Election Security: Perception and Reality

Voters' trust in elections comes from a combination of the mechanisms and procedures we use to record and tally votes, and their confidence in election officials' competence and honesty. Electronic voting systems pose considerable risks to both the perception and reality of trustworthy elections.

David Evans and Nathanael Paul
University of Virginia

Democracy depends on losers accepting the results of elections. Citizens' confidence that votes are cast anonymously and without coercion, and that the reported election results accurately reflect the collective will of the voters is essential. Although trustworthy elections are essential to democracy, achieving them requires balancing security, cost, and convenience. Voting technologies have a substantial impact on both the actual and perceived security of elections.

With a plain paper-based voting system, voters can rely on some aspects of the process based solely on their own actions and observations. Voters know that the ballot they cast accurately reflects their intent because they can examine that ballot themselves. Furthermore, they know that a physical record of their vote exists. That record cannot be destroyed, lost, or tampered with without leaving some physical evidence. (See the "Trust in the Election Process" sidebar for a review of trust issues in small- and large-scale elections.)

Voting systems that do not produce a physical record, such as mechanical-lever and electronic-voting machines, create additional trust issues. We lose transparent verifiability and must trust that the machines function correctly. This expands the scope of trust from the local election officials to include the manufacturers who make those machines as well as the people and processes used to inspect, maintain, and operate them.

Electronic voting also increases the potential for large-scale fraud. If many voting machines run the same software, and no mechanisms exist for voters to verify their votes are recorded correctly or for election officials to conduct a meaningful recount, an intentional or accidental flaw in that software can irrevocably affect an election's outcome.

Voting technologies

Plain paper ballots are becoming rare in the US—approximately 1 percent of the population used a paper ballot in 2000.[1] The alleged problem with hand-marked and hand-counted paper ballots is the time required to count them. As ballots become more complicated with multiple offices and propositions in a single election, humans find it more difficult to reliably and quickly count hand-marked paper ballots.

To enable mechanized counting, alternatives to paper require voters to cast ballots so that officials can easily tabulate votes. Punch cards gained prominence in the late 1960s because of their speed of tabulation and low costs.[2] Voters use a punch card by punching holes through the card with a punch device. Votomatic and datavote are two styles of punch-card ballots. In the votomatic system, the punch card contains many rows of holes that correspond to particular choices. The voter inserts the card into a slot in the punch-card holder, which contains several pages indicating the actual issues. In the datavote system, the voter receives one or more punch cards with the choices printed on the card itself. The votomatic system sacrifices voter verifiability for reduced cost. Although a voter can examine the punch card and attempt to determine if her vote is correctly recorded, it is difficult to do so. The holes on the card do not have labels that identify candidates directly; to verify the punch card, the voter must map the card holes back to the ballot. With the datavote system, voters can more easily verify their choices because the candidate names are printed on the punch card. However, this requires printing custom punch cards for each election. Furthermore, a typical election will require more than one punch card per voter, which leads to

Trust in the Election Process

Secret ballot elections must satisfy the competing goals of anonymity and integrity. Anonymity requires that voters can't be associated with the votes they cast. Strong integrity guarantees would let all participants verify the final tally accurately reflects all legitimate votes cast. In addition, to limit the possibility of vote coercion and selling, it should be impossible for voters to prove to others how they voted.

On a small scale, a simple election process satisfies all these requirements:

1. All eligible voters gather in a room. Voters determine each others' eligibility to vote based on personal acquaintance.
2. A collection of identical ballots and identical markers is produced. A ballot box is opened to show all voters that it is empty and then closed.
3. Each voter selects one ballot and marker. Ballot distribution is done in plain view of all the voters to ensure that no voter takes more than one and to prevent any association between voters and ballots.
4. Each voter marks the ballot in a clear (but indistinguishable) way to indicate her vote. Voters mark their ballots behind a privacy curtain and then conceal their vote before exiting the voting booth. Voters are observed to ensure that they do not display their vote to anyone else.
5. All voters deposit their ballots in the ballot box. The depositing is observed by the other voters to ensure that voters don't examine, remove, or tamper with other ballots when they deposit their own ballots and to prevent voters from depositing more than one ballot into the ballot box.
6. In plain view, each ballot is removed from the ballot box, presented to all the voters to view and recorded in the tally, which is maintained in plain view of all the voters.

This process gives voters a high degree of confidence that their votes are anonymous because the ballots and markers being identical and the ballots are dissociated from the voters when they are mixed in the ballot box. Even with such a process, though, voters still rely on trust to some degree. For example, voters believe no hidden cameras are watching them mark their ballot, no pressure sensors under the table on which they mark their ballot, and the ballots are not analyzed for fingerprints after the election. Of course, if these are serious concerns, countermeasures could be employed against each of these threats such as burning all the ballots after they are counted. Furthermore, this process provides only minimal protection against vote coercion—a voter could uniquely mark a ballot so that it would be recognizable to the coercer when revealed for counting. If coercion is a serious concern, it would be possible to require voters mark ballots in a standard, simple way, and only ballots that follow the prescribed marking rules are counted.

Voters can trust the count because of the transparency in the process. Every voter can observe that only eligible voters deposit ballots in the ballot box, the ballot box was not tampered with during the process, the votes counted correspond to the votes cast, and the counting is done correctly.

This process doesn't scale to large elections when all the voters don't know each other, can't fit in the same room, and can't observe the counting process. We can mimic the secure small-scale election process in large-scale elections with a voter experience like this: Go to the local precinct, present documents for voting authorization and identification, receive a ballot, walk to an enclosed voting booth and vote, drop the ballot in a ballot box, and then leave.

With such a process, instead of observing the entire election, voters must rely on trusted election officials to ensure the election's integrity. For example, the confidence that only eligible voters can vote and that no eligible voters are prevented from voting is now based on a complex registration and authentication process. Voters trust that the poll workers will only permit eligible voters to vote and that all eligible voters will be permitted to vote. There are well-known cases where this has not been the case. Mayor Richard J. Daley of Chicago reputedly influenced the result of the Illinois vote in the 1960 presidential election for John F. Kennedy by using such tactics as having local precinct captains register and vote for dead people. Before the 2000 presidential election, Florida hired a private firm to remove convicted felons from the voter rolls. It is suspected that over 50,000 nonfelon eligible voters (disproportionately minorities) were incorrectly removed.

For anonymity, voters' trust that all the ballots are initially identical and that there is no way to associate their ballot with their identity. Generally, anonymity at the level of voting precincts is satisfactory. If voters require anonymity beyond their voting precinct, they are trusting that ballots cast at different precincts are mixed up before votes are revealed. Voters trust that the ballot box into which they deposit their ballot initially contained no ballots, only eligible voters were permitted to deposit cast ballots into the box, each voter could deposit only one ballot, the ballots in the box are counted accurately at the end of the election, and the correct results are reported. Voters trust the election officials and the procedures they follow to ensure all these properties.

logistical difficulties and new possibilities for error.

The problems with punch-card ballots became well-known after the state of Florida's 2000 US Presidential election. Because voters might not completely remove punch-card holes, it can be unclear from a punch card what the voter intended. Unlike permanent markings on paper, punch-card ballots are susceptible to accidental tampering if loosely attached chads (the holes voters punch out) are shaken free. Although we can count punch cards automatically, quickly, and cheaply, the results of those counts are not necessarily accurate. In fact, two tabulations from a punch-card machine rarely produce exactly the same count.[2]

Regardless of whether punch cards can record votes

accurately, voters have lost faith in them, which makes them unacceptable. After the Florida elections served to destroy voter confidence in punch-card systems, the US government passed a law encouraging states to replace their punch-card and mechanical-lever systems. The Help America Vote Act of 2002 (HAVA) allocated US$3.86 billion for election upgrades. According to the HAVA act, US states that accept funds must replace their existing punch-card and mechanical-lever voting machines.

A newer type of voting system that attempts to combine the clarity and transparency of paper ballots with the mechanized counting of punch cards is the optical mark-sense voting system. Twenty-eight percent of US voters used this system in 2000.[1] Instead of making a physical hole in the paper as in a punch-card-type system, voters indicate their choice by filling in an oval or completing an arrow. Because voters mark the ballot in such a way that clearly relates to their actual vote, they can easily examine their ballot and determine if the ballot reflects their intentions.

Nevertheless, voters commonly mark a ballot in a way that cannot be scanned correctly. If the ballot is scanned at the precinct before the voter deposits it, the scanner can warn the voter of any undervotes (the voter did not indicate a choice for all elections) or overvotes (the voter selected more choices than permitted), and give the voter an opportunity to change or replace the ballot before depositing it.

However, scanning is often done at a central site after the voter has left the poll site. Hence, the voter trusts the scanning machine to correctly interpret the ballot but knows that a human could also examine the ballot. Manual recounts with optical-scan machines, punch cards, and paper ballots require subjective judgment to ascertain the voter's intent. All these paper-based solutions distribute a voter's trust among a group of poll workers and election observers. Because a physical record of the vote exists, it is likely to require many corrupt election officials to perpetrate widespread fraud. Although some possibility exists that the physical record does not accurately reflect the voter's intent (especially with votomatic-style punch cards), voters' ability to personally examine their vote and deposit it in a secure ballot box makes large-scale fraud difficult. The election's integrity depends on the security of the ballot box and the counting process' accuracy, and multiple election officials and independent observers are responsible for ensuring this.

Paperless Voting

With paperless technologies, trust becomes more centralized. Because no physical record of votes is kept, more opportunity exists for votes to be lost, altered, or inserted without detection.

Mechanical-lever machines, first made in 1892, were used by a majority of US voters by the 1960s but by only 18 percent of voters in 2000.[1] With these machines, a voter pulls a lever associated with a candidate, and the lever causes a counter wheel to rotate inside the machine. Election officials read and record the counter wheels' tallies at the close of the election. The machines offer good privacy— a handle opens and closes the booth's curtain and resets all levers to the "off" position when leaving the booth. Anonymity is guaranteed as long as many voters use each lever machine and no one can observe the levers in the booth or the counter wheels prior to, during, or just after the time of the vote. The voters' trust that their vote is recorded is based largely on the satisfying thunk the machine makes when they pull the lever.

The problem with mechanical-lever machines is that after the voter leaves the booth, no physical record of the voter's intent exists. If the voting machine worked correctly, the values of the counters associated with the voter's selections all increased by one, and the values of all other counters remained the same. If not, there is no way to perform a manual recount. In practice, mechanical voting machines have been notoriously inaccurate. The Caltech/MIT voting technology report[1] found that mechanical-lever machines had significantly higher residual vote numbers than any other voting technology in Senate and gubernatorial elections, although the residual vote counts in presidential elections were similar to those with other technologies.

The electronic equivalent of a mechanical-lever machine is a direct recording electronic (DRE) voting machine. Many of the reasons for the increased adoption of DRE machines include accessibility and prevention of voter mistakes. DRE machines can use audio interfaces to let the visually impaired vote without assistance. We can program DRE machines to prevent overvotes and to warn voters about undervotes. Furthermore, DRE machines save precincts the costs associated with producing and securing paper ballots.

On the other hand, DRE machines raise serious security concerns. They make the election process less verifiable and greatly expand the aspects of an election for which voters must rely solely on trust. As with a mechanical-lever machine, no physical record of a voter's intent exists. Unlike a mechanical-lever machine, however, the mechanism for recording a vote is hidden in the code for the machine, which vendors keep secret.

Voting interfaces

When it comes to voting, usability and security are closely intertwined. An election's integrity depends on the recorded votes accurately reflecting the voter's intent. This could be compromised either by tampering with the recording of voter intent or by interfaces that increase the probability that the recorded votes will not accurately reflect the voter's intent.

A notorious example is the butterfly ballot design used in Palm Beach County, Florida, in the 2000 US presidential election, which made it easy for voters to mistakenly record their intent. This design used the votomatic punch-card ballot, which makes it difficult for

voters to verify that their ballot reports their intent. One report estimated that 2,000 votes in Palm Beach that were intended for Democratic candidate Al Gore were mistakenly recorded for Republican candidate Pat Buchanan.[3]

A voter interacts with a DRE machine through a user interface, often using a touch-screen display. The goal of minimizing voter error can increase the voting machine software's complexity and conflict with ease of use. For example, DRE machines can require a voter to confirm an undervote, but this requires an additional step from the voter.

Complex interfaces also make pre-election ballot reviews more difficult. With paper ballots, it is easy to print and publicly review sample ballots before an election. With DRE equipment, it is harder to review the ballot presentation because it is in the form of a complex user interface. A series of screenshots can't capture an interface fully, and ballot issues might not be apparent without conducting test votes using the DRE machines (which exposes the machines and raises other security issues). These issues were apparent in the recent California gubernatorial recall election in which DRE machines in Alameda County were programmed so that voters could not view the instructions after they began voting. This might have increased the number of voters who voted against the recall but did not cast a vote for a replacement candidate, even though they were allowed to do so.

Recent studies conducted in Georgia and Maryland concluded that although most voters can use DRE machines without difficulty, a significant proportion of voters, especially older ones, required assistance. The Carl Vinson Institute of Government conducted a public opinion telephone survey to study voter confidence in DRE machines in Georgia.[4] They found that fewer than 2 percent of responders reported difficulties in using DRE touch-screen machines. A University of Maryland study[5] conducted an exit poll on voters using Diebold's AccuVote-TS touch-screen DRE machines in two counties in Maryland. Three percent of voters encountered technical problems with the machines, 7 percent reported that they were not easy to use, and 9 percent asked for assistance using the machines. Difficulty with the interface was correlated with age and education. Twenty-one percent of the voters 65 years or older asked for help; the lowest age group asking for help was those 35 to 49, who asked for help 5 percent of the time. The youngest voters, ages 18 to 24, were second highest in asking for help at 16 percent, but this might be largely due to inexperience with voting in general. Of those with no college experience, 18 percent asked for assistance. Voters with a four-year degree or some college experience asked for help 9 percent of the time, and only 8 percent of voters with graduate school education asked for help.

The amount of assistance required does play a role in voter trust in a voting system because that help will usually come from a poll-site worker. Voters who ask for help risk compromising their anonymity, and voters who need assistance might be reluctant to ask for it because of this or just personal embarrassment. This study's results indicate that many voters who did not ask for help received help anyway. This likely indicates that these voters were closely observed by poll-site workers trying to help, which some voters might interpret as a violation of privacy.

Vote recording

Given a user interface that voters believe lets them enter their vote without error, DRE machines' trustworthiness depends on how accurately the recorded vote reflects the entered vote. The trust citizens place in DRE machines depends on their experience using them as a voter and their understanding (or misunderstanding) of how the machines and the surrounding process works.

One of the reasons voters trust DRE machines is their surface resemblance to ATMs. After all, if we can trust an electronic machine to count money, surely we can trust it to count votes. The fallacy in this argument is the difference in accountability. With an ATM machine, the user receives a paper receipt as well as a monthly bank statement. If any discrepancies exist, the customer can dispute the statement with the bank—in the US, it is the bank's responsibility to prove the transaction record is correct. With a DRE machine, there is no receipt, no transaction statement, and no way for a voter to dispute the recorded results.

The Maryland study[5] asked voters if they felt confident that their vote was recorded according to their intent, and 10 percent of respondents did not feel confident that their vote was accurately recorded. The study also asked voters if they trusted the mechanical-lever or punch-card system used in previous elections. Compared to 90.7 percent of voters who trusted the DRE machines used in the election just conducted, only 70.5 percent of voters trusted the mechanical-lever or punch-card system they used in previous elections.

Always vote for principle, though you may vote alone, and you may cherish the sweetest reflection that your vote is never lost.

—John Quincy Adams

The Georgia phone survey reported that 70 percent of Georgians in 2002 were "very confident" their vote was accurately recorded (23 percent were "somewhat confident") compared to 56 percent in 2001.[4] The trust citizens place in particular voting technologies correlates significantly with race; in the same survey 79 percent of whites were very confident their vote was accurately counted

while only 40 percent of blacks were very confident.

Both studies show consistently higher voter confidence with DRE equipment than previous systems, although one in 10 voters (according to the Maryland study, or worse for the Georgia study) lacks confidence that their vote is recorded accurately. Both studies were conducted prior to recent media coverage regarding problems with electronic voting, including security flaws found in commonly used DRE machines.[6] Anecdotal evidence from the recent California gubernatorial recall election suggests that many voters lack confidence in DRE machines.[7]

DRE machines' actual trustworthiness depends on many properties that are invisible to the voter. Unlike paper-based voting systems where voters can personally examine the physical record of their vote and deposit it in a secure ballot box, voter trust in DRE equipment depends on trusting the voting machine hardware and software in combination with the people and procedures designed to safeguard it.

Increasing trust

Several mechanisms have been proposed to provide voters with increased confidence that their vote is cast as intended and that the votes are tabulated correctly. The procedures protecting the voting process as well as the actual process can enhance trust. A particular design decision's impact on security is often different from the impact perceived by typical voters.

Measures that increase the perception of security often do not significantly increase actual security. To perform a recount electronically for votes that only exist as data on a computer, the trusted official can push a recount button, but because this is merely recounting the electronically recorded votes, it provides little actual benefit. Independently stored audit records that record each individual vote can provide somewhat more trust enhancement but only against computational counting mistakes or careless fraud that modifies only the vote totals and not the audit records. These recounts have little security benefit against accidental or malicious programming errors in the vote recording process because they can affect both the counts and audit records.

On the other hand, measures that might increase actual security might not increase perceived security. Voting systems that depend on cryptographic techniques will improve voter trust only if officials present those cryptographic mechanisms to the public clearly and convincingly.

Verification procedures

We can use verification procedures to increase confidence in voting machines. In paperless voting systems, the voters depend entirely on the voting machine to record their votes correctly. Although officials are rapidly phasing out mechanical-lever machines, it is instructive to consider the procedures necessary to adequately verify a mechanical-lever machine and compare those to the difficulties of verifying DRE equipment.

When a ballot is cast, a mechanical-lever voting machine should increment by one all counters corresponding to the voter's selections and not change any other counter. A physical inspection of the machine mechanics should reveal any indirect connections between levers and counters. It is reasonable to require a straightforward, independent path between each lever and its corresponding counters. If the mechanics are hidden, however, it would be necessary to try all possible combinations of lever settings and counter values to ensure there are no unexpected interactions. A typical lever machine has over 100 levers and each lever has two positions, so testing all possible combinations of lever settings ($2^{100} \approx 10^{30}$) is infeasible even without considering possible interactions between lever settings and counters. (In fact, there are reasons to be suspicious of the mechanical counters because they require more force to increment counters when many wheels must turn such as when the count advances from 999 to 1,000.) If we can determine mechanically that all levers and their corresponding counters are independent, the verification process is reduced to independently checking that each lever and counter behave appropriately.

Verifying DRE machines is substantially more complex than testing mechanical-lever machines. The state space for a mechanical-lever machine is fairly well defined by the lever and counter settings. By contrast, the state space for a DRE machine includes all election data recorded by the machine as well as all internal data. DRE machines often run on top of complex operating systems such as Windows CE, so the operating system's internal state might also affect the machine's behavior. Usability demands complex interfaces that require many lines of code, often involving multiple threads. In addition, the DRE machine's state also includes the ballot definition, which is a complex input file specifying election parameters and candidates. We can program a DRE machine to act differently based on that description, unlike a mechanical-lever machine where the ballot is affixed to the surface.

Unlike mechanical-lever machines, it is impossible to examine a DRE machine and have high confidence that components will not interact unexpectedly. Mechanical interactions are physically apparent, but with software systems, the interactions between components are often subtle and unexpected, even to the system's programmers. Yet subtle programming mistakes can impact an election's result, especially if the same software is deployed on voting machines in thousands of precincts. A programmer could deliberately introduce a subtle bug into the software running the DRE machine where the first few hundred votes are counted correctly and some randomly selected small fraction of later votes are intentionally misrecorded to vote for a different party than the voter intended.

In many states, companies known as independent testing authorities (ITAs) certify voting machines, following US Federal Election Commission guidelines.[8] This certification process includes inspecting source code that the voting sys-

tem vendor produces. The inspection guidelines specify coding standards (such as using dynamic memory management, checking array bounds and pointer references, and naming conventions) and design guidelines (small modules that can be tested independently). A thorough inspection can increase confidence that code is correct, but it does not prove the absence of intentionally malicious logic. Because the coding standards let DRE software be written in type unsafe languages such as C, it is especially difficult to inspect code for correctness and noninterference.

A clever programmer can construct buggy code that would pass even a thorough code inspection but would behave incorrectly in a particular situation. For example, one of the inspection guidelines (guideline 5.4.2t) specifies, "code should use explicit comparisons in all `if()` and `while()` conditions. For instance, `if(flag)` is prohibited, and shall be written in the format `if (flag == TRUE)`." This is highly questionable because in C/C++ the value 0 is interpreted as false and all other values are interpreted as true. Hence, the `== TRUE` comparison will make the wrong decision if nonzero values other than `TRUE` are used. A subtle attack would exploit this, possibly taking advantage of components that return nonzero values other than `TRUE`. Importantly, the inspection guidelines exclude commercial-off-the-shelf (COTS) compilers used to produce the voting machine executable and operating systems running on the voting machine. The vendor has no control over the COTS components, but any software in the machine can adversely affect the voting machine. Conspiracies in which a programmer at Microsoft designs an operating system function specifically to tamper with votes in a possible voting machine application are possible, but unlikely. Unintentional bugs with underlying operating systems are more likely, and they might lead to flaws that someone could exploit or to dysfunctional machines.

The value of any certification process depends on procedures to ensure that the machines used are identical to the machines that were certified. With mechanical-lever machines, we can use a tamper evident seal to ensure that the machine's mechanics cannot be tampered with without detection. With DRE machines, however, this is a much harder problem. Unlike mechanical tampering, software changes are not obviously apparent. Because ballot definition files must be loaded into the machine, there must be a way of changing what is in the machine's memory and any changes to the voting machine code will not be apparent. Although election procedures are designed to limit access to voting machines to trusted election officials, voting machines are often kept in insecure locations. Procedures in Alameda County for the recent California recall election left voting machines unattended in insecure poll sites with ballots loaded for several days before the election.[9]

In addition to certification, officials test voting machines at poll sites. These tests, known as logic-and-accuracy tests typically involve poll workers casting a series of test votes on voting machines and then checking that the reported tally is consistent with the entered votes. Testers might make mistakes when entering numerous votes, so if the reported tally is off by one, this might not raise alarm. Because the number of test votes is limited by the testers' patience, it is unlikely that testers would detect malicious voting software that records the first few hundred votes correctly and then incorrectly counts later votes.

Another approach uses automated testing. This increases the number of votes possible and reduces the chances of testing mistakes. However, it makes the testing easily distinguished from normal use by the software and lets a clever programmer inject a bug that appears only in normal use—some DRE machines even specify a test mode, so that the tests run different code from production use. Testing procedures specify when officials should do the logic and accuracy tests, generally before and after the election, but not during it. Malicious software could be programmed to count votes correctly at all times except during the middle of election days.

Local testing is important to detect flawed equipment, but it is easily thwarted by a sophisticated attacker. Any confidence it provides to voters and election officials is largely misguided. Inspection provides somewhat greater confidence but only in conjunction with procedures that ensure the voting machines' actual code matches the inspected code and that any uninspected components included come from trustworthy sources. Because private companies conduct the inspections and results are not public, public confidence in inspections depends largely on trust in those companies' integrity and expertise.

Voter-verifiable ballots

One way to decrease the trust voters must place in voting machine software is to let voters physically verify that their intent is recorded correctly. Rebecca Mercuri has proposed a method for voter-verifiable ballots. After a voter has finished making selections using a DRE machine, the machine prints out a paper ballot that contains the voter's selections for each choice. The printed ballot is kept behind a window to prevent voters from having any opportunity to tamper with it. Voters can examine the ballot and confirm that it accurately reflects their selections. If voters approve the ballot, they press a button to confirm their vote and observe the printed ballot drop into an opaque ballot box. If voters do not approve the ballot, they must consult a poll worker to void the ballot and vote again. This process provides voters with a high degree of confidence that their intended vote was accurately recorded. The paper ballots also provide a mechanism for validating results reported by the electronic voting machine. They can be manually counted or electronically counted to confirm the results if there is a dispute regarding the election results. Some elections

should also be randomly selected for counting paper ballots. Recently, some jurisdictions including California have adopted requirements that DRE machines produce a voter verified paper audit trail (see http://www.ss.ca.gov/executive/press_releases/2003/03_106.pdf).

Voter-verifiable paper ballots eliminate the need to trust the voting machine software to correctly record the voter's intent. Furthermore, they provide voters with substantially increased confidence that their intended vote will be counted. They are not, however, without cost. The need to support printing and collection of paper ballots increases the maintenance costs and election complexity for the poll workers. Voters must perform two steps to complete their vote: the first confirmation prints the ballot, and the second confirmation casts the vote and deposits the ballot in the ballot box. With a fully integrated voting booth, it would be possible to include a curtain that opens only after the second confirmation is complete. Without this, voters might assume they are finished after the first confirmation and the next voter would have an opportunity to examine the previous voter's ballot and decide whether to cast it or void it.

It's not the people who vote that count. It's the people who count the votes.

—Joseph Stalin

Another issue that must be considered carefully is that there are now two potentially contradictory records of the election: the data stored in the DRE machine and the paper ballots. Any discrepancies between the results reported by the DRE machine and the paper audit count must be examined carefully and would likely lead to controversies. Except in cases where evidence of tampering with the paper ballots exists, the paper ballots should be considered the official record of the election because the voters had an opportunity to confirm that they recorded their intent correctly.

Voter-verifiable results

In the US, citizens have generally accepted election results because they trust that the process of tallying ballots is done securely and transparently. They also trust that representatives of both major political parties and independent organizations, such as the League of Women Voters, monitor the process. Voters trust the results without obtaining any direct evidence that their vote was included correctly.

In emerging democracies especially, but also in the US as elections are increasingly scrutinized, more explicit mechanisms might be necessary to guarantee acceptance of the reported results. The strongest mechanisms would provide individual voters with the ability to know that their vote is accurately included in the tally. We must do this without compromising the voter's anonymity or providing voters with a means to prove to someone else how they voted. A simple, but unsatisfactory way to do this would be to provide each voter with a unique, randomly assigned ballot number, and then to publish a transcript of each ballot number and vote after the election. Voters would be able to check the published transcript to verify that their vote was included and recorded correctly. In the event that it was not, however, the voter could not prove to anyone else that his vote was not included. A voter could, however, prove to a coercer how she voted by revealing her ballot number to the coercer before the election transcript is published. Furthermore, there is no proof that forged votes were not added to the tally.

Hence, we must provide voters with a receipt that lets them check that their vote is included correctly or prove otherwise to an election official if it was not, without enabling a voter prove to a coercer how she voted.

Two systems that propose to provide this are VoteHere's VHTi[10] and David Chaum's SureVote.[11] (See page **XX** for Chaum's "Secret Ballot Receipts and Transparent Integrity" article.) Both systems provide voters with a coded receipt that reflects their vote but does not reveal it to anyone else. The systems distribute trust among a group of election trustees. In Chaum's scheme, the voting machine prints the receipt in two layers that are encoded using visual cryptography. When laminated together, they reveal the voter's selections, but each layer independently is meaningless dots. The voter examines the laminated receipt before approving the vote and then selects one of the two layers to retain as a receipt. The machine then prints the remainder of the receipt, including a barcode that contains a receipt signature on the layer the voter will keep. The other layer is surrendered to a poll worker and shredded. People find visual cryptography intrinsically satisfying and reassuring—although they have no way to really know the independent layers are meaningless, the random pattern on each layer is convincing to most individuals.

In VoteHere's scheme, the receipt is a printed receipt—but instead of containing the voter's selections in plain text, it prints only code numbers for each selection. The code numbers are based on a code book that is generated uniquely for the voter's randomly assigned ballot sequence number, which is stored on a key card given to the voter. The code books are generated using a code book key that a group of election trustees generate; that key must be securely distributed to and stored on each voting machine. As long as the key is uncompromised, the code numbers are meaningless to anyone other than the voter, who sees the code numbers displayed by the voting machine. Voters can check that the code numbers printed on their receipt match the displayed code numbers for their selections. Integrity depends on the voting machine displaying the correct code numbers for each ballot. VoteHere proposes ensuring this by using observers at randomly selected times

throughout the election who act as voters until the final casting step and audit the displayed code books .

After the election is complete, both SureVote and VoteHere publish a transcript of the election receipts. Any voter can verify that her receipt is included in the published transcript. Additional procedures must also ensure that additional votes have not been added to the receipt transcript. Next, a series of steps is taken to decrypt the published ballot receipts in a way that ensures that the decrypted votes correspond to the published encrypted votes while also ensuring that a particular encrypted vote cannot be tied to its corresponding decrypted vote. Both systems use verifiable mixes for this process in which each trustee partially decrypts each ballot using its share of the key and then scrambles the ballots' order. It must be evident that each output ballot corresponds to a unique input ballot without revealing the mapping between input and output ballots.

Chaum's scheme does this by grouping the trustees in pairs where the second trustee takes the output ballots of the first trustee as its input ballots. After both trustees have completed the steps, half the input ballots are randomly selected, and the first trustee must verify that each ballot corresponds to one of the output ballots. Then, the second trustee verifies that each of its input ballots that do not correspond to one of the verified output ballots of the first trustee corresponds to one of its output ballots. Thereby, each trustee has a 50 percent chance of getting caught if any ballot is tampered with. However, there is no way to connect any input ballot to the corresponding output ballot of the second trustee in the pair. VoteHere's system uses validity proofs to show that the decryption calculations were done correctly.

Cryptographically verifiable results provide the potential for high confidence in the entire election process with much less dependence on trust in people and procedures than other processes. For individual voters, it is not as clear that confidence in complex cryptographic systems is better than confidence in more transparent processes such as protecting paper ballots. The proposed schemes also add substantially to the complexity of the voter experience. With Chaum's scheme, in addition to the normal voting process, a voter must enter a receipt layer selection, separate the receipt into its layers, and give one layer to a poll worker for shredding. This increases the time required to vote and the number of poll workers needed. VoteHere's scheme requires less change to the voting process. However, the receipts it produces do not provide an obvious representation of the voter's selections. As with the votomatic punch cards, a voter would need to map the numbers printed on the receipt to those displayed on the voting machine to know the receipt correctly reflects her selections.

Electronic voting increases the potential for large-scale fraud. It is essential that citizens perceive elections as secure, and the lack of transparency and auditability in DRE machines threaten that perception regardless of whether the machines themselves are secure.

The tradeoffs between security, cost, and convenience are political, but those decisions much be made with a clear understanding of the risks that they are accepting. Without providing voter-verifiable paper audit trails, with today's technology and testing procedures, DRE machines pose an unacceptable risk for the potential benefits they provide. □

References

1. V*oting: What is, What Could Be*, Cal Tech–MIT Voting Technology Project Report, Jul. 2001; www.vote.caltech.edu/Reports/.
2. F. Fessenden, "Counting the Vote: The Machine; New Focus on Punch Card System," *New York Times*, 19 Nov. 2000, p. 40.
3. J.N.A. Wand et al., "The Butterfly Did It: The Aberrant Vote for Buchanan in Palm Beach County, Florida," *American Political Science Review*, vol. 95, no. 4, 2001, pp. 793–810.
4. Peach State Poll, Georgians Express Confidence in New Electronic Voting System, Carl Vinson Inst. of Government, Feb. 2003; www.cviog.uga.edu/peachpoll/.
5. P. Herrnson, B. Bederson, and O. Abbe, *An Evaluation of Maryland's New Voting Machines*, Center for American Politics and Citizenship, Univ. of Maryland, 2002; ftp://ftp.cs.umd.edu/pub/hcil/Reports-Abstracts-Bibliography/2002-25html/2002-25.pdf.
6. T. Kohno et al., Analysis of an Electronic Voting System, tech. report TR-2003-19, Johns Hopkins Information Security Inst., July 2003.
7. R. Konrad, **"**Some Voters Skeptical of E-voting Systems," *Washington Post*, 8 Oct. 2003.
8. Federal Election Commission, *Voting System Standards*, www.fec.gov/pages/vssfinal/vss.html.
9. K. Zetter, "Time to Recall E-Vote Machines?," *Wired News*, 6 Oct. 2003, http://www.wired.com/news/politics/0,1283,60713,00.html.
10. C.A. Neff and J. Adle, *Verifiable e-Voting: Indisputable Electronic Elections at Polling Places*, Aug. 2003, www.votehere.net/vhti/documentation/verifiable_e-voting.pdf.
11. D. Chaum, "Secret-Ballot Receipts and Transparent Integrity," *Proc. Palo Alto Workshop on Information Dynamics in the Networked Economy*, 2002.

David Evans *is an assistant professor at the University of Virginia's Department of Computer Science. His research interests include program analysis, securing sensor networks, and biological approaches to programming and security. He has a SB, SM, and PhD in computer science from the Massachusetts Institute of Technology. Contact him at evans@cs.virginia.edu.*

Nathanael Paul *is a second-year PhD student at the University of Virginia's Department of Computer Science. His research interests are electronic voting, user authentication, and applied cryptography. He has a BS in computer science from Bob Jones University and an MS in computer science from Clemson University. Contact him at nrpaul@cs.virginia.edu.*