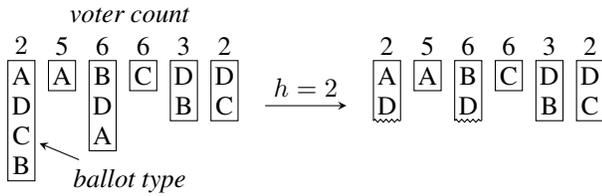


Figure 1: On the left, an example profile with  $k = 4$  candidates A, B, C, D and  $n = 24$  voters of 6 types with partial ballots. Ballots are listed top-down, with the number of voters of each type above each ballot. On the right, the profile is truncated to ballot length  $h = 2$ .

texts other than IRV—for instance, research on the Boston school choice mechanism found that limiting the number of schools parents could rank to five resulted in undesirable strategic behavior (Abdulkadiroglu et al. 2006). There has also been research on ballot length in approval voting from a learning theory angle, seeking to recover a population’s preferences efficiently (Garg et al. 2019).

### Preliminaries

An IRV election consists of  $k$  candidates labeled  $1, \dots, k$  and  $n$  voters. Each voter  $j$  has a preference ordering over a subset of the candidates denoted by the ordered subset  $\pi_j$ , which we refer to as a *ballot*. At any point down the ballot,  $\pi_j$  can terminate, at which point the voter is indifferent over the remaining options. If  $\pi_j$  includes all candidates, we call it *full*, otherwise we call it *partial*. We call a collection of ballots a *profile*. Unless otherwise specified, a profile may contain partial ballots.<sup>2</sup> If multiple voters have identical ballots, we say they are of the same *type*. Given a profile, IRV proceeds by eliminating the candidate with the fewest ballots ranking them first and removing them from all ballots. Ballots that have all their candidates eliminated are *exhausted*. Eliminations continue until only one candidate remains, who is declared the winner (equivalently, one can terminate when one candidate has the majority of votes from non-exhausted ballots). Ties can be broken as desired (for instance, by coin-flip), although they are unlikely in large elections.

In many real-world elections, the number of candidates a voter can rank is limited to  $h < k$ , which we call the *ballot length*. We assume that if the ballot length is  $h$ , voters submit the length  $h$  prefix  $\pi_j(1, \dots, h)$  of their ideal ballot  $\pi_j$ . Voters who would have submitted a ranking listing  $h$  or fewer candidates are unaffected. Thus, we say that ballots are *truncated* to the ballot length  $h$ . See Figure 1 for an example of a profile with partial ballots truncated to  $h = 2$ . Note that there is no difference between running IRV with ballot length  $k$  and  $k - 1$ , since only one candidate remains after the  $(k - 1)$ th elimination.

The main question we focus on is how ballot length affects an election. For instance, how many different candidates can win as the ballot length varies for a fixed profile? In order to address this question, we make some assumptions

<sup>2</sup>All 168 elections in the PrefLib data have partial ballots.

about the lack of consequential ties, since in trivial cases such as zero voters, any candidate can win depending on tie-breaks. We say that a profile is *consequential-tie-free* if tie-breaks do not affect the winner under any ballot length  $h$ . We say it is *elimination-tie-free* if a tie for last place never occurs when running IRV for any ballot length  $h$ . Finally, we say it is *tie-free* if no two candidates ever have tied vote counts when running IRV at any ballot length  $h$ . We note that the problem of determining if a given candidate could win under some tie-breaking sequence is known to be NP-complete (Conitzer, Rognlie, and Xia 2009).

### Worst-Case Analysis of Ballot Truncation

We say a profile has  $c$  *truncation winners* if  $c$  different candidates can win depending on the ballot length. Previous simulation work found up to  $k - 2$  truncation winners for  $k = 4, 5$ , and 6 (Kilgour, Grégoire, and Foley 2020). One of our main results is that up to  $k - 1$  truncation winners are possible for any  $k$ . We note that it is impossible to have all  $k$  candidates win under different ballot lengths, since lengths  $k$  and  $k - 1$  behave the same way. All proofs omitted for space can be found in the extended version of the paper (Tomlinson, Ugander, and Kleinberg 2022).

First, we establish an exact lower bound on the number of voters required in order to achieve  $k - 1$  truncation winners in consequential-tie-free profiles. Our voter lower bound is based on the observation that the winner at  $h = 1$  (the plurality winner) must be eliminated second under ballot lengths  $\geq 2$  for  $k - 1$  truncation winners to occur. In order for the plurality winner to be eliminated second, the first elimination must redistribute enough votes for every other candidate to overtake the plurality winner.

**Theorem 1.** *For any  $k > 3$ , a consequential-tie-free profile requires at least  $2k^2 - 2k$  voters in order to produce  $k - 1$  truncation winners. For  $k = 3$ , the lower bound is  $k^2 = 9$ .*

Our main theoretical result is a construction matching this lower bound, showing that  $k - 1$  truncation winners can occur for any  $k \geq 3$ . Our construction can not only produce  $k - 1$  truncation winners, but *any* sequence of winners over ballot lengths  $1, \dots, k - 1$ , provided that a candidate has not yet been eliminated.

**Theorem 2.** *Let there be  $k > 3$  candidates, labelled  $1, \dots, k$  in their full-ballot IRV elimination order. Fix any sequence of candidates  $w_1, \dots, w_{k-1}$  such that  $w_h \in \{h + 1, \dots, k\}$  for all  $h \in [k - 1]$ . There exists a consequential-tie-free profile with  $2k^2 - 2k$  partial ballots whose sequence of truncated IRV winners from  $h = 1, \dots, k - 1$  is  $w_1, \dots, w_{k-1}$ . For  $k = 3$ , such a profile exists with 9 ballots. Any sequence where  $w_h \leq h$  for some  $h \in [k - 1]$  is impossible to realize as the sequence of truncated IRV winners for any consequential-tie-free profile.*

The idea behind the construction is to maintain a tie for second place among all candidates but two: the candidate about to be eliminated, in last, and the candidate next in the winner sequence, in first. Each elimination redistributes ballots to move the next candidates into first and last place. By carefully designing ballots, they become exhausted at just

the right moment to freeze the order once we reach step  $h$  of IRV, causing the candidate currently in first to win. The example in Figure 1 uses this construction for  $k = 4$  to achieve different winners at ballot lengths 1, 2, 3 (namely, A, B, C). Note that the full-ballot elimination order labeling of candidates A, B, C, D is 2, 3, 4, 1, which makes the truncation winner sequence 2, 3, 4 feasible. In contrast, the sequence 2, 2, 4 would not be feasible since the candidate eliminated second under full ballots cannot win at ballot length 2. Intuitively, a winner sequence with elimination order labeling is feasible if it is element-wise at least 2, 3,  $\dots$ ,  $k$ .

### Restrictions on Profiles

Since IRV can behave very erratically across ballot lengths for general profiles, we might hope that imposing restrictions on the space of profiles makes IRV more well-behaved. We consider three classic profile restrictions from voting theory, single-peaked (Black 1948; Arrow 1951), single-crossing (Gans and Smart 1996), and 1-Euclidean preferences (see (Elkind, Lackner, and Peters 2022) for a survey of preference restrictions). A profile is *single-peaked* if there exists an order  $<$  over the candidates such that, for every ballot  $b$  ranking  $i$  first, if  $j < k < i$  or  $i < k < j$ , then  $j$  is not ranked above  $k$  in  $b$ . A profile is *single-crossing* if there exists an ordering  $L$  of the ballots such that for every ordered pair of candidates  $(i, j)$ , the set of ballots ranking  $i$  above  $j$  forms an interval of  $L$ . Finally, a profile is *1-Euclidean* if there exist embeddings of the voters and candidates in  $[0, 1]$  such that if voter  $b$  is closer to candidate  $i$  than to candidate  $j$ , then voter  $b$  ranks  $i$  above  $j$ .

Intuitively, single-peaked profiles arise when there is a political axis arranging candidates from left to right and voters prefer candidates closer to their ideal point on the axis (each voter can have their own ideal point). Single-crossing preferences arise when voters are arranged on an ideological axis and each candidate is most appealing to voters at a certain point on this axis. While the definitions appear similar, neither condition implies the other. 1-Euclidean profiles are both single-peaked and single-crossing—but there are profiles that are both single-peaked and single-crossing, but not 1-Euclidean (Elkind, Faliszewski, and Skowron 2014).

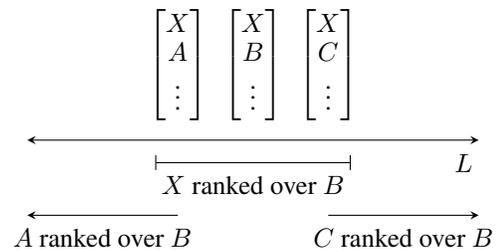
In contrast to general profiles, where  $k - 1$  truncation winners can occur, we show that such cases are impossible under either single-peaked or single-crossing preferences (and therefore 1-Euclidean profiles).

**Theorem 3.** *With  $k \geq 5$  candidates, no consequential-tie-free single-peaked profile has  $k - 1$  truncation winners.*

*Proof.* Suppose for a contradiction that a single-peaked profile has  $k - 1$  truncation winners ( $k \geq 5$ ). We know the candidate eliminated first cannot win under any ballot length. In order for the candidate eliminated second ( $h \geq 2$ ) to win at some ballot length, it must be at  $h = 1$ —i.e., the plurality winner must be eliminated second under  $h \geq 2$ . Thus, they must be overtaken by at least three candidates (for  $k \geq 5$ ) when the first eliminated candidate  $X$ 's ballots are redistributed. But the second place on ballots listing  $X$  first can only be the candidate to the left or right of  $X$  in the single-peaked ordering, making this impossible.  $\square$

**Theorem 4.** *With  $k \geq 5$  candidates, no consequential-tie-free single-crossing profile has  $k - 1$  truncation winners.*

*Proof.* As in the proof of Theorem 3, we'll show that the first candidate eliminated,  $X$ , can only redistribute ballots to two candidates. Suppose for a contradiction that they redistribute ballots to at least three candidates. Call these candidates  $A, B$ , and  $C$  in the order in which they first appear as second choices in the ballots ranking  $X$  first in the single-crossing order  $L$ . By the single-crossing property, all ballots to the left of ballots starting  $X$ ,  $A$  must rank  $A$  above  $B$ , since a ballot to its right ranks  $B$  above  $A$ , namely those starting  $X, B$ . Moreover, all ballots to the right of ballots starting  $X, C$  must rank  $C$  above  $B$  by symmetric reasoning. But this means  $B$  cannot have any ballots ranking them first, contradicting that  $X$  (who does have ballots ranking them first) is the first eliminated. See below for a visual depiction of this argument:



$\square$

Although the upper bound on truncation winners is strictly lower for single-peaked profiles than for general profiles, the number of achievable truncation winners still grows with  $k$ . In particular, we can show that  $\Omega(\sqrt{k})$  truncation winners are possible in a consequential-tie-free single-peaked profile with  $\Theta(k)$  voters.

**Theorem 5.** *With  $k = \kappa(\kappa + 1)/2$  candidates ( $\kappa \geq 3$ ), there is a single-peaked consequential-tie-free profile with  $3\kappa(\kappa + 1)/2$  partial ballots that results in  $\kappa$  distinct truncation winners.*

The exact upper bound on the number of truncation winners for single-peaked (and single-crossing) preferences remains an open question—it could be as large as  $k - 2$ . Additionally, we do not know a non-trivial lower bound on the number of achievable truncation winners for single-crossing or 1-Euclidean profiles.

### Restrictions on Ties

Since our main theorem allows ties (albeit only ties that do not affect the winners), one might be concerned that the large number of truncation winners is a byproduct of these ties. In the following results, we show that even if no vote counts are ever tied, there can still be arbitrary truncation winner sequences. We can therefore get any feasible winner sequence regardless of the tiebreaking rule. As before, we start by establishing lower bounds on the number of voters required for  $k - 1$  truncation winners and then provide a matching construction for tie-free profiles achieving any truncation winner sequence.

**Theorem 6.** For any  $k \geq 3$ , an elimination-tie-free profile must contain at least  $(k^3 - 3k)/2$  voters in order to produce  $k - 1$  truncation winners.

**Theorem 7.** For any  $k \geq 3$ , a tie-free profile must contain at least  $(2k^3 - 5k^2 + 3k)/2$  voters in order to produce  $k - 1$  truncation winners.

Note that for consequential-tie-free profiles, the lower bound on voters for  $k - 1$  truncation winners is  $\Omega(k^2)$ , but  $\Omega(k^3)$  for elimination-tie-free and tie-free profiles.

**Theorem 8.** Given the same setup as in Theorem 2, there exists a tie-free profile with  $(2k^3 - 5k^2 + 3k)/2$  ballots whose sequence of truncated IRV winners from  $h = 1, \dots, k - 1$  is  $w_1, \dots, w_{k-1}$ .

The constructions for consequential-tie-free and tie-free profiles both use  $\Theta(k^2)$  distinct ballots. However, only  $\Theta(k)$  distinct ballots are required to produce  $k - 1$  truncation winners. This is asymptotically tight, since each candidate who wins at some ballot length needs at least one ballot type listing them first.

**Theorem 9.** Given  $k > 3$  candidates, there is a tie-free profile producing  $k - 1$  truncation winners with  $\Theta(k^3)$  voters of  $\Theta(k)$  types.

## Full Ballots

So far, all of our constructions have relied on partial ballots. For profiles with full ballots, a simple extension of our constructions using filler candidates allows us to achieve up to  $k/2$  truncation winners, and in fact any feasible sequence of winners in the first half of ballot lengths.

**Corollary 1.** Let  $k = 2\kappa$  for some  $\kappa > 3$ . Label the candidates  $1, \dots, 2\kappa$  in order of their elimination under full ballots. Fix any sequence  $w_1, \dots, w_{\kappa-1}$  such that  $w_h \in \{\kappa + h + 1, \dots, 2\kappa\}$  for all  $h \in [\kappa - 1]$ . There exists a full-ballot consequential-tie-free profile with  $2\kappa^2 - 2\kappa$  voters and a full-ballot tie-free profile with  $(2\kappa^3 - 5\kappa^2 + 3\kappa)/2 + \kappa(\kappa - 1)/2$  voters whose sequences of truncation winners from  $h = 1, \dots, \kappa - 1$  are  $w_1, \dots, w_{\kappa-1}$ .

While we have not found a general construction with full ballots and  $k - 1$  truncation winners, we have found full-ballot elimination-tie-free profiles with  $k - 1$  truncation winners up to  $k = 10$  using a linear-programming-based search (described at the end of this section). Full ballots make intuitive constructions more challenging, but do not appear to prevent a large number of truncation winners. However, how a full ballot requirement does or doesn't change our main result remains an open question.

If instead of requiring ballots to be full, we require them to all have length at least  $k/2 - c$ , we can improve the above extension of our constructions and get an additional  $c$  ballot lengths at which we can specify the winner.

**Corollary 2.** Let  $k = 2\kappa$  for some  $\kappa > 3$ . Suppose we require ballots to have length at least  $\kappa - c$  for  $c < \kappa$ . Label the candidates  $1, \dots, 2\kappa$  in order of the elimination under full ballots. Fix any sequence  $w_1, \dots, w_{\kappa+c-1}$  such that  $w_h \in \{\kappa - c + h + 1, \dots, 2\kappa\}$  for all  $h \in [\kappa - 1]$ . There exists a consequential-tie-free profile with  $2\kappa^2 - 2\kappa$  ballots and a

$k$	# trunc. winners	ballot types	voters	voters lower bound (Theorem 6)
4	3	7	29	26
5	4	12	55	55
6	5	23	99	99
7	6	36	161	161
8	7	57	974	244
9	8	85	1759	351
10	9	122	4855	485

Table 1: LP full-ballot constructions. We used different search strategies for  $k \leq 7$  and  $k \geq 8$ , leading to profiles farther from the voter lower bound for  $k \geq 8$ .

tie-free profile with  $(2\kappa^3 - 5\kappa^2 + 3\kappa)/2 + (\kappa - c)(\kappa - c - 1)/2$  ballots whose sequence of truncated IRV winners from  $h = 1, \dots, \kappa + c - 1$  is  $w_1, \dots, w_{\kappa+c-1}$ .

Given that explicit full-ballot constructions appear quite challenging, we turn to a computational approach to investigate whether full-ballot profiles can produce  $k - 1$  truncation winners. Using a linear programming (LP) search, we identified elimination-tie-free profiles with full ballots and  $k - 1$  truncation winners for  $k = 4, 5, 6, 7, 8, 9, 10$  (the sizes of these profiles are shown in Table 1). Moreover, this approach was able to find instances with voter counts matching the exact lower bound in Theorem 6 for  $k = 5, 6, 7$ . Our approach was not able to match the voter lower bound for  $k = 4$ . For  $k \geq 8$ , we faced runtime constraints since the number of variables is exponential in the number of candidates, leading us to restrict the search space (described in further detail below). We consider elimination-tie-free profiles since they are easiest to encode as an LP, where we use constraints to enforce unambiguous eliminations.

The idea behind the search is to construct possible elimination orders across all  $h$  that could result in  $k - 1$  winners, express these as conditions on the sums of counts of every full ballot type in  $S_k$  (the set of permutations on  $k$  elements), and then use an LP to find a feasible real-valued solution of ballot type counts that result in that elimination order. We round these fractional ballot counts to be integers and check if the resulting profile has the desired elimination order. If not, we can try another possible elimination order or increase the gaps in the constraints so that rounding is less likely to make a solution infeasible. For  $k \leq 7$ , we tested all possible elimination orders, but only tested a single elimination order for  $k \geq 8$  due to runtime constraints. The exact LP formulation can be found in the extended version (Tomlinson, Ugander, and Kleinberg 2022).

## Ballot Length in Simulation

Our theoretical results show that the winner of an IRV election can change dramatically as the ballot length varies. Here, we ask how likely these changes are through simulated profiles. Such simulation analysis was previously conducted for  $k = 4, 5, 6$  (Kilgour, Grégoire, and Foley 2020). We extend these simulations up to  $k = 40$  (our real-world IRV data has examples of elections with up to  $\approx 30$  candidates).

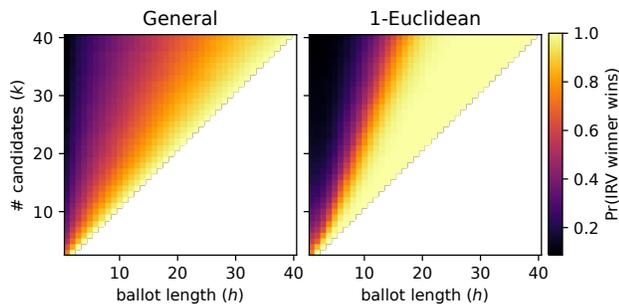


Figure 2: Probability that truncated ballots produce the full IRV winner for candidate counts  $k = 2, \dots, 40$  and ballot lengths  $h = 1, \dots, k - 1$ . (Left) For general preferences, the probability of producing the IRV winner increases smoothly with the ballot length  $h$ . (Right) For 1-Euclidean preferences, there is a sharper transition around  $h = k/2$ .

We simulate two different types of profiles: general profiles with rankings sampled uniformly at random and 1-Euclidean profiles with voters and candidates embedded in one dimension. For the general profiles, we fix 1000 voters. For 1-Euclidean profiles, we simulate an infinite voter population uniformly distributed over  $[0, 1]$ , where the number of first-place votes a candidate  $i$  has is the size of the interval of  $[0, 1]$  containing points closer to  $i$  than any other candidate.

For both general and 1-Euclidean profiles, we simulate both full and partial ballots to gauge the effect of forced truncation with and without voluntary truncation. For general profiles with partial ballots, we independently and uniformly perform voluntary truncation on each voter’s preferences before applying forced truncation in the form of ballot length. For 1-Euclidean partial ballots, we do the same with each ballot type.

In Figure 2, we show the probability that the full-ballot IRV winner is selected with each ballot length  $1, \dots, k - 1$  for  $k = 3, \dots, 40$  with initially full ballots (the heatmaps were qualitatively identical for partial ballots (Tomlinson, Ugander, and Kleinberg 2022)). For general preferences, the probability of selecting the full-ballot IRV winner increases smoothly as ballot length increases. Additionally, for any fixed ballot length, the probability of selecting the IRV winner decreases as the number of candidates increases. For instance, for  $h = 3$ , the probability of selecting the IRV winner first dips below 50% at  $k = 12$ . For 1-Euclidean preferences, small ballot lengths are even less likely to produce the full IRV winner: for  $h = 3$ , the probability first drops below 50% for  $k = 9$ . On the other hand, the probability rapidly increases around  $h = k/2$ . For ballots longer than  $k/2$ , uniform 1-Euclidean preferences almost always produce the full IRV winner.

In Figure 3, we visualize the same data in a different way. We plot the mean and maximum observed numbers of truncation winners across ballot lengths (the figure also includes PrefLib winner counts described in the next section). While the difference between general and 1-Euclidean profiles was pronounced in the previous heatmaps, they result in almost

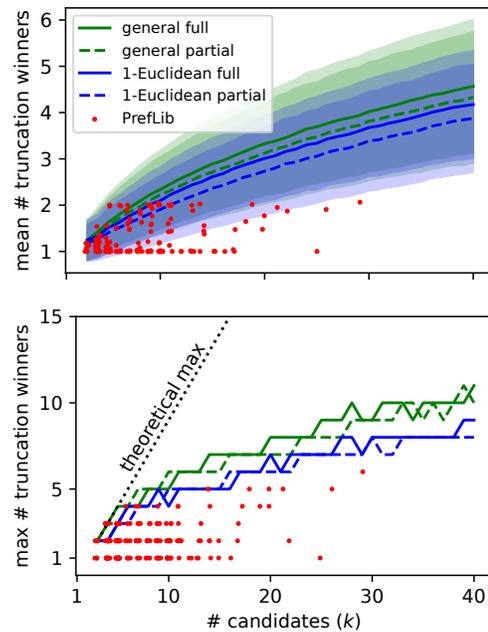


Figure 3: Mean (top) and maximum (bottom) number of truncation winners in 10000 synthetic ballot simulations and 10000 PrefLib resampling trials. We simulated uniform general and 1-Euclidean preferences for  $k = 3, \dots, 40$ . The shaded regions show standard deviation across trials. To simulate partial ballots, each ballot is voluntarily truncated at a random length between 1 and  $k$ . While up to  $k - 1$  truncation winners are possible, the mean number of truncation winners only reaches 4 around  $k = 40$ . 1-Euclidean profiles and profiles with partial ballots tend to produce slightly fewer truncation winners. For the PrefLib data, each point represents a single election, with horizontal jitter added for legibility. Real elections tend to produce even fewer truncation winners, although it is not rare to have more than 1.

the same number of truncation winners on average. Additionally, these simulated profiles tend to have a small number of truncation winners relative to the theoretical maximum. On average for  $k \leq 10$ , there are around two truncation winners, while the theoretical maximum is nine. Additionally, the maximum observed number of winners in 10000 simulated trials was well below the theoretical maximum, especially for larger  $k$ : we only began generating any profiles with 10 truncation winners around  $k = 40$ .

Intuitively, these simulation results therefore indicate that profiles with large numbers of truncation winners are very rare in the space of profiles, at least under these (uniform) measures. However, they do not appear to be significantly rarer among 1-Euclidean profiles than among general profiles, as one might have expected given the increased structure of 1-Euclidean profiles. On the other hand, profiles in which there are more than one winner across ballot lengths are very common. Thus, while truly extreme cases with  $k - 1$  truncation winners might be rare, cases where ballot length has an effect occur readily in simulation.

## Truncating Real-World Election Data

Given that many truncation winners are theoretically possible, we now ask how often multiple truncation winners occur in real-world election data. To this end, we analyze voter rankings from 168 elections in PrefLib (Mattei and Walsh 2013). This collection includes 12 American Psychological Association (APA) presidential elections (Regenwetter et al. 2007) ( $h = 5$ ), 14 San Francisco local elections ( $h = 3$ ), and 21 Glasgow local elections ( $h = k$ ), among others. The number of candidates in these elections ranges from 3–29 and the number of voters from tens to hundreds of thousands. See the extended version for additional summary statistics of these elections (Tomlinson, Ugander, and Kleinberg 2022). Some of these PrefLib datasets included a small number of ballots with multiple candidates listed at the same rank (0.5% of all ballots), which we omit.

In order to evaluate the impact of ballot length, we truncate the rankings at each possible shorter ballot length than the election actually used. We then run IRV on the truncated ballots. We assume that if ballots had been shorter, voters would have reported the same ranking, but truncated to the ballot length. It is possible that voters would express their preferences differently depending on the ballot length, so our approach should be seen as an approximation to this counterfactual scenario.

In 41/168 elections, there were two different winners across ballot lengths, and in one election, there were three different winners. Overall, 25% of the elections were sensitive to ballot length. Among the elections with ballot length  $h \leq 5$ ,  $12/85 = 14\%$  of them had two different truncation winners; for elections with  $h > 5$ ,  $29/83 = 35\%$  of elections had two or more different winners. In order to better understand the landscape of possible outcomes in each election, we also performed resampling of ballots. Given a collection of  $n$  ballots, we resample a collection of  $n$  ballots with replacement to simulate another possible election outcome with the same pool of voters. We then truncate those collections of votes to assess the impact of ballot length. In 10000 resampling trials, we observed up to six different truncation winners across the elections, but the expected number of truncation winners under resampling was between one and two for all elections (see Figure 3). In Figure 4, we also use ballot resampling to visualize the sequence of truncation winners in two PrefLib elections. The 2009 Burlington Mayoral election famously had a different plurality winner (Kurt Wright) than the elected IRV winner (Bob Kiss), but our visualization reveals that at ballot length  $h = 2$ , the election was a complete toss-up and could have gone either way with only a small change in ballot counts. In the right subplot, we visualize the sequence of truncation winners in the one PrefLib election that had three distinct truncation winners. Not only does this election have three truncation winners, but the sequence of winners flips back and forth, as we proved theoretically possible.

The smaller number of truncation winners in real data is likely due to the small number of front-runners in real-world elections, in contrast with the uniform preferences in our synthetic data. Our observations here are in line with the finding that ballot truncation is less likely to change the win-

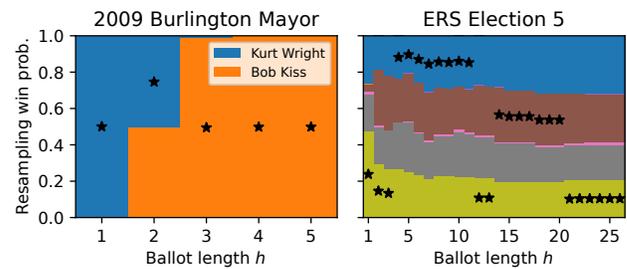


Figure 4: Two elections in the PrefLib data, the infamous 2009 Burlington mayoral election (left,  $k = 6$ ,  $n = 8974$ ) and an anonymous intra-organization election from the Electoral Reform Society (right,  $k = 26$ ,  $n = 104$ ). The stacked bars show the probability candidates have of winning at each ballot length under ballot resampling. Stars indicate the winners at each ballot length with actual ballot counts.

ner in the Mallows model when preferences are more tightly clustered around the central ranking (Ayadi et al. 2019).

## Discussion

Our theoretical results are fairly pessimistic: IRV election outcomes can change dramatically with ballot length. Our analysis of real and simulated data, on the other hand, presents a more mixed picture: ballot length regularly has an effect on the identity of the winner even in real elections, but the extreme changes between winners that are theoretically possible rarely occur, which may be cause for some degree of optimism. Nonetheless, changes in ballot length by truncation can often result in two or three different winners, even when the ballot length is short.

There are a number of open theoretical questions around ballot length. First, is it possible to achieve every feasible truncation winner sequence with complete ballots? We suspect the answer is yes, but an explicit construction has proved elusive. Second, are more than  $O(\sqrt{k})$  truncation winners possible for single-peaked ballots? How many truncation winners are possible with single-crossing ballots? Similar questions could be asked for other profile restrictions, such as 1-Euclidean preferences.

Our interest in IRV is due to its increasing popularity of IRV in United States local elections, but one could also investigate the effects of ballot length in other ranking-based voting systems such as Borda count or Copeland’s method. Additionally, we do not address what ballot length should be used in practice, which requires making a tradeoff between competing desires. Finally, it would be interesting to understand when elections are close enough for ballot length to affect the winner. There has been research on calculating the margin of victory for IRV (Sarwate, Checkoway, and Shacham 2013; Blom et al. 2016; Magrino et al. 2011), defined as the number of votes which would need to be altered to change the winner, which is NP-hard to compute (Xia 2012). A notion of margin of victory that relates to winners across different ballot lengths would be valuable.

## Acknowledgments

This work was supported in part by ARO MURI, a Simons Investigator Award, a Simons Collaboration grant, a grant from the MacArthur Foundation, the Koret Foundation, and NSF CAREER Award #2143176. We thank the anonymous reviewers for their helpful feedback.

## References

- Abdulkadiroglu, A.; Pathak, P. A.; Roth, A. E.; and Sonmez, T. 2006. Changing the Boston School Choice Mechanism. *NBER Working Paper*, (w11965).
- Arrow, K. J. 1951. *Social Choice and Individual Values*. John Wiley & Sons.
- Ayadi, M.; Amor, N.; Lang, J.; and Peters, D. 2019. Single transferable vote: Incomplete knowledge and communication issues. In *AAMAS*.
- Baumeister, D.; Faliszewski, P.; Lang, J.; and Rothe, J. 2012. Campaigns for lazy voters: truncated ballots. In *AAMAS*, 577–584.
- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy*, 56(1): 23–34.
- Blom, M.; Teague, V.; Stuckey, P. J.; and Tidhar, R. 2016. Efficient Computation of Exact IRV Margins. In *ECAI*, 480–488.
- Burnett, C. M.; and Kogan, V. 2015. Ballot (and voter) “exhaustion” under Instant Runoff Voting: An examination of four ranked-choice elections. *Electoral Studies*, 37: 41–49.
- Chevalyere, Y.; Lang, J.; Maudet, N.; and Monnot, J. 2010. Possible winners when new candidates are added: The case of scoring rules. In *AAAI*.
- Coll, J. A. 2021. Demographic disparities using ranked-choice voting? Ranking difficulty, under-voting, and the 2020 Democratic primary. *Politics and Governance*, 9(2): 293–305.
- Conitzer, V.; Rognlie, M.; and Xia, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *IJCAI*.
- Elkind, E.; Faliszewski, P.; and Skowron, P. 2014. A characterization of the single-peaked single-crossing domain. In *AAAI*.
- Elkind, E.; Lackner, M.; and Peters, D. 2022. Preference Restrictions in Computational Social Choice: A Survey. *arXiv:2205.09092*.
- FairVote. 2022. Details About Ranked Choice Voting. <https://www.fairvote.org/rcv>. Accessed: 2023-03-15.
- Fishburn, P. C.; and Brams, S. J. 1984. Manipulability of voting by sincere truncation of preferences. *Public Choice*, 44(3): 397–410.
- Florida Legislature. 2022. State of Florida Chapter No. 2022-73, Senate Bill No. 524. <http://laws.flrules.org/2022/73>. Accessed: 2023-03-15.
- Gans, J. S.; and Smart, M. 1996. Majority voting with single-crossing preferences. *Journal of Public Economics*, 59(2): 219–237.
- Garg, N.; Gelauff, L. L.; Sakshuwong, S.; and Goel, A. 2019. Who is in your top three? Optimizing learning in elections with many candidates. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, 22–31.
- Hoffman, C.; Kauba, J.; Reidy, J.; and Weighill, T. 2021. Proportionality in multi-winner RCV elections: A simulation study with ballot truncation. *Available at SSRN 3942892*.
- Kamwa, E. 2022. Scoring rules, ballot truncation, and the truncation paradox. *Public Choice*.
- Kilgour, D. M.; Grégoire, J.-C.; and Foley, A. M. 2020. The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice*, 184(1): 197–218.
- Konczak, K.; and Lang, J. 2005. Voting procedures with incomplete preferences. In *IJCAI Multidisciplinary Workshop on Advances in Preference Handling*, volume 20.
- Langan, J. P. 2004. Instant runoff voting: A cure that is likely worse than the disease. *Wm. & Mary L. Rev.*, 46: 1569.
- Lewyn, M. E. 2012. Two cheers for instant runoff voting. *Phoenix L. Rev.*, 6: 117.
- Magrino, T. R.; Rivest, R. L.; Shen, E.; and Wagner, D. 2011. Computing the Margin of Victory in IRV Elections. In *Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE)*.
- Mattei, N.; and Walsh, T. 2013. PrefLib: A library for preferences. In *ADT*, 259–270. Springer.
- Menon, V.; and Larson, K. 2017. Computational aspects of strategic behaviour in elections with top-truncated ballots. *Autonomous Agents and Multi-Agent Systems*, 31(6): 1506–1547.
- Narodytska, N.; and Walsh, T. 2014. The Computational Impact of Partial Votes on Strategic Voting. In *ECAI*, 657–662.
- Neely, F.; and Cook, C. 2008. Whose votes count? Undervotes, overvotes, and ranking in San Francisco’s instant-runoff elections. *American Politics Research*, 36(4): 530–554.
- Neely, F.; and McDaniel, J. 2015. Overvoting and the equality of voice under instant-runoff voting in San Francisco. *California Journal of Politics and Policy*, 7(4).
- Regenwetter, M.; Kim, A.; Kantor, A.; and Ho, M.-H. R. 2007. The unexpected empirical consensus among consensus methods. *Psychological Science*, 18(7): 629–635.
- Saari, D. G.; and Van Newenhizen, J. 1988. The problem of indeterminacy in approval, multiple, and truncated voting systems. *Public Choice*, 59(2): 101–120.
- Saltsman, M.; and Paxton, R. 2021. Start Spreading the News—Ranked-Choice Voting Is a Mess. *The Wall Street Journal*.
- Sarwate, A. D.; Checkoway, S.; and Shacham, H. 2013. Risk-limiting audits and the margin of victory in nonplural-elections. *Statistics, Politics, and Policy*, 4(1): 29–64.
- Swensen, S. 2021. Minutes of the Political Subdivisions Interim Committee. <https://le.utah.gov/MtgMinutes/>

publicMeetingMinutes.jsp?Com=INTPOL&meetingId=17750, 2:06:27. Accessed: 2023-03-15.

Tennessee Legislature. 2022. State of Tennessee Public Chapter No. 621, Senate Bill No. 1820. <https://publications.tnsosfiles.com/acts/112/pub/pc0621.pdf>, note=Accessed: 2023-03-15.

Tomlinson, K.; Ugander, J.; and Kleinberg, J. 2022. Ballot Length in Instant Runoff Voting. *arXiv:2207.08958*.

US Congress. 2019. H.R.4464 - Ranked Choice Voting Act. <https://www.congress.gov/bill/116th-congress/house-bill/4464>. Accessed: 2023-03-15.

Xia, L. 2012. Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 982–999.

Xia, L.; and Conitzer, V. 2011. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41: 25–67.