

A Deeper Look: Rebutting Shamos on e-Voting

Arthur M. Keller, Ph.D.¹

Edward Cherlin²

David Mertz, Ph.D.³

Table of Contents

1) Abstract.....	2
2) A Taxonomy of Error.....	3
3) Philosophical Errors.....	3
3.1) Ignoring Voting Procedures’ Political Natures.....	3
Misunderstanding Paper’s Benefits.....	4
3.2) The Corner of Unreasonably Difficult Proofs.....	5
Misplacing the Burden of Proof.....	7
3.3) It Seems to Have Worked So Far, So It’s Fine.....	8
3.4) Comparing Voting Systems to Flight Control Systems.....	9
3.5) Comparing Voting Systems to Financial Systems.....	10
A Better Parallel: Gambling Systems.....	12
4) Implementational Errors.....	13
4.1) Overestimating the Effort Required to Cheat and Underestimating What Cheats Can Accomplish.....	13
Cheating With Communications Devices.....	14
Cheating with Hardware – The Malware Loader.....	15
4.2) Touting Separation of Candidate Names As a Panacea.....	16
4.3) Overestimating What Testing Can Accomplish.....	17
4.4) Misunderstanding the Nature of Open-Source Software.....	19
4.5) Overestimating the Benefits and Practicality of Independent Audit Devices.....	19
4.6) Systematically Misunderstanding Paper Trails and Ballots.....	20
5) Miscellaneous Errors.....	22
6) Conclusion.....	22
7) Acknowledgements.....	23

¹ Arthur M. Keller, Ph.D., is a researcher of Technology and Information Management at the Baskin School of Engineering of the University of California, Santa Cruz. He also advises startups and serves as expert witness on patent infringement cases, especially in the areas of databases and e-commerce, as managing partner of Minerva Consulting.

² Edward Cherlin has been a mathematician, a Peace Corps Volunteer in South Korea, a Buddhist monk, a computer programmer, a high-tech market analyst, Managing Editor of a computer magazine, leader of anti-spam and Free Software organizations, and on one occasion a precinct voting inspector. He is currently working on software and hardware products, including voting software, designed to be affordable in poor villages around the world, and to meet the other stringent requirements for helping to end extreme poverty.

³ David Mertz, Ph.D. is Chief Technology Officer of the Open Voting Consortium (“OVC”). He is a well-known author on programming—and sometimes political—topics. He feels that procedural democracy requires that the technical instruments of governance be open for public inspection, every bit as much as it requires the legal acts of government remain so open. David has written extensively about OVC’s design principles.

1) Abstract

In his widely cited paper *Paper v. Electronic Voting Records—An Assessment*,⁴ Michael Ian Shamos⁵ surveys a variety of objections to Direct Recording Electronic (“DRE”) voting systems (herein “*An Assessment*”). While acknowledging and validating some of the most pressing objections, he breezily dismisses many others, often by packaging them as straw men or by impugning objectors’ maturity, reasoning ability, or thoughtfulness. Given Professor Shamos’ justified reputation as a respected elections expert, it is disappointing that *An Assessment* sidesteps not only key technical issues, but also important issues of transparency, accountability, and the nature of the American democratic republic.⁶

Our paper identifies *An Assessment*’s most significant errors, and also considers electronically-based voting systems⁷ security more broadly, especially compared to that of other electronic systems, such as financial systems and gambling devices. We focus mainly on the possibility of vendor-sponsored fraud, since vendors’ access to and knowledge of their voting systems, and their ability to keep their inner workings secret by force of law, gives them unique power over how votes are solicited, recorded, and counted.⁸

We hope that our paper dispels some of the unjustified trust currently placed in opaque, unverified, and unverifiable e-voting systems.

⁴ Michael Ian Shamos, “Paper v. Electronic Records – An Assessment,” April 2004, <http://euro.ecom.cmu.edu/people/faculty/mshamos/paper.htm>. This article has been re-published in numerous places, including by NIST at http://vote.nist.gov/threats/papers/paper_v_electronic_records.pdf.

⁵ In endnote 1 of his paper, Professor Shamos describes himself thusly:

The author is Distinguished Career Professor in the School of Computer Science at Carnegie Mellon University and an attorney admitted to practice in the Commonwealth of Pennsylvania and before the United States Patent and Trademark Office. From 1980-2000 he was statutory examiner of electronic voting systems for the Commonwealth of Pennsylvania. From 1987-2000 he was the designee of the Attorney General of Texas for voting system certification. During those years he personally examined more than 100 different computerized voting systems for certification purposes. In the 2000 election, machines for which he participated in certification (which did not include Florida) were used to count more than 11% of the popular vote of the United States.

He is widely cited, and his paper has been quite influential.

⁶ A “democratic republic” (or simply “republic”) usually is defined as a form of government in which citizens elect representatives, who actually conduct the affairs of state. In a pure “democracy,” in contrast, citizens conduct the affairs of state directly.

⁷ This class includes not only DREs, but also any system that uses electronic devices to present selections to voters, record, cast, or tabulate votes. Thus it includes most existing punchcard and optical-scan systems as well as Automark-type systems.

⁸ The same poorly- or un-reviewed development process that permits a dishonest vendor to cheat also permits vendors’ errors to go undetected, which may allow others, like hackers or dishonest elections officials, to cheat.

2) A Taxonomy of Error

An Assessment's errors are of two main types: philosophical and implementational.

Within the first class of errors, *An Assessment* seems to misconceive the role of citizens in a democratic republic. That role should not be to accept the proclamations of government officials (or of the vendors they choose) on faith, but instead to treat them with a healthy, sustained skepticism. This role is particularly important on issues concerning the maintenance of the republic itself. And no issue is more central to that maintenance than the honest and accurate administration of elections, since they are citizens' chief means of control over the republic's course. Yet *An Assessment* would make citizens prove that such systems are inaccurate and/or dishonest, rather than requiring their proponents⁹ to prove their accuracy and honesty by opening them to total public review and constant verification. This burden of proof inverts the proper role of the citizen from skepticism of government to faith in government, establishes a defective heuristic for evaluating systems or policies whose malfunctions can cause significant harm, and makes it unreasonably difficult to prove election fraud.

An Assessment also appears confused about the essential nature of democratic republics themselves, writing

...I believe I and the republic will survive if a president is elected who was not entitled to the office....¹⁰

It should be unnecessary to state that a "republic" whose Presidency is occupied by someone not entitled to the office is, at least for that term of office, no longer a republic. It is a soft dictatorship run by those who maneuvered the illegitimate President into office.¹¹ How such a "republic" can recover its legitimacy and standing via future elections is best left in the realm of theory; in the real world, an illegitimate Presidency is a disaster.

An Assessment also makes many errors of the second, implementational, class, chiefly, it appears, through lack of imagination and through failure properly to consider motive. It variously underestimates the misdeeds a crooked vendor can accomplish and the importance of vendor error, fails to account for advances in technology that make (or will make) misdeeds easier to perpetrate, overestimates the efficacy of its security prescriptions, and, in one case, even proposes an "ultimate" solution that actually makes the problem substantially worse.

The following sections describe *An Assessment's* errors in more detail. Section 3 discusses its philosophical errors, and section 4 its errors concerning voting technology's implementation. Additionally, section 5 explores a variety of miscellaneous errors.

3) Philosophical Errors

3.1) Ignoring Voting Procedures' Political Natures

The *purpose* of voting is, of course, political, not technical. But, while the *procedure* of voting is technical—a ballot is prepared, cast, and counted—it is also political. It is not enough for the technical requirements (such as accuracy, privacy, and accessibility) to be satisfied. The political requirements must also be satisfied. These requirements are (1) sufficient transparency to inspire a solid assurance that (2)

⁹ Generally these would be vendors and public officials.

¹⁰ Shamos at §1.1.

¹¹ Given the small Presidential election margins of late, even a tiny amount of cheating radically could change the nation's direction.

every eligible vote is properly counted, and only eligible ones are counted. *An Assessment* falls down by considering the problem of voting as merely technical. It advocates the use of systems that, from the voter's viewpoint, are black boxes: the voter does something on a touchscreen, goes home, and views the election results on her television. While this arrangement might –with sufficient effort–reach an acceptable level of technical performance and security, it has not currently come close.¹² Further, its opacity prevents voters from effectively supervising the voting process and (justifiably) undermines many voters' confidence in its correctness.¹³

Misunderstanding Paper's Benefits

In advocating DREs, *An Assessment* misses the fundamental purposes of paper ballots,¹⁴ which are not simply technical accuracy, but also transparency and citizen participation. Paper ballots serve the same underlying democratic purpose as open meeting laws. A voter can fully examine and understand her ballot. Depending upon the surrounding system, she might also be able to understand the procedures by which it is counted, and meaningfully participate in the counting or auditing herself.¹⁵ This kind of transparency cannot be achieved with DREs, since only a tiny proportion of voters can competently evaluate their security, even if vendors were to disclose their source code and permit their hardware to be inspected.

Unaided, paper ballots do not prevent fraud—there exists a long and ignoble history of falsified paper ballots, ballot box stuffing, and the like. But they can help limit fraud. Paper ballots must be used as part of carefully designed systems with chain-of-custody controls and trustworthy tabulation systems, whether hand-tabulated or computer tabulated. For example, an electronic ballot printer (“EBP”) that produces both voter-verifiable paper ballots (“VVPB”)¹⁶ and corresponding electronic ballot images makes it harder to cheat than DREs without a voter-verified paper audit trail (“VVPAT”), since doing so requires manipulating both sets of records.¹⁷ As in double-entry bookkeeping, the ability to compare the two sets improves security. Further, the paper ballot permits the voter to determine whether the system accurately has represented her votes.¹⁸ The existence of an electronic audit trail and cryptographic marks on electronically-printed ballots¹⁹ (where used) can make box stuffing more difficult than unaided paper ballots, but with the added risk of presentation fraud.²⁰ VVPB are like hand-marked paper ballots in that

¹² Even Shamos's best suggestions would not raise e-voting security to the “acceptable” level.

¹³ Contrary to some skeptics, authors believe that computer experts from the general public can provide adequate supervision if proper procedures are enacted and enforced and the systems themselves incorporate suitable security, reliability, and auditability measures, such as the use of paper ballots.

¹⁴ Potentially including ballots produced by an electronic ballot printer, but verifiable by the voter.

¹⁵ Hand counts are conceptually easy for the average voter to understand – and to effectively supervise if law and custom allow it. Machine vote presentation, printing, and counting are much more difficult for average voters to understand and to effectively supervise. However, hand recounts of a statistically-significant set of randomly-selected precincts, or hand checks of the tabulation of a statistically-significant, randomly-chosen set of ballots in every precinct, would improve both security and transparency.

¹⁶ Such ballots constitute the definitive statements of their casters' intent.

¹⁷ Or, manipulating the totals after the records have been cross-checked and tabulated.

¹⁸ It is, however, possible for a crooked EBP to encode invisible marks on the ballot directing the tabulator to count it differently than it appears to the voter.

¹⁹ Shamos casually dismisses such marks; though it is true that voters cannot transparently verify them, and that a dishonest machine could print false marks in an attempt to void ballots containing votes it wishes to cancel.

²⁰ A presentation fraud attempts to deceive the voter into making a different choice than she would absent the fraud, by modifying the presentation of choices or the interpretation of selections. For example, a fraud might occasionally omit a candidate from the ballot, or put her name at the bottom of the ballot, or make it more difficult to select her name, or provide a “default” choice for voters who might otherwise choose not to cast a vote for a given office.

they require the tabulation of all the paper *ballots*, rather than the statistical spot-checks required by VVPAT. We must note that using voter-verified paper ballots with an electronic audit trail does not by itself eliminate the potential for error or fraud. How the electronically-printed paper ballots are tabulated and compared against the electronic audit trail also has potential for error or fraud, which may be mitigated by such techniques as fully open source, random hardware inspections, and parallel testing.

Finally, voting fraud is not, as *An Assessment* seems to presume, either present everywhere or absent everywhere; fraud comes in degrees and increments. A malicious DRE, created and distributed by one vendor to hundred of thousands of polling places, systematically can falsify millions of votes. It is fraud on a wholesale level. Stuffing a ballot box, in contrast, works at a retail level. A tamperer, however malicious and skilled, can stuff only as many ballots as might plausibly be cast at the polling place. While an organized group could stuff multiple ballot boxes, malicious DRE software could affect far more votes. Risks must be evaluated in terms of their potential harms, not in a vacuum as *An Assessment* does.

3.2) The Corner of Unreasonably Difficult Proofs

An Assessment's main argument hinges upon the bold assertion that “[T]he United States has been using direct-recording electronic voting equipment for well over 20 years without a single verified incident of successful tampering.”²¹ The paper seemingly moderates this stance farther on, noting that “hacking has been advancing at a[n] alarming rate, and new attacks are constantly being discovered, so we are entitled only to a small bit of comfort from DRE history.”²² However, the underlying theme of *An Assessment* remains the idea that e-voting systems are generally safe and that many security concerns are over-inflated, unjustified, or just plain invention.

This argument has a variety of defects. First, it hinges upon what a “verified incident of successful tampering” is. As explained below, *An Assessment* appears to define this so as to be extraordinarily difficult to prove. Second, the argument implicitly inverts the appropriate burden of proof. The potential ill effects of dishonest voting equipment, like the potential ill effects of bad drugs, can be grave. This argues that vendors should have to prove that their systems are safe and effective, rather than skeptics having to prove that they are not. Third, as shown in section 4, *An Assessment's* argument ignores a variety of very practical cheating methods.

What would *An Assessment* accept as a “verified incident of successful tampering”? He does not spell out the criteria, but a reasonable inference from the paper as a whole would appear to be:

- (1) a confession by a person who engaged in or who knows of such cheating; or
- (2) direct proof of the existence of cheating code (e.g., its presence in a vendor-supplied source listing) combined with anomalies in election results²³ in which that code was used that are consistent with its expected effects; or possibly
- (3) an election result that differs so substantially from expectations that no other explanation suffices.

Alternative (1) requires a tamperer, or at least someone who knows of tampering, to blow the whistle. Perhaps a dishonest vendor's employee would be motivated to do so by a guilty conscience or by an intense attack of patriotism, or a disgruntled employee by revenge. However, by and large, those involved in wrongful or criminal enterprises remain silent. We cannot reasonably rely upon whistle-blowers to

²¹ Shamos at §1. This argument also ignores numerous, serious, well-documented instances of DRE (and other e-voting system) errors. See “Malfunctions and Miscounts, Sorted by Vendor,” *VotersUnite.org* <http://www.votersunite.org/info/messupsbyvendor.asp>, and also infra §3.2.

²² Ibid.

²³ How would Shamos require anomalies to be proven? Does he accept, for example, the validity of exit polls conducted to existing standards?

A Deeper Look: Rebutting Shamos on e-Voting

disclose tampering.

Alternative (2) is very difficult to prove because e-voting vendors generally (always?) keep their source code secret,²⁴ and only rarely allow others to inspect their hardware, and even then only under severely restricted conditions that exclude the public.²⁵ For example, while there is not currently direct evidence of software cheating, there is the suspicion that software may have been involved in the 18,000 undervotes in the closely contested 2006 Sarasota, Florida race.²⁶ Yet the Florida trial and appellate courts have ruled that unless there is clear evidence that the software *was* involved in the cheating, the vendor's trade secret rights trump the rights to a fair and free election. This creates a classic "Catch-22" situation since an expert will likely need to inspect the software in order to determine whether it was involved in the cheating.

In any case, a dishonest vendor is exceedingly unlikely to disclose source containing its cheating code, as we just discussed. Instead, such a vendor, if somehow forced to disclose source, would disclose similar source showing no trace of cheating. Of course, nothing guarantees that the disclosed source was actually used to create the executable application used in the questioned election,²⁷ but the general lack of understanding of computer security on the part of the public (and of many elections officials) would almost ensure that this critical point would remain unaddressed.²⁸ Further, a malware loader²⁹ can be used to corrupt even a fully-reviewed, honestly-implemented voting application.

Thus, the only apparently acceptable remaining way to prove tampering is alternative (3): where an election result differs so substantially from expectations that no other explanation suffices. In one place, *An Assessment* calls this "[t]he smell of irregularity" and claims that it is "sufficient to set off alarms resulting in investigations and recounts."³⁰ Since he wrote his paper in early 2004, he must have known of the 2002 Georgia senate race. In that contest, Max Cleland polled 49% to Saxby Chamblis' 45% a few days before the election,³¹ then lost 53% to 46%: an 11-point swing. The election was Georgia's first using Diebold DREs,³² and at least one uncertified "patch" was installed prior to the election.³³ Yet no official investigation of this anomaly ever seems to have been launched. Apparently Shamos believes that this contest does not present even a whiff of the "smell of irregularity," since he declines to mention it.

²⁴ See infra §4.4 for why this matters.

²⁵ See infra §4.1.

²⁶ Marc L. Songinin, "Court prohibits access to touch-screen source code", *ComputerWorld*, June 19, 2007, http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9025252&source=rss_topic17Computerworld.

²⁷ Such guarantees could be provided by the proper use of encryption and verification procedures.

²⁸ To his credit, Shamos, in his §1.2, partially recognizes this issue, and proposes using code escrow and encryption procedures (for code verification) to address it. His solution, however, is incomplete (e.g., vendors could still put cheating code into escrow, or use hardware to cheat), and, to our knowledge, no vendor has implemented such encryption procedures.

²⁹ See infra §4.1.

³⁰ Shamos at §1.1.

³¹ Welch, William M., "Control of Senate riding on tightest contests," *USA Today*, 11/3/2002, http://www.usatoday.com/news/politicselections/2002-11-03-state-polls-usat_x.htm

³² Dugger, Ronnie, "How They Could Steal the Election This Time," *The Nation*, 7/29/2004, at p.7, <http://www.thenation.com/doc.mhtml?i=20040816&c=7&s=dugger>.

³³ Cohen, Adam, "The Results Are in and the Winner Is . . . or MaybeNot," *New York Times*, 2/29/2004, <http://theory.lcs.mit.edu/~rivest/voting/press/nyt/2004-02-29%20NYT%20The%20Results%20Are%20in%20and%20the%20Winner%20Is%20.%20.%20.%20or%20Maybe%20Not.pdf> (reprint of <http://www.nytimes.com/2004/02/29/opinion/29SUN3.htm>, which is now available only by subscription).

A Deeper Look: Rebutting Shamos on e-Voting

Perhaps a 20-point swing is required?³⁴

Further, many elections are decided by margins small enough that they cannot, by definition, differ enough from expectations to trigger *An Assessment's* “smell” test. For example, the 2004 Presidential race hinged upon 2% of the vote in Ohio,³⁵ and the 2000 Presidential race upon 0.01% of the vote in Florida. Such races can be thrown by very modest cheating, yet nothing in *An Assessment's* criteria would allow us to detect it.

Misplacing the Burden of Proof

Under *An Assessment's* criteria, then, it is extremely difficult to prove e-voting cheating. But what is the rational response? Should we trust that cheating will not occur, as *An Assessment* seems implicitly to advocate? Or should we discard the criteria that would force us to do so? Perhaps an analogy will help us make this decision.

Once upon a time, but still within living memory, pharmacological drugs were largely unregulated. The implicit societal presumption was that drugs were safe and effective unless proven otherwise. In those days, however, many commonly available drugs were either harmful, ineffective, or both. For example, an arthritis remedy purchased at a local drug store might have contained radium,³⁶ and a digestive remedy (or a Prohibition-circumventing beverage, depending upon your viewpoint) might have contained a neurotoxic industrial chemical.³⁷ This state of affairs persisted for many years, until the “miracle drug” Elixir Sulfanilamide – one of whose ingredients was the antifreeze relative diethylene glycol – killed over 100 people in 15 states.³⁸ The ensuing outrage forced Congress to enact the 1938 Food, Drug, and Cosmetic Act, which, for the first time in American history, required drug manufacturers to show that their drugs were safe and effective before marketing them.³⁹

Today, protected by this burden of proof, we largely take drug safety for granted. We know that manufacturers must submit to a comprehensive scheme of pre-marketing trials and inspections, and even to some post-marketing surveillance.⁴⁰ And, by and large, drug safety is quite good – and far better than it was during the patent-medicine era.

Now, drugs are preparations used to maintain the health of, or to treat diseases of, the human body – the

³⁴ An interesting analysis of this contest, including a discussion of puzzling regional vote patterns, appears in the second half of Smith, Van, “Future Vote,” *Baltimore City Paper*, 12/11/2002, <http://www.citypaper.com/news/story.asp?id=3381>. Further, even if an event like this occurs through innocent error, failure to investigate it can itself constitute a form of cheating.

³⁵ “CNN.com Election Results,” *cnn.com*, 11/2/2004, <http://www.cnn.com/ELECTION/2004/pages/results/states/OH/P/00/>.

³⁶ Macklis, R. M., “Radithor and the era of mild radium therapy,” *Journal of the American Medical Association*, Vol. 264 No. 5, 8/1/1990, <http://jama.ama-assn.org/cgi/content/abstract/264/5/614>.

³⁷ Segel, Lawrence, “Ginger Jake Blues,” *The Medical Post*, Vol. 38, Iss. 34, 9/24/2002, <http://www.medicalpost.com/mpcontent/article.jsp?content=/content/EXTRACT/RAWART/3834/41A.html>

³⁸ Ballentine, Carol, “Taste of Raspberries, Taste of Death: The 1937 Elixir Sulfanilamide Incident,” *FDA Consumer Magazine*, 6/1981, <http://www.fda.gov/oc/history/elixir.html>.

³⁹ “The History of the FDA: The 1938 Food, Drug, and Cosmetic Act,” *U.S. Food and Drug Administration*, <http://www.fda.gov/oc/history/historyoffda/section2.html>.

⁴⁰ “Update from the Office of Postmarketing Drug Risk Assessment,” *U.S. Food and Drug Administration*, <http://www.fda.gov/cder/present/dia-62000/opdra/>; “Post-Marketing Surveillance,” *U.S. Food and Drug Administration*, <http://www.fda.gov/cder/handbook/postmark.htm>. The recent problems related to certain COX-2 inhibitor drugs have raised calls to strengthen post-marketing surveillance, e.g., Tanne, Janice Hopkins, “FDA will increase postmarketing surveillance of drugs,” *British Medical Journal*, 11/20/2004, <http://bmj.bmjournals.com/cgi/content/extract/329/7476/1203-a>.

“body personal.” Voting systems are, likewise, preparations (though electronic instead of chemical) used to maintain the health of, or to treat diseases of, the “body politic” – that is, the American democratic republic. Since the events of Florida 2000, e-voting systems have been touted as cure for a variety of electoral ills, from reducing tabulation errors, to eliminating over-voting and curbing unintentional under-voting, to improving accessibility for the disabled, to making voting faster and easier for everyone. Under something of a congressional mandate,⁴¹ jurisdictions around the nation have raced to adopt (usually DRE-based) e-voting systems, all without any significant assurance against vendor-based cheating.

As even *An Assessment* admits, e-voting systems have failed numerous times.⁴² But the failures have not been limited, as he implies, mainly to denials of service. They sometimes have included, for example, failure to record cast votes,⁴³ the subtraction of votes from a candidate’s legitimate total,⁴⁴ inability to cast the desired vote,⁴⁵ “miscounts,”⁴⁶ failure to show some candidates and showing others incorrectly,⁴⁷ the attribution of one candidate’s votes to another,⁴⁸ and many others.⁴⁹ Perhaps all of these well-documented failures are innocent errors. And perhaps, hidden beneath *An Assessment*’s weighty burden of proof, some actual frauds lurk. But even accepting, for the sake of argument, that none do, dishonest individuals or entities fraudulently can use otherwise-innocent failures by ignoring those favoring their chosen candidates or causes, and correcting those that do not.⁵⁰ Thus even innocent failures imperil voting security.

Clearly we have been dosing the body politic with unsafe “drugs.” Perhaps it is time to reverse the burden of proof that gave us those “drugs,” and force their proponents to show their safety (security) and effectiveness (proper operation) *before* we use them.

3.3) It Seems to Have Worked So Far, So It’s Fine

An Assessment’s argument about “without a single verified incident of successful tampering” has another aspect that the burden-of-proof discussion did not explore. And that is the implied argument that because something has (appeared) to operate correctly so far, it will continue to operate correctly. In some contexts, particularly those involving continuous systems, this argument has validity. For example, it is

⁴¹ “Help America Vote Act of 2002,” *Federal Election Commission*, http://www.fec.gov/hava/law_ext.txt. Section 301(a)(3) of this statute, in particular, has motivated the widespread adoption of DRE-based e-voting systems. Interestingly, this section requires only that each polling place provide “at least one” voting machine accessible to the disabled. Probably because of the (very legitimate) desire to buy, maintain, and use only a single system, this section has become a lever forcing the use of DRE-based systems by *all voters*, whether or not they need the accessibility they provide.

⁴² Shamos at §1.5.

⁴³ “Diebold in the News – A Partial List of Events,” *VotersUnite.org*, <http://www.votersunite.org/info/Dieboldinthenews.pdf> at p.6 (San Diego).

⁴⁴ *Ibid.* at p.2 (Florida).

⁴⁵ *Ibid.* (Maryland).

⁴⁶ *Ibid.* at p.3. (Kansas).

⁴⁷ *Ibid.* at p.11. (Georgia).

⁴⁸ *Ibid.* at p.3-4 (California).

⁴⁹ “Malfunctions and Miscounts, Sorted by Vendor,” *VotersUnite.org*, <http://www.votersunite.org/info/messupsbyvendor.asp>

⁵⁰ One potent technique for making partisan use of ‘random’ error is for dishonest elections officials to distribute reliable voting machines to precincts favoring the candidates they favor, and less-reliable (but still legally-compliant) machines to other precincts. If failures mainly impede vote casting, this technique will cause a greater proportion of votes to be lost in the disfavored precincts, thus tilting the election in favor of the dishonest officials’ candidates.

reasonable to say, “This bridge has stood for 15 years. It’s safe.” Why? Because, in 15 years, a bridge successfully will have withstood the vast majority of the stresses it is likely to encounter on any given day.⁵¹ Further, the bridge responds continuously—and therefore predictably⁵²—to new stresses. Software, on the other hand, is a kind of discrete system. That one has exercised the “it works OK” operational path trillions of times says nothing whatsoever about the existence of an error operational path⁵³—or a “steal the election” operational path. It may be that the right trigger for that path has never been received—or that, as previously discussed, our threshold for detecting cheating is set so high that we are unable to perceive its effect.

Finally, unlike bridges, software is subject to constant revision. Even assuming that vendors have been strictly honest so far, nothing prevents them from cheating tomorrow, the next day, or whenever vigilance wanes and the opportunity presents itself.⁵⁴ Real security arises from procedures that prevent cheating, such as total public software and hardware review, not from mere trust.

3.4) Comparing Voting Systems to Flight Control Systems

An Assessment argues there is no “engineering difference that allows us to entrust our lives to aircraft but would impel us to avoid voting machines.”⁵⁵ This argument ignores at least four crucial distinctions between these two applications of technology.

First, the motivations are quite different. The motivation to “cheat” (that is, to murder people) with flight software can only be that of a terrorist, a murderer, or an insane person. It is overwhelmingly likely that such a motivation will be harbored by at most a single person on a project, whose chicanery is thus likely to be caught by others. The motivation to cheat with voting software, however, encompasses not only individuals, but also – more importantly – whole organizations. A vendor might, for example, wish to further its business or political objectives by promoting the political success of a particular party or candidate. Its control of a voting system would allow it to do so illegitimately. This incentive has no parallel for manufacturers of flight control systems.⁵⁶

Second, aircraft use redundancy to protect against certain types of failure. Typically they employ three physically separate flight control systems, the proper operation of any one of which supports proper

⁵¹ This does not, of course, mean that the bridge cannot fail, only that it is very unlikely to do so.

⁵² Continuous systems are not necessarily predictable; some (e.g., planetary climate) may exhibit chaotic behavior. Bridges, however, are engineered to exhibit predictable characteristics throughout (and significantly beyond) their intended range of operation. Probably the most famous bridge collapse was the Tacoma Narrows Bridge, which collapsed quickly after its construction (see, e.g. <http://www.lib.washington.edu/specialcoll/exhibits/tmb/> and <http://www.wsdot.wa.gov/TNBhistory/Connections/connections3.htm>). Failures, such as the Nimitz Freeway collapse of 1989, now occur rarely, and usually are caused by the violation of established engineering standards. See, e.g., Wearne, Phillip, “Collapse: When Buildings Fall Down.” *TV Books, L.L.C.*, 2000. Violations of appropriate computer security standards, and failure to treat irregularities seriously, likewise endanger elections’ security, but appear to be far from rare.

⁵³ This is a major reason that all significant software contains bugs. The possible number of paths through any significant application is astronomical. An engineer cannot possibly consider them all. Proper engineering practices, such as encapsulation, can reduce the effective number of paths that need to be considered by grouping well-tested operations together, but they only reduce the problem’s magnitude without solving it.

⁵⁴ Vendors sometimes have included uncertified code in “updates.” See, e.g., “Staff Report on the investigation of Diebold Election Systems, Inc.,” *California Secretary of State*, 4/20/2004, http://www.ss.ca.gov/elections/ks_dre_papers/diebold_report_april20_final.doc.

⁵⁵ Shamos at §1.1.

⁵⁶ It is true that a flight system vendor might take shortcuts to reduce costs, thus reducing quality and potentially causing accidents. But that does not amount to the intent to cause an accident, and thus has little parallel to a dishonest vendor’s intent to defraud.

A Deeper Look: Rebutting Shamos on e-Voting

operation of the aircraft.⁵⁷ These separate systems cooperate to identify faults among themselves. Existing voting systems contain nothing like this; they are open-loop, non-redundant systems that provide no assurance of proper operation.

Third, commercial aircraft are operated by highly skilled flight crews, who continuously monitor their proper operation and can compensate for many (though not all) flight control system failures. Existing voting systems have essentially no such check upon them, particularly with respect to cheating,⁵⁸ and voting staffs typically employ lightly-trained volunteers.

Fourth, flight software is subject to *far* more rigorous specification, implementation, review, and testing procedures than voting software.⁵⁹ Flight software vendors take great pains to assure proper operation. For example, of the roughly \$5 billion Boeing spent to develop the B777 airliner, an estimated \$1 billion (20%) was spent *solely on software testing*.⁶⁰ This kind of attention to detail cannot be compared to that given by voting system vendors, particularly in light of existing voting systems' numerous, serious documented failures.⁶¹

3.5) Comparing Voting Systems to Financial Systems

In his §2.1, *An Assessment* implicitly argues that voting is similar to financial transactions (such as those conducted through a bank's check-clearing system), and, since financial transactions successfully have been computerized, voting successfully can be computerized as well – presumably through the use of DREs. However, he ignores seven key differences between financial transactions and voting with existing e-voting systems – differences that erode the core of his argument.

First, financial transactions (other than small cash transactions) are almost always traceable to the participating individuals or entities, while votes *must not* be traceable to the voters who cast them. Permitting vote traceability would encourage voter coercion and vote selling – maladies once so severe that they motivated the nearly-universal adoption of the secret ballot.⁶²

Second, financial transactions form a closed loop whose correctness can be verified in its entirety – such as by the comparison of receipts and periodic statements, double-entry accounting, the comparison of inflow and outflow records between multiple entities, and other forms of auditing.⁶³ Voting, as commonly implemented and as *An Assessment* advocates, is an open-loop process in which the voter must simply trust that her vote is properly counted.⁶⁴

⁵⁷ Cortellessa, Vittorio, et al., “Certifying Adaptive Flight Control Software,” <http://www.isacc.com/presentations/3c-bc.pdf> at p.3.

⁵⁸ See supra §3.2.

⁵⁹ See, e.g., the FAA's DO178-B standard, <http://www.linuxworks.com/solutions/milaero/do178-b.php3>, which governs all aviation software.

⁶⁰ Cortellessa, Vittorio, et al., “Certifying Adaptive Flight Control Software,” <http://www.isacc.com/presentations/3c-bc.pdf> at n.1.

⁶¹ See note 8 supra.

⁶² See, e.g., *Burson v. Freeman*, 504 U.S. 191, 200-05 (1992), <http://caselaw.lp.findlaw.com/scripts/getcase.pl?navby=case&court=us&vol=504&page=191>

⁶³ These auditing techniques almost always are based upon the comparison of two or more data sets describing the same transactions. See “Generally Accepted Accounting Principles,” <http://www.fasab.gov/accepted.html>, which discusses auditing standards generally applicable to business transactions.

⁶⁴ The proper use of voter-verified paper trails (“VVPAT”) or ballots (“VVPB”) can help reduce risks during vote casting by providing an independent audit trail that can be compared against matching electronic records or against official totals, possibly even by citizens themselves. VVPB may be more secure than VVPAT because it requires the tabulation of all the paper *ballots*, rather than simply the statistical spot-checks ...continued on 11

A Deeper Look: Rebutting Shamos on e-Voting

Third, each participant in a financial transaction has a strong incentive to verify its correctness. The incentive to verify correctness in voting is generally substantially weaker. While some highly-motivated voters always will wish to verify whether their votes properly are counted, many others will not.⁶⁵ Further, the incentive to verify correctness is likely to erode with the passing of each (apparently) uneventful election.

Fourth, there is strong legal recourse for financial fraud, including restitution, financial penalties, and frequently-enforced criminal sanctions. Further, evidence of fraud usually causes an in-depth investigation to be launched.⁶⁶ In contrast, there is little real legal recourse for vote fraud, and officials often are reluctant to investigate. Generally election contests must be made within extraordinarily short time frames, and must overcome high evidentiary hurdles. For example, Ohio's 2004 Presidential recount began the day the election effectively became irrevocable,⁶⁷ and the Supreme Court cited a similar deadline as justification for ending the 2000 Presidential litigation.⁶⁸

Fifth, financial attacks generally are waged by single hackers or by trusted insiders, not by the banks that conduct the transactions. Banks that condone fraud can look forward to quick bankruptcy, and their officers to lengthy prison sentences. The concern with computer-related vote fraud, on the other hand, is not so much with hackers as with vendors.⁶⁹

Sixth, financial transactions are repeatable, money often can be reclaimed, and a certain level of loss can be recovered by raising the prices charged to conduct the transactions. Elections, in contrast, are not

...10 required by VVPAT. Both VVPAT and VVPB are subject to presentation attacks (see note 4) and vote recording attacks, such as when the voter does not actually verify the paper (although how would the system know when the voter will not verify the paper?) and when the paper is not actually counted. Other technologies, like Neff's (<http://www.votehere.com>), may also be helpful, though even this technology is vulnerable to vendor-based attacks. For one thing, it also does not prevent certain presentation attacks. For another, it can be understood only by those possessing an advanced understanding of cryptography, and thus cannot effectively be supervised even by most computer scientists, let alone by ordinary citizens. These systems are still at risk of error or fraud in the process of aggregating the individual voting machine or precinct totals. Elimination of the potential for fraud or error in the aggregation process is an area requiring further study.

⁶⁵ Particularly at busy polling places or after waiting in a long line.

⁶⁶ One exception is retail-level credit card fraud, which card issuers have tended to write off, concluding that the cost of investigation is higher than the probable restitution. The recent trend, is, however, toward prevention and more thorough investigation. See, e.g., MacDonald, Jay, "Internet fraud: credit crooks and the new economy," <http://www.bankrate.com/brm/news/cc/20010417a.asp>

⁶⁷ David Cobb, the Green Party candidate for President, asked for a statewide recount on 11/19/2004, 17 days following the election. It was not granted until 12/13/2004, when Ohio's electors were already voting. "Cobb Testifies Before Congressional Forum in Ohio," *Cobb-LaMarche 2004*, <http://www.votecobb.org/press/2004/dec/pr2004-12-13b.php>

⁶⁸ See *George W. Bush, et al., v. Albert Gore, Jr., et al.*, case 00-949(12/12/2000) at II(B), <http://supct.law.cornell.edu/supct/html/00-949.ZPC.html>:

...That statute, in turn, requires that any controversy or contest that is designed to lead to a conclusive selection of electors be completed by December 12. That date is upon us, and there is no recount procedure in place under the State Supreme Court's order that comports with minimal constitutional standards. Because it is evident that any recount seeking to meet the December 12 date will be unconstitutional for the reasons we have discussed, we reverse the judgment of the Supreme Court of Florida ordering a recount to proceed.

⁶⁹ Or, possibly, with officials or hackers exploiting vendors' malware or security flaws.

repeatable, votes usually cannot be reclaimed, and vote loss cannot be compensated through other means.

Seventh, the software involved in financial transactions generally is very carefully reviewed and controlled. Banks and other financial institutions strictly control how insiders can use their systems. In contrast, the software involved in voting has, to date, often been handled in an extraordinarily sloppy manner.⁷⁰

A Better Parallel: Gambling Systems

An Assessment is not entirely off the mark in comparing financial and voting systems: one kind of electronic financial system is closely analogous to e-voting systems. And that is *electronic gambling systems*.⁷¹ As with existing e-voting systems, gambling systems are opaque, open-loop systems. The gambler puts her coins in the machine, pulls the lever, and gets a chance at a payoff; the voter makes her selections, clicks “Cast Vote,” and gets a chance at having her vote counted. Neither the voter nor the gambler knows what is inside the system, and neither can check whether she receives the outcome she is due.

Nevada – the gambling capital of the nation – long has recognized the need to ensure gambling devices’ honesty, and has established an agency to address it, whose

goal and objective...is to ensure that gaming is conducted honestly and competitively and the public has confidence and trust in the gaming industry. In other words, we want the games to be fair and patrons to have a good time.⁷²

To further these goals, Nevada requires electronic gambling systems to undergo both software and hardware inspections, including random inspections during ordinary use.⁷³ In fact, an entire division of the Nevada Gaming Control Board – the Electronic Services Division – is charged with “inspect[ing] gaming devices in its laboratory and in the field to ensure continued integrity, and assist[ing] in resolving gaming patron disputes through analysis of device electronics and software.”⁷⁴

The Gaming Control Board has detected a variety of vendor frauds. Perhaps the most prominent involved American Coin, which had rigged its poker machines to unfairly reduce the probability of royal flushes, and its keno machines to unfairly reduce the probability of top jackpots. The Board pulled the company’s license and fined it \$1 million.⁷⁵ In another case, Universal Distributing had programmed its slot machines to show clear misses as “near misses” (where two symbols were on the line and one very close) to entice gamblers to continue playing. The Gaming Control Board eventually outlawed this practice.⁷⁶ There was even a fraud perpetrated by a Gaming Control Board member himself, who programmed some

⁷⁰ See, e.g., “Electronic Voting Systems: the Good, the Bad, and the Stupid,” <http://www.acmqueue.org/modules.php?name=Content&pa=showpage&pid=219&page=2> (Diebold).

⁷¹ Others have made this comparison, e.g., “Gambling on Voting,” *New York Times* editorial, 6/13/2004, <http://query.nytimes.com/gst/abstract.html?res=F60F17FD38540C708DDDAF0894DC404482&incamp=archive:search>.

⁷² “Slot Machine Malfunctions: Frequently Asked Questions,” *Nevada Gaming Commission and State Gaming Control Board*, http://gaming.nv.gov/slot_machine_malfunctions.htm.

⁷³ Nevada Revised Statutes ch. 463.670, <http://leg.state.nv.us/NRS/NRS-463.html>.

⁷⁴ “State Gaming Control Board Information Page,” *Nevada Gaming Commission and State Gaming Control Board*, http://gaming.nv.gov/about_board.htm#esd.

⁷⁵ Bourie, Steve, “Are Slot Machines Honest?” *American Casino Guide* (1999), <http://www.americancasino.com/Tips/Slots-Honest.shtml>. The programmer who disclosed the cheating subsequently was murdered. O’Connell, Peter, “Pair sentenced for roles in killing of witness in slot-rigging case,” *Las Vegas Review-Journal*, 12/22/1999, http://www.lvrj.com/cgi-bin/printable.cgi?lvj_home/1999/Dec-22-Wed-1999/news/12605224.html.

of the slot machines he inspected to pay off when coins were inserted in a certain sequence.⁷⁷

This shows not only that vendors and others in positions of power can cheat, but also that even intrusive government regulation (though valuable) can be insufficient to uncover or discourage cheating. It thus argues for total public review of voting hardware and software, and for a mechanism that guarantees that the reviewed software is actually used on Election Day.⁷⁸

4) Implementational Errors

4.1) *Overestimating the Effort Required to Cheat and Underestimating What Cheats Can Accomplish*

An Assessment consistently overestimates the amount of effort and mechanism required for a dishonest vendor to install and operate cheating software. For example, in §1.2, *An Assessment* considers what would be required for a vendor to cheat if it were constrained to provide each precinct in the nation with identical software, and concludes that “[t]he practical possibility of such a scheme is nil.”⁷⁹ In reality, it is “nil” only if we accept his invalid assumptions about how such cheating would be conducted and how its effects would appear to elections officials and to the public.

For example, he says that “[i]t is not possible to move a constant fraction of votes from one party to another in each jurisdiction without it being obvious that manipulation is going on because the political demographics of the precincts are too individualistic and distinctive.” There is little evidence for this proposition. A cheat that moves 1-2% of the votes from one major party’s candidates to the other’s would be very unlikely to be detected in the vast majority of precincts,⁸⁰ as long as it can discern tests from actual voting.⁸¹ Why? First, opinion polls routinely fluctuate by small single digits from day to day,⁸² and elections routinely are decided by such margins. Second, even large variations from expected results seem not to trigger inquiries into potential fraud. For example, the 11-point, 3-day swing in the 2002 Cleland/Chambliss race⁸³ has not, to the authors’ knowledge, triggered any official investigation. Shamos himself seems to believe that no such investigation is warranted, since he fails even to mention this race, though it was the jurisdiction’s first using DREs.⁸⁴

⁷⁶ Ibid.

⁷⁷ Robison, John, “Ask the Slot Expert,” *Casino City Times.com*, April 29, 2002, <http://robison.casinocitytimes.com/articles/417.html>; Bourie, Steve, “The World’s Greatest Slot Cheat?” *American Casino Guide* (1999), <http://www.americancasinoguide.com/Tips/slot-cheat.shtml>

⁷⁸ See note 19 *infra* for such a procedure.

⁷⁹ Shamos at §1.2.

⁸⁰ Even precincts that most heavily favor a given major party plausibly could be manipulated by small amounts. For example, in the 2004 Presidential race, San Francisco precinct 3823 polled 95.7% for John Kerry and 1.8% for George Bush. <http://web.sfgov.org/site/uploadedfiles/election/Guides/SOV041102.pdf> A fraud probably could have doubled Bush’s tally without raising alarms. Reducing Bush’s tally would have been more chancy, since the precinct recorded only 12 votes for him. Generally, moderately raising a candidate’s totals where she does poorly can affect the overall result without being noticed, as can moderately lowering her totals where she does well.

⁸¹ This need not be very difficult. See §4.3 *infra*. Further, tests often are conducted incorrectly, or not at all; as Shamos notes, “The administrative procedures concerning the handling of DRE machines and materials are usually not spelled out at all, or, if spelled out, then not circulated and not followed.” Shamos at §3.3.

⁸² See, e.g., <http://www.electoral-vote.com/2004/info/state-graphs.html> (Day-to-day tracking polls of 2004 Presidential election); <http://pollingreport.com/BushJob.htm> (Bush job ratings; look especially at the Gallup and CNN/USA Today/Gallup results, which cover long, continuous spans of 3-4 day sampling intervals).

⁸³ See *supra* §3.2.

A Deeper Look: Rebutting Shamos on e-Voting

Further, a cheat need not be so unsophisticated as to move a constant fraction of votes. It could observe the voting and move votes only when required, using some sanity checks (e.g., never move more than 5% of the votes, never reduce a candidate's total to below some small number, etc.) to avoid triggering suspicions. Other, more sophisticated cheats using concealed communication devices could be even more flexible, and are discussed below.

Weaving another link into its chain of error, *An Assessment* write,

Therefore the software would have to be distributed with a database telling it how to alter the vote for each relevant candidate in each precinct. The database would have to contain at least the names of political parties and possibly candidates and would have to know in advance the precise hours during which all future elections are to be conducted so the machine would know when to behave properly.

Even accepting, for the sake of argument, that it is correct about the impracticality of moving a constant fraction of votes from one major party's candidates to the other, a dishonest vendor need not go to such lengths. For example, it could routinely distribute an "update" to each jurisdiction prior to each election. Even an honest vendor might well use such a procedure to distribute bug fixes and enhancements.⁸⁵ A dishonest one would simply also include information about the upcoming election. This easily could be accomplished by providing each precinct with its own "login" username and password, which it would use to download the updates; the usernames would tell the vendor's website which election information to include. Or, if the voting system uses Shamos' "ultimate protection against malicious code,"⁸⁶ the vendor's image-generation program surreptitiously could include the election information.

Finally, if none of these approaches were open to a dishonest vendor, it could still cheat by combining the appropriate software with concealed communication devices. This is actually the most flexible kind of cheat, since it permits the vendor easily and surreptitiously to update both data and code.⁸⁷

Cheating With Communications Devices

Many modern computers contain wireless ("WiFi") devices. These usually operate over ranges of 100-200 feet.⁸⁸ The newest wireless innovation, called WiMax, has a range of up to 30 miles.⁸⁹ In addition, broadband over power lines ("BPL"), which uses ordinary power lines to carry its signals, could bring Internet connectivity to every power outlet in the nation.⁹⁰ Communications devices using these approaches are small and easily concealed within any computer, and thus within any computer-based voting system.

⁸⁴ Ibid.

⁸⁵ The law requires such updates to be certified, but certification testing, like any other form of black-box testing, is very unlikely to discover a competently-designed cheat. And vendors sometimes have included uncertified code in updates; see supra note 9.

⁸⁶ See infra §4.2.

⁸⁷ It also would be possible for a vendor to program a cheat to be activated, or to be notified of the appropriate cheating procedure, via a surreptitious sequence of keystrokes or ballots. Using such a cheat requires, however, the cooperation of more individuals than does an ordinary cheat.

⁸⁸ Zyren, Jim, et al., "802.11g Starts Answering WLAN Range Questions," *CommsDesign.com*, 1/14/2003, <http://www.commsdesign.com/story/OEG20030114S0008>.

⁸⁹ Vaughan-Nichols, Steven J., "WiMax is Coming, and It's Going to Be Big," *The Channel Insider*, 6/2/2004, <http://www.thechannelinsider.com/article2/0,1759,1605972,00.asp>.

⁹⁰ "BPL – Broadband over Powerline," *InfoCellar.com*, <http://www.infocellar.com/networks/new-tech/BPL/BPL.htm>.

A Deeper Look: Rebutting Shamos on e-Voting

Once a voting system contains a communications device, it becomes simple for a dishonest vendor to use it to modify the system's code or data.⁹¹ The device provides a standard Internet connection, which can be used to transfer anything. With short-range wireless devices, the vendor would have to establish a station to supply cheating information close to each polling place, which would be impractical in rural areas, but quite feasible in cities – mobile stations placed in ordinary cars would suffice. With WiMax and BPL, far fewer cheating stations would be necessary. As this technology (and other, more advanced communication technology) is deployed, this form of cheating will become practical nationwide.

One might try to thwart this kind of cheating by assiduous application of open-source techniques, including full public review, the proper and consistent use and checking of cryptographic wrappers, and so forth. These techniques are valuable, but even they are not proof against hardware-based cheating.

Cheating with Hardware – The Malware Loader

What is “hardware-based cheating”? It is a kind of cheating that uses special firmware⁹² in concert with a communications device. Its most obvious form is a “malware⁹³ loader.”^{94,95} Such a loader can operate even if the originally-loaded voting application and the operating system are completely free of cheating code. It makes only one demand on the voting application: that it make periodic calls to a function whose address is fixed and known to the firmware⁹⁶ (the “periodic function”). An application might make such a call for a variety of legitimate purposes, such as to update the time on the display, make an audit entry, flush data to disk, animate a logo, and so forth, so its presence would not trigger any suspicion during review. The loader does not care what the function does; it could simply return. It is only necessary that the application call it periodically.

Cheating proceeds thusly: The vendor places the malware loader's code in some piece of firmware. Assume, for this argument, that the vendor chooses the video BIOS. That is, it modifies the BIOS code that writes to the video display to also call the malware loader. When it is called,⁹⁷ the loader monitors the system's concealed communications device for a “cheating signal.” When it detects this signal, it disables interrupts on the CPU, which prevents anything (including the operating system) from interfering with its task. Next, it loads a “malware bootstrap” into an unused area of memory.⁹⁸ When it has done so, it modifies the first instruction of the periodic function to be a jump to the beginning of the malware bootstrap. Then it purges the CPU's caches and instruction pipeline to ensure that the CPU can see the malware bootstrap, re-enables interrupts, and returns to its caller.⁹⁹

⁹¹ This cheat works whether the communications device is concealed (e.g., wireless) or patently present (e.g., Ethernet), as long as the device provides access to a network (e.g., the Internet) that is also accessible to the vendor.

⁹² “Firmware” is software built into some pieces of computer hardware. The most familiar kind of firmware is a PC's BIOS (“Basic I/O System”), which provides a standard interface between the operating system and the hardware, such as keyboards and disks. A somewhat less familiar kind is the video BIOS, which usually resides on a video card and provides a similar interface for video operations. But computing devices often contain other, less well-known kinds of firmware, including that present in FPGAs (“Field-Programmable Gate Arrays”) and ASICs (“Application-Specific Integrated Circuits”).

⁹³ “Malware” is a common term for software that performs a malicious function.

⁹⁴ Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten, “Security Analysis of the Diebold AccuVote-TS Voting Machine”, <http://itpolicy.princeton.edu/voting/>.

⁹⁵ State of Maryland Affidavit of Curtis Feeney, http://www.buzzflash.com/alerts/04/12/images/CC_Affidavit_120604.pdf, executed on 12-06-04.

⁹⁶ This restriction can be removed by appropriate use of versioning information.

⁹⁷ (presumably when the application asks the operating system to write something to the video display)

⁹⁸ The address and size of this area would be predetermined and fixed for all versions of the voting system.

A Deeper Look: Rebutting Shamos on e-Voting

At this point the malware bootstrap is in memory and the first instruction of the periodic function contains a jump to it. Eventually the voting application receives control, and eventually it calls the periodic function. When it does so, the malware bootstrap takes over. It knows all about the application and the communications device, and it uses the device to load a cheating application into memory, overwriting portions of, or even the entire, original application.¹⁰⁰ When it has finished loading, it gives control to the cheating application, which does whatever is necessary to accomplish its task.¹⁰¹

Note, again, what the malware loader accomplishes. It starts with a communications device (possibly concealed, like WiMax or BPL), a small amount of cheating firmware, and an application and operating system that contain *no cheating code whatsoever*. It requires only that the application periodically call a function whose address is known: something that has many legitimate uses. Using these resources, and the appropriate information available via the communications device, it loads a fully functional cheating application into a voting machine's memory and gives it control.¹⁰²

What does this mean? At the very least, it means that we must inspect voting system hardware to ensure that it contains no concealed communications devices. Further, we must ensure that such systems have no access to the Internet, since such access could be used in place of a concealed communications device to broadcast the cheating signal and to allow access to the cheating application. Unfortunately, technological advances will make communications devices ever smaller and harder to detect. It is even now not beyond the realm of possibility for a vendor to place such a device inside an otherwise-innocent ASIC, perhaps one that also houses a video or disk controller.

Finally, the malware loader's practicability also shows that even experts—like Shamos—cannot anticipate all the ways in which a dishonest vendor might corrupt its voting system. As always, the best defenses are humility and continual vigilance.

4.2) Touting Separation of Candidate Names As a Panacea

In §3.6, *An Assessment* proposes defeating vendor-based frauds by removing the knowledge of candidate names from voting stations. In this proposal, the candidates' names would be represented as graphic files, each corresponding to a certain ballot position. Shamos calls this “the ultimate protection against malicious code.”¹⁰³

Not only does this proposal fail to accomplish its objective, it can be used to assist the functioning of malicious code, and even to assist its installation. First, as even *An Assessment* admits, OCR (optical character recognition) technology would enable the voting station to read and interpret the images.

⁹⁹ Which is probably the operating system.

¹⁰⁰ The cheating application is, of course, available via the communications device.

¹⁰¹ A malware loader need not reside in firmware. It could also be embedded in the operating system or in one of its device drivers. The authors have chosen to depict the firmware-based loader because it illustrates the need for thorough hardware inspection, and shows that even the most subtle spiders – or expert computer scientists – sometimes leave loose threads. Note that it would be easier to implement an OS-based loader than a firmware-based one, because the operating system knows more about applications than does the firmware, and because the programming environment is more flexible and forgiving.

¹⁰² Note that the malware bootstrap itself remains in memory, and can supervise further cheating. For example, it could save the cheating application in an unused (and unlikely to be used) area of the system's hard disk or flash memory, thus ensuring that it is available for use (by the malware loader) in succeeding elections, even if the communications device becomes unavailable or if the cheating signal is not broadcast.

Further, the malware bootstrap could avoid detection by even highly intrusive audits by replacing the cheating application with the original one when the polls close, and then deleting itself. This step is arguably unnecessary, since it is very difficult to audit the contents of a running system's memory.

¹⁰³ Shamos at §3.6.

Second, the graphic files presumably would be generated either directly by the vendor, or by using a vendor-provided program. Either method would allow the vendor to insert hidden data into the image, either via proprietary headers, via steganography,¹⁰⁴ or possibly via other less-obvious means. A dishonest vendor would, at a minimum, include in this data the mapping between the image files and candidates' names and parties, thus utterly defeating Shamos' scheme. But a craftier vendor might also include data indicating when to cheat, how much to cheat, how to detect when an anti-cheating test is being conducted, and so forth. Further, a vendor could include executable code in the images, which the voting station could then load and execute.¹⁰⁵

Finally, this approach does not impede the malware loader described in section 4.1, nor does it address frauds during tabulation.

4.3) Overestimating What Testing Can Accomplish

In his §1.2.1 and §3.5, *An Assessment* advocates something Shamos calls “parallel testing.” In this approach

a set of examiners is empowered to enter any polling place at the start of voting and commandeer any voting machine for test purposes. No actual voters cast votes on the selected machine. No change whatsoever is made to the test machine – it is not even moved from its position (to counter the argument that it might contain a motion sensor to warn that it was under test). The examiner votes a number of predetermined ballots comparable to the number that would be voted on a typical machine in that precinct. Of course, manual entry of votes by a human is an error-prone process, so a video camera is used to capture his actual vote entries. At the normal close of polls, the votes on the test machine are tabulated and compared with the expected totals.¹⁰⁶

Of this he writes

If any software is present that is switching or losing votes, it will be exposed.¹⁰⁷

This assertion is incorrect. While this test is valuable, and should be used during every election, it can fail to achieve its purpose in a number of ways.

First, the testing conditions, including the vote stream, must be truly indistinguishable from those that obtain during regular voting. If the voting station is a networked device, it might not be possible to do this without cueing it into the test, since it is necessary, of course, to prevent the “test” votes from being included in the actual election results – and this might be accomplished by disconnecting the voting station from the network. Further, a cheat might observe the vote stream and notice that votes are being cast much more rapidly than during regular voting, or that a block of votes was cast during the first hour of the election, then none for the remainder of the day. Or the voting station might “helpfully” provide a “test mode” – most elections officials would have no idea that this could be a cue to record votes

¹⁰⁴ Johnson, Neil F. and Jajodia, Sushil, “Steganography: Seeing the Unseen,” *IEEE Computer*, February 1998: 26-34, <http://www.jjtc.com/pub/r2026.pdf>

¹⁰⁵ The vendor could even regularly update the information added by the image generation program. One approach would be to host the program on its website and to require elections officials to login, which would identify them – and the elections they administer. Another would have the program surreptitiously download information from the vendor's website, and still another would have the vendor “helpfully” provide regular “updates” for downloading.

¹⁰⁶ Shamos at §3.5.

¹⁰⁷ Ibid.

A Deeper Look: Rebutting Shamos on e-Voting

honestly.¹⁰⁸ Lastly, the testers might use something – like a special login or logout procedure, or a voting card with a known, constant test key – that would cue the voting station into the test. It is very difficult to design a test that, even when executed perfectly, outwits a determined adversary.^{109,110}

Second, a communications device (or the security hole created by *An Assessment's* “ultimate protection against malicious code”¹¹¹) could be used to tell the voting station how to detect the latest testing procedures.

Third, the test must actually be performed diligently and in a statistically-significant set of precincts. Yet *An Assessment* itself criticizes both elections procedures and the manner in which they commonly are performed, saying

The administrative procedures concerning the handling of DRE machines and materials are usually not spelled out at all, or, if spelled out, then not circulated and not followed.¹¹²

Given this, it seems likely that parallel testing often will fail of its purpose. This failure will be particularly acute if the first few tests “find nothing,” since that will create a natural tendency to view the process as pointless.

Fourth, elections officials must actually believe, and act properly upon, the suggestion or discovery of cheating. Far too often officials treat anomalies (such as the one in which 600 voters somehow cast 3,900 votes for a Presidential candidate in Ohio—using DREs¹¹³; or 18,000 empty electronic ballots Sarasota, Florida in 2006¹¹⁴) as “glitches,” and merely remove the offending machine(s) from service. But does a more reasonable response exist? It is Election Day, the machines are cheating, and what is an official to do? Shut down the polls? There really is no practical way to save that election, though proper exposure of such anomalies would help to prevent future cheating.¹¹⁵

¹⁰⁸ And, of course, there might be no good way to do a “dry run” except by use of “test mode,” since officials have to ensure that the “test” votes are not included in the actual election results.

¹⁰⁹ E.W. Dijkstra, in “On the Reliability of Programs,” says of a similar issue that “program testing can be used very effectively to show the presence of bugs but never to show their absence.”
<http://www.cs.utexas.edu/users/EWD/transcriptions/EWD03xx/EWD303.html> at p.3.

¹¹⁰ See also note 14 supra for cheats that can be activated via specific key or ballot sequences. Testing is extremely unlikely to discover such cheats.

¹¹¹ Shamos at §4.1.

¹¹² Ibid. at §3.3.

¹¹³ McCarthy, John, “Voting machine error gives Bush 3,893 extra votes in Ohio,” *San Francisco Chronicle*, 11/5/2004, <http://sfgate.com/cgi-bin/article.cgi?f=/news/archive/2004/11/05/politics1149EST0515.DTL>.

¹¹⁴ Kim Zetter, “Docs Point to E-Voting Bug in Contested Race,” *Wired Magazine*, March 17, 2007, <http://www.wired.com/politics/onlinights/news/2007/04/evotinganalysis>.

¹¹⁵ It would be desirable to establish strong standards for such investigations, possibly along the lines of those used by the National Transportation Safety Board to investigate commercial aviation disasters.

4.4) *Misunderstanding the Nature of Open-Source Software*

In §3.2, *An Assessment* endorses open source requirements¹¹⁶ for “the ballot setup, display, tabulation and reporting sections of voting system code.” This is a step forward, but it is immediately followed by an expression of sympathy for security by obscurity for other, unspecified types of code, on the basis that such obscurity makes hackers’ jobs more difficult.

This “partial open source” approach essentially nullifies open source’s security benefits, because the undisclosed code can modify the operation or the results of the “open source” code. Since it is undisclosed, we have no way to know what it does.¹¹⁷ One might be tempted to conclude that a smaller amount of undisclosed code (judged, perhaps, by executable file size) is more secure than a larger amount, but only a tiny amount of code is necessary to implement a malware loader.¹¹⁸ Further, while security by obscurity may, indeed, make a hacker’s job more difficult, dishonest vendors are a significantly greater concern.

4.5) *Overestimating the Benefits and Practicality of Independent Audit Devices*

In §3.1, *An Assessment* proposes separating vote-casting stations into three devices, each made by a different vendor. In this scheme, the touchscreen displays choices and accepts voter input, then sends the voter’s choices to both an audit device and a tabulation device. The audit device has its own screen on which the voter can verify the correct recording of her votes. If the voter accepts the choices, both the

¹¹⁶ A canonical statement of the procedure necessary to ensure that the publicly reviewed (“open”) source actually makes it into the machines on Election Day was written by co-author David Mertz. We present it with slight modifications:

- (1) Date 0: The “build maven” releases code and build instructions to the wider software and elections community. This includes the instruction: “This code, when compiled/assembled/linked/processed should hash to...”
- (2) Dates 0 through N: The community checks the “release candidate” code. Such checks include both the mechanical check of crypto stuff (hashes and/or other things, public-key, etc.) and examination of the underlying code.
- (2a) Dates 0 through N: If problems are encountered (bugs, failed hashes, etc.), restart the process.
- (3) Date N+1: The hash codes for the final “this year’s election” code are published in the relevant newspapers, websites, etc.
- (4) Date N+M: Hold the election. Poll workers compare the hashes on the CDs they receive to those widely published and accepted by the community of evaluators.

The poll workers need to be trained to do something like the following:

- a) Insert the CD into an independent, separate computer (i.e., not running the election software CD itself).
- b) Type something like `sha voting-station`, where “sha” is a non-vendor-coded program that uses the appropriate (publicly-known) algorithm to compute the CD’s hash.
- c) Hold the local newspaper that pre-published the correct hash in their hand.
- d) Look at the hash displayed on screen.
- e) Make sure the screen looks like the newspaper.

¹¹⁷ Shamos makes a similar error in his §1.4, in which he implies that public review of “the basic loop that interrogates portions of a touchscreen and interprets them as votes” is a sufficient security guarantee.

¹¹⁸ See *supra* §4.1.

A Deeper Look: Rebutting Shamos on e-Voting

audit device and the tabulation device record her votes. When the polls close, the records are compared. Of this arrangement, *An Assessment* asserts

So long as there is no collusion between the audit device manufacturer and the tabulation manufacturer, no amount of tampering with either machine will go unremedied.

This might be correct as far as it goes,¹¹⁹ but it ignores the possibility of presentation frauds on the touchscreen device. The touchscreen might, for example, place its preferred candidate first, or show her name in bigger, bolder text than the others, or enlarge her selection area (and shrink those of the disfavored candidates) to make it easier to vote for her, or even occasionally omit the names of disfavored candidates. These cheats might not deter a determined voter, but they significantly could influence which choices less-decisive voters make.

An Assessment's §3.1 also requires an audit procedure that may often not be followed correctly, blithely assumes that an audit mismatch will result in "an investigation [being] launched," and further assumes that such an investigation will be effective. But, as *An Assessment* notes in §3.3, "The administrative procedures concerning the handling of DRE machines...are usually not spelled out at all, or, if spelled out, then not circulated and not followed." Why, then, should we assume that the §3.1 audit procedure will be followed correctly? Further, even if this procedure is followed correctly, and an investigation of mismatches is launched, it is quite unlikely to do anything to repair the election that prompted it. At best it would help prevent a future fraud.¹²⁰

Finally, elections officials are unlikely to buy the system *An Assessment* hypothesizes, since doing so requires dealing with three unrelated vendors (messy, time-consuming, and possibly in violation of legal requirements concerning dealings with suppliers) and most officials (who are not trained in computer security) will not see the need.

4.6) Systematically Misunderstanding Paper Trails and Ballots

While *An Assessment*'s critique of paper ballots¹²¹ cites some real deficiencies of paper-only systems (e.g., that they can allow ballot-box stuffing), it also conflates machine-printed paper ballots and paper "trails"¹²² and includes a considerable amount of incorrect material. This section discusses some of the most important points.

First, *An Assessment* asserts that statutes requiring paper records to govern over electronic records when they mismatch "make[] the insecure paper record paramount over the secure electronic one."¹²³ This, of course, assumes that the electronic records are, in fact, secure. But, as we show elsewhere, existing e-

¹¹⁹ No obvious security flaw suggests itself, but that does not rule one out. Of course, nothing prevents the three vendors cited from merging into one after many of these systems have been deployed. Mergers and acquisitions are hardly unknown in the elections services industry. For example, Diebold acquired Global Election Systems in 2002. Kropko, M. R., "Dented Diebold: Ohio firm finds voting machine business stormy," *The Cincinnati Post*, 5/8/2004, <http://www.cincypost.com/2004/05/08/diebold05-08-2004.html> (sidebar).

¹²⁰ See supra §4.3.

¹²¹ Shamos at §2.4.

¹²² A machine-printed "paper ballot" is a human-readable paper recording of votes. It acts as their authoritative record if there is a conflict between it and any other medium on which the same voter's votes are recorded. A machine-printed paper ballot is intended to be verified by the voter for accuracy before the voter casts it, which she ordinarily does by placing it in a ballot box. Such paper ballots may be counted by hand or by a computerized tabulator. A "paper trail," commonly referred to as a "voter-verified paper audit trail" or "VVPAT," is a paper recording of votes that the voter verifies for accuracy, but never handles or "casts." It is maintained for auditing or recount purposes, but generally the corresponding electronic record is tabulated first. In practice, only a small number of VVPAT records are actually tabulated, and usually manually as part of a statistical test.

A Deeper Look: Rebutting Shamos on e-Voting

voting systems exhibit no real security, and even some of the proposed fixes do little or nothing to improve the situation.¹²⁴ Further, existing e-voting systems are insecure in a way that hand-counted paper systems are not. Since one program might handle tens of millions of votes, it might also subvert them. Paper ballots, on the other hand, must be subverted a ballot (or ballot box) at a time – a much more time-consuming and difficult procedure, and one much more likely to be discovered.

He goes on to assert that paper trails increase the likelihood of DRE failure because “the presence of the mechanism increases the load on the machine’s power supply and processor.”¹²⁵ With respect to the power supply, he neglects to note that the vendor will simply use a larger power supply to handle the additional load. With respect to the processor load, the authors are unaware of any credible information indicating that the small additional processor load required to print a paper trail (or paper ballot) would in any way reduce the processor’s reliability.

An Assessment asserts that the disabled cannot readily view paper trails,¹²⁶ but ignores that fact that ballot scanners, when equipped with headphones or other assistive technologies, can permit the disabled to verify printed ballots unassisted and confidentially.¹²⁷ Of course, such scanners introduce their own security issues, but these are smaller than the security issues introduced by ordinary DREs, since they will affect, at most, a small percentage of the ballots cast.

An Assessment further asserts that the statutory changes that would be required to accommodate paper trails or machine-printed paper ballots are “of great legal and, in some states, constitutional significance.”¹²⁸ Of course, there is nothing about defining the term “ballot,” nor about dictating what is counted in what sequence and under what conditions, that is beyond the capabilities of any state legislature or of the Congress. If there were, DREs never could have been adopted anywhere, since their introduction radically redefined what it means to cast, represent, and count a vote.

Finally, and most oddly, *An Assessment* seeks support in:

A report of the Caltech-MIT Voting project concluded that the presence of paper trails actually decreases confidence in the voting system [34].¹²⁹

Given this introduction, one might expect the “report”¹³⁰ to reveal significant research on the question of paper trails and voter confidence, possibly including experiments, surveys conducted during actual elections, and the like. In reality, however, the report contains precisely one sentence on this topic, for

¹²³ Shamos at §2.4 point 4.

¹²⁴ See, e.g., supra §4.2 (separation of candidate names); supra §4.4 (partial use of open source).

¹²⁵ Shamos at §2.4 point 6.

¹²⁶ *Ibid.* at point 9.

¹²⁷ See, e.g., “The Open Voting Consortium FAQ (Frequently Asked Questions),” http://www.openvotingconsortium.org/modules.php?name=FAQ&myfaq=yes&id_cat=8&categories=Access+for+Voters+with+Disabilities. See also “State of California Standards for Accessible Voter Verified Paper Audit Trail Systems In Direct Recording Electronic (DRE) Voting Systems,” http://www.ss.ca.gov/elections/ks_dre_papers/avvpatt_standards_6_15a_04.pdf, which describes standards for accessible paper trails (“AVVPAT”).

¹²⁸ Shamos at §2.4 at point 7.

¹²⁹ Shamos at §2.4 point 10.

¹³⁰ Cited by Shamos at endnote 34 as ‘Selker, Ted. et al, “The SAVE System: Secure Architecture for Voting Electronically: Existing Technology, with Built-in Redundancy, Enables Reliability,” CalTech/MIT Voting Project VTR Working Paper, Oct. 22, 2003, revised January 4, 2004.’ This paper is available at http://www.vote.caltech.edu/media/documents/vtp_WP7r.pdf

which it cites no authority:¹³¹

Adding paper to electronic voting system [sic] undermines the public confidence in electronic voting systems.

The authors confess to being stumped by this unsupported assertion.

5) Miscellaneous Errors

In §1.3, *An Assessment* claims that “Congress has no power to determine the manner in which presidential electors are chosen other than to specify the time and date of their election.” This, and its following assertion of unconstitutionality, are very likely incorrect. Section 5 of the 14th Amendment explicitly empowers Congress to enforce its provisions,¹³² one of which happens to be the equal protection clause.¹³³ This was the very clause that the Supreme Court invoked to stop the Florida recount in *Bush v. Gore*¹³⁴—a presidential contest. Thus it is likely that the Court would uphold Congressional action aimed at improving equality in vote counting.¹³⁵ This criticism also applies to a similar assertion concerning the “Protecting American Democracy Act of 2003” in §2.1.

In §2, *An Assessment* implicitly blames the invalidity of 337,297 ballots cast in Taiwan’s 2004 presidential election upon the use of paper ballots, saying, “Surely if the voters could rely on the paper ballots to be counted properly this result could not have occurred.” He provides, however, no breakdown of the reasons behind the declarations of invalidity, and does not address the potential contribution of the “Million Invalid Ballot Alliance,” which asked voters to reject both candidates by spoiling their ballots.¹³⁶

Finally, in §1.5, *An Assessment* implicitly argues that DREs are suitable for America because India has chosen to use them nationwide. The authors fail to understand Shamos’ intent here. Perhaps he will issue an updated version of his paper explaining his argument more clearly.

6) Conclusion

We face a serious problem. Existing e-voting systems easily might be subverted by their vendors, or their faults used by others to cheat.¹³⁷ To put our heads in the sand does not solve the problem, but permits it to fester, to our republic’s peril. We must face the problem squarely, with the full array of intelligence, wisdom, and humility at our command.

Some believe that this problem compels us to use only election systems that can effectively be supervised by members of the general public possessing ordinary intelligence, schooling, and common sense. Since

¹³¹ Selker at §5.1.

¹³² <http://caselaw.lp.findlaw.com/data/constitution/amendment14/>. §5 reads, “The Congress shall have power to enforce, by appropriate legislation, the provisions of this article.”

¹³³ <http://caselaw.lp.findlaw.com/data/constitution/amendment14/>.

¹³⁴ [George W. Bush, et al., v. Albert Gore, Jr., et al., case 00-949 \(12/12/2000\), http://supct.law.cornell.edu/supct/html/00-949.ZPC.html](http://www.supct.law.cornell.edu/supct/html/00-949.ZPC.html).

¹³⁵ For example, four of the five Supreme Court Justices who upheld the Voting Rights Act Amendments of 1970’s lowering of the voting age to 18 in national elections (including Presidential elections) relied upon section 5 of the 14th Amendment to do so. *Oregon v. Mitchell*, 400 U.S. 112, 112-16 (1970), <http://caselaw.lp.findlaw.com/scripts/getcase.pl?court=us&vol=400&invol=112>

¹³⁶ Eyton, Laurence, “Taiwan polls: Off the streets, into the courts,” *Asia Times Online*, 4/2/2004, <http://www.atimes.com/atimes/China/FD02Ad01.html>; “Q&A: Taiwan election dispute,” *BBC News*, 7/8/2004, <http://news.bbc.co.uk/1/hi/uk/3560355.stm>.

¹³⁷ See the end of §3.2 supra.

A Deeper Look: Rebutting Shamos on e-Voting

e-voting systems can effectively be supervised only by the tiny subgroup of citizens that is well-versed in computer security, some conclude that we should not use them. They instead recommend that we use hand-marked paper ballots and count them by hand, with citizen labor and under citizen supervision, in the precincts in which they are cast. The totals thus computed would be posted outside each precinct and reported to the media, thus committing each precinct to its citizen-computed and citizen-supervised totals and heading off central aggregation frauds.

Others, including the authors, believe that computer experts from the general public can provide adequate supervision if proper procedures are enacted and enforced and the systems themselves incorporate suitable security, reliability, and auditability measures, such as the use of voter-verified paper ballots, full public disclosure and review of all software and hardware, and proof of secure operation before deployment. The security, reliability, and auditability measures must include not only the in-precinct vote casting and tabulation, but also centralized tabulation systems at each county and statewide.

Finally, we must understand the full range of techniques that may be used to subvert voting systems—from vote casting to tabulation to reporting—and maintain vigilance for the appearance of new techniques, such as in irregularities in voter registration and purging. Whatever course we take, if we are willing vigorously to police electronic gambling systems, we should be willing to do the same—and more—for voting systems. Our votes—and our republic's future—are at stake.

7) Acknowledgements

We would like to thank Ron Crane and Alan Dechert to their contributions to this paper. We would also like to thank the anonymous reviewers for their suggestions on improving this paper.