



ARTICLE

Votes Can Be Confidently Bought in Some Ranked Ballot Elections, and What to Do about It

Jack R. Williams¹, Samuel Baltz²  and Charles Stewart III³ 

¹Democracy Works, Brooklyn, NY, USA; ²Department of Political Science, Massachusetts Institute of Technology, Cambridge, MA, USA; ³Department of Political Science, Massachusetts Institute of Technology, Cambridge, MA, USA

Corresponding author: Samuel Baltz; Email: sbaltz@umich.edu

(Received 19 May 2023; revised 13 November 2023; accepted 11 January 2024; published online 6 May 2024)

Abstract

We show that, in some ranked ballot elections, it may be possible to violate the secret vote. There are so many ways to rank even a handful of candidates that many possible rankings might not be cast by any voter. So, a vote buyer could pay someone to rank the candidates a certain way and then use the announced election results to verify that the voter followed through. We examine the feasibility of this attack both theoretically and empirically, focusing on instant runoff voting (IRV). Although many IRV elections have few enough candidates that this scheme is not feasible, we use data from San Francisco and a proposed election rule change in Oakland to show that some important IRV elections can have large numbers of unused rankings. There is no evidence that this vote-buying scheme has ever been used. However, its existence has implications for the administration and security of IRV elections. This scheme is more feasible when more candidates can be ranked in the election and when the election results report all the ways that candidates were ranked.

Keywords: instant runoff voting; ranked choice voting; vote buying; election security; electoral systems; election fraud detection

Edited by: Lonna Atkeson

1. Introduction

A central obstacle to vote buying and voter coercion in contemporary democracies is the secret ballot. However, simply removing voters' names before announcing the votes is not always sufficient to maintain voter privacy (Adler and Hall 2013; Bernhard *et al.* 2017; Kuriwaki 2020; Kuriwaki, Lewis, and Morse 2023). When jurisdictions are weighing a change in their voting technologies, or a different way of announcing election results, a major question is whether this change could compromise voter secrecy (Castelló 2016). But rarely is this question applied to the electoral system itself. Can changing the ballot format threaten the secrecy of ballots?

Different political institutions are known to be more or less susceptible to illegitimate electoral activities like vote buying (Birch 2007; Hicken 2007), but among cases that use the secret ballot, the ballot format rarely has a *direct* bearing on the feasibility of identifying how somebody voted (Rae 1967). One reason is that, in most widely used electoral systems, voters communicate roughly the same amount of information with their votes. This is not true of ranked ballot elections,¹ and we will argue that it is an especially acute issue in instant runoff voting (IRV). IRV is a historically rare electoral system that has rapidly gained traction in several countries and has especially surged in popularity as an alternative electoral system for American elections (Santucci 2022; Tolbert and Kuznetsova 2021). IRV, which has

¹By “ranked ballot,” we mean any ordinal voting system, in which voters can provide an ordered ranking of candidates.

long been used in Australian elections, now elects federal and state representatives in two U.S. states (Alaska and Maine), and is the system of municipal elections in two of the country's largest cities (New York and San Francisco). IRV is a superficially straightforward system, in which a winner is chosen by iteratively eliminating candidates and redistributing the eliminated candidates' votes to whomever was ranked after them on each ballot. However, the intuitive simplicity of IRV masks deep complexities (Atsusaka 2023; Baltz 2022b; Xia 2012).

We will demonstrate that the combinatorial properties of IRV can make it possible to verify how someone voted with their cooperation, which in turn could make it possible to buy their vote.² The problem is that a voter may be able to arrange the candidates on their ballot in an explosively large number of different ways. When there are half a dozen candidates on the ballot, common vote-counting rules make it possible to generate thousands of ways to order those candidates, and a vote buyer could assign each sequence to a different voter. If the election administrator reports the number of times that each sequence was cast in a reporting unit of (say) a few hundred people, then the majority of possible sequences will likely not be cast, so a vote buyer could be confident that each assigned sequence which appears in the election results was cast by the voter to whom it was assigned. This opportunity to leave a unique “fingerprint” on your ballot—and to thereby sell your vote—is baked into the rules of IRV.

We examine the feasibility of this scheme both theoretically and using data from real IRV elections. The problem we identify is an *abstract statistical* problem that could hypothetically be used in vote-buying activity. Vote buying is vanishingly rare in many democracies, including the cases where IRV is most relevant like Australia and the United States, but we will argue that the problem is nevertheless substantively important to understand. The core measure in our analysis is the proportion of distinct rankings cast in an IRV election. When this proportion is close to 0, a very small share of the possible rankings will be cast, and when the proportion is near 1, nearly all the possible rankings will be cast. In most real IRV elections, voters may not have a realistic opportunity to cast an identifiable ballot. Of the 36 IRV elections for public office available in the PrefLib library (Mattei and Walsh 2013), the median proportion of possible sequences cast is almost 0.98, while the mean is around 0.8. However, we identify several contests for prominent municipal offices where the expected proportion leaves at least a third of the possible sequences available, and we find some examples where most sequences are not expected to be cast.

These analyses make three main contributions: a theoretical contribution to the study of election security and electoral systems, a methodological contribution to the study of election rules, and a substantive contribution to the administration of IRV elections. Theoretically, researchers and administrators commonly consider election security when examining new voting technologies, new ways of handling voter data, or new ways of announcing vote totals. While substantial research has focused on the relationship between various electoral institutions, the secret ballot, and vote buying, we are not aware of another case where a security issue is baked into the abstract rules of the ballot format itself.


Methodologically, we introduce new ideas for how to study such vulnerabilities. We identify several rules for handling rankings that voters leave blank, and we derive the number of sequences that can be cast in ranked ballot elections under each of these rules. We then adapt methods from ecology and information science to model the expected number of sequences that will be cast in a ranked ballot election under very conservative assumptions. Using election result data, we demonstrate that our estimated proportions of unused sequences are even smaller than the number of unused sequences in some real IRV elections.

Substantively, we outline a vulnerability in a prominent and suddenly popular electoral system, and we suggest ways that election administrators can mitigate the statistical possibility of IRV votes being

²From now on, we will refer only to vote buying (offering a reward in exchange for a vote) and omit the idea of voter coercion (punishing somebody for how they voted). This is just for the sake of brevity. Every conclusion in this article applies just as much to voter coercion as it does to vote buying.

Pre-Election:

6 candidates are running for office.



The vote-buyer creates a random sequence using the candidates' names.




The vote-buyer promises a reward to a specific voter if they rank the candidates in that order.



Election + Post-Election:

The voter follows the specified order on their secret ballot.



The election administrators announce the list of ballots cast.

1	A	F	C	D	B	
2	A	D	F	B	C	E
3	A	D	F	B	C	E
4	A	E				
5	A	F	E	D	C	
6	A	C	B	E	D	F

The vote buyer checks for the sequence in the list of cast votes.



If anyone ranked the candidates in the specified order, the vote-buyer rewards the voter.

Figure 1. A step-by-step example of the vote-buying scheme.

identifiable. Our analysis suggests the following risk factors: (a) releasing a list of all of the ways that voters ranked the candidates, (b) including the specific spots skipped on each ballot, (c) allowing a large number of candidates to be ranked, and (d) reporting those rankings at a granular level.

Finally, it is worth emphasizing three arguments that we emphatically *do not* make. First, we should not be misread as opposing IRV. Second, we have no evidence of vote buying in any countries that use IRV. Third, we do not argue that switching to IRV alone would suddenly cause vote buying to proliferate in places where it is not currently common.

2. The Scheme and Its Potential Scope

Consider a contest between κ candidates, where voters are allowed to rank $L \leq \kappa$ of those candidates. Ballots are private, and there is no record explicitly connecting cast votes to individuals. We do not even require that it is known *whether* a specific person voted in the election. We only assume that when the election results are reported, they include a list of all the orders in which candidates were ranked. Figure 1 shows a step-by-step example of a vote-buying scheme that could exploit this situation.³

There are two reasons to focus on IRV instead of (a) other ordinal voting systems or (b) Cast Vote Records which report the votes that someone cast in a series of single-vote elections. The core reason is that the signaling is especially likely to be cost-free in IRV, because candidates ranked beneath the beneficiary of the scheme will only receive an extra vote if that candidate has already been eliminated. In contrast, a vote buyer who purchases a sequence of votes on a Cast Vote Record is buying random votes in unrelated contests, while a candidate who uses our scheme in some other ranked ballot systems (e.g., Borda Count) would actually be *buying votes for their competitors*. We expand on this reasoning in Section 1 of the Supplementary Material, where we also establish that the number of rankings in IRV is

³We are deeply grateful to Claire DeSoi for the visual design of this figure.

dramatically larger than the number of rankings on a Cast Vote Record or multi-mark ballots (Maloy 2019).⁴

Closely related schemes have been concocted before. In an online post, Quinn (2004) identified and briefly discussed the possibility of casting identifiable ballots in the related system of Single-Transferable Voting, and the idea was also mentioned in a report by an Irish commission on electronic voting (Commission on Electronic Voting 2004, 67). Benaloh and Tuinstra (1994) write that a similar scheme may have actually been used in Italian villages.⁵ However, to the best of our knowledge, the conditions under which this scheme is feasible have never been formally studied.

Could this scheme really be carried out? Our primary goal is to understand the conditions under which the scheme is *theoretically* feasible. We will focus on how many candidates need to run, and how detailed election result reporting needs to be, for the scheme to be statistically plausible. To motivate why its theoretical properties are of interest, though, we will first draw on political science theory to establish that there are narrow areas in which the scheme may be *substantively* feasible. These substantive claims should be read with two major caveats in mind. First, *we have no evidence that our scheme has ever been used, and we do not expect that current IRV elections are in imminent danger from our scheme.* Second, *this vulnerability can be easily solved*, and we will make specific suggestions for how to render it infeasible.

The first substantive point is that our scheme is only relevant in secret ballot elections. The worst violation of a secret ballot is when someone's vote choice can be revealed without their consent, but it is also crucial that "voters should not be able to prove how they voted to anyone, *even if they wish to do so*" (Bernhard *et al.* 2017, 2). Although "uncertainty about whether voters actually vote the way they say they will" is the essential firewall that prevents vote buying (Hicken 2011, 293), there are many countries which have both secret ballots and thriving vote-buying operations (Cruz 2018; Nwankwo 2018). These operations often involve making imperfect inferences about how people voted (Brusco, Nazareno, and Stokes 2004). While we are not aware of any constituency that uses IRV and has living practices of vote buying, or any democracy where our scheme would be legal, the fact that IRV could add another tool to the voter buyers' toolbox should be a reason for caution in introducing IRV without the safeguards that we will identify.

It is particularly pressing to better understand IRV as it rapidly proliferates across local and regional American elections, and we will show that, in some IRV jurisdictions, the richness of the available election data sets the stage for votes to be, at least in principle, identifiable. This is only a hypothetical problem, since the secret ballot long ago effectively eliminated vote buying in American elections (Cox and Kousser 1981). However, this equilibrium is maintained by a thicket of political institutions, including tight restrictions on what is allowed in polling places (Fitz 2022), policies regarding absentee voting and ballot selfies (Koutsoulias 2018), regulations restricting overly identifiable cast vote data (Kuriwaki *et al.* 2023), and scrutiny on new voting technologies (Castelló 2016).

A change in electoral systems is a major overhaul in political institutions, and it is not outlandish to suggest that new institutions could, over many years, enable illicit activity on a non-negligible scale. In fact, a voter has *already* exploited identifiability in IRV to identify *his own* vote in a real election in Aspen, Colorado, and, by combining this security issue with others, he was able to confidently de-anonymize other peoples' votes (Zimet 2009, 3). Mass vote buying does not occur in contemporary American elections, but this is partly because of careful attention to how to mitigate risks in various political institutions.

Our scheme may also be more plausible in small-population elections for nongovernmental offices. IRV is used in various organizations, from the leadership contests of major political parties to faculty senates. It is not hard to imagine our scheme being used to sway, say, an election that shapes the governance of a university.

⁴There may be other systems, perhaps, for example, Score Then Automatic Runoff, where we suspect the scheme may be statistically even more feasible than in IRV. However, IRV is rapidly gaining ground as a way to fill important public offices, and the relative simplicity of IRV makes exposition clearer.

⁵We are grateful to Johan Ugander, and in turn to Jon Kleinberg, for bringing this to our attention.

Finally, even if no vote buying actually takes place, a particularly pressing concern is the *perception* of insecurity. Trust in American elections is a newly polarized issue, with Republican voters trusting elections less than they have in recent history (Stewart III 2022). Indeed, Atkeson *et al.* (2023) have shown that many voters do not believe that the ballot was actually secret in recent American elections. If steps are not taken to mitigate our scheme, its feasibility could become one more reason for people to doubt that elections are secure.

3. Ballot Identifiability in IRV

The plausibility of the scheme hinges on the ability of a vote buyer to confidently infer that, because someone ranked the candidates in a particular order in the election, it was likely the voter to whom that sequence was assigned. So, how confident can the vote buyer be that, if they find a certain sequence among the election results, it is because a specific voter cast that sequence? This depends on how many sequences can be cast, which in turn depends on how votes are counted and reported. The key question is whether a voter may decide not to rank any candidate in a certain position (say, rank someone first, skip the second spot, and then rank someone third), and whether the position that they skipped is reported.⁶

The strictest rule is to discard any incomplete or repeated rankings, so that every voter must rank someone first, someone else second, and so on. We are not aware of a large election for public office using this “No Blanks” rule, but it is a useful lower bound. Under the No Blanks rule, the number n of sequences with κ candidates on a length L ballot is the number of length- L permutations of candidates:

$$n = \frac{\kappa!}{(\kappa - L)!}.$$

A more common way to handle blank rankings is to skip them. In Maine, for example, “If you skip a single ranking, the ranking after the single skipped ranking will be moved up and counted” (State of Maine 2018).⁷ This is equivalent to allowing voters to leave any number of spaces blank, but only at the end. Under this “Blanks Last” rule,⁸

$$n = \sum_{i=0}^{L-1} \frac{\kappa!}{(\kappa - L + i)!}.$$

While “Blanks Last” is a common rule for *counting* votes, it does not seem to be the most common rule for *reporting* votes. In Section 2 of the Supplementary Material, we identify the most detailed reporting in several recent IRV elections for public office. The most common reporting method actually lists the rankings that were submitted, including which spots each voter left blank. This “Any Blank” ruleset describes the main election reporting files in Alaska and Maine⁹ and actually understates the identifiability of ballots in New York City and San Francisco. The major exceptions are Australia and Papua New Guinea, where ballots are much less identifiable. The number of possible ballots in the “Any Blank” ruleset, previously derived by Baltz (2022a, 77), is given by

$$n = \sum_{i=0}^{L-1} \binom{L}{i} \frac{\kappa!}{(\kappa - L + i)!}.$$

We will consider all three rules. “Any Blank” is closest to how votes are reported in several real IRV elections. “Blanks Last” is how votes are often counted. “No Blanks” is the lowest bound on how many

⁶Why should voters ever be allowed to leave spots blank? In some ranked ballot systems, blank spots might make a difference in vote counting and could correspond to deliberate voter decisions (e.g., ranking one’s favorite candidates first, one’s least-favorite last, and skipping the middle spots). Also, many voters empirically do leave spots blank, so to forbid it may be equivalent to discarding or curing a large share of ballots.

⁷Setting aside the detail that, in Maine, votes after two consecutive skips are not counted.

⁸We assume that the voter does not leave *every* spot blank.

⁹Maine counts votes using roughly Blanks Last but reports using Any Blank.




Figure 2. The number of possible distinct sequences on a log-linear scale, when the number of rankings equals the number of candidates, under each of the three rulesets: when voters may not leave any of the spaces on their ballot blank, when they may leave any number of spaces blank at the end of their ballot, and when they can leave any combination of spaces blank.

sequences could exist under any counting rule. So, what sorts of numbers do these equations produce in realistic elections?

We will focus on ballot lengths from 2 candidates up to 10, which is the maximum number of candidates supported by some common voting machines.¹⁰ In that range, the number of possible sequences under each of these rulesets is shown in Figure 2. Just for the sake of visualization, we restrict our attention to the situation in which the number of candidates in the election equals L , though we will relax this restriction later.

Figure 2 shows that, for elections with up to 10 candidates, there can be millions or tens of millions of distinct sequences formed by reordering the candidates, and the number of possible sequences grows exponentially in the number of candidates. When there are four candidates or fewer, the number of rankings is only in the tens or hundreds for every ruleset, and the risk of a security issue seems extremely minimal. When there are six candidates, any ruleset generates enough sequences that a unique sequence could be assigned to every voter in many American precincts (Baltz *et al.* 2022). Once the number of candidates rises to about 10, with the No Blanks ruleset, a sequence could be uniquely assigned to every voter in Colorado. If the Any Blank rule is used, then after assigning a unique sequence of 10 candidates and blank spots to every voter in the United States, there would still be tens of millions of sequences left over.

Because of the sheer number of possible sequences that can be cast in an IRV election, it will often be possible to assign a unique sequence to every voter, so vote-buying activity would not be constrained by the number of sequences that can be cast. However, the vote buyer would also need to assess how certain it is that the assigned voter is the only person casting a specific sequence. To estimate that, we need to consider the probability of *collisions*, that is, when a sequence is assigned to one voter and then coincidentally cast by another. How confident can a vote buyer be that, if they assign a sequence to a voter, the sequence will be cast by *only* that voter?

3.1. Estimating the Number of Uncast Sequences

The main quantity of interest is the number of rankings that will not be cast in the election. This problem has been extensively studied in cryptography and ecology, in situations where individuals in a population each belong to some “type,” a sample is taken from the population, and the question is the expected number of types represented in the sample. This is also closely related to the coupon collector’s problem (Adler and Ross 2001). Good (1953) derived that in a population with S different types, where

¹⁰For example, the Democracy Suite System by Dominion Voting (Dominion Voting Systems 2023).

individuals have a probability p_i of belonging to type i , the expected number of types in a sample of size m is¹¹

$$E[S_m] = S - \sum_{i=1}^S (1 - p_i)^m.$$

For example, if a niche contains animals of S different species, and m animals are observed, this equation estimates the number of species we expect to find among those m animals. In our application, the types are the possible sequences. Everyone in the electorate has a probability of ranking the candidates a certain way, and the election is a realization of those rankings.¹² The challenge is that the expected number of types in the sample depends on the proportion p_i of voters who will cast each ranking, which requires a strong assumption about the distribution of support for the candidates.

We must make some assumption about expected vote choices, so what is the most scientifically conservative assumption we can make? The vote-buying scheme depends on a large pool of unused votes, so the more distinct rankings that are cast in the election, the harder the scheme will be. Therefore, the assumption that will most strongly play against our claim that vote buying may be possible in IRV is whichever assumption *maximizes* the expected number of distinct sequences cast.

Maximizing $E[S_m]$ as a function of p_i is an optimization problem, constrained by the fact that the sum of the proportions is one. In Section 3 of the Supplementary Material, we derive the single critical point $p_i = \frac{1}{S}$ for all i , and we solve the associated Lagrangian equation to show that this critical point maximizes the expected number of types in the sample. So, the most conservative possible assumption is that all of the types are equally likely to appear in the population. The intuitive explanation for this result is that each type has the best chance of appearing in the sample when as many individuals as possible are of that type, and the best way to simultaneously maximize the number of people in each type is to split them equally between each type.

In our application, this means that people are equally likely to cast each possible ranking of the candidates. Is this a reasonable model of vote choice? Of course not. The point is that, if we estimate the number of unused rankings available for the vote buyer under this assumption, then under *any* reasonable model of vote choice, there will be *at least* as many rankings available to the vote buyer.

Table 1 uses this assumption to answer the question of how many available sequences a vote buyer should expect *not* to be cast in an electorate of a certain size, under our middle-of-the-road rule (Blanks Last). The table shows the expected number of available (un-cast) sequences for ballot lengths ranging from 3 to 10 and populations from 10 voters up to 100,000,000 voters (the latter being the order of magnitude that represents the largest electorates in the world, e.g., in Indonesian presidential elections). Section 4 of the Supplementary Material includes the analogous tables for the No Blanks and Any Blank rulesets, and shows how both the possible and expected numbers of sequences vary by the number of candidates that can be ranked on a ballot.

Table 1 shows that there is ample opportunity in many realistically sized electorates for vote buyers to take advantage of sequences that will not be cast while also clearly demonstrating that having a small number of candidates in an IRV election eliminates any opportunity for our vote-buying scheme. We have not addressed *which* sequences to buy, which depends on the probability that a given sequence will be cast, but when there is a very large pool of unused sequences even a randomly purchased sequence is likely not to be cast by anyone else. We will next use the tools we have developed to examine the theoretical feasibility of a vote-buying scheme in a numerical example drawn from a real election.

¹¹ We follow the notation of Chao and Jost (2012, 2535).

¹² This matches the classic modeling paradigm in which candidates make guesses about the types of voters in an electorate (Coughlin 1992, Section 1.6).

Table 1. The expected number of sequences that will *not* be cast in an election with a certain ballot length and in an electorate of a given population, using the Blanks Last vote-counting method. We take the maximally conservative assumption that all sequences are equally likely to be cast and focus on the case where the number of candidates is equal to the length of the ballot.

Ranks	3	4	5	6	7	8	9	10
Population								
10^1	8	55	315	1,946	13,689	109,590	986,399	9,864,090
10^2	0	13	239	1,858	13,599	109,500	986,309	9,864,000
10^3	0	0	15	1,173	12,735	108,605	985,410	9,863,100
10^4	0	0	0	12	6,602	100,043	976,460	9,854,105
10^5	0	0	0	0	9	44,010	891,311	9,764,605
10^6	0	0	0	0	0	12	357,914	8,913,118
10^7	0	0	0	0	0	0	39	3,579,147
10^8	0	0	0	0	0	0	0	390

4. Numerical Example

Consider a proposed change to how the city of Oakland conducted its IRV election for mayor in 2022. Ten candidates ran, and about 125,000 people voted. In the election, voters were only allowed to rank five candidates, but after the election, there has been serious discussion concerning whether voters should actually have been allowed to rank all 10 candidates (Mukherjee 2023).

Now imagine a hypothetical scenario that, to the best of our knowledge, *did not* happen in Oakland that year. Suppose that a supporter of some candidate c decided to buy votes for that candidate. The vote buyer consults a poll and decides to purchase 1,500 votes. The challenge confronting the vote buyer is how many votes they have to buy in order to obtain 1,500 verifiable sequences (so, how many votes must be bought in order to expect that 1,500 votes will be cast by only the voter from whom the vote was bought?). Let us also, for the sake of illustration, imagine that the votes were reported using the Any Blank ruleset—Oakland actually releases much less information about the rankings submitted, but as we showed in Section 2 of the Supplementary Material, Any Blanks is common across many other IRV elections. We will first consider the case where 5 candidates can be ranked and then the case where all 10 candidates can be ranked.

The vote buyer begins by generating 1,500 different permutations of the candidates. Because the vote buyer is trying to boost the vote total of candidate c , all of those permutations must have c in the first position, so there are now nine candidates who can be assigned to four ballot positions. There are 5,508 possible sequences that can be cast with candidate c in the first position, and the number of ways that all 10 candidates can be arranged in the five ballot positions is 63,590. The expected number of unique sequences cast, under our conservative assumption, is

$$E[S_m] = S - \sum_{i=1}^S (1-p)^m,$$

$$E[S_m] = 63,590 - 63,590 \left(1 - \frac{1}{63,590} \right)^{125,000},$$

$$E[S_m] \approx 54,684,$$

meaning that only about 8,906 possible sequences will not be cast in the election. The vote-buyer wants to buy 1,500 votes that will *only* be cast by the voter from whom they bought the vote, but each sequence they buy has some chance of colliding with a sequence that another voter casts for legitimate reasons. In Section 5 of the Supplementary Material, we suggest several ways to estimate the probability that

a bought sequence will collide with a legitimate vote. Under the simplest approach, which naïvely models the probability that a bought vote collides with a legitimate vote as a Bernoulli trial, in order to expect to obtain 1,500 sequences, the vote buyer needs to buy the following much larger number of votes:

$$1,500 \times \frac{63,590}{63,590 - 54,684} \approx 10,710.$$

What happens if, instead, city officials determine that voters should actually be able to rank all 10 candidates? How does the vote buyer's calculus change? When all 10 candidates can be ranked in Any Blank ruleset, there are 234,662,230 possible sequences overall, of which 17,572,113 can be assigned. Then

$$E[S_m] = 234,662,230 - 234,662,230 \left(1 - \frac{1}{234,662,230}\right)^{125,000},$$

$$E[S_m] \approx 124,966.$$

Now how many votes should the vote buyer purchase in order to expect that 1,500 of them will be unique? Under the simple Bernoulli approximation, the vote buyer can now expect to secure 1,500 unique sequences by buying the following much smaller number of votes:

$$1,500 \times \frac{234,662,230}{234,662,230 - 124,966} \approx 1,501.$$

5. Identifiability in Real IRV Elections

Thus far, all of our analysis has been theoretical. We now check how many sequences were actually cast in San Francisco's 2019 IRV contests and compare it to our theoretical estimates.¹³

Before focusing on one example, we should emphasize that there are many IRV elections in which identifiability is neither a practical nor a theoretical concern. There are two ways that identifiability is effectively mitigated in real IRV elections. The first is reporting practices. For example, it is not unusual for more than half a dozen candidates to contest elections to the Australian House of Representatives,¹⁴ so we might expect that voters have the opportunity to cast identifiable ballots. However, as discussed in Section 2 of the Supplementary Material, there is no obvious way for the public to obtain the sorts of granular lists of rankings that would enable the scheme we outline in this paper.¹⁵ In the American case, one way to limit the identifiability of ballots while still producing reasonably transparent election result data may be to withhold precinct identifiers, or at least not reveal the location of votes from small or split precincts, while in other countries, the same might apply to polling place-level data (Kuriwaki *et al.* 2023).

A second mitigation strategy is to limit the number of candidates who can be ranked. In Alaska, granular lists of all the rankings cast are available, but because only four candidates can contest IRV elections there, the maximum number of unused rankings is very small. In Papua New Guinea, both precautions are present, since only three candidates can be ranked *and* there is no public release of the lists of every ranking submitted.

The following discussion of San Francisco's 2019 elections should be read with the caveat that it may focus on a slightly unusual case among IRV elections, since San Francisco combines slightly larger numbers of candidates than most IRV elections with extremely granular election result reporting.

¹³We stress that we are *not* searching for evidence of vote buying. We are not aware of any evidence that this sort of activity has ever taken place in a real IRV election.

¹⁴In a recent sample ballot from the Australian Electoral Commission, voters can rank eight candidates (Australia Electoral Commission 2019).

¹⁵We are grateful to the Australian Electoral Commission and to Campbell Sharman for pointing us to the information that comes closest to a list of the sequences cast in Australian preferential voting elections.

Table 2. The number of candidates, possible sequences, and total ballots cast in San Francisco's 2019 municipal contests (San Francisco 2023).

Contest	Candidates	Possible sequences	Possible sequences	Ballots Cast
		Blanks Last	Any Blanks	
Mayor	6	1,956	13,326	180,254
District Attorney	4	64	208	193,746
Supervisor (District 5)	4	64	208	23,718
City Attorney	1	1	1	161,464
Public Defender	1	1	1	158,621
Sheriff	1	1	1	146,693
Treasurer	1	1	1	152,762

In Section 6 of the Supplementary Material, we estimate the share of sequences cast in every IRV election available in the PrefLib library (Mattei and Walsh 2013), and we find that the San Francisco mayoral election (for example) has a larger proportion of available sequences than the median contest does, but it is not a total outlier, and indeed it has a smaller estimated proportion of available sequences than the mean.

San Francisco's 2019 municipal elections included seven IRV contests, of which three were contested (San Francisco 2023). Votes were counted using approximately the Blanks Last ruleset, so we conservatively focus on Blanks Last in this section, but they are actually reported in a way that is even more identifiable than Any Blanks, so we supply the corresponding information for Any Blanks in Section 7 of the Supplementary Material. Table 2 summarizes the number of possible sequences under both rules and the number of ballots cast across these elections. In the contested races, the fact that votes are reported using Any Blanks generates many times more possible sequences than if they were reported using Blanks Last.

What proportion of possible sequences, $\frac{S_m}{S}$, were cast in these contests? Figure 3 shows the proportion of possible sequences cast in each precinct for each contested IRV election, alongside the proportion cast across the whole contest (the solid horizontal line), and the expected number of sequences cast in a precinct of a given size when we conservatively set $p = \frac{1}{S}$.¹⁶ The increased identifiability caused by reporting votes at the precinct level can be observed by comparing a given dot (the proportion of sequences cast in some precinct) to the horizontal line (the proportion of sequences cast overall). The corresponding figures using the Any Blanks ruleset in Section 7 of the Supplementary Material show that vote reporting method results in a much larger share of unused sequences.

Comparing the real number of sequences cast to the expected proportion, we find that the assumption was indeed conservative, in that it overestimated the proportion of sequences that would be cast. The difference between the mayoral contest and the others in Figure 3 also demonstrates how an increase of as few as two candidates can greatly increase the proportion of available uncast sequences.

In Section 9 of the Supplementary Material, we simulate a vote-buying scheme in the mayoral election by randomly sampling “bought” sequences and counting the number of times that a simulated bought vote would have coincided with a legitimate one. We find that over 95% of randomly purchased sequences do not match a legitimate vote in the real election, and would therefore be identifiable.

¹⁶In Section 8 of the Supplementary Material, we show that this proportion is equal *under our assumptions* to the variable that ecologists have studied under the name “sample coverage” (Chao and Jost 2012, 2535). Also note that the Board of Supervisors subfigure in Figure 3 has fewer precincts than the others because this contest took place within district 5.




Figure 3. The proportion of all possible sequences cast in three San Francisco IRV races, using the Blanks Last ruleset. Each dot represents the proportion cast in a precinct, with the total number of ballots cast in that precinct on the x-axis. The solid horizontal line is the proportion of possible ballots cast across the whole election, and the dashed curve is the expected number of ballots cast in a precinct of a given size under our conservative assumption $\rho = \frac{1}{5}$.

6. Conclusion

We have revealed a security flaw in some ranked ballot elections. The risk is larger in elections using IRV, where there is a public list of all the ways that voters ranked the candidates, that list includes the specific spots that voters skipped, many candidates contest the election and many candidates can be ranked, and the results are reported in small-population areas. A promising area for future study is to examine the normative costs and benefits of how to approach these risk factors.

Under these conditions, voters can send uniquely identifiable signals that allow a third party to verify that the voter cast a particular vote. We examined the conditions under which this weakness is most important, by computing the expected number of votes that can be uniquely identified in simple IRV election setups. Although there is no evidence of vote-buying activity in countries where IRV has been implemented, and it would require more than just a change in ballot format for such activity to proliferate, we have shown that this system can make it statistically possible to identify votes from a large pool of cooperating voters.

Some of the risk factors can be mitigated, but the trade-offs deserve careful consideration. Changing the number of candidates who can be ranked or limiting candidate entry would reduce identifiability, but this is strong medicine indeed, since these approaches could also change election results (Tomlinson, Ugander, and Kleinberg 2023). Likewise, administrators could avoid releasing detailed or granular election data, but transparency is also an important value (Kuriwaki *et al.* 2023). Importantly, attempts to combat vote-buying schemes might also require transparent election result data. Election forensics techniques could be adapted to this problem, for example, by checking for very large numbers of sequences that rank the same candidate in first place and are each cast by just one person, which could be a distinctive trace of our scheme (Hicken 2011).¹⁷ One particularly promising remedy may be that, if votes are counted using the Blanks Last reporting scheme, there is no reason to announce the specific spots that were skipped on peoples' ballots.

Data Availability Statement. Replication code for this article has been published in Code Ocean, a computational reproducibility platform that enables users to run the code (Williams, Baltz, and Stewart III 2024). That code can be viewed interactively at <https://codeocean.com/capsule/1544948/tree/v1>.

Supplementary Material. For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2024.4>.

References

- Adler, E. Scott, and Thad E. Hall. 2013. "Ballots, Transparency, and Democracy." *Election Law Journal* 12 (2): 146–161. <https://doi.org/10.1089/elj.2012.0179>
- Adler, Ilan, and Sheldon M. Ross. 2001. "The Coupon Subset Collection Problem." *Journal of Applied Probability* 38 (3): 737–746.
- Atkeson, Lonna Rae, Eli McKown–Dawson, M. V. Hood III, and Robert Stein. 2023. "Voter Perceptions of Secrecy in the 2020 Election." *Election Law Journal* 22 (3): 234–253. <https://doi.org/10.1089/elj.2022.0064>
- Atsusaka, Yuki. 2023. "Causal Inference with Ranking Data: Application to Blame Attribution in Police Violence and Ballot Order Effects in Ranked-Choice Voting." Preprint, [arXiv:2207.07005](https://arxiv.org/abs/2207.07005).
- Australia Electoral Commission. 2019. "Practise Voting—House of Representatives." https://www.aec.gov.au/voting/how_to_vote/practise/practise-house-of-reps.htm.
- Baltz, Samuel. 2022a. "Computer Simulations of Elections, with Applications to Understanding Electoral System Reform." PhD Thesis, University of Michigan.
- Baltz, Samuel. 2022b. "The Probability of Casting a Pivotal Vote in an Instant Runoff Voting Election." Preprint, [arXiv:2210.01657](https://arxiv.org/abs/2210.01657) [cs].
- Baltz, Samuel, Alexander Agadjanian, Declan Chin, John Curiel, Kevin DeLuca, James Dunham, Jennifer Miranda, et al. 2022. "American Election Results at the Precinct Level." *Scientific Data* 9: 651. <https://doi.org/10.1038/s41597-022-01745-0>
- Benaloh, Josh, and Dwight Tuinstra. 1994. "Receipt-Free Secret-Ballot Elections." In *STOC '94: Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, 544–553. New York: Association for Computing Machinery. <https://doi.org/10.1145/195058.195407>
- Bernhard, Matthew, Josh Benaloh, J. Alex Halderman, Ronald L. Rivest, Peter Y. A. Ryan, Philip B. Stark, Vanessa Teague, Poorvi L. Vora, and Dan S. Wallach. 2017. "Public Evidence from Secret Ballots." Preprint, [arXiv:1707.08619](https://arxiv.org/abs/1707.08619) [cs].
- Birch, Sarah. 2007. "Electoral Systems and Electoral Misconduct." *Comparative Political Studies* 40 (12): 1533–1556. <https://doi.org/10.1177/0010414006292886>
- Brusco, Valeria, Marcelo Nazareno, and Susan C. Stokes. 2004. "Vote Buying in Argentina." *Latin American Research Review* 39 (2): 66–88.
- Castelló, Sandra Guasch. 2016. "Individual Verifiability in Electronic Voting." PhD Thesis, Universitat Politècnica de Catalunya.
- Chao, Anne, and Lou Jost. 2012. "Coverage-Based Rarefaction and Extrapolation: Standardizing Samples by Completeness Rather Than Size." *Ecology* 93 (12): 2533–2547. <https://doi.org/10.1890/11-1952.1>
- Commission on Electronic Voting. 2004. "First Report of the Commission on Electronic Voting on the Secrecy, Accuracy, and Testing of the Chosen Electronic Voting System." https://web.archive.org/web/20050523021828/http://www.cev.ie/htm/report/first_report/pdf/05Part.pdf.
- Coughlin, Peter J. 1992. *Probabilistic Voting Theory*. Cambridge: Cambridge University Press.
- Cox, Gary W., and J. Morgan Kousser. 1981. "Turnout and Rural Corruption: New York as a Test Case." *American Journal of Political Science* 25 (4): 646–663. <https://doi.org/10.2307/2110757>

¹⁷We are grateful to Walter R. Mebane, Jr. for interesting discussions on this point.

- Cruz, Cesi. 2018. "Social Networks and the Targeting of Vote Buying." *Comparative Political Studies* 52 (3): 382–411. <https://doi.org/10.1177/0010414018784062>
- Dominion Voting Systems. 2023. "RCV Brochure." <https://www.dominionvoting.com/optional-solutions/>.
- Fitz, Rebecca M. 2022. "Peering into Passive Electioneering: Preserving the Sanctity of Our Polling Places." *Idaho Law Review* 58: 270–287.
- Good, I. J. 1953. "The Population Frequencies of Species and the Estimation of Population Parameters." *Biometrika* 40, no. 3/4: 237. <https://doi.org/10.2307/2333344>
- Hicken, Allen. 2007. "How Do Rules and Institutions Encourage Vote Buying?" In *Elections for Sale: The Causes and Consequences of Vote Buying*, edited by Frederic Charles Schaffer, 67–87. Boulder, Colorado: Lynne Rienner.
- Hicken, Allen. 2011. "Clientelism." *Annual Review of Political Science* 14: 289–310. <https://doi.org/10.1146/annurev.polisci.031908.220508>
- Koutsoulias, Isidora. 2018. "Ballot Selfies: Balancing the Right to Speak Out on Political Issues and the Right to Vote Free from Improper Influence and Coercion." *Journal of Law and Policy* 26: 349–394.
- Kuriwaki, Shiro. 2020. "The Administration of Cast Vote Records in U.S. States." Preprint. Open Science Framework. <https://doi.org/10.31219/osf.io/epwqx>
- Kuriwaki, Shiro, Jeffrey B. Lewis, and Michael Morse. 2023. "The Still Secret Ballot: The Limited Privacy Cost of Transparent Election Results." Working Paper. <https://arxiv.org/abs/2308.04100>.
- Maloy, J. S. 2019. *Smarter Ballots: Electoral Realism and Reform*. New York: Palgrave Macmillan.
- Mattei, Nicholas, and Toby Walsh. 2013. "PrefLib: A Library for Preferences." <http://www.preflib.org/> In *Algorithmic Decision Theory*, edited by Patrice Perny, Marc Pirlot, and Alexis Tsoukiás. Lecture Notes in Computer Science, 259–270. Berlin–Heidelberg: Springer. https://doi.org/10.1007/978-3-642-41575-3_20
- Mukherjee, Shomik. 2023. "Why Didn't Oakland's Ranked Choice Ballot Follow City Charter?" The Mercury News.
- Nwankwo, Cletus Famous. 2018. "Vote Buying in the 2018 Governorship Election in Ekiti State, Nigeria." *Open Political Science* 1 (1): 93–97. <https://doi.org/10.1515/openps-2018-0005>
- Quinn, Ciaran. 2004. "Vote Buying, Intimidation of Voters—The Unintended Consequences of Electronic Voting in Ireland." <https://web.archive.org/web/20041215021221/http://election.polarbears.com/art0037.htm>.
- Rae, Douglas. 1967. *The Political Consequences of Electoral Laws*. New Haven: Yale University Press.
- San Francisco. 2023. "Past Election Results—Department of Elections." <https://sfelections.sfgov.org/past-election-results>.
- Santucci, Jack. 2022. *More Parties or No Parties: The Politics of Electoral Reform in America*. Oxford: Oxford University Press.
- State of Maine. 2018. "Instructions for Voters: Marking a Ranked-Choice Voting Contest." <https://www.maine.gov/sos/cec/elec/upcoming/pdf/HORZ.boothposterFINALDRAFT.061218.pdf>.
- Stewart III, Charles. 2022. "Trust in Elections." *Daedalus* 151 (4): 234–253. <https://doi.org/10.1162/daed'a'01953>
- Tolbert, Caroline J., and Daria Kuznetsova. 2021. "The Promise and Peril of Ranked Choice Voting." *Politics and Governance* 9 (2): 354–364. <https://doi.org/10.17645/pag.v9i2.4385>
- Tomlinson, Kiran, Johan Ugander, and Jon Kleinberg. 2023. "Ballot Length in Instant Runoff Voting." *Proceedings of the AAAI Conference on Artificial Intelligence* 37 (5): 5841–5849. <https://doi.org/10.1609/aaai.v37i5.25724>
- Williams, Jack R., Samuel Baltz, and Charles Stewart III. 2024. "Replication Code for 'Votes Can Be Confidently Bought in Some Ranked Ballot Elections, and What To Do about It.'" Version 1. <https://codeocean.com/capsule/0741918/tree>.
- Xia, Lirong. 2012. "Computing the Margin of Victory for Various Voting Rules." In *Proceedings of the 13th ACM Conference on Electronic Commerce—EC '12*, 982–999. Valencia: ACM Press. <https://doi.org/10.1145/2229012.2229086>
- Zimet, Millard J. 2009. "Complaint to Aspen Election Commissioners." <https://web.archive.org/web/20230308184415/http://static1.1.sqspcdn.com/static/f/270139/4001698/1251676428717/ZimetComplaintToElecComm.pdf?token=RMASy0Jxc2LnNZmZX47xYtzGwE%5C%3D>.