PAPER TRAILS, CRYPTOGRAPHY AND OTHER APPROACHES TO VOTE VERIFICATION

Paul Herrnson

Let me start out by thanking Shalyn for the invitation, and everyone else their warm hospitality. I haven't been to Albany in a while. I'm originally from New York; born in the Bronx, raised on Long Island, and it's good to have my feet on home turf.

The title of my talk is *Paper Trails, Cryptography and Other Approaches to Vote Verification*. The first thing to ask is where did the demand for vote verification systems come from? And the answer is: the introduction of electronic voting systems has reinvigorated some traditional concerns.

As you know, Kim Brace has been a leader in documenting the adoption of new voting systems in the U.S. Direct Recording Electronic (DRE) voting systems, sometimes referred to as computerized systems, touch-screen systems, or ATM-style systems, were in use prior to the 2000 presidential election and the passage of the Help America Vote Act, but those two events stimulated their use. Once they became widely used, some people began to have some concerns about these systems. One of them is security.

There has always been a concern in United States elections that ballots could be altered or stolen. It is probably easier to do this with paper ballots or punch cards, or even maybe mechanical systems, but the stealing of large numbers of paper ballots is pretty detectable, that is unless they are replaced by fraudulent ones. The concern with the direct recording electronic voting systems is theft may not be as detectable.

Then there's a concern with reliability. Some systems, particularly electronic ones, or lever ones, can break down. And so there is concern that we need to have a backup. With DRE and mechanical systems, there is no backup, but the presence of the paper itself, in terms of paper ballot systems, is very reassuring.

It's a physical medium. With a DRE system, should it crash or should there be some other problem, people would want reassurance that their votes were counted. At least that is what many would argue. Now as of mid-2007, roughly 27 states required all voting systems, including computerized ones, to have some kind of paper back up.

So this brings us to the study of vote verification. And the questions that are generally raised, and ones we have addressed in our study, with Dick Niemi as a co-PI, are: Do verification systems improve the voting experience? Do they increase voter satisfaction? Are people more confident that their votes are accurately recorded, that they will be accurately counted, and that their privacy will be protected? These issues are very important to election administrators who are considering mating some kind of verification system to a DRE system.

Another set of questions we addressed is: Does the verification system detract from the voting process? Does the addition of vote verification systems make it more difficult to vote? Do they distract voters from what they are doing? Do they increase the need for help? Such need could have the potential to violate someone's privacy because an election official would become involved in the casting of that person's ballot. Does the addition of a verification system increase the time needed to vote, which is of obvious concern to election officials and voters standing in line? And finally, do vote verification systems improve voters' abilities to cast their votes as intended? What good is a voting machine if you go to cast your ballot on it, and you push a bunch of buttons thinking you're voting for candidate A, and you mistakenly vote for candidate B? Ask Al Gore, voting accuracy is very important.

There are many different types of voter verification systems on the market or in various stages of development. We tested several of them. Recall that these systems are added on to an existing voting system to make it possible for the voter to verify that the buttons he or she has pushed on the voting system to select their preferred candidates are going to be recorded as votes on their ballot. In addition to assuring voting accuracy they are designed to reassure people. They are also designed to make it possible for recounts, in many cases, so an election can be audited.

The vote verification systems we tested are illustrative of the design features of all commercially available systems and several prototypes. The first is a system that has a voter verifiable paper audit trail, a VPAT, or VVPAT, or sometimes called a receipt.

The second system, and you can see pictures of the system in your packet, is one that has a separate monitor. The system does two things: First, after you vote it shows your selections, which you can compare to those cast of the voting system; Second, while you are voting, should the selection you make on the voting system fail to match what is being recorded by the voting system's memory and on the vote verification system, the verification system will alert you and election officials that there is something wrong with the voting system's programming.

The next system is a basic audio system, a tape recorder and headphones. The last system is a cryptographic system, which has a computer in it. It has the same vote-checking feature as the system with the monitor. That is, if you make a candidate selection and the vote recorded is different from the selection you made, the system will alert you. But the key feature of this system is that it gives each voter a receipt with a bunch of codes they can use to check whether their votes have been counted using a toll-free number or the internet.

We tested these four verification systems in comparison with a control system, that is, a voting system that had no vote verification system on it. All of the verification systems were connected to an interface that was either the Diebold AccuVote-TS DRE System, or a system that mimicked that system. The AccuVote-TS was also the voting system that had no verification system. Thus, the basic voting system paired with the different verification systems was the same. You can see pictures of these in the handout.

The Diebold AccuVote-TS is a DRE system that is widely used, and as you may recall, has been severely criticized for security and the politics of the company's president. The next slide shows a picture of the model of the AccuVote-TS with the AccuVote View Printer Module. A voter uses this system just like any other touch-screen or ATM-style DRE system. When the voter has cast his or her ballot, the system prints a paper under a glass cover that shows who the person voted for. Voters can use that paper to verify that their votes were recorded as cast. That paper is meant to reassure voters.

One basic problem with the paper verification system is that the paper comes on a roll. Thus, the order in which you cast your ballot is kept intact. If someone kept track of who voted on each system, they could violate your privacy. Another problem is the paper spool is difficult to change and is prone to jam. Also, as is

the case with paper ballots, the rolls of paper recording individuals' votes can be stolen.

The next verification system, which uses the extra computer monitor, is the Scytl Pnyx. After someone finishes voting, their selections are presented on the monitor, one at a time, enabling the voter to compare the selections on the screen with those they believed they made. One of the shortcomings of this system, we tested a prototype, is it did not provide the name of the write-in candidate. It just said write-in. Other shortcomings included the size of the extra monitor: It was too small for some voters to read, but large enough for some individuals beside the person voting to look at.

The next system, the audio system with the little tape recorder and the headphones, was developed at MIT. It is a very simple and inexpensive system. When a voter selects a candidate, such as Joe Smith, the system informs the voter in a computer-simulated voice that he or she voted for Joe Smith. If the voter decides to change their selection and deselects the candidate, the system states you deleted your vote for Joe Smith. Then, when a new selection is made, say for Paul Herrnson, the simulated voice will state that you voted for Paul Herrnson. The tape serves as a backup that can be used for election audits. The biggest drawback to this system is that the simulated voice can be very irritating and it can butcher ethnic names and long names. Another shortcoming is all votes are kept in the order in which they were cast. This is the same problem associated with the Deibold DRE with the printer module.

The last system we tested, the boxy-looking cryptographic system in the hand out, is the Vote Here Sentinel. Explaining how this system works would take many hours, and I still may not be able to explain its operation to some people's satisfaction. In short, once a voter casts a ballot, it gives them a receipt with some codes on it. Voters can take that receipt home and dial in to an 800 telephone number, or check on the internet, communicate their codes, and learn whether their vote was counted. That is the system's standard verification protocol. There is also an advanced verification system that can be used, usually in the polling place, to let a voter confirm the selections they made for each office.

We used two techniques to study these vote verification systems. First, we did what is referred to as an expert review. This involved employing human-computer interaction experts, the type of researchers that make and test consoles for airplanes, spaceships, and cars to examine how voters interacted with the vote

verification systems. Some of the experts we hired had expertise with voting systems. All of them assessed the strengths and weaknesses of each verification system.

Second, we conducted a field study. That is what I'll focus on today. The field study consisted of 850 voters who put the machines through their paces. We did the testing in a variety of locations in Maryland in order to involve a diverse sample of participants. We recruited voters from office buildings, shopping malls, senior citizen centers, and various locations on universities. We gave them a ballot that was relatively short but was normal in other respects.

Our protocol described some candidates and asked voters to choose among them based on their preferences. It also directed them in some situations to choose one specific candidate from a few. We asked voters to vote for more than one candidate in a race because in several jurisdictions, including in Maryland, voters are asked to vote for more than one person for state delegate and some other offices. Finally, we asked voters to change a vote, which is something that some individuals do in the voting booth. We put together a questionnaire with several core questions that applied to all of the vote verification systems and some specific questions designed to assess the unique aspects of several of them.

Our field study worked as follows. We recruited individuals to participate in it. The first thing we did was give the individual a voter an information sheet, similar to a mock ballot distributed in newspapers or by election officials, and ask him or her to circle their choices. Using this information sheet, the voter would vote on a voting system coupled with a vote verification system or the control system. Then the voter would answer a questionnaire about the system. Then they would repeat this process on each system until they had voted on them all. When they were done voting and evaluating all of the vote verification systems, we gave the voters a questionnaire to record their background characteristics and prior voting experiences.

One of the things we did was to give each person participating in the study a unique name to use for their write-in vote. Using the write-in vote, we could match the votes a person cast on each system to the intended votes they had circled on their information sheet. We were also able to link their intended votes and actual votes to the information they provided on their demographic voting experience questionnaires. This enabled us to learn whether the vote they intended to cast was the same as the vote they actually

cast, and to assess whether voter characteristics were systematically related to their evaluations of the vote verification systems or the accuracy of their votes.

What are the findings of the study? The chart with the bar graphs presents measures of the average voter satisfaction with each of the verification systems and the control system. The satisfaction measure includes such things as voter confidence that their votes were accurately recorded, confidence that their votes were counted properly, and confidence that their ballot was cast privately. The standard for assessing the vote verification system is the control system, the Diebold AccuVote-TS with no verification system. It scored about 5.65 on a scale of 1 to 7, where 1 was the lowest possible rating and 7 was the highest possible rating. The only system to score higher than the control system was the Diebold AccuVote-TS with the AccuView Printer Module - the paper trail system. The other vote verification systems, which are intended to improve confidence in these aspects of voting, actually scored worse than the system with no verification unit attached to it.

The next finding shows the average ratings for ease of use. It is important to note that we measured only the ease of use of the vote verification systems, not the control system or the voting system itself. Once again, the paper trail system is rated better than the other verification systems. The Vote Here Sentinel (the cryptographic) system was also rated fairly well. The system with the independent computer monitor and the audio systems were not judged as easy to use.

The next finding concerns how distracting, on average, voters found each verification system. Again, the paper system did quite well. It was judged as the least distracting. The cryptographic and the computer monitor system were judged slightly more distracting. The MIT audio system was considered to be the most distracting. Why? Because after each vote is cast, the voting process is interrupted while the voter is told who he or she voted for.

The next item we tested was the need for help. Workers on the study recorded whether a study participant needed help using the vote verification system. So the percentages of voters asking for help in the slide refer only to those who needed help using a vote verification system. This can be thought of as the amount of help a voter needs above and beyond the help they might receive when using the basic voting system. When people vote on a voting

system without a verification system, they may need help. If a printer, a cryptographic system, independent monitor system or an audio verification system is added to a voting system, the number of voters who request assistance is between 4.5 and 7 percent age points greater than if the voters had used a voting system that had no vote verification system. This increase has tremendous implications for the administration of elections, particularly the staffing of polling places.

The next finding concerns the accuracy of voting that occurred when voters used each vote verification system versus the control system. Accuracy refers to a voter's ability to cast his or her votes as intended. The findings for this criterion are, generally speaking, that the addition of a vote verification system did not significantly improve voter accuracy. When using the system with the paper trail, the cryptographic system, or the system with the independent monitor, voters made slightly fewer errors on than on the control system with no verification unit. On the margin, these systems may have improved voter accuracy slightly. The audio system, which was the most intrusive in terms of the voting process, led to the commission of more voting errors.

I'll make a couple more points before I wrap up. First, in terms of administrative issues, electronic voting systems are complicated in and of themselves. When a paper trail, a computer monitor, a cryptographic unit, a tape recorder, or some other verification system is added, the number of steps that are needed to set up, maintain, or take down a voting system increases significantly. The Scytl Pnyx was especially difficult to program and integrate into the voting system. A computer programmer worked for days to accomplish this. The paper roll on the Diebold AccuVote-TS with AccuVote View Printer Module was very difficult to change. One might think that dropping a paper spool into a printer would be simple. Well, it is not simple. Complicating the process was the fact that the printer would only print one side of the paper.

Second, some of the vote verification systems crashed while in use. Think about this. These systems are designed to increase voter confidence, yet while a voter is checking a ballot using them they stopped working. Election officials (actually, stand-ins who staffed our field study) came over to investigate the problem while someone was voting. This was not very reassuring.

To conclude, in testing a variety of vote verification systems, we learned a number of things. First, voters do not rate them as significantly better than voting systems having no voter

verification units. Indeed, all but the paper system was rated worse. Second, they resulted in more voters needing help while casting their ballots. Third, some verification systems improved voting accuracy slightly; one made voting less accurate. Fourth, all posed significant administrative challenges. On balance, these systems did not improve the voting process.

Now, there is currently a great deal of political pressure on political decision makers to adopt voting systems with paper trails or to add verification systems to DRE voting systems. Some of the pressure comes from a vocal group of computer scientists; some of it comes from advocates. There's also some resistance from other computer scientists and from election officials that have to purchase and deploy voting systems.

New York has to make some challenging decisions in terms of the voting system it selects, and whether or not the system will have some form of independent verification system, such as those tested in our study. There will be pressure to adopt such systems, but adopting systems may introduce more problems than are solved. Thank you.