

Testing Voters' Understanding of a Security Mechanism used in Verifiable Voting

Morgan Llewellyn¹, Steve Schneider², Zhe Xia³, Chris Culnane², James Heather², Peter Y.A. Ryan⁴ and Sriramkrishnan Srinivasan²

¹IMT Lucca Piazza S. Ponziano, 6, 55100 Lucca, Italy

²University of Surrey, Guildford, Surrey, GU2 7XH, UK

³Wuhan University of Technology, China

⁴University of Luxembourg, Luxembourg

Abstract

Proposals for a secure voting technology can involve new mechanisms or procedures designed to provide greater ballot secrecy or verifiability. These mechanisms may be justified on the technical level, but researchers and voting officials must also consider how voters will understand these technical details, and how understanding may affect interaction with the voting systems. In the context of verifiable voting, there is an additional impetus for this consideration as voters are provided with an additional choice; whether or not to verify their ballot. It is possible that differences in voter understanding of the voting technology or verification mechanism may drive differences in voter behaviour; particularly at the point of verification. In the event that voter understanding partially explains voter decisions to verify their ballot, then variance in voter understanding will lead to predictable differences in the way voters interact with the voting technology.

This paper describes an experiment designed to test voters' understanding of the 'split ballot', a particular mechanism at the heart of the secure voting system Prêt à Voter, used to provide both vote secrecy and voter verifiability. We used a controlled laboratory experiment in which voter behaviour in the experiment is dependent on their understanding of the secrecy mechanism for ballots. We found that a two-thirds majority of the participants expressed a confident comprehension of the secrecy of their ballot; indicating an appropriate level of understanding. Among the remaining third of participants, most exhibited a behaviour indicating a comprehension of the security mechanism, but were less confident in their understanding. A small number did not comprehend the system. We discuss the implications of this finding for the deployment of such voting systems.

1 Motivation

In response to problems with the counting and certifying of ballots in countries such as the United States of America and Iran, academics and politicians have worked together to create and implement new voting technologies. Many of these technologies provide new mechanisms or procedures that are designed to provide greater ballot security and opportunities for voters to verify their ballot. While technically sophisticated, there can arise a disconnect between technical understandings of the threats and an understanding of how voters will accept and interact with the technology. This disconnect is in some sense surprising as one motivation for technical work is the lack of voter confidence in some current voting technologies.

Typical concerns of the technical variety include distinguishing between trusted and untrusted agents [7], security and access to publicly posted ballot data, threats stemming from the inclusion of proprietary software, and the lifetime of the encryption. While valid, these concerns primarily consider threats in the context of the voting technology and only tangentially consider the impact of voter understanding of the selected voting technology. It is likely that most voters will avoid the technical deliberation and alternatively focus on the evaluation of non-technical questions such as, “Can I use it?”, “Is my vote choice secret?” and “Will my vote count?” [5, 2, 6]. If voters believe a new security or verification mechanism endangers ballot secrecy or accuracy, these beliefs may influence voter interaction with the technology; such as not voting or choosing not to verify their vote.

For instance, a loss of confidence in the butterfly ballot following the 2000 U.S. Presidential election, is partially responsible for its being driven into oblivion in U.S. elections [1]. In India, election officials have sought to increase public confidence in the election process by adopting electronic voting machines as a method to overcome “booth capturing” which is often associated with paper ballots. If voters believe, rightly or wrongly, that a voting technology leaks ballot information or that the technology provides an advantage to one political party, these voters may move to block the use of the technology or under some regimes alter their voting behaviour. In fact, research on American voters suggests individuals without voting experience may alter their turnout decision based upon their perceptions of ballot secrecy [4]. In light of the role voter confidence can

have on the voting process as a whole, we wish to evaluate voter understanding of ballot secrecy and the subsequent impact on verifiable voting technologies.

2 Verifiable Voting

One basic requirement for a verifiable voting technology is to provide a voter with the ability to check that his/her vote has been counted and verify that the election tally is correct. While straightforward in the absence of a secret ballot, verifiable voting is more difficult to achieve under a secret ballot. Despite the creation of several voting schemes which allow for verifiable voting with a secret ballot, it is unclear how voters will response to these “sophisticated” voting technologies. In particular, it is an open question how skeptical voters, those who do not understand the underlying security and secrecy mechanism(s), will react when given the choice to verify a ballot. One possibility is skeptical voters may accept election official guarantees or verify out of interest. Another possibility is those skeptical voters may simply refrain from verifying their ballot, or worse exit the electoral process. This interaction between a voter’s understanding of the voting technology’s security mechanisms and their decision to verify their ballot are at the heart of any serious debate seeking to implement a verifiable voting technology.

In order to evaluate how voters may behave during a verifiable election, we evaluate the rationality of individual decisions to verify their ballot through the lens of their understanding of ballot secrecy. The verifiable voting technology used in this context is a basic implementation of Prêt à Voter [8]. In the version considered here, Prêt à Voter represents a paper-based voting system that allows voters to verify their ballot was properly recorded in the vote tally. Prêt à Voter was chosen due to the simplicity of the mechanism that ensures the “receipt” or posted information contains no information. The Prêt à Voter scheme requires that individuals who wish to verify their ballot retain a portion of their ballot.

In our design, voters possessing strong beliefs over ballot secrecy during the election and verification procedures will have an incentive to post their receipt to a public bulletin board.¹ Whereas, voters that possess significant doubts over

¹The posting of the ballot to a public bulletin board is a step found in many variants of verifiable voting.

either the ballot or verification system’s ability to maintain ballot secrecy will possess an incentive to refrain from posting their receipt. We suspect that those participants with significant uncertainty over ballot secrecy may fluctuate between posting and not posting their ballot.

3 Prêt à Voter

We now give a high level overview of the classic Prêt à Voter system [8, 3, 9]. The ballot form consists of two columns with a perforation vertically down the middle. The left hand side (LHS) lists the candidate names in a random order and the candidate ordering varies from ballot to ballot. The voter can use the right hand side (RHS) to mark her choice, and at the bottom of the RHS, there is an encrypted value. When it is decrypted, the corresponding candidate ordering on the LHS can be retrieved. An example of the ballot form is shown in Figure 1.

Bob	
Echo	
Crystal	
David	
Alice	
	7q3Kyr

Figure 1: A Prêt à Voter ballot form example

On the election day, each authenticated voter will be provided with such a ballot form, in secret, for example in a sealed envelope. She takes it into a voting booth, and marks her choice on the RHS against her preferred candidate. After that, she separates the ballot form into two halves along the perforation and shreds the LHS which contains the candidate names. Then the voter takes the remaining RHS to the election officials who will scan it into the election system. Finally, the RHS will be returned to the voter, and she can take it home as her receipt.

After the election day, the election system will publish all the received votes onto a public bulletin board, which may be understood as like a newspaper: once information is published, it cannot be removed and it will be available to the

public. The voter can now check whether her receipt is correctly displayed on the bulletin board. If not, she can use her receipt as a proof of her vote to challenge the election.

The key innovation of the Prêt à Voter system is that each voter is provided with a receipt. The receipt contains the voter’s vote, but only in encrypted form. Hence it does not reveal how the voter has voted. Meanwhile, thanks to the receipt, the voter does not need to trust the election system to correctly include her vote, since the receipt can be used to enforce this through verification.

4 Methodology

Whether voters understand the security of Prêt à Voter and how differences in understanding affect voter interaction and behaviour is an open question. We hypothesize that voters will understand the security mechanism of Prêt à Voter. We developed a simple game to test our hypothesis where the actions of the participants will vary depending upon their understanding. In particular, in the context of Prêt à Voter we are interested in evaluating whether voters understand the choice of making public their receipt will not reveal their vote choice.

To test our hypothesis we slightly modified the Prêt à Voter protocol and designed a game theoretic experiment where participant decisions to verify their ballot truthfully reveal their understanding of the security mechanism. Monetary rewards are used to motivate the subjects to behave truthfully and to take actions that are in the subject’s perceived best interest. The game works as follows: within a group of 12 subjects each subject casts a vote in a fictitious election. When marking their choices, each subject will also need to select whether they wish to “post” their receipt (anonymously) so that all subjects can see it. In this case, a receipt will be similar to that shown in Figure 2. The incentives are structured such that if a subject chooses to post her receipt, she will receive a reward of £1 (the amount is denoted as \mathcal{A}). Otherwise, she will receive nothing. After all subjects have cast their votes, all those who selected to post receipts will have their receipts made public. Once the receipt is made public, each subject makes a guess of every participant’s vote choice, whether or not that participant published her receipt. This subject will receive a reward of £.50 for each correct guess and zero otherwise. Moreover, she will lose an amount of £.50 for every

Post:	<input checked="" type="checkbox"/>
Not post:	<input type="checkbox"/>
X	
7q3Kyr	

Figure 2: A receipt example

participant who correctly guesses her vote choice. Note that the game design and reward values are chosen so that the experiment can separate those who believe that the receipt divulges no information from those who believe that the receipt divulges some information.

The game's economic incentives will induce subjects to self-reveal their understanding of the security mechanism. As the receipt contains no information, the dominant strategy is for subjects to post their ballot in each and every round. Given the receipt provides no information over vote choice and there is a reward for posting one's ballots, voters who possess a full or high level of understanding about the security mechanism will choose to post their ballot. However, a voter who does not fully understand the security mechanism may be afraid that the publication of her receipt will allow the other voters to guess her correct vote choice. Hence she may choose not to publish her receipt.

5 Subjects & Location

A total of five experiments were run at the University of Surrey, UK. The dates of the experiments were the 28th and 30th of June and the 1st of July, 2011; with two experiments on each of the 30th and 1st. Three experiments were conducted during lunchtime, with an additional afternoon experiment on each of the 30th and 1st. Each experiment contained twelve subjects and lasted about one and a half hours.

Subjects were recruited during the week prior to the experiment on the University of Surrey campus via fliers and email. Prospective subjects were informed they would be paid between £10 and £20. While some university employees did express interest in participating, the majority of responses to the call for subjects came from the student body.

The experiment location was the seminar room within the Department of Computing at the University of Surrey. This location was chosen as it was both large enough to allow the installation of the voting equipment and was equipped with the proper technology necessary for the experiment. A designated voting area was built within the seminar room containing a voter registration table, ballot box, and three voting booths. The voting booths were equipped with walls to ensure privacy, a table for marking the ballot and a shredder for destroying the left-hand portion of the ballot.

In addition to the designated voting area, a staging area was constructed where subjects read the experiment instructions (reproduced in Appendix A), waited for their turn in the voting booth, and participated in the game following voting. The staging area consisted of 12 rectangular tables facing a whiteboard. These twelve tables were spaced such that no individual seated at one desk could read a piece of paper on any other desk.

As subjects arrived they were assigned an individual table in the staging area. The subject was handed a copy of the instructions and told that when all the subjects arrived the instructions would be read aloud to the entire group.

Before the start of each experiment, the instructions were read aloud to the entire group. Following the instructions, participants answered two questions to verify they understood the purpose of the experiment, to assess understanding of the voting technology, and to verify each participant understood how they would be compensated. Next the randomized ballot ordering was explained and displayed to each participant. This demonstration consisted of displaying multiple ballots simultaneously on an overhead projector and discussing how the placement of a candidate's name on any one ballot varied randomly between ballots. After all subjects made a vote choice over all the ballots, individual payoffs for the round were averaged for those who chose to post their ballot and those subjects who did not. The averages of these two groups were then publicly posted for the subjects to

observe. After the posting of the average payoffs, the next round of voting started.

Principal steps in the experiment:

1. Each subject receives a ballot and takes it into the polling booth;
2. The subject fills in her vote and chooses whether to post her receipt;
3. The subject separates the ballot into two halves, shreds its LHS and drops its RHS into a ballot box;
4. The receipts are made public for those who choose to post their receipts;
5. Each subject attempts to identify the vote choice of every participant, including herself;
6. Average payoffs for the posting group and the non-posting group are announced;
7. Start a new round from Step 1 if the experiment is not finished.

6 Experiment Results

In total there were 29 rounds of voting, with 5 rounds in the June 28th experiment and 6 rounds in each subsequent experiment. Participant earnings ranged between £10 - 20. On average the experiments took approximately 1.5 hours from start to finish.

Given the game design and the careful choice of payoffs, participants who fully understood the security mechanism possessed an incentive to always post their ballot—as participants were paid £1 for posting their ballot. However, participants possessing beliefs that viewing the right-hand side of the ballot helps others guess their vote choice should be less likely or unwilling to post their ballot. This behaviour was expected as a result of the game’s payoffs; participants were penalized for each participant who correctly guessed their vote choice. Thus, we expected that individuals who possessed significant doubts or a high degree of uncertainty over the secrecy of the election technology should either refrain from posting or engage in a mixing strategy; switch between posting and not posting.

Table 1: Ballot Posting by Round: All Experiments

	Did	
	Not Post	Posted
Round 1	13%	87%
Round 2	18%	82%
Round 3	15%	85%
Round 4	15%	85%
Round 5	7%	93%
Round 6	8%	92%

Of the 348 votes cast, in 87% of the votes cast the voter chose to publicly post the right-hand side of their ballot. The breakdown of posting by round is displayed in Table 1. In each of the six rounds, the majority of subjects chose to post their ballot. On average, there were 1 to 2 people who did not to post their ballot in any one round, out of a possible 12 people. We anticipated that by the later rounds all participants would post their ballot. While Table 1 does indicate that the proportion of participants who post their ballot is increasing over time, the size of this statistic is insignificant.² We conclude that even in the later rounds, round five or six, there is one individual, on average, who does not post their ballot. This implies that even after observing multiple occurrences of the voting process, posting of results, and ballot guessing game, there remains some doubt over the amount of information contained in the right-hand side of the ballot.

Analyzing the posting behaviour between the five experiments reveals similar rates of non-posting activity. In each of the five experiments, there was little variance in the total number of non-posted ballots; between 7 and 10 ballots. While posting behaviour between rounds varied slightly across experiments, there was remarkably little variance in the aggregate non-posting behaviour at the experiment level. The consistency of the posting behaviour across the five experiments lends credibility to the robustness of the result; i.e. more experiments are unlikely to yield different results.

²A simple logistic regression of round on subject posting behaviour reveals a positive but insignificant coefficient for the variable round.

Table 2: Posting behaviour by Participant Type

	High Certainty	Low Certainty
High comprehension	Always Post	Usually Post
Low comprehension	Never Post	Usually not Post

While in the aggregate the security mechanism appears fairly well understood, as 87% of ballots were posted, we analyze individual decisions to evaluate two dimensions of participant perceptions over the voting system. Table 2 separates participant understanding of the security mechanism along two dimensions; comprehension and certainty. Participants may either comprehend the security mechanism or not, and these participants may possess either a high or low degree of certainty over this assessment. The comprehension dimension approximates a participants posting behaviour within any one round, while the certainty dimension helps explain participant behaviour over multiple rounds. Participants with a high degree of comprehension and high level of certainty will exhibit a behaviour consistent with always posting; this person will post in each and every round. On the other-hand, participants with a low degree of comprehension and a high level of certainty will never post their ballot. While individuals with a high degree of certainty will play the same strategy across rounds, we expect that individuals with a low level of certainty will mix between posting and not posting in an effort to reduce their exposure to the wrong strategy. Thus, individuals with a low degree of certainty over their comprehension will be observed mixing between the two posting behaviours; where high certainty types will always, or never, post.

Analyzing the individual posting behaviour, the majority of subjects posted their ballot in every single period. Out of a total 60 subjects, 62% (37 participants) posted their ballot at every opportunity. These results indicate that a majority of participants possessed a high comprehension of the security mechanism and possessed a high degree of certainty over this assessment. This data suggests that a majority of individuals took actions consistent with those of someone who understands that the right-hand side ballot contained little or no information.

Of the 23 participants who did not post a ballot in at least one round, we find evidence of only two subjects who made decisions consistent with an individual

possessing a low level of comprehension and a high level of certainty over this assessment. That is out of a total of 60 participants only two subjects never posted a ballot. This behaviour is consistent with a profile that does not understand the security mechanisms provided by candidate order randomization.

The remaining 21 participants are classified as low certainty individuals, due to their posting behaviour deviating in at least one round. Of the 21 participants classified within the low certainty group, no individual posted in less than 50% of possible opportunities. These results indicate that while roughly one-third of participants possessed some uncertainty over their comprehension of the security mechanism, in no case was this assessment so low as to warrant that individual posting in fewer than 50% of the rounds. These results indicate that even among uncertain participants the overall level of understanding was fairly high.




Figure 3: Posting Behaviour by Subject

Despite the posting of average payoffs following each round, we did not observe wild swings in behaviour between rounds as a result of the previous round's outcome. Due to the low number of participants who did not post their ballot, there were several incidents where the individual who did not post their ballot received a higher average payoff than those who posted their ballot. Despite reporting this result, we did not observe a large change in the posting behaviour during subsequent round. Additionally, despite showing that on average

individuals who posted their ballots tended to receive higher payoffs, even in the sixth round we observed two participants choosing not to post their ballot. We conclude that aggregate and outside information may have an affect on individuals with higher levels of uncertainty, but these information flows may not affect those who are either rightly or wrongly are certain of their understanding.

7 Conclusion

The majority of participants took actions consistent with those of a voter fully understanding the security mechanism of Prêt à Voter. However, we did show that differences in understanding can directly affect voter interaction with the voting technology. Approximately one-third of participants took actions consistent with a voter expressing a high level of comprehension over the security mechanism, but a corresponding low level of certainty over this assessment. Our findings present initial evidence that voter interaction with a verifiable voting technology may significantly vary by the individual voter’s understanding of the technology and their confidence in this understanding. This finding raises the additional concern that, if implemented, a Prêt à Voter style voting system may be vulnerable through indirect attacks via voter beliefs over the secrecy of the voting process. While any such attack may be reduced via public information campaigns and party support, additional research is needed to determine the size and significance of these strategies.

Given the high educational level of the participants, detailed instructions, repeated nature of the experiment, and financial incentives, we hypothesize that in a general election setting the level of understanding will likely be lower than that which we observed. However, in a general election it is also uncertain to what degree understanding of the security mechanism and will simply be replaced with a voter’s notion of trust. Additional research is needed to better understand the interaction of voter understanding of the election technology and voter trust in that technology. While results indicate the transfer of information has a limited affect on individuals ”certain” in their understanding of the security mechanism, it is necessary to further study this behaviour.

Acknowledgements

We are grateful to Vincent Koenig, and to the anonymous referees, for comments on an earlier draft of this paper. This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/G025797 (Trustworthy Voting Systems) and the FNR (National Research Fund) Luxembourg under project SeRTVS-C09/IS/06.

References

- [1] R. Michael Alvarez, Thad E. Hall, and Morgan H. Llewellyn. Are americans confidence their ballots are counted? *The Journal of Politics*, pages 754–766, 2008.
- [2] R. Michael Alvarez, Gabriel Katz, Ricardo Llamosa, and Hugo E. Martinez. Assessing voters’ attitudes towards electronic voting in latin america: Evidence from colombia’s 2007 e-voting pilot. *E-Voting and Identity*, pages 75–91, 2009.
- [3] David Chaum, Peter Y. A. Ryan, and Steve A. Schneider. A practical voter-verifiable election scheme. *Proceedings of the 10th European Symposium on Research in Computer Science (ESORICS’05)*, pages 118–139, 2005. LNCS 3679.
- [4] Alan S. Gerber, Gregory A. Huber, David Doherty, and Conor M. Dowling. Is there a secret ballot? Ballot secrecy perceptions and their implications for voting behaviour. *British Journal of Political Science*, 43(1), 2011.
- [5] Anne-Marie Oostveen and Peter Van den Besselaar. Internet voting technologies and civic participation: the users’ perspective. *The Public*, pages 61–78, 2004.
- [6] Anne-Marie Oostveen and Peter Van den Besselaar. Users’ experiences with e-voting: a comparative case study. *International Journal of Electronic Governance*, pages 357–377, 2009.
- [7] Ronald L. Rivest. Electronic voting. *Proceedings of Financial Cryptography (FC’01)*, pages 243–268, 2001. LNCS 2339.

- [8] Peter Y. A. Ryan, David Bismark, James Heather, Steve Schneider, and Zhe Xia. Prêt à Voter: a Voter-Verifiable Voting System. *In IEEE Transactions on Information Forensics and Security (Special Issue on Electronic Voting)*, 4(4):662–673, 2009.
- [9] Peter Y. A. Ryan and Steve A. Schneider. Prêt à Voter with re-encryption mixes. *Proceedings of the 11th European Symposium on Research in Computer Science (ESORICS'06)*, pages 313–326, 2006. LNCS 4189.

A Voting instructions

The following instructions were provided to the participants to read when they arrived at the experiment, and were read out when they were all present, before the start of the experiment. The questions at the end were used as a self-test for the participants to check that they understood the instructions.

University of Surrey Voting Experiment Instructions

General. You are about to participate in a voting process experiment in which you will cast a ballot for one of numerous alternatives. The purpose of the experiment is to gain insight into your understanding of the voting technology and features of the ballot form. The instructions are simple. You will be paid at the conclusion of the experiment. You have the right to withdraw from the study at any time and all identifiable data and information will be confidential. This study has received a favourable ethical opinion from the University of Surrey Ethics Committee. If you have any complaints or concerns about any aspect of this experiment please contact Professor Steve Schneider, s.schneider@surrey.ac.uk, 01483 689637.

Overview. You are about to participate in a voting experiment. Your compensation for this experiment will depend upon the decisions you make. While some aspects of the voting process may resemble those you have encountered in the past, there are some differences. We therefore ask you to follow these instructions and ask any questions that may arise during the course of these instructions. We kindly ask you to refrain from conversation during the experiment.

The voting procedure you will take part in tests a ballot form where the left-hand side of the ballot contains a random ordering of candidate names. The experiment consists of several rounds and the random candidate ordering is independent within and across rounds. That is in each round, the ordering of the candidates on your ballot is unrelated to the ordering of candidates on the other participants' ballots both within and across all previous rounds.

Instructions to Participants.

At the beginning of the experiment you will be allocated £10. This experiment will consist of a set number of rounds. Each round will consist of two PHASES:

PHASE I

1. First you will vote for a candidate. To vote for a candidate you will enter the voting booth and place a mark in the appropriate box to the right of the candidates name.
2. Second, you must choose whether to publicly post your ballot. To publicly post your ballot mark the box that says "Post". To decline to post your ballot mark the box that says "Don't Post". In each round, if you post publicly post your ballot you will receive £1.00. If you do not publicly post your ballot in a round, you will not receive the £1.00.
3. Next, you will separate the ballot along the perforated edge. You will then shred the left-hand side of the ballot; the portion containing the candidates' names.
4. Deposit the right-hand side of your ballot in the ballot box provided.

PHASE II

1. Next you will select the vote choices of the other participants using the electronic handset provided. If an individual chose to publicly post the right-hand side of their ballot, you will view the right-hand portion of their ballot prior to selecting their vote choice. For individuals who chose not to post the right-hand side of their ballot, you will be asked to select their vote choice but will not view the right-hand side of the ballot.

2. You will win £0.50 for each vote choice you correctly identify. You will lose £0.50 for each individual who correctly identifies your vote choice.

The experimenters will keep track of payments and obtain the totals to pay participants. You will learn your total earnings only after the completion of the experiment.

After each round some information will be publicly posted. The first piece of information is the average payoff for the group that publicly posted their ballot. The second piece of information is the average payoff for the group that did not publicly post the ballot.

Are there any questions? We kindly ask you to complete the following questions as these should help you understand the instructions.

Questions

- 1) The purpose of this experiment is to study which item?
 - a) Who you will vote for in the next election.
 - b) Your social interaction in an election setting.
 - c) Your understanding of the voting technology and features of the ballot form.
- 2) What three activities comprise your compensation?