Buying random votes is as hard as buying no-votes

Stefan Popoveniuc and Jonathan Stanton*

April 9, 2007

Abstract

In voting systems where a mark in a fixed position may mean a vote for Alice on a ballot, and a vote for Bob on another ballot, an attacker may coerce voters to put their mark at a certain position, enforcing effectively a random vote. This attack is meaningful if the voting system allows to take receipts with them and/or posts them to a bulletin board. The coercer may also ask for a blank receipt. We analyze this kind of attack and prove that it requires the same effort as a comparable attack would require against any voting system, even one without receipts.

1 Introduction

Privacy in a voting system means that a vote cannot be bought or sold. Trying to buy votes has been a practice for many years and at every election there are accusations of vote buying or coercion. A coercer will typically pay some amount of money if a voter delivers a vote for a certain candidate, or will try to prevent a voter from being able to cast a ballot. The design of some of the state of the art voting system do not allow people to prove how they voted, even if they would like to. We say that those voting systems provide involuntary privacy.

Some of these new voting systems use ballots in which either the order of the candidates is randomized [CRS05] or the order of the candidates is fixed, but there is another level of indirection to vote for the desired candidate [PH06]. In these systems the voter gets to keep a receipt with the marks that she made inside the booth. An example of how the ballots and receipts appear is provided in Figure 1. In such a system, the receipt does not allow a voter to show evidence of who they actually voted for, so traditional vote buying, where a vote for a specific candidate is purchased, is not possible. However, a coercer may force her to put the mark at a certain position, effectively enforcing a vote for a random candidate.

Most of the time, the ballot does not contain an explicit "No Vote" for each contest, but the voter is allowed to "not vote" in that contest. So an alternative attack is for the coercer to demand to see no mark in that contest on the receipt, which will force the voter to abstain from that contest.

What we show in this paper is that these sorts of attacks can be conducted on basically any voting system, that the specific "random receipt" buying attack does not influence the election results any more then a voter being disenfranchised by being prevented from voting, and that the format of encrypted receipts does not facilitate these types of attacks.

*Authors (poste@gwu.edu, jstanton@gwu.edu) are with the George Washington University, Department of Computer Science, Washington DC 20052

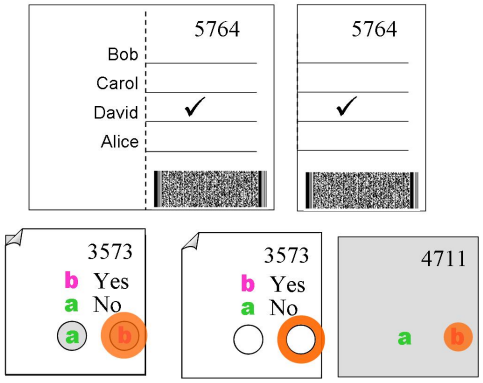

Figure 1: On the left are ballots and on the right are receipts. The top example is from Pret-A-Voter, while the bottom is PunchScan

2 Voter Disenfranchisement

Election systems have numerous ways in which voters may be disenfranchised[BBJM01, Cra01, Jon05]. They range from technical flaws in the voting technology which cause votes to be lost[Bol04, WRM04], to registration systems which prevent valid voters from being registered, to voting processes that make it harder to reach a polling place on election day, to malicious destruction of ballots or actively preventing voters from reaching a polling place or getting their absentee ballot. In this paper we recognize that many of these methods are outside our expertise and can occur no matter what voting technology is used for the casting of ballots. So we focus on the ways in which the voting system itself, and specifically the provision of ballot receipts affects voter disenfranchisement and the results of the election.

In the United States, at least, active voters' lists are a matter of public record. Often times election staff buy those lists from the election officials to conduct targeted campaigns. Although whether or not having public voter lists is good public policy is not a question we discuss, in practice they are available, even if such availability is limited to some subset of the public (such as registered political parties) the result is that interested groups (such as the parties) can know whether each individual voted in a particular election or not.

The availability of voter lists allows an attacker to verify participation in any particular election, therefore they do have the ability to coerce someone to "not vote" by paying them to not show up to the polls or threatening them if they do. Such attacks are individual in scope as each voter must be contacted and bought or coerced not to vote.

The overall process of forcing a voter to abstain is equivalent in a ballot receipt based system, as each voter must be contacted both before and after the election – before to obtain their agreement to the voting rule (abstain or random vote 3) and after to view their receipt and verify that they did as instructed. The receipt does allow some flexibility about selecting which groups of voters to coerce as the receipt is *portable* and can be shown to someone at work, in a neighborhood, at a group or association meeting, etc. While with traditional verification methods, the voter lists from each precinct must be checked to verify voter participation which is easy for neighborhood based disenfranchisement, but harder for work-based or arbitrary association based where the members do not live in the same area and may vote in hundreds of different precincts.

So although the process of forcing a random vote or abstention may be a bit different in receipt based vs traditional voting systems, the overall difficulty and requirements do not appear different.

3 Random votes

Assume we have a voting system where a mark in a fixed position may mean a voter for any candidate, depending on the serial number of the ballot. Several of the currently proposed systems support such an assumption. A coercer can ask a voter to bring a receipt with a mark on a certain position in a contest. For example Mallory promisees 10$ to anyone that shows her a receipt with the first position marked for some contest. Since the positions are randomly assigned to candidates, this would effectively mean a random vote in that contest.

We prove that this attack is as costly as coercing someone not to vote at all. More precisely, we show that a random vote does not influence the results of the election any more then an abstention, which as discussed in Section 2 is possible in many ways.

Assume that in an honest election, the results would be:

Candidate 1 - x_1 votes
Candidate 2 - x_2 votes
...
Candidate n - x_n votes

Let $X = x_1 + x_2 + ... + x_n$ be the total number of votes. Note that X is the number of votes, not the number of voters. In the case of a choose n out of m race, the total number votes is n times the total number of voters (assuming that all the voters do choose n candidates). The mathematical demonstration is only concerned with the number of votes, and therefore is valid for binary contests, choose one out of n or choose n out of m contests. The discussion is not valid for rank voting since there are multiple ways of "counting the votes", deciding the candidate ranking, and determining the winner when it comes to rank voting.

So for victor-take all elections, after an eventual rearranging, let $x_1 > x_2 > ... > x_n$, then assuming a plurality election candidate one would win. In the case where a candidate must also receive a majority of the votes, then $x_1 > X/2$ as well as x_1 having the most votes must occur. In this case, although we show the order of the candidates will not change, in some cases a candidate may be able to overcome the majority condition if voters are forced to randomly vote or are disenfranchised, but would not overcome it if they were not. This can be considered affecting the election, but since it can occur either by random votes or disenfranchisement it is not a problem caused by the random votes enabled by receipts.

Some elections, especially important referendums, require 50%+1 of the total number of possible votes, not just from the total number of expressed votes. An attacker may choose to coerce people to stay at home or to show and cast random ballot, depending on her interest, but in this particular situations one attack may have an advantage over the other one.

3.1 Forcing uniform random votes

Assume Mallory buys a fraction $0 \leq f \leq 1$ of the total votes and forced them to become random votes. Then candidate k will:

- Loose $x_k \times f$ votes, because these votes will now be random instead of in favor of x_k
- Win some votes from all the other candidates. From candidate i, $x_i \times f \times \frac{1}{n}$ will now be in favor of candidates k, for a total of:

$$x_1 \times f \times \frac{1}{n} + x_2 \times f \times \frac{1}{n} + \cdots + x_n \times f \times \frac{1}{n} = \frac{f}{n} \times (x_1 + x_2 + ... + x_n) = \frac{f}{n} \times X \text{votes.}$$

In general x'_k, the number of votes for candidate k after buying a fraction f of random votes, is $x'_k = x_k - x_k \times f + \frac{f}{n} \times X$

For Mallory's favorite candidate, k, to win it would have to have the highest number of votes after the coercion attack. So, $\forall j, 0 \leq j \leq n, j \neq k, x'_k > x'_j$,

$$x_k - x_k \times f + \frac{f}{n} \times X > x_j - x_j \times f + \frac{f}{n} \times X \quad \Rightarrow \quad (x_j - x_k) \times f > x_j - x_k$$

If k would have been the winner even without the attack,($x_j < x_k$), the k would still be the winner after the attack (the inequations becomes $f < 1$, which is true). If $x_j > x_k$, thus k, would not be the winner before the attack, the inequality is a contradiction, because it would require $f > 1$, thus k cannot become the winner after the attack. In both cases k would not change the election results by doing this sort of attack.

3.2 Forcing uniform random silence

We prove that choosing random voters and forcing them not to vote at all would not result in an attack that would change the order of the candidates, just as above. Moreover, we prove that the mathematical equations end up being the same.

When buying uniform random silence candidate k will loose $x_k \times f$ votes, because these voters will be forced to stay at home.

Thus $\forall k, x'_k = x_k - x_k \times f$. For Mallory's favorite candidate, k, to win it would have to have the highest number of votes after the coercion attack. So, $\forall j, 0 \leq j \leq n, j \neq k, x'_k > x'_j$,

$$x_k - x_k \times f > x_j - x_j \times f \quad \Rightarrow \quad (x_j - x_k) \times f > x_j - x_k$$

. Thus is the exact same equation as above. The same discussion applies.

3.3 Forcing nonuniform random votes

We now focus on an attack that would succeed: assume that Mallory determines which voters will not vote for Mallory's desired candidates. She decides to buy random votes only from these *non-supporing* voters. The result is that Candidate k does not loose any votes from his loyal voters (since none are targeted for the attack), but actually wins some votes from each of their opponents, because they pick up some of the random votes.

Thus, $\forall j, 0 \leq j \leq n, j \neq k$

$$\begin{aligned} x'_j = \quad & x_j - x_j \times f + x_1 \times f \times \frac{1}{n} \\ + \quad & x_2 \times f \times \frac{1}{n} + \dots \\ + \quad & x_{k-1} \times f \times \frac{1}{n} \\ + \quad & x_{k+1} \times f \times \frac{1}{n} + \dots \\ + \quad & x_n \times f \times \frac{1}{n} \\ = \quad & x_j - x_j \times f + \frac{f}{n} \times (X - x_k) \end{aligned}$$

So each candidate J loses some votes from their own voters being forced to randomly vote, and gains some votes from other candidates voters being force to randomly votes, but does not gain any

votes from *k's* voters as they are not targeted in the attack. So Candidate k only gains from this attack, and will end up with x'_k votes.

$$x'_k = x_k + \frac{f}{n} \times (X - x_k)$$

Generally, for any two candidates different from k, their rank will be preserved by this attack. If i was before j without the attack, i would still be before j after the attack. Mathematically, $\forall i, j$ such that $i > j, i \neq k, j \neq k$, if $x_i < x_j$ then

$$\begin{aligned} x'_i < x'_j &\Leftrightarrow x_i - x_i \times f + \frac{f}{n} \times (X - x_k) > x_j - x_j \times f + \frac{f}{n} \times (X - x_k) \\ &\Leftrightarrow (x_i - x_j) > (x_i - x_j) \times f \\ &\Leftrightarrow 1 > f(TRUE) \end{aligned}$$

Therefore, for k to win the elections, it would have to have more votes then candidate 1. Thus,

$$\begin{aligned} x'_k > x'_1 &\Rightarrow x_k + \frac{f}{n} \times (X - x_k) > x_1 - x_1 \times f + \frac{f}{n} \times (X - x_k) && (1) \\ &\Rightarrow x_1 - x_k < x_1 \times f && (2) \\ &\Rightarrow 1 - \frac{x_k}{x_1} < f && (3) \end{aligned}$$

3.4 Buying nonuniform random silence

Assume that Mallory coerces a random fraction of voters known not to vote with k, to not vote at all.

All the candidates other then k loose some a fraction of their votes: $\forall j, 0 \leq j \leq n, j \neq k, x'_j = x_j - x_j \times f$, while candidate k keep all his votes: $x'_k = x_k$

Just as in the previous cases, to have candidate k win, we need

$$\begin{aligned} x'_k > x'_1 &\Rightarrow x_k > x_1 - x_1 \times f \\ &\Rightarrow x_1 \times f > x_1 - x_k \\ &\Rightarrow f > 1 - \frac{x_k}{x_1} \end{aligned}$$

same as equation 3.

Therefore, it takes the same effort to make k win, regardless if voters are paid to vote randomly or not to vote at all.

Note: If the fraction f is not the same for each candidate, the result is similar: $f_j > 1 - \frac{x_k}{x_j}$ for every $j < k$

4 Conclusions and Acknowledgments

We have proved that the ability for receipt based voting systems to support purchase of a verifiable random vote, does not make it easier for a coercer to influence an election. A random vote has a similar effect to someone buying abstentions from voters, which can be verified by active voter lists or other known methods. Therefore, voting systems that allow receipts to be posted on a bulletin board or that have marks on them from the voting process that indicate a vote for *some* candidate do not weaken the systems protection from vote buying.

References

[BBJM01] Henry E. Brady, Justin Buchler, Matt Jarvis, and John McNulty. Counting all the votes: The performance of voting technology in the united states. Technical report, Berkeley: University of California, 2001.

[Bol04] Erica Bolstad. Votes from 134 residents were not counted. *The Miami Herald*, January 7 2004.

[Cra01] L. F. Cranor. Voting after Florida: No easy answers. *Ubiquity*, 47, February 13-19 2001.

[CRS05] David Chaum, Peter Y. A. Ryan, and Steve Schneider. A practical voter-verifiable election scheme, 2005.

[Jon05] Douglas W. Jones. Threats to voting systems. *NIST Workshop on Threats to Voting Systems*, October 2005.

[PH06] Stefan Popoveniuc and Ben Hosp. An introduction to punchscan. In *IAVoSS Workshop On Trustworthy Elections (WOTE 2006)*, Robinson College, Cambridge UK, June 2006. Also GWU-CS Technical Report.

[WRM04] Jr. Walter R. Mebane. The wrong man is president! overvotes in the 2000 presidential election in florida. *Perspectives on Politics*, 2004. Forthcoming.