ThreeBallot in the Field

Harvey Jones, Jason Juang, Greg Belote
Instructor: Ronald L. Rivest

December 13, 2006

Abstract

Voting systems have been the subject of much recent controversy. Due to the difficulty of securing and auditing electronic voting systems, a variety of different paper-based cryptographic voting schemes have emerged. Ronald Rivest has proposed a purely paper-based system called ThreeBallot, which strives to achieve the same level of security as cryptographic systems without using cryptography. Although ThreeBallot has been subject to academic criticism, it has not been tested in the field. This paper describes a paper-based and a computer-aided implementation of ThreeBallot. Any successful voting system must be usable, must be secure, and must preserve the secret ballot. To test usability, we held mock elections and observed voters. To test security and privacy, we executed attacks against these mock elections.

In one mock election, 20% of voters successfully sold their vote. One student, when given control of tallying the votes, was able to throw the entire election. In our usability studies we confirmed voter difficulty in using ThreeBallot. We found that about 10% of voters didn't understand ThreeBallot well enough to check another's ballot, and in one mock election more than 30% of voters failed to cast a valid ballot on their first try.

1 Introduction

Voting systems fail for many reasons. They can be too complex for the average voter to use, they can be run by corrupt voting officials, or they can compromise voter privacy. A modern voting system must avoid these pitfalls under the pressure of a real-world election if it is to be a feasible replacement for current voting technology. Professor Ronald Rivest has proposed ThreeBallot [1], a paper-based voting system that claims to be private and secure, with a small usability trade-off.

ThreeBallot has been subject to academic criticism, for example by Appel [2] and Strauss [3], but it has not been used in an actual election environment. With the goals of usability, security, and privacy in mind, we conducted several mock elections designed to test ThreeBallot with actual voters, and to discover where the system succeeds and fails.

We tested the usability of the system by running elections open to the public in Building 32 and the East Campus dormitory. We anticipated usability problems with the vanilla ThreeBallot system, so we implemented a computer-based ThreeBallot machine with an interface similar to a touch-screen electronic voting machine. We recorded our interactions with voters, the number of inaccurate ballots that were submitted, and users' understanding of the system as measured by a quiz after their vote was tallied. We found that voters were initially skeptical about the system. We found that a substantial number of voters experienced difficulty using paper-based ThreeBallot for the first time, with more than thirty percent of voters having their ballots rejected on submission. After voters had successfully completed one election, they were familiar enough with the system to distinguish correct and incorrect ballots.

In order to test the system's robustness against election-stealing, we ran a mock election for students familiar with the system, and provided incentives for any student able to produce substantially skewed election results that were not subsequently repudiated by the class. One student successfully mounted an attack, combining analysis of all submitted ballots with spying on other voters' receipts. His attack was able to change the winner of all three races in the in-class election.

In order to test ThreeBallot's robustness against privacy compromises, we provided incentives to voters in the in-class election for proving to us how they voted. We also implemented a proposed privacy attack, which we discuss in Section 2.4.4, and attempted to prove how voters voted, assuming we were able to coerce them into giving us their receipts. We were able to reconstruct one-third of the ballots cast in the in-class election and five-sixths of the ballots cast at the Stata Center election, but could not reconstruct any of the ballots cast in the East Campus election.

2 Background

2.1 Design Criteria

Seemingly contradictory criteria must be satisfied by any successful voting system. The ideal system is transparent enough to inspire complete public trust, yet hides enough information to ensure the secret ballot requirement. Designs that satisfy both of these requirements typically sacrifice usability and/or scalability.

Privacy A voter's choice must be kept secret, even if he or she may wish to divulge it. If the secret ballot requirement is not held, voters may sell their votes, or the powerful may force other voters to select a particular candidate.

Verifiability A voting system should be publicly verifiable. Voters should be able to determine, to a very high probability, that their votes were accurately recorded and tabulated. An audit trail should be available so that the count can be verified.

Usability A voting system should be usable by the voting population. Complicated, hard to use, or error-prone systems may disenfranchise some voters.

Scalability The voting system should scale well to large numbers of voters. This includes not only handling a large number of total voters, but also a large number of *simultaneous* voters, to maintain a short waiting times throughout the election day. Ideally, any extra resources required to handle a larger number of voters should be cheap and easy to procure, and large elections should not compromise election security.

2.2 Cryptographic voting schemes

An ideal voting scheme allows an election to be auditable by any voter, but protects the security of every voter's ballot. Several ideas have been proposed that use cryptography and zero knowledge proofs to provide privacy and verifiability.

David Chaum proposed such a scheme [4] employing visual cryptography and mixnets. His scheme allows a voter to take home a receipt. The election official publishes copies of all the receipts on a web site, where voters can check that their receipt appears and is identical to the physical receipt. This receipt does not reveal who the voter voted for. In fact, the receipt is information-theoretically secure.

The tabulation process then begins with these receipts, and proceeds through stages of decryption, each of which is auditable by the public to ensure that there is no foul play.

Chaum's scheme suffers from the fact that the decryption process is difficult for the average voter to understand. Chaum uses an analogy with Russian nesting dolls to help clarify his explanation, but even so, it is not a particularly simple system. The voter's receipt is also somewhat difficult to verify, as it consists of a very large number of "pixels," and the voter must check that all pixels match the published receipt.

Chaum's "Punchscan" [5] and Adida and Rivest's "Scratch & Vote" [6] provide further examples of recently-proposed cryptographic voting schemes. With regards to the complexity of the involved cryptography, Chaum asserts, "Just like the little padlock in the corner of browser windows, users benefit without having to understand the inner workings of the crypto, especially where the software is public as with Punchscan" [7]. Furthermore, Chaum believes that the cryptography behind Punchscan is much simpler than that of the scheme we described earlier, and that it may be simple enough to be taught in advanced high school classes.

2.3 ThreeBallot

Rivest's ThreeBallot system [1] is an attempt to satisfy the design criteria we discussed earlier, protecting the secret ballot in a cryptographically secure manner, but without the use of cryptography. This arguably simplifies the process, increasing its transparency, and therefore bolstering voter confidence in the system.

2.3.1 The Procedure

To vote using ThreeBallot, each voter receives a *multiballot*. The multiballot consists of three ballots. This may be implemented any number of ways, but we will imagine, as Rivest does, that a multiballot is a single piece of paper with perforations, allowing it to be separated into three parts. Each of the three ballots is identical, except for a unique ID number. (See Appendix A for a sample multiballot.)

To vote for a candidate, the voter marks that candidate's name on exactly two of the three ballots. The two may be chosen arbitrarily. Otherwise, the voter marks the candidate's name on exactly one of the three ballots, again chosen arbitrarily. The voter may choose not to vote for any candidate by simply marking every candidate's name exactly one time. Once finished voting, the voter runs the multiballot through a *checker*, which verifies that the ballots have been filled out according to these constraints.

After the checker has validated the multiballot, the voter is given the option of keeping a copy of *one* of the three ballots (including the ID number) as a receipt. Nobody except the voter should know which one was copied. Once the voter has gotten a receipt, and checked to ensure that it is, in fact, a copy of one of the ballots, the multiballot is separated into the three separate ballots, all of which are then cast. From this point on, the ballots have no association to each other; nobody should be able to determine which three ballots were originally part of one multiballot. As a result of the voting constraints, the voter has cast one vote for every candidate, and (possibly) a second vote for some candidate.

There are now $3n$ ballots in the ballot box, where n is the total number of voters in the election. Election officials tally the votes from all $3n$ ballots, then subtract n votes from each candidate to obtain the final tallies. All $3n$ ballots are posted publically on the *public bulletin board*, so that anyone may verify the official tallies.

2.3.2 Properties

Because all ballots are available to the public, any individual or group may tally the election results and audit the official totals.

Each voter can use his or her receipt to audit the election and verify that his or her vote was counted correctly. To do so, the voter looks up his or her ballot ID number on the public bulletin board. The ballot displayed on the bulletin board should match the voter's receipt. If it does not, somebody has tampered with the election in some way, and the voter complains to the appropriate authorities.

However, the receipt does not prove how any voter voted, as the other two ballots in the voter's multiballot could have been filled out in any pattern, independently of the contents of the receipt. Thus, the receipt does not violate the secret ballot requirement.

ThreeBallot, as presented, has the same requirements, in terms of voting machinery, as an optical scan machine; all of the selection and recording of votes is done by paper. As a result, ThreeBallot is scalable; higher throughput can be achieved by increasing the number of voting booths and writing implements available. Additional checkers may be added to avoid bottlenecks at the ballot box, but the number of checkers needed is low, since the checking process should be relatively quick.

2.3.3 Refinements and concerns

Stateless Checker The machine that verifies that submitted multiballots are legal sees all three ballots together, and therefore "knows" the voter's intent. It also will "know" which ballot the voter requests a receipt for. Retaining this information would allow corrupt election officials to subvert the public verifiability of the election. So, it is vitally important that the checker be designed to purge any information about the prior checked ballot before accepting the next ballot.

Vote-selling attack Rivest notes that voters may sell their votes by simply remembering all three ID numbers from their multiballot, and revealing them to another party after casting their ballot. This links a specific set of three ballots to the voter, thus violating the voter's privacy and allowing vote-selling.

A potential solution to this problem is to use the "Shamos checker," described in Rivest's paper. In this case, no ID numbers are printed on the blank ballots. Instead, the checker prints the numbers on the ballots after validating the voter's multiballot, and then creates the voter's receipt and deposits the ballots into the ballot box without allowing the voter to see the printed ID numbers (except for the one on the receipt, of course).

2.4 Criticisms of ThreeBallot

2.4.1 Complex and unfamiliar interface

To a first-time user, ThreeBallot can seem intimidating. It requires that the voter make many marks on the ballot, following a strict set of rules. The system, being unlike most other voting systems, is outside the experience of the typical voter. The benefits gained from the added complexity are not immediately obvious to the untrained eye, and as such, some voters may be skeptical that the added complications to their voting experience are necessary.

2.4.2 Inability to handle write-in votes

Charlie Strauss [3], of Verified Voting New Mexico, asserts that write-in votes are difficult, if not impossible, due to the design of ThreeBallot. He notes that straightforward implementations of write-in votes are likely to violate the law in many jurisdictions. Voters could cast overvotes for write-in candidates in jurisdictions where voter intent laws require that written-in names be counted as votes, regardless of whether or not the bubble is filled in.

Additionally, there is the problem of how to represent write-in votes on the public bulletin board without facilitating vote-selling. Ballots containing write-in votes for uncommon or peculiar names can be easily identified on the public bulletin board, compromising ballot secrecy.

2.4.3 Cannot accomodate other election systems

Rivest acknowledges that ThreeBallot cannot be easily extended to preferential voting systems, where voters rank their preferences on the ballot from highest to lowest. While most elections in the United States still use first-past-the-post, plurality-wins balloting, some localities, such as Cambridge, MA, use Single Transferable Vote systems for some races. The desire to keep these arguably superior systems would hinder the adoption of ThreeBallot.

2.4.4 Strauss's privacy compromise

Strauss [8] has contended that due to the constraints posed by the ThreeBallot system, it may be possible to reconstruct complete triples of ballots if the number of cast ballots is far fewer than the number of legal ThreeBallot votes. If a third party could coerce voters into revealing their receipts, he could then determine how the voter voted, and potentially use intimidation to influence elections.

Each ballot has a number of bubbles equal to the total number of candidates running in all elections. Thus, a ballot can be represented as a bit string, where “1” is a filled-in bubble and “0” is an empty bubble. Strauss found that when 100 voters were aggregated onto a bulletin board, 25-bit ballots were sufficient to reconstruct 90% of cast votes. At an aggregation level of 1,000, 35-40 bits was the threshold, and at 10,000 voters, 45-50 bits were sufficient. This poses significant problems, as aggregating beyond 1,000 voters is likely to be difficult, asne can only aggregate over sets of identical ballots. This sets a hard limit on the level of integration at the level of the state legislative district. Strauss also found that party affiliation tended to exacerbate the problem; straight-ticket voters' ballots were cracked far easier than random voters.

3 End User Testing

We conducted and observed sample elections in order to test the usability of the ThreeBallot system. We observed voter interactions with both computer and paper-based systems. We noted how many times the “scanner” rejected ballots, and we took notes on which parts of ThreeBallot were the most difficult to explain. Overall, we found that although ThreeBallot was the subject of suspicion initially, most voters were able to competently use the system after one election.

3.1 Methodology

Users approached the voting booth and were assigned to the computer voting station or paper ballots, depending on which station was free at the time.

Computer Implementation We implemented a ThreeBallot voting application and ballot renderer using Tkinter and Python. Voters received a quick explanation of the ThreeBallot system and were told that the computer would generate a valid multiballot for them. Voters selected their candidates much as they would on a touch-screen machine. Voters then reviewed their ballot, and either cast it, or discarded it and started over. The application then generated a LaTeX file and converted it to PostScript, which was displayed on the screen. The voter was given the opportunity to review his or her ballot one more time, before sending a print command. At this point, a person simulated the checker, and filled out a receipt for a column of the user's choice. In this case, and in the paper ballot case, voters were made aware that in an actual election, the checker and receipt printer would be implemented by machines.

Paper Implementation The paper voters were given a similar description of the system, and handed a paper multiballot (a blank version of the one depicted in Appendix A). After

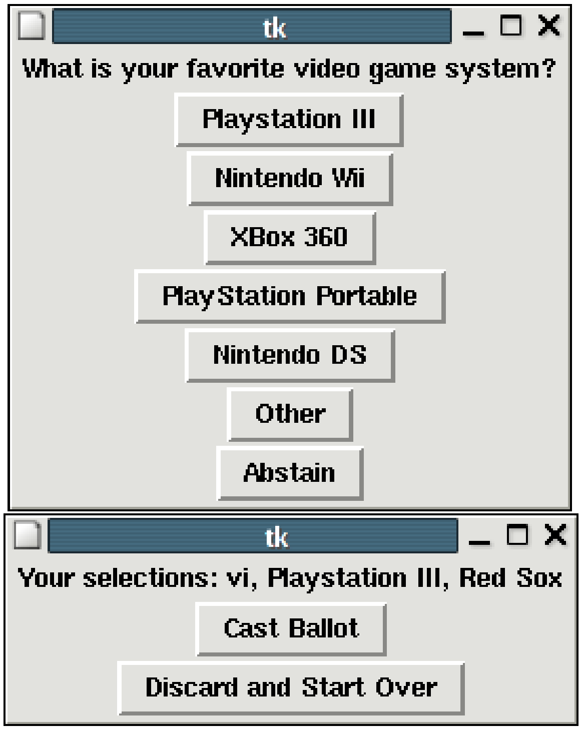

What is your favorite video game system?

Your selections: vi, Playstation III, Red Sox

Figure 1: User interface of computer-aided ThreeBallot

they completed the multiballot, they handed it to a person who was simulating the checker machine. After the checker machine validated the multiballot, the person filled out a receipt by hand and cast the multiballot.

Survey We surveyed each voter after his or her ballot was cast. The survey asked voters about their voting experience, and measured their understanding of the ThreeBallot system. It also asked voters how confident they were in the security of ThreeBallot, and whether or not they would approve of its use in a federal election. The back side of the survey contained a quiz that asked voters to interpret multiballots. Our quiz contained one valid multiballot, one overvoted ballot (all three Pat Buchanan bubbles were filled out), and one abstention (all candidates had one bubble filled out). Please see Appendix B for the full text of the survey.

3.2 Comparison of Systems

Accuracy By and large, voters were able to cast votes for their intended candidate(s). Our post-vote survey asked voters to write down what they intended to vote for in each race. Most complied, and there were no cases where a voter's ballot did not match their survey response. Computer-based and paper-based balloting were both accurate systems in our study.

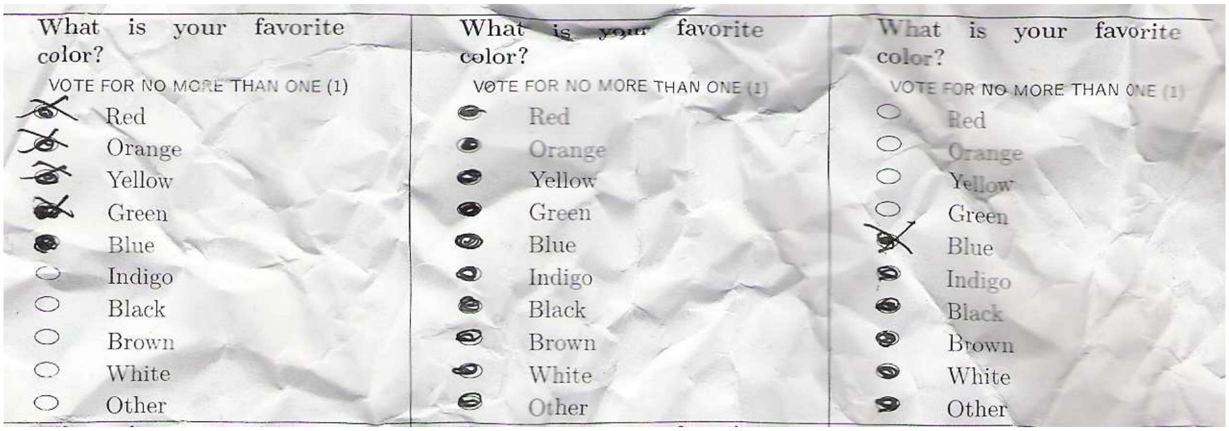

What is your favorite color?
VOTE FOR NO MORE THAN ONE (1)
☒ Red
☒ Orange
☒ Yellow
☒ Green
● Blue
○ Indigo
○ Black
○ Brown
○ White
○ Other

What is your favorite color?
VOTE FOR NO MORE THAN ONE (1)
● Red
● Orange
● Yellow
● Green
● Blue
● Indigo
● Black
● Brown
● White
● Other

What is your favorite color?
VOTE FOR NO MORE THAN ONE (1)
○ Red
○ Orange
○ Yellow
○ Green
☒ Blue
● Indigo
● Black
● Brown
● White
● Other

Figure 2: One voter crumpled up his voided ballot and threw it at the poll workers before leaving in frustration

Efficiency For paper balloting, there was a large incidence of the checker initially rejecting ballots. In the two elections run in Building 32 and East Campus, fifty-one voters voted on paper ballots. Of these, sixteen ballots were initially rejected by the checker. Twelve of the sixteen voters got it right on the second try. One voter submitted to the checker four times before finally submitting a valid ballot, and another voter became frustrated after his second rejected submission, ultimately leaving without casting a vote. For computer voters, one voter mis-clicked and had to re-do her ballot. The computer did not print out any receipts that were rejected by the checker. Thus, the computer was more efficient in terms of votes cast per poll worker time and in terms of total time taken to vote.

Transparency Unfortunately, using computers to assist voters negates many of the advantages of paper-based ThreeBallot. Computers are very susceptible to tampering. In order to attain the same level of security as paper-based ThreeBallot, computer-aided ThreeBallot relies on each voter to read and check his or her receipt for voter and computer error (or malice). Thus, the hope that computers will ease the learning curve for ThreeBallot must be tempered with the knowledge that the security of the system still requires every voter to understand the underlying principles behind ThreeBallot.

Computer-based systems may also exacerbate throughput issues. Whereas optical-scan OneBallot and paper-based ThreeBallot require only paper and pens at the bottleneck point in the voting process, computer-aided ThreeBallot requires one computer per voter throughout the entire voting process.

Prior experience with existing direct-recording electronic voting systems may confuse users of this computer-based ThreeBallot implementation. One voter escaped the polling place with a ballot, which we discovered 24 hours later when one of us noticed a complete computer-generated multiballot away from the polling place. When we asked the voter why she had the ballot, she said that she thought that she only had the receipt, and that the computer had recorded her tally.

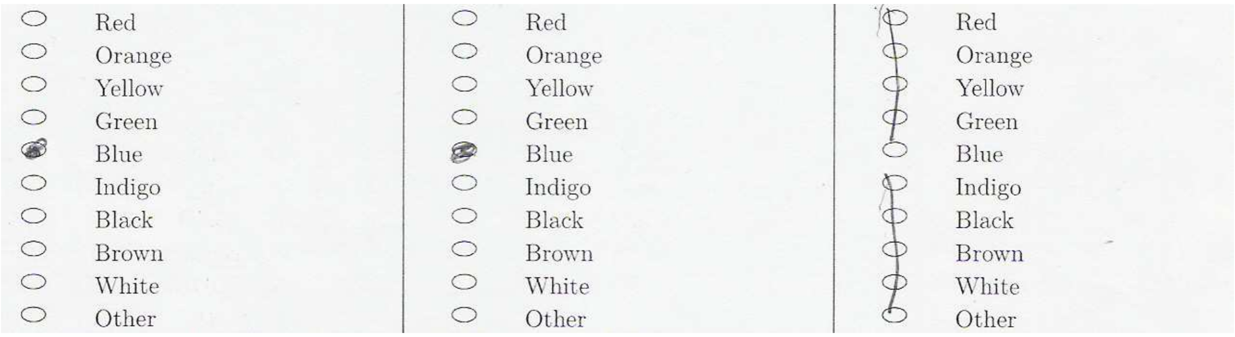

○	Red	○	Red	⦶	Red
○	Orange	○	Orange	⦶	Orange
○	Yellow	○	Yellow	⦶	Yellow
○	Green	○	Green	⦶	Green
●	Blue	●	Blue	○	Blue
○	Indigo	○	Indigo	⦶	Indigo
○	Black	○	Black	⦶	Black
○	Brown	○	Brown	⦶	Brown
○	White	○	White	⦶	White
○	Other	○	Other	⦶	Other

Figure 3: An actual bullet-voted multiballot from our mock elections

3.3 Implementation details

Terminology can be a problem; during the Building 32 election, the ballots instructed voters to “vote for at most one candidate”. Many voters interpreted this to mean that they should only fill in one bubble per race per ballot. When we removed this instruction, voters still treated the election as “vote for at most one”, but were able to fill out the paper ballots much more easily.

In both of the mock elections we conducted, two voters received receipts but left them at the polling place. An adversary that obtained these receipts could have changed these votes with impunity, in the worst case adding a vote for the preferred candidate and subtracting one from the leading candidate in every race. Because the election officials at the polling site have the easiest means of accessing the ballots, it is particularly worrisome when voters leave their receipts at the polling place, within easy reach of potentially corrupt officials. We imagine that dumpster diving near the polling place would also be rewarding, although we did not attempt it during our mock elections.

3.4 User Reactions

The first reaction we received from many voters, upon explaining ThreeBallot for the first time, was a look of disbelief, usually followed by a comment that the ballot was more complicated than they were expecting. Others questioned the need for such a complex system. In order to help address these concerns, we made sure to explain the motivation for using ThreeBallot. In particular, we made sure to mention that each voter would get to keep a copy of part of the ballot as a receipt, and to show each voter a blank receipt as an example.

Despite these efforts, many voters still appeared to be overwhelmed by the complexity of the instructions on the ballot itself, as well as by the shorter, more intuitive directions given orally by poll workers.

Further evidence of voter frustration lies in the large number of observed "bullet votes"—races for which the voter marked just the desired candidate in two of the columns, and marked all the remaining candidates in the third column (see Figure 3 for an example), or voted in a similar pattern.[1] Several voters became impatient with the large number of marks to be made and simply drew vertical lines through entire columns.

In the East Campus mock election, fourteen multiballots, out of thirty-eight paper ballots cast, exhibited a bullet vote in at least one of the two races with five or more candidates.

Many voters were noticeably relieved when told, after having listened to our explanation of ThreeBallot, that they would be voting with computer assistance, although a few insisted that they would have done just fine with a paper ballot.

Question	Yes	Total	Percentage
Secret	61	81	75%
Secure	53	82	65%
Would Share Receipt with a Third Party	69	84	82%
Would Use in Federal Election	44	82	54%

Figure 4: Results of survey regarding voter confidence in ThreeBallot

About forty percent of survey respondents said they would not feel comfortable with instituting ThreeBallot as the primary voting method in a federal election. We did not ask these voters to explain their reasons for saying this, but many volunteered a response. The most popular reason, given by about a quarter of these voters, was the feeling that ThreeBallot would be too confusing for many voters, especially the elderly. The next most popular volunteered reason was a lack of confidence that ThreeBallot has been thoroughly tested and examined, since it is so new.

Three-quarters of survey respondents were confident that ThreeBallot preserves the secrecy of a voter's ballot.

3.5 User Understanding

The results from the "pop quiz" we gave to voters are summarized in Figure 5. It appears that most voters had a basic understanding of ThreeBallot after using it once; the vast majority were able to recognize a valid ballot and a basic overvote. However, more than half of voters did not recognize that marking each candidate's name exactly once is a legal way

[1] More specifically, we defined a "bullet vote" for ThreeBallot to be one of two voting patterns. One case is described and depicted above. The other case is where a voter voted for every candidate on one ballot, voted for no candidates on the second ballot, and chose a single candidate on the third. Furthermore, the term only applies to races with four or more candidates, since in cases with three or fewer candidates, these patterns are very common.

Question	Accurate Responses	Total	Percentage
Identified Valid Ballot	58	64	91%
Identified Triple Vote	59	64	92%
Identified Abstention	19	49	39%

Figure 5: Pop quiz on multiballots

of simply not casting any vote in that parcticular race. [2] Voters failed to recognize this, despite an explicit instruction on the ballot explaining how to abstain from voting in a race, as well as a sample ballot at the polling site demonstrating how to do so.

This indicates that voters can quickly pick up the basics of ThreeBallot and begin using it effectively. It also indicates, however, that voter intuition will not help the voters when presented with novel ThreeBallot scenarios.

Further anecdotal evidence that users can quickly adapt to the system comes from our third mock election; one non-expert had gained enough familiarity with the system by his second time voting to explain the system and check the ballots of five voters.

3.6 Recommendations

Sample Ballots Our experience strongly suggests that sample ballots are very effective in helping voters to understand how to correctly vote using ThreeBallot. However, even when given access to a sample ballot, some voters still had trouble casting a valid ballot unless a poll worker walked through the sample ballot with them and pointed out how each vote on the ballot was cast.

In creating our sample ballot, we simply took a blank ballot from the election and filled it out, creating a valid multiballot. However, in an actual election, this would be illegal in many, if not most, jurisdictions (including the state of Florida [9]), as the mere mention of a candidates name at a polling place is considered to be electioneering. While this problem is easily solved by creating a mock-up of an actual ballot containing false names, it is an additional concern that election administrators should keep in mind.

Clear Terminology Ballot designers and elections officials must be very careful to avoid any confusing terminology. Among the problems we encountered during our mock election was the multiple senses of the word "vote." Some interpreted a single vote as being the act of casting a multiballot and registering one vote for each race, while others interpreted each mark on the ballot as a vote, and so a multiballot contained at least one vote every

[2] We only took the East Campus results for this survey question. After the Building 32 election, it was clear that there was a substantial amount of voter confusion regarding abstention. We modified the sample ballot presented in the East Campus election to contain an "abstain" vote, and we added "Abstain" to the computerized voting machine, in order to make voters aware that abstention was legal.

candidate, and two total votes for some candidates. As a result, directions such as "Vote for at most one candidate" caused confusion, and we imagine that similar instructions in the context of multi-seat elections, such as city council elections, would only result in further confusion.

Shading Rows A common error made by voters was accidentally skipping a row while filling out the ballot, thus failing to mark any bubbles for that candidate. To solve this ballot design issue, one voter suggested shading every other row to make it easier for voters to check over the ballot at a glance.

Training Sessions for Poll Workers While it should be common sense that poll workers need to be trained for the introduction of any new voting system, we want to stress the importance of allowing election officials to operate a mock election, in order to help ensure that everything moves smoothly on election day. Despite our familiarity with ThreeBallot, our second and third mock elections ran much more smoothly than the first, demonstrating the power of just a couple hours of experience.

ThreeBallot is a more complex system than the current "OneBallot" system, and so we would expect that most voters will require more time to vote. In this sense, it is particularly important that poll workers have had some practice in order to avoid any further delays in the system.

Fortunately, our experience tells us that it should not be too difficult to train poll workers. As mentioned above, one voter had gained sufficient experience to act as a poll worker after having experienced the voting process only twice. While this is an exceptional case, we believe that it indicates that training is not excessively difficult.

Gradual Introduction of System On a similar note, any introduction of ThreeBallot should begin with a smaller-scale election, such as a primary or an off-year when no high-profile national races are on the ticket, and be accompanied by a large voter education movement. The gradual introduction of the system over the course of a couple years should ease the transition.

Handling Unwanted Receipts There will be some voters who do not want receipts. Perhaps the voter trusts the system, or is apathetic towards the voting process. Whatever the reason, these voters should not simply leave their receipts, unguarded, at the polling site, since this allows corrupt election officials to tamper with votes, as mentioned earlier.

Several courses of action could be taken here. The voter may be given the option of simply not receiving a receipt at all. The receipt is either never produced, or is destroyed immediately after it is produced. A similar strategy would be to provide a paper shredder at the polling site, so that voters may easily destroy unwanted receipts. Our preferred solution is to use one of these two strategies, and additionally to strongly encourage voters to surrender their receipt to a trusted third party, in lieu of destroying it.

4 6.857 Test

We challenged students in 6.857 to compromise the security and privacy properties of ThreeBallot. Students were able to prove how they voted, opening the door to vote selling and voter coercion. One student acting as a corrupt voting official was able to tamper with the results of the election.

4.1 Methodology

We held a mock election in the 6.857 lecture. Each student was encouraged to try to compromise ThreeBallot in one of two ways. Students could sell their votes, by proving to us how they voted before the complete tally was posted. The use of digital cameras was forbidden.

We also allowed students to play the role of a corrupt election official by submitting to us a fraudulent public bulletin board. We then published the fraudulent bulletin boards and allowed other students to try to use their receipts, or other means, to prove that each bulletin board had been tampered with. We considered such an attack to be successful if the student changed 20% of the ballots cast.

The procedure was as follows. On Sunday afternoon, we e-mailed the class informing them of the mock election, and of these procedures. On Monday morning, at the beginning of lecture, we held the election, then publically posted the results at midnight, 12 hours after the lecture. On Wednesday evening, we sent the fraudulent bulletin boards to the class. Proofs of votes and repudiations of faked results were due by Friday at 11:59 PM. All attempts to sell votes or submit fake bulletin boards were to be sent to a special e-mail address that we created for this purpose.

4.2 Vote Selling

Because the ballot IDs were pre-printed, attackers could memorize their ballot ID numbers (or write them down) and provide them to a third party before vote totals are posted. Thus, the third party can reconstruct the three ballot in question, and be reasonably certain that the vote they are paying for (or are coercing) was accurately cast.

Four students successfully sold their votes using this attack. One suggested that voters write

ID	Race 1	Race 2	Race 3
746159405	+1 Red, -1 O, G, I, Bk.	+1 W, -1 PS3.	+1 Brf, -1 Bxr, Buch.
905778831	+1 O, G, Bl, Br.	+1 PS3.	+1 Bxr
285952325		+1 XBox, PSP, DS.	+1 Brf.
765500665	+1 Red.	+1 W, -1 XBox, PSP, DS, Ot	-1 Buch.
908907354	+1 Red. -1 Bl.	-1 XBox.	
997495774		+1 Wii.	
534252307		+1 Wii.	+1 Brf, -1 Buch.

Figure 6: Compromised Ballots

their ballot numbers somewhere on their body, not only to make detection more difficult, but also to prevent election officials from confiscating the evidence.

4.3 Election Throwing

One student (Yoyo Zhou) combined Strauss's reconstruction techniques with spying on others' receipts: during a break in the lecture, he wandered around the classroom, spying on students and copying down their receipt ID numbers. He then cross-correlated known triples with known receipt numbers, yielding pairs of ballots which he knew to be uncheckable. In other cases he was only able to establish a unique pair, but knew which ID was issued for the pair, and yielded a single uncheckable ballot. He then modified the uncheckable ballots, secure in the knowledge that no receipt existed. By doing so, he was able to swing every race in the mock election.

This attack is particularly worrying, given Rivest's intent that helper organizations be permitted to have copies of the receipt, so that large-scale receipt verification may be performed by a trusted third party [1]. Over eighty percent of voters in our post-vote survey indicated willingness to provide a third party, such as the League of Women Voters, with a copy of their receipt. An attacker posing as the League of Women Voters, or other well-known, trustworthy organization, could easily collect receipt numbers with which to carry out this attack.

Yoyo was able to identify six unique triples, including his own vote, and he ended up changing seven ballots total. By adding votes to his preferred candidate and subtracting votes from the leading candidate, he was able to throw the election without leaving behind evidence of ballot-box stuffing.

In total, Yoyo added three votes for Red, subtracted one from Blue and Indigo, added a vote to Brown, added four votes for the Wii, subtracted one each from the XBox and Other, added three votes to Briefs, and subtracted three from Buchanan. This was sufficient to change the winner of all three contests. We were reconstructing ballots at the same time, and found the triples that he had. We ran our reconstruction code against his results to see if there was evidence of impossible vote totals (more votes cast than voters, ballots with no possible complements, etc.) We couldn't find any evidence of tampering using these methods. Yoyo added votes, but he also subtracted them to keep each race's vote total under 18. He also modified more votes than he needed to, in order to ensure that each ballot still formed part of a valid triple.

4.4 Recommendations

Race-wise split The more information that is encoded onto one ballot, the smaller the probability that a voter's intent will be disguised by other similar ballots that will form a valid three-ballot triple. We suggest splitting the ballots apart, and then assigning each individual race an ID number and splitting those. This vastly increases the difficulty of reconstructing triples, as the number of bits one will have to work with is the number of candidates in a given race. This method will make verifying receipts more tedious, but as the bulk of this work will be done by receipt aggregators in close elections, the impact will be minimal.

No pre-printed IDs Pre-printed ID numbers led to 25% of the class successfully selling their votes in the in-class mock election. Unfortunately, it will be virtually impossible in practice to prevent information from leaking out of the voting booth, given that one would not only have to ban camera phones, but writing utensils as well. Long or complex IDs have troubling implications from a transparency perspective; ThreeBallot is a system designed to be simple enough for voters to understand and verify. A barcode or encrypted ID number will add complexity to the checker, and make it impossible for voters to check the functionality of the checker/receipt printer. We believe that the Shamos checker, or a similar mechanism, is an absolute necessity for a practical ThreeBallot election.

Guarding receipts closely While voters should still be encouraged to share their receipts with trusted third parties, so that large-scale receipt verification can take place, they should also be warned against sharing their receipts with suspicious persons. In particular, voters should be cautioned about leaving their receipts in public places where untrusted parties can view them, or even steal them.

5 Reconstruction

We implemented a ballot matcher based on Strauss's reconstruction attack and ran it on the Building 32, in-class, and East Campus results. Our implementation used the naive $O(n^3)$ algorithm. For each pair of ballots, the matcher searched through all other ballots to see if the resultant triple is a legal multiballot.

For the Building 32 election, we were able to reconstruct 25 of 30 ballots, each of which contained 3 races, with a total of 18 candidates. For the in-class election, we were able to reconstruct 6 of 18 ballots, each of which contained 3 races, with a total of 19 candidates. We also discovered that one voter submitted a multiballot that contained a "zero vote" for a candidate, i.e. the voter had not marked any of the three ballots for that candidate. Somehow, this ballot had slipped through the checker unnoticed. For the East Campus election, 58 ballots were cast, each with 3 races and a total of 14 candidates. We were unable to reconstruct any valid triples, although we were able to narrow down several ballots to 4 possible companions.

The elections we ran consisted of very small ballot sizes (three races, with at most nineteen total candidates). ThreeBallot bulletin boards would be much more reconstructable for elections using longer ballots.

The size of the ballot for our in-class election was similar to that of the Building 32 election, and fewer ballots were cast in class. We would have expected that ballots from the class election would be more easily reconstructed. This turned out to be false. We believe that our relative difficulty in reconstructing in-class results is due to the prevalence of "bullet voting"; the numerous bullet votes resulted in many ballots with similar patterns being cast, making it more difficult to find combinations which had to be unique. One possible explanation for the high number of bullet votes is that every voter had been encouraged to try to sell his or her vote, and voting in an easy-to-remember pattern facilitated the most common vote-selling attack.

6 Receipt Verification

Providing voters with receipts is useless if the voters fail to verify them against the posted election results. In order to determine if voters do, in fact, verify their receipts, we analyzed the webserver logs from our public bulletin board website.

For our first election, in the Stata Center, there were twenty-nine ballots cast. The only way to check a receipt for this election was to type a ballot number into a form on the results website. The logs from the election results server show that only two receipts were checked in the two days following the election.

For the other two mock elections, voters had the option of either typing their ballot number into the form, or simply looking at the list of all ballots and finding their ballot number manually.

Out of eighteen ballots cast in our in-class mock election, three receipts were checked using the form. However, the logs reveal that form submissions came from only two unique IP addresses, so either two students shared a computer to access the website, or the form was used by an attacker who had stolen receipt numbers from other voters. We believe that the former is more likely, since many students taking the class know each other, or even live together.

Seven unique IP addresses, including the two mentioned above, accessed the full listing of all ballots for the in-class election. This was expected, since we encouraged the class to create fraudulent bulletin boards, and doing so with any success would almost certainly require access to the original unspoiled results.

More interesting results come from the mock election held at the East Campus dormitory. Fifty-eight ballots were cast in this mock-election. Seven unique IP addresses checked a total of eight different ballot numbers within the two days after the election was held. Two additional IP addresses accessed the full ballot list without submitting a form, so we estimate that a total of ten voters checked their receipts online. [3]

Certainly, in a real election where voters are motivated to vote and have a vested interest in the outcome, a somewhat higher fraction of voters would take the time to verify their ballots, especially in situations where the outcome may be in question.

Additional measures to increase voter motivation could include offering a bounty for receipts proving election fraud, but as Appel [2] notes, this permits insiders to make a lot of money by changing ballots, even if they can't change the election results. A random drawing for monetary rewards may also encourage voters to check their ballots, but voters may be put off by the idea. An initiative proposal for such a voter reward in Arizona was defeated by a nearly 2-to-1 margin [10] in the November 2006 general election, with many critics arguing that rewarding voters would cheapen democracy's most sacred institution.

7 Conclusion

In this project, we provided a proof-of-concept for computer-based ThreeBallot. We tested the demands of ThreeBallot on poll workers in a real-world election and found that users require several minutes of interaction with a poll worker before they are comfortable with the system. We assessed voters' attitudes towards ThreeBallot, and found that they were suspicious of the complexity, but understood the desirability of secret verifiable elections. We tested users' understanding of the system after having voted, and found that greater than 90% of voters understood ThreeBallot well enough to check another person's ballot. We tested ThreeBallot's robustness against vote selling and election tampering, and found that pre-printed ID numbers and long ballots will pose significant challenges.

Overall, there are reasons for both optimism and pessimism with regards to ThreeBallot's viability in real-world elections. Although there is a definite start-up cost to introducing

[3] It should be noted that we were working on analyzing the mock election and writing this paper in dormitory lounges, in front of several voters. It is quite possible that this prompted several of them to check their results.

ThreeBallot, the learning curve does not appear to be as steep as we feared; once voters understood the system, they were able to accurately cast their ballot in the vast majority of cases. ThreeBallot's usefulness will strongly depend on how much of an over improvement the current system it is, weighed against the problems adopting it will cause. Ultimately, ThreeBallot is a method that trades usability and voter throughput for secrecy and verifiability.

The need to make such a tradeoff will vary depending on local conditions, but it is notable that the three biggest US election controversies in recent years have been regarding usability (2000 Palm Beach ballot), throughput (2004 Ohio elections) and usability/voting machine issues (2006 Florida-13). Although corrupt election officials exist and have been at the forefront of previous election controversies, recent American voting irregularities have centered on subtler issues. Voting officials should be very wary of introducing complex systems without evidence that they are resolving a significant problem.

References

[1] Ronald L. Rivest. The ThreeBallot Voting System. `http://theory.lcs.mit.edu/~rivest/Rivest-TheThreeBallotVotingSystem.pdf%`, October 2006.

[2] Andrew W. Appel. How to defeat Rivest's ThreeBallot Voting System. `http://www.cs.princeton.edu/~appel/papers/DefeatingThreeBallot.pdf`, October 2006.

[3] Charlie Strauss. The trouble with triples: A critical review of the triple ballot (3ballot) scheme. `http://www.cs.princeton.edu/~appel/voting/Strauss-TroubleWithTriples.pd%f`, October 2006.

[4] David Chaum. Secret-ballot receipts: True voter-verifiable elections. *Security & Privacy Magazine, IEEE*, 2(1):38–47, Jan.-Feb. 2004.

[5] Punchscan. `http://punchscan.org/`.

[6] Ben Adida and Ronald L. Rivest. Scratch & vote: self-contained paper-based cryptographic voting. In *WPES '06: Proceedings of the 5th ACM workshop on Privacy in electronic society*, pages 29–40, New York, NY, USA, 2006. ACM Press.

[7] Punchscan FAQ. `http://punchscan.org/faq-general.php`.

[8] Charlie Strauss. A critical review of the triple ballot voting system, Part 2: Cracking the triple ballot encryption. `http://www.cs.princeton.edu/~appel/voting/Strauss-ThreeBallotCritique2v%1.5.pdf`, October 2006.

[9] Florida Statutes, Section 102.031. `http://www.leg.state.fl.us/Statutes/`, accessed December 2006.

[10] `http://www.azsos.gov/results/2006/general/BM200.htm`, November 2006.

[11] Lawrence Norden, et al. The machinery of democracy: Usability of voting systems. Technical report, Brennan Center for Justice at NYU School of Law, August 2006.

[12] Benjamin B. Bederson, Bongshin Lee, Robert M. Sherman, Paul S. Herrnson, and Richard G. Niemi. Electronic voting system usability issues. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 145–152, New York, NY, USA, 2003. ACM Press.

Appendices

A Sample Multiballot

INSTRUCTIONS

Each candidate or choice on this ballot is duplicated three times. For each office or ballot issue:

- To vote **for** a candidate: fill in **exactly two** of the three ovals next to that candidate's name.
- Fill in **exactly one** of the three ovals for every other candidate in that race.
- To **abstain** fill in **exactly one** of the three ovals for *every* candidate in that race.

Be sure to fill in ovals completely, and avoid making any stray marks outside the oval.

WARNING! You **must** mark at least one and at most two ovals in **every** row. A ballot on which any row has zero or three marks **is not valid** and will not be accepted.

Once you have finished voting, bring your ballot to the Checker to verify that your ballot has been filled out correctly, and to obtain a receipt. You may use this receipt to prove that you voted, and to ensure that your vote was correctly recorded and counted.

—CUT ALONG THIS LINE WHEN DONE VOTING—

What is your favorite color?	**What is your favorite color?**	**What is your favorite color?**
● Red	○ Red	○ Red
○ Orange	○ Orange	● Orange
○ Yellow	● Yellow	○ Yellow
● Green	● Green	○ Green
○ Blue	● Blue	○ Blue
○ Indigo	○ Indigo	● Indigo
○ Black	● Black	○ Black
○ Brown	○ Brown	● Brown
○ White	● White	○ White
● Other	○ Other	○ Other
What is your favorite video game system?	**What is your favorite video game system?**	**What is your favorite video game system?**
○ Playstation III	● Playstation III	○ Playstation III
● Nintendo Wii	○ Nintendo Wii	○ Nintendo Wii
○ XBox 360	● XBox 360	○ XBox 360
○ PlayStation Portable	○ PlayStation Portable	● PlayStation Portable
○ Nintendo DS	● Nintendo DS	○ Nintendo DS
● Other	● Other	○ Other
Best baseball team?	**Best baseball team?**	**Best baseball team?**
○ Yankees	○ Yankees	● Yankees
● Red Sox	● Red Sox	○ Red Sox
○ Pat Buchanan	● Pat Buchanan	○ Pat Buchanan
`Ballot ID: 212010336`	`Ballot ID: 240232886`	`Ballot ID: 623451381`
END OF BALLOT	**END OF BALLOT**	**END OF BALLOT**

B Survey

Post-vote survey

What did you vote for?

Would you be comfortable with giving a copy of your receipt to a third party, such as the League of Women Voters?

Yes No

Are you confident that this system preserves the secrecy of your ballot?

Yes No

Are you confident that this system is secure and accurate?

Yes No

Would you feel comfortable if this voting system was used in your precinct for a federal election?

Yes No

How well do you feel you understand how this system works? (1 = Not at all, 10 = Perfectly)

1 2 3 4 5 6 7 8 9 10

Do you have additional comments about your voting experience?

Who is your favorite AFC East team?	**Who is your favorite AFC East team?**	**Who is your favorite AFC East team?**
VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)
● Buffalo Bills	○ Buffalo Bills	○ Buffalo Bills
○ Miami Dolphins	○ Miami Dolphins	● Miami Dolphins
○ New England Patriots	● New England Patriots	● New England Patriots
○ New York Jets	● New York Jets	○ New York Jets
END OF BALLOT	**END OF BALLOT**	**END OF BALLOT**

Is this a valid ballot? If so, for whom is it a vote?

Yes No

Who is your favorite AL East team?	**Who is your favorite AL East team?**	**Who is your favorite AL East team?**
VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)
○ Boston Red Sox	○ Boston Red Sox	● Boston Red Sox
○ New York Yankees	● New York Yankees	○ New York Yankees
● Pat Buchanan	● Pat Buchanan	● Pat Buchanan
END OF BALLOT	**END OF BALLOT**	**END OF BALLOT**

Is this a valid ballot? If so, for whom is it a vote?

Yes No

What are your favorite colors?	**What are your favorite colors?**	**What are your favorite colors?**
VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)	VOTE FOR NO MORE THAN ONE (1)
● Blue	○ Blue	○ Blue
○ Green	● Green	○ Green
● No, blue!	○ No, blue!	○ No, blue!
○ African	● African	○ African
● European	○ European	○ European
○ I don't know	● I don't know	○ I don't know
END OF BALLOT	**END OF BALLOT**	**END OF BALLOT**

Is this a valid ballot? If so, for whom is it a vote?

Yes No