

DAVID JEFFERSON

COMPUTER SCIENTIST, LAWRENCE LIVERMORE NATIONAL LABORATORY¹
BOARD CHAIRMAN, VERIFIED VOTING FOUNDATION
BOARD OF DIRECTORS, CALIFORNIA VOTER FOUNDATION
d_jefferson@yahoo.com

If I can shop and bank online, why can't I vote online?²

There is widespread pressure around the country today for the introduction of some form of Internet voting in public elections that would allow people to vote online, all electronically, from their own personal computers or mobile devices. Proponents argue that Internet voting would offer greater speed and convenience, particularly for overseas and military voters and, in fact, any voters allowed to vote that way.

However, computer and network security experts are virtually unanimous in pointing out that online voting is an exceedingly dangerous threat to the integrity of U.S. elections. There is no way to guarantee that the security, privacy, and transparency requirements for elections can all be met with any practical technology in the foreseeable future. Anyone from a disaffected misfit individual to a national intelligence agency can remotely attack an online election, modifying or filtering ballots in ways that are undetectable and uncorrectable, or just disrupting the election and creating havoc. There are a host of such attacks that can be used singly or in combination. In the cyber security world today almost all of the advantages are with attackers, and any of these attacks can result in the wrong persons being elected, or initiatives wrongly passed or rejected.

Nonetheless, the proponents point to the fact that millions of people regularly bank and shop online every day without apparent problems. They note that an online voting transaction resembles an ecommerce transaction, at least superficially. You connect your browser to the appropriate site, authenticate yourself, make your choices with the mouse, click on a final confirmation button, and you are done! All of the potential attacks alluded to above apply equally to shopping and banking services, so what is the difference? People ask, quite naturally, *"If it is safe to do my banking and shopping online, why can't I vote online?"*

This is a very fair question, and it deserves a careful, thorough answer because the reasons are not obvious. Unfortunately it requires substantial development to explain fully. But in brief, our answer is in two-parts:

1. It is *not* actually "safe" to conduct ecommerce transactions online. It is in fact very risky, and more so every day. Essentially all those risks apply equally to online voting transactions.
2. The technical security, privacy, and transparency requirements for voting are structurally different from, and actually much more stringent than, those for ecommerce transactions. Even if ecommerce transactions *were* safe, the security technology underpinning them would not suffice for voting. In particular, the voting security and privacy requirements are unique and in tension in a way that has no analog in the ecommerce world.

¹ *Analyses and views stated herein are drawn from my expertise as a computer scientist working on national security applications and are my own. They are not to be ascribed to my employer, Lawrence Livermore National Laboratory, which takes no position on these issues.*

² © David Jefferson, 2011

The rest of this essay expands upon these two points in order.

Ecommerce transactions are not, in fact, “safe”

Why do security experts say that ecommerce transactions are not safe when millions of people do them every day, mostly without problems? The question needs to be refined: “Safe for whom?” and “What degree of safety is required”?

Ecommerce transactions may be relatively safe for consumers, but they certainly are not safe for financial institutions or merchants. Banks, credit card companies, and online merchants lose billions of dollars a year in online transaction fraud³ despite huge investments in fraud prevention and recovery. People have the illusion that ecommerce transactions are safe because merchants and banks don't hold consumers financially responsible for fraudulent transactions that they are the innocent victims of. Instead the businesses absorb and redistribute the losses silently, passing them on in the invisible forms of higher prices, fees, and interest rates. Businesses know that if consumers had to accept those losses personally most online commerce would collapse. Instead, they routinely hide the losses, keeping the magnitude secret so the public is generally unaware. It's a good business strategy.

There are many techniques for ecommerce fraud that are directly applicable to online voting. A common pattern starts with theft of credentials, e.g. names, account numbers, credit card numbers, passwords, or the answers to personal challenge questions. The theft can be initiated through phishing scams, drive-by malware installation, or other means, and such tricks can just as easily be used to steal online voting credentials as well. Recently a new botnet named Zeus has been in the news that installs malware on PCs.⁴ Zeus is specifically designed to wait until you connect to your bank and then it steals your bank account number and password or PIN as you type them into your browser. The botmasters use those credentials to transfer money out of your accounts and to fake your online financial statements to hide the theft (for a while at least). It makes no difference that you have a “secure” connection to your banking site because the malware operates inside your computer and can see and modify everything you type in the clear, before it is encrypted for transmission down the “secure” connection. There are now illicit businesses that help other people set up Zeus botnets, or rent time on a botnet already created.⁵ Most people, however, are completely unaware of these threats.

Zeus exemplifies what could just as easily happen if online voting becomes widespread. Eventually someone, perhaps a partisan political operative or a foreign intelligence agency, will deploy a similar botnet to infect thousands of voters' computers and modify their votes invisibly as they are being transmitted. Again, having a “secure” connection to the remote election server will make no difference. There is no effective way to prevent such an attack, and no effective recovery. Banks, online merchants, and high tech companies that do business online have huge security budgets to defend themselves against cyber attacks, and even so they are frequently victimized. If these organizations with such great expertise and capability in computer and network security can be successfully attacked, then no voting system vendor or local election administration has any realistic chance of successfully defending against similar threats.

We have to recognize that the cost to the attacker of conducting a remote online attack has declined drastically over the last few years as various programming templates, libraries, and toolkits for malware production have become widely available. One recent study demonstrated that it was possible to duplicate even very sophisticated attack vectors like Stuxnet, the malware that did great

³ See <http://www.mcafee.com/us/resources/reports/rp-financial-fraud-int-banking.pdf>, p. 4

⁴ See [http://en.wikipedia.org/wiki/Zeus_\(trojan_horse\)](http://en.wikipedia.org/wiki/Zeus_(trojan_horse))

⁵ See http://threatpost.com/en_us/blogs/new-service-helps-attackers-get-zeus-botnet-ground-011011

damage to Iranian nuclear facilities, in about two months time for under \$20,000.⁶ We are now in a very different threat environment than we were even a few years ago.

What level of security is sufficient to protect elections? The scale of fraud that ecommerce and electoral systems can tolerate are very different from one another. In the ecommerce world if one out of every thousand ecommerce transactions is lost or is fraudulent it is not really a vital concern. Banks, merchants and purchasers routinely deal with online revenue losses over 10 times higher than that⁷, and have many tools to deal with the loss. As unjust and frustrating as it may be, no catastrophic global consequence ensues from a small ecommerce fraud rate. Ecommerce markets are relatively *robust*, i.e. not overly sensitive to small-scale losses. But in the voting world we are all familiar with the cases where, within about one decade, a senator, a governor, and a U.S. president were all elected by margins much smaller than one vote in a thousand. Small changes in vote totals sometimes have very big, even global consequences, and can push a whole city, state or nation in a new direction. Elections outcomes are thus *very sensitive to small errors or frauds* in a way that ecommerce systems simply are not. Election security is thus a matter of *national security*, and the security standards have to be designed to reliably prevent, detect, and correct even very small problems and attacks. That level of security and reliability is neither needed nor cost effective for ecommerce systems.

Voting security, privacy, and transparency requirements are structurally different from those for ecommerce transactions

The second point of our argument is that the security, secrecy, and transparency requirements for online voting transactions are structurally very different from, and generally much stricter than, those for ecommerce transactions. The security mechanisms that make ecommerce transactions relatively safe (for consumers at least) are not sufficient to guarantee the safety of online voting.

The first major distinction is that we can at least eventually *detect* ecommerce errors and fraud, but we may never even know about online election fraud.⁸ In the ecommerce world problems are reliably detected because of such practices as receipts, double entry bookkeeping, and financial audit records kept by both sides of every major transaction. But in the online election world there are no receipts, no double entry bookkeeping, and no meaningful audit trail information. Security experts routinely call for an independent, end-to-end audit trail that can be used to verify that the electronic ballots received by election officials are identical to those the voters sent, and that none were forged, lost, or modified in transit. But the only reliable way to accomplish this with current technology is for voters to send paper copies of their ballots back to their local election officials along with a signed attestation, and for the officials to use those copies in a formal risk limiting audit procedure.⁹ That would solve most of the security problems associated with online voting (though not the privacy problems). But most advocates of Internet voting oppose such a paper-based audit requirement because of the additional burden on voters to mail back paper copies of their ballots and signed attestations, which is essentially equivalent to sending an ordinary paper absentee ballot. Yet without a meaningful end-to-end audit trail a well-constructed attack may lead to the attackers' choice of candidates being elected and *there may well be no way to know that anything happened at all*. Even if there is suspicion of a problem there will be no way to prove or disprove it. And because of

⁶ See http://hosted.ap.org/dynamic/stories/U/US_TEC_HACKING_CONTROL_SYSTEMS?SITE=AP&SECTION=HOME&TEMPLATE=DEFAULT&CTIME=2011-10-23-08-23-54

⁷ See <http://www.mcafee.com/us/resources/reports/rp-financial-fraud-int-banking.pdf> , p. 4

⁸ See <http://servesecurityreport.org/paper.pdf>

⁹ For papers on audit procedures for elections a good place to start is <http://statistics.berkeley.edu/~stark/Vote/#papers>

ballot secrecy even if there were strong evidence that *particular* persons cast illegal ballots, or their ballots were tampered with, officials cannot know *which* ballots to remove from the count. Hence, fraudulent online voting will most often be *undetectable*, and almost certainly *uncorrectable* even if detected.

Vote fraud is much less manageable than ecommerce fraud. There is no election analog to the natural business practice of “spreading the cost” or “spreading the risk”. There is no way to pass on to other voters the “losses” due to illegal ballots cast by ineligible voters or attackers, or to recover votes changed by malicious software. There is no “insurance” that one can buy to cover those losses. There is just no way to compensate for damage done to an election.

There are several ways in which the security requirements for voting are strictly stronger than those for financial transactions. Eligibility checking is one. In the ecommerce world essentially anyone including criminals, non-citizens, and minors, is allowed to buy and sell online. Non-human entities, e.g. corporations, government agencies, and estates, are free to engage in ecommerce transactions as well. And there are usually no residency requirements for ecommerce transactions. But all such factors play a role in determining eligibility to vote.

Then there is the issue of proxy transactions. In the ecommerce world you can freely authorize someone else to act as your agent for purchases or funds transfers, or you may authorize others to spend funds from your accounts simply by giving them your credit card number and security code, and/or your PIN or password. By doing so you take responsibility for the consequent risk. For larger transactions you can accomplish the same thing by setting up a joint bank account, signing a contract, appointing a trustee or guardian, granting power of attorney, etc. But in the voting world you are *never* permitted to transfer your right to vote to anyone else, at least not in the U.S. No one is legally allowed to act as your proxy to vote for you, not even your spouse, and not even with your written permission.

The prohibition of double voting is a third election security requirement that has no analog in the ecommerce world. A person is free to engage in as many ecommerce transactions as he pleases but the rule of one person, one vote is fundamental. The double vote check is actually complex because it has to cover not just voting a second time online, but also voting a second time by paper absentee ballot or in person at the polls.

Because of the need for eligibility checking, proxy vote prevention, and double vote prevention we are required to *verify the actual identity of voters*. In contrast, for an ecommerce transaction we only have to verify that the person doing the transaction is *authorized to use a suitable financial account*, which is a much lower requirement. We need a strong identity verification procedure for online voting because if an attacker can figure out how to cast one illegal vote online through a weakness in the identity verification, then he can probably automate that attack to cast thousands of phony votes. But reliably verifying the actual identity of a potential voter remotely through the Internet is a difficult and unsolved problem in the U.S. The U.S. does not issue national identity cards with private keys embedded in them, and even if it did today’s computers and mobile devices are not equipped with devices to read them securely. Nor do election jurisdictions keep a database of faces, fingerprints, or other biometric data about registered voters, and once again even if they did computers today are not equipped to capture and transmit them securely. It is not sufficient for the voter to just present a PIN number or password or to answer to a challenge question (e.g. “What city were you born in?”). Any such data might be given away, guessed, stolen, or sold, and thus does not constitute sufficient proof of identity because of the danger of *automated* online buying and selling or stealing of such voting credentials.

In most states voters prove their eligibility to vote initially when they register and then provide an ink signature sample for use later use. Voters prove their identity again when they vote, either at the polls or via paper absentee ballot, by duplicating that ink signature on record in the registration database. Some states are now going even further and requiring voters to provide photo ID

documents at the time of voting. But we cannot get a wet ink signature from a voter through the Internet to compare against the registration records, nor can the voter present his or her face along with a matching photo ID or passport. *As of now there is no reliable infrastructure in place to verify over the Internet the actual identity of a person sitting at a PC or holding a mobile device.*

There is no comparable identification requirement for ecommerce transactions. All that is really required for an online transfer of funds out of your bank account is knowledge of the name, account number, and password or pin associated with the account. *There is no check of the actual identity of the person doing the transaction.* As another example, consider that when you sign up for an ecommerce account, e.g. at Amazon.com, they ask for your name and address, but they do not ask for a picture, or an ink signature, or your driver's license, or passport, or other proof of identity. They never really check those, and they have no way to do so. After creating an Amazon account all that is *really* required to make a purchase is reasonable evidence that you are in possession of some (any!) valid credit card. You demonstrate that by giving the name on the card, and the account number, security code, expiration date, and password or pin, but you do not need to present a drivers license, passport, fingerprint, or picture ID. Merchants do like to know who their customers are for marketing purposes, but they are not rigid about it. They will sell to anyone who can type in numbers for a valid credit card. If the credit card turns out later to have been stolen, the problem will be sorted out after the fact.

The privacy requirements for ecommerce and voting transactions are also fundamentally different. An ecommerce transaction is generally *symmetric* between buyer and seller, with both parties in theory fully aware of all the details of what is being bought and sold, for what price, with what warranties, and who has what rights to void the transaction, etc. For larger transactions there is usually an exchange of official paper receipts with names, dates, prices, conditions, and other transaction details so that in case of a dispute either the buyer or seller can *prove* to a third party (e.g. a court) exactly what the transaction was supposed to be so the dispute can be resolved.

But it cannot be the same with voting transactions. While the voter of course knows the details of his votes, election officials must not. Officials know the names of those who voted, and the contents of the cast ballots, but they are never supposed to know exactly who cast which ballot. This is a requirement for partial blindness on the part of one side in the transaction that has no analog in the ecommerce world. Furthermore, although each voter knows how he personally voted and is free to tell anyone, he is not allowed to have any *proof* of how he voted that could convince a third party. This inability to prove how you voted is the most powerful protection we have against the threat of vote selling and vote coercion, and is a requirement unique to voting. I know of no other security situation in which people are completely free to *disclose* a fact that they know (how they voted), but are not permitted to have any *proof* of that fact that can convince someone else that they are telling the truth. In this respect voting privacy requirements are almost the *opposite* of ecommerce privacy expectations in which both sides generally insist on possessing proof of the details of a transaction.

The unusual vote privacy rules have strong consequences that we cannot avoid. As noted earlier, if for some reason officials learn after the fact that a particular person has succeeded in casting an illegal ballot there is no way to find it to remove it from the count. In the U.S. and most other countries once a voting transaction is complete it cannot be undone even in principle because the information needed has been deliberately lost. In that sense a voting transaction is *irreversible*. In the ecommerce world, however, we go to some lengths to make sure most transactions are reversible in case they are found to be faulty in some way, or fraudulent, or sometimes even if one party simply has second thoughts. Merchandise can be returned, and records can be corrected, and a monetary transfer transaction can be reversed with another transaction. For that reason people feel free to take prudent risks with online financial transactions based on the reputation of the merchant or the credit history of the buyer. But there is no concept of "reputation" or "credit worthiness" in the election world to help manage risk. These differing vulnerabilities to failures and fraud lead to very

different security approaches in online transaction software. For election security there is a very strong imperative for *up front, absolute prevention of errors and fraud*. For ecommerce there is usually much reduced need for strong security barriers up front because problems can usually be corrected later and those that cannot can be absorbed.

The flip side of privacy is *openness* or *transparency*. Once again, the requirements are completely different for ecommerce and for online voting. In the ecommerce world a person buying something online is entitled to know everything about his particular transaction, but nothing about other people's transactions. A buyer is not entitled to know how many other transactions there are, what the merchant's revenues or profits are, who else the merchant sells to, or what price others pay for the same goods or services, and he has no right to audit the books of the merchant he is dealing with.

In the voting world, however, most of this is reversed. Complete election information is (or should be) open to all. Election officials report not just the names of the winners, but also exactly how many votes were cast and how many each candidate received down to the precinct level. The list of exactly who voted is also usually public, and in some jurisdictions so are the original ballot images. In principle *all* information bearing on the outcome of an election that does not compromise vote privacy is (or should be) public. Candidates, parties, and the public are entitled to participate in open audits, challenges, and recounts so that everyone, especially losing candidates, can be satisfied that the election was conducted according to law and the votes were counted accurately. Election officials are thus accountable to candidates and voters for the integrity of every relevant detail of an election, whereas merchants are usually accountable only to buyers, and then only for each buyer's own transactions.

The pattern of motivation for fraud is profoundly different between the commercial and electoral worlds. In an ecommerce situation all transactions are essentially independent. A buyer has no particular incentive to spoil or tamper with another buyer's online purchase since two buyers rarely have conflicting interests. In any case the problem would almost certainly be detected and corrected. And it is hard to imagine a motive for another nation to bother messing with many Americans' ecommerce transactions. But the situation is completely different with voting. There is a powerful partisan incentive to block or change other people's votes, especially if it can be done without detection. The motivation to *automate* that process to affect thousands of online votes is that much greater. Such attacks can be done for tens of thousands of dollars or less, while the monetary value of changing the outcome of an election can be hundreds of millions of dollars or more, and the nonmonetary value can be even more immense. With *Internet* voting the danger is actually much worse because anyone on Earth, including foreign governments, could derive great benefit from tampering with with U.S. elections, especially since it is unlikely they will be caught or brought to justice. Online voting fraud is thus a *national security risk* in a way that ecommerce fraud simply is not.

The sum of all of these considerations is simple. The security, privacy, and transparency requirements for online voting are much more complex and stringent than they are for ecommerce transactions. The acceptability of small losses and the strategies for managing risk are very different between the two. And it is hard to grasp the full implications of the fact that online elections might be compromised and the wrong people elected via silent, remote, automated vote manipulation that leaves no audit trail and no evidence for election officials or anyone else to even detect the problem, let alone fix it. These points are all pretty basic, and they are not going to change for the better any time soon. While there is plenty of research going on in the computer security community to try to deal with the security and privacy problems of Internet voting, there is no technology on the horizon that is going to resolve them all any time soon, or even in the foreseeable future. For the time being we simply cannot provide satisfactory security for online voting even though we can for online commerce.