Ethical Issues in E-Voting Security Analysis [*]

David G. Robinson[1] and J. Alex Halderman[2]

[1] Information Society Project, Yale Law School, `david.robinson@yale.edu`
[2] The University of Michigan, `jhalderm@eecs.umich.edu`

Abstract. Research about weaknesses in deployed electronic voting systems raises a variety of pressing ethical concerns. In addition to ethical issues common to vulnerability research, such as the potential harms and beneifts of vulnerability disclosure, electronic voting researchers face questions that flow from the unique and important role voting plays in modern democratic societies. Should researchers worry that their own work (not unlike the flaws they study) could sway an election outcome? When elected officials authorize a security review, how should researchers address the conflicted interests of these incumbent politicians, who may have powerful incentives to downplay problems, and might in principle be in a position to exploit knowledge about vulnerabilities when they stand for re-election? How should researchers address the risk that identifying specific flaws will lead to a false sense of security, after those particular problems have been resolved? This paper makes an early effort to address these and other questions with reference to experience from previous e-voting security reviews. We hope our provisional analysis will help practicing researchers anticipate and address ethical issues in future studies.

1 Introduction

Over the past seven years, computer security researchers have conducted more than a dozen significant studies of vulnerabilities in fielded electronic voting systems (e.g., [2–5, 7, 11, 12, 15, 18, 21–23]). Like many computer security studies, these projects have focused on identifying concrete technological problems and solutions. Yet voting occupies a special place in democratic public life—its integrity is of common concern to all citizens—and security analyses of voting systems can shape a democratic state's actual—and perceived—legitimacy. In this paper, we seek to identify, and describe, some of the ethical choices that are inevitably relevant to security analysis of e-voting systems.

We begin in Section 2 by considering high-level questions: whether researchers should perform such studies at all, and whether, in so doing, they should be concerned with the political consequences of their findings. In Section 3, we consider some of the quandaries that arise when obtaining access to voting systems through means such as leaks, anonymous sources, and direct government

---

[*] Appeared in *2nd Workshop on Ethics in Computer Security Research* (WECSR 2011), Castries, St. Lucia, March 4, 2011.

authorization of studies. In Section 4, we consider the potential for collateral damage during the process of studying real systems. In Section 5, we consider issues that arise after the research is complete, such as whether, when, and how to publicly disclose the findings. We conclude, in Section 6, that there is ample room for further inquiry into the ethical issues surrounding voting machine security research. We also suggest that the computer security community might achieve more, in the future, by becoming more involved in public policy debates at an earlier stage—before, rather than after, potentially vulnerable technologies have been adopted.

2 High-Level Questions

2.1 Whether to Perform Such Studies?

Researchers who aim to improve e-voting security must consider whether experimental evaluation of the security of deployed systems actually advances this goal. Many researchers believe that paperless electronic voting machines are inherently insecure, because they lack the transparency or verifiability necessary to prevent attacks by dishonest insiders. Those who accept this view might argue that empirically determining that a particular paperless system is insecure teaches us nothing.

Some in the field, such as Rebecca Mercuri [16], have argued that evaluations that point out specific security problems can actually make the general problem worse. These studies allow officials or vendors to correct some of the immediate problems, then claim that the systems have been tested and fully secured. Moreover, where an evaluation fails to find problems, such a negative result might be hailed by officials or vendors as confirmation of the system's security. Of course, while negative results are a favorable indicator of voting system security, this kind of analysis *cannot* definitively establish that a system is secure: adversaries could always be smarter, luckier, or better funded than testers, and find problems they did not.

Can researchers overcome these objections? One rationale for participating in the e-voting security evaluations is that demonstrating specific security problems may be more persuasive than arguing about abstract architectural weaknesses. Another is that if policymakers, having already heard the arguments about architectural weaknesses, still insist on using the machines, discovering vulnerabilities can provide new information with which to assess the machines' suitability. It may also allow the specific problems to be corrected before they can be maliciously exploited, although in some cases machines have been used in elections with documented vulnerabilities unpatched.

Empirical security evaluations help close the gap between theory and practice, by providing case studies in *how* security fails in practice, in addition to confirmation that it does. Security vulnerabilities remain dangerous even if a voting system provides a paper record of each vote and audits this record to detect fraud—even when they are detected, security or integrity breaches in real elections can still

compromise privacy or disrupt elections. By better understanding the kinds of vulnerabilities that arise in deployed systems and seeking their underlying causes, we can hope to strengthen future voting systems, both paperless and not.

2.2 Whether to Consider Near-Term Political Consequences?

The principal goal of electronic voting security research is to ensure high-integrity elections. But in the near term, disclosing findings also has the potential to distort the fortunes of political actors, and the course of political debate, in the places that use these systems. By the same token, the decision to remain silent about known problems may have important political results. Electronic voting vulnerabilities, once detected, may place their discoverers in an inherently political position. Disclosing problems may not only increase the chance they will be remedied (or, the risk that they may empower attackers), but also have immediate and profound effects on voter confidence and turnout.

Even when they are careful to avoid making any claims about the actual integrity of past elections, researchers who identify security concerns in an incumbent voting technology do give voters reason to doubt the legitimacy of that technology, and of the results it has produced. Recent studies of voting system integrity—because they have tended to find major flaws, rather than to offer support for the security or integrity of field-deployed systems—have tended to offer at least implicit or indirect reinfrocement for the electroal integrity concerns of losing candidates and their supporters.[3]

Given these factors, choices made by electronic voting researchers might at the margin change who wins and who loses an election. They could, for example, influence whether a U.S.-aligned political faction in another country does or does not prevail over its domestic rivals. Whatever one's normative views about war and peace, welfare policy, and all the other important choices made by elected officials, these secondary effects of the voting research could easily be the work's most important near-term impact.

It might be tempting for electronic voting researchers to attempt to anticipate, and tailor their actions around, these potential collateral impacts of their work. But we believe this would be a mistake. Political prediction is notoriously difficult even for its foremost practitioners, and effects that help or hurt political incumbents will recede in importance, over time, as different parties trade off in power. Unintended consequences could cut in any number of directions, so researchers could only speculate about what the second-order effects of their work will be.

More broadly, accurate democratic representation based on an honest count of votes is a worthwhile goal in its own right. Voting itself represents a kind

---

[3] See, *e.g.*, [19] (In a Democratic Senate primary, the losing candidate describes the "well-documented unreliability and unverifiability of the voting machines used in South Carolina."); [6]("At last week's hearing, [losing candidate] Rawl trotted out a parade of forensic, academic and computer experts who pointed to security, software and statistical irregularities.").

of epistemic modesty that denies in principle that any one political actor can know what is best for the system as a whole. To speculate about the second-order impact of increased democratic integrity—let alone basing one's actions on such impact—would itself be an anti-democratic choice. Where the release of reserach findings would create a significant risk of physical violence or other clear and concrete harm, researchers might reasonably decide to keep their results temporarily private. But such a choice should be the exception, rather than the rule, and we believe researchers should not pay condition their disclosures on the routine ebb and flow of electoral politics.

3 Obtaining Access

Researchers typically aim to provide a security evaluation that is independent of vendor and official influence. When access is limited, and cooperation with these parties enables otherwise impossible research, researchers must be vigilant to retain as much independence as is feasible—and transparent about the extent to which their end product is informed or shaped by other actors. If vulnerabilities are found, there is a further ethical question about disclosure: Is it ethical for researchers to bind themselves not to disclose such vulnerabilities to the public? On the other hand, how should researchers approach the ethical problems that can arise when their access comes through channels that are not officially approved?

In a typical e-voting security evaluation, researchers analyze a system, design specific attacks against it, and then attempt them in a demonstration or testing environment that mirrors the conditions under which system is actually used. This requires detailed technical information about how the system functions. In practice, researchers obtain this information by analyzing the system's source code or, where source code is not available, by reverse engineering voting machines. Obtaining the necessary access to voting machines or source code is one of the major prerequisite challenges of e-voting research, since vendors and system developers have historically been reluctant to support independent security reviews [17].

There are three main ways researchers have obtained such access: through leaks and anonymous sources (e.g., [11, 15, 22]), through government-sponsored studies (e.g., [4, 21]), and by purchasing government-surplus machines (e.g., [2, 7]).

3.1 Leaks and Anonymous Sources

Leaks and anonymous sources provided access for some of the earliest studies. In 2003, Kohno et al. [15] analyzed source code for components of the Diebold voting system software; this code had been posted to the company's public FTP site, where it was discovered and retrieved by e-voting activist Bev Harris [14]. In 2007, Feldman et al. [11] studied a Diebold AccuVote-TS paperless DRE machine after they were given unrestricted hands-on access to the machine by a nongovernmental source, who provided the machine on condition of anonymity. In 2010, Wolchok et al. [22] analyzed the electronic voting machines used in India

by studying a machine given to coauthor Hari Prasad by a government source under condition of anonymity.

Working with leaks and anonymous sources raises several concerns. One is legality: Is the source lawfully permitted to provide the machine? Do intellectual property protections preclude reverse engineering or working with obtained source code? Honoring promises of anonymity may create further risks; for instance, Indian researcher Hari Prasad spent over a week in a Mumbai jail and faces an ongoing legal battle to protect the identity of his source [20]. Researchers should consider whether legal risks may limit their ability to thoroughly evaluate system or disclose their findings.

Another concern is the source's motives. Researchers should questions whether sources that offer to provide or leak material have political motivations. We have argued that researchers have at most a limited duty to predict the secondary political effects of e-voting analyses, but they should be wary about the integrity and authenticity of the machines under study. Sources could hypothetically tamper with them to make them appear more vulnerable or plant evidence of past tampering, jeopardizing the integrity of the study's results. This is particular a concern when working with unknown sources or sources that request anonymity, since readers of the subsequent study will not be able to judge for themselves whether the source is trustworthy. In any case, it creates an extra duty of care for researchers, and it may necessitate clear disclaimers about the provenance of the study material.

3.2 State-Sponsored Studies

Studies sponsored by government entities raise another set of concerns. State-sponsored studies, such as the California secretary of state's Top-to-Bottom Review [21] and the Ohio secretary of state's Project EVEREST [4], provided researchers access to hardware and software for multiple e-voting systems. States' can often compel voting system vendors to provide source code access (for example, California threatened vendors with decertification if they did not), which simplifies the technical aspects of these studies and removes some kinds of legal risk for the researchers; however, cooperating with elected officials leads to other quandaries.

Working with government sources requires clear ground rules about how the study will be performed and how the results will be disclosed. These ground rules often take the form of a legal agreement between the researchers and public officials. Researchers need to ensure that these rules allow them to maintain their independence. If researchers are asked to sign a nondisclosure agreement, they should ensure that the terms allow them to disclose problems they might find, and do not overly restrict their ability to perform future work.

Researchers may also be asked to allow the government to review the findings prior to making them public, or to grant the government the ability to designate certain findings as confidential and prevent public disclosure. Disclosing vulnerabilities to officeholders who were elected (and may face reelection) using the same insecure technologies is deeply troubling, particularly if the vulnerabilities are

not fully disclosed to the public and if these officeholders have the authority to decide whether the election technology will continue to be used. Apart from the opportunity to exploit security flaws, public officials may have a strong incentive to downplay information that could cast doubt on the legitimacy of their own past or future elections. Researchers should ensure from the outset that there are clear rules that set appropriate conditions for disclosure, and that set a definitive deadline for all results to become public.

Official studies may require researchers to operate within constraints not applicable to real-world attackers. For example, researchers may be asked to operate within tighter time constraints, while real attackers have potentially unconstrained time to complete their attacks. Such constraints magnify the risk that the study may fail to uncover the full extent of problems. Where conditions are imposed, researchers must decide whether or not it is on balance worthwhile to proceed. In any event, researchers who agree to conduct limited analyses of voting systems should disclose these limitations in their reports, and should emphasize that their findings *cannot* establish that the systems under study are secure, since real-life attackers need not play by similar rules.

3.3   Government-Surplus Equipment

Government-surplus equipment has been obtained by researchers in a number of cases. When Buncombe County, North Carolina replaced its Sequoia AVC Advantage DREs in 2007, Princeton professor Andrew Appel purchased a lot of five machines for $82 [1]; these machines were the subjects of studies by Appel et al. [2] and Checkoway et al. [7]. In 2009, researcher Jeremy Epstein and colleagues purchased two Sequoia AVC Edge DREs for $100 after they were sold by Williamburg, Virginia after the state banned paperless DREs. Halderman and Feldman [13] performed a brief analysis of one of these devices and showed that they could easily alter its software (reprogramming it to play Pac-Man).

In many ways, government-surplus equipment raises fewer concerns that materials from other sources. Such machines often carry less legal encumbrance, and, depending on the chain of custody between government use and the researchers, they may raise fewer doubts about whether the machines under study are the same as the machines actually in use. However, other concerns can arise in later research phases when working with machines that have been used in real elections and may still contain real vote data.

4   Accidents During Analysis

Once researchers have obtained access to machines or source code, the process of security analysis consists of understanding the behavior of the system, identifying vulnerabilities, conceiving attacks, constructing attack demonstrations, and performing experiments to confirm that the attacks work. A number of ethical issues can arise during these efforts as a result of accidental access to data and to other systems.

4.1 Accessing Confidential Voter Information

One concern that may arise when analyzing voting equipment that has been used in real elections is that confidential voter information may remain present on the machines. Whether machines are provided by government or nongovernmental sources or purchased government-surplus, the sources may fail to completely sanitize the storage before turning the equipment over to researchers.

Machines obtained in the India voting study [22] and the AVC Advantage investigations [2, 11] (among several instances), contained vote data from the last elections in which they were used. In several cases the researchers discovered attacks that could deanonymize votes based on this data. Protecting the confidentiality of voters' ballots in instances like these requires researchers to take special precautions to prevent the data they recover from the machines from being publicly disclosed. Researchers may be ethically obligated to erase the data as soon as they discover it (especially if they cannot ensure its security), though this may be complicated by legal requirements for election data retention.

4.2 Risks of Collateral Damage

Other issues arise when testing *Internet* voting systems, such as in the recent public trial of a web-based voting system orchestrated by the Washington, D.C. Board of Elections and Ethics [9]. Researchers from the University of Michigan who participated in the trial [23] (including the second coauthor of this paper) encountered several unexpected ethical quandaries.

The D.C. election officials organized a mock election prior to the start of real voting. They claimed this system was disconnected from unrelated election facilities and promised not to take legal action against well-intentioned efforts to demonstrate security problems in the system. The researchers were able to penetrate the system and take control of the election server, changing votes and compromising ballot secrecy. They were also able to penetrate several other pieces of network infrastructure (routers, switches, and a terminal server) located on the subnet that election officials had initially designated for testing. The researchers noticed that officials were still in the process of configuring this equipment, but they continued their attack on the belief that the insecure components were being prepared for use in D.C.'s real voting system. After the public trial concluded, the researchers learned that these devices were in fact unrelated to the voting trial, and were being prepared for use elsewhere in the D.C. government network—they had discovered a critical security breach, but arguably one outside the intended scope of the testing. Researchers have ethical duty to limit potential for unintentional damage to unrelated equipment like this.

4.3 Risks of Unintentionally Disrupting Real Elections

The Michigan researchers made another unexpected discovery during the D.C. voting trial: they found that election officials, in preparing and testing the system, had uploaded to the test system the credentials that were to be used by real voters

following the trial. The researchers only discovered that these credentials were real after they had downloaded them over an insecure connection and transferred them to their own systems.

The researchers in this case had no intention of interfering with a real election, but, had D.C. officials not decided as a result of the trial to refrain from using the system in the real election, the researchers' access to these credentials would likely have constituted an unrecoverable security breach: Any party with access to the credentials would have been able to cast votes on behalf of the real voters, and the researchers had inadvertently exposed these voters to risks of being by malicious third parties. Issuing new credentials was impractical, since they had to be delivered to voters by postal mail, and it was too late to send a new batch. Researchers should weigh the risk of unintentionally disrupting real elections against the potential benefits of participation, and take steps to minimize such risk.

5   Disclosure

After the technical work of evaluation has been completed, researchers need to document their findings and decide what to disclose, to whom to disclose it, and when and how to disclose it. Some of the ethical considerations involved are common to other kinds of security vulnerability disclosures, and others are particular to e-voting security research.

5.1   What to Disclose?

In deciding what to disclose, security researchers must balance the need to convincingly convey the dangers they have found against the potential for making those problems worse by providing details that could aid real attackers. In e-voting research, this problem is complicated by the nature of the decision-making process involved. Researchers could choose to describe certain problems or details only to election officials and vendors, in an effort to limit the potential for misuse. However, election officials and voters have different incentives—officials suffer adverse results from the *appearance* of problems with the voting systems, whereas voters suffer from the existence of such problems, whether visible or not. Achieving greater security sometimes requires convincing voters that the system is vulnerable, which argues for wider disclosure.

Convining the public that security problems exist does not necessarily require full disclosure of the details of those problems. In practice, researchers often choose to illustrate the problems by creating demonstration attacks that they can perform for officials and journalists and convey to the broader public on video (e.g., [7, 11, 22]).

Researchers have rarely chosen to release full source code for these demonstration attacks. Instead, demonstration videos typically provide evidence that the problems exist without conveying all the technical details required to exploit them. While this practices makes its somewhat more difficult for malicious parties

to carry out the attacks, it cuts against the academic norm that research results should be shared in a form allowing reproduction, and it requires voters to take the researchers' word about the findings.

Advocates of full disclosure (and of "responsible disclosure", which gives vendors some time to rectify the problem before making all details public [8]) argue that these practices put added pressure on vendors to produce timely fixes, since they ensure that real attackers will have the information needed to exploit the problem. This argument seems less compelling in the e-voting context, where effectively securing systems like paperless DREs may require replacing them entirely. The governments that own these machines often lack the funds or political will to do so—and, as a result, systems with known vulnerabilities remain in widespread use. For example, Maryland continues to use the Diebold AccuVote-TS DRE that was discredited by researchers in 2005 [11]. This would probably still be the case even if the authors of that study had made attacks even easier by publishing their voting machine virus source code.

5.2   Disclosing Negative Results

If researchers examine a voting system in secret and are unable to discover a way to attack it, should they publicize this fact? On one hand it seems intellectually dishonest to suppress results like this, and it may lead other researchers to waste effort attempting the same thing. On the other hand, as we discussed earlier, this kind of negative says very little about the security of the system, and it may be misrepresented by others to argue that the system has been tested and found to be *fully* secure. We know of no instance where credible researchers have announced a negative e-voting test result.

5.3   When to Disclose?

Unlike most systems studied in security research, election systems are generally used only a few times a year, during elections that are scheduled long in advance. This schedule significantly impacts decisions about when to disclose vulnerabilities. Revealing problems so soon before an election that there is not time to implement any effective remedies would create risks without significant countervailing benefits. On the other hand, if researchers know about problems and there is sufficient time to mitigate them, they may have an obligation to publicly disclose them. Balancing these factors requires, in part, reasoning about what remedies can be practiacally achieved in time.

Researchers might consider giving election officials or voting system vendors advance notice about their findings prior to public disclosure, to allow them to begin implementing mitigations. Though sometimes beneficial, this approach is problemmatic. Researchers who studied systems without authorization may run the risk of political retribution or lawsuits attempting to suppress publication of their results. The mitigations that are implemented may be weaker in the absense of public pressure from voters. The risk of insider attacks, one of the most important categories of threats against voting systems, is certainly not reduced

by disclosing new attacks only to insiders. For these reasons, researchers often choose not to disclose problems to officials and vendors in advance of publication.

5.4 Attacking Real Elections

One course of action that is clearly unethical is for researchers to exploit vulnerabilities they have discovered to attack real elections. People outside the research community sometimes suggest that researchers should change an election outcome to an obviously incorrect result in order to demonstrate conclusively that the system is vulnerable in practice. Not would not only criminal, but also a subversion of the democratic process that this body of research serves.

6 Conclusions

We have tried to articulate the scope of ethical concern for electronic voting security researchers, and to describe some of the issues that arise within that scope. Our map of this ethical terrain is far from perfect, but we hope it can be useful—both to researchers facing ethical quandaries, and to the lay public as it considers the value and impact of security research into electronic voting.

This paper explores ethical choices that actually confront today's researchers. Arguably, however, the most important ethical lesson of the electronic voting experience is about what might have been. The troubled modern history of electronic voting owes a great deal to the 2002 passage of the Help America Vote Act, [10] which gave states time-limited funds to purchase computerized voting equipment without setting meaningful standards for its security. Policymakers assumed, or allowed themselves to be persuaded, that widely sold paperless electronic voting machines were as secure as their manufacturers claimed. The vulnerabilities that have since been found may surprise Congress and the public, but they are much less surprising to experts in the field. HAVA's deep flaws reflect our research community's failure to intervene effectively in the public policy debate. In the future, as legislatures consider computerized approaches to emerging challenges in healthcare, defense, and other areas, computer security researchers should do all they can to get out ahead of possible security problems, and to dissuade policymakers from indulging in the kind of wishful thinking that generated the electronic voting morass of the last eight years.

References

1. Andrew W. Appel. How I Bought Used Voting Machines on the Internet. Feb 7, 2007. http://www.cs.princeton.edu/~appel/avc/.
2. Andrew Appel, Maia Ginsburg, Hari Hursti, Brian W. Kernighan, Christopher D. Richards, Gang Tan, and Penny Venetis. The New Jersey Voting-Machine Lawsuit and the AVC Advantage DRE Voting Machine. *Proc. Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE),* 2009.

3. Adam Aviv, Pavol Cerný, Sandy Clark, Eric Cronin, Gaurav Shah, Micah Sherr, and Matt Blaze. Security Evaluation of ES&S Voting Machines and Election Management System. *Proc. USENIX/ACCURATE Electronic Voting Technology Workshop (EVT),* 2008.

4. Jennifer Brunner, et al. Ohio Secretary of State's Evaluation & Validation of Election-Related Equipment, Standards & Testing (EVEREST). December 2007.

5. Kevin Butler, William Enck, Harri Hursti, Stephen McLaughlin, Patrick Traynor, and Patrick McDaniel. Systemic Issues in the Hart InterCivic and Premier Voting Systems: Reflections on Project EVEREST. *Proc. USENIX/ACCURATE Electronic Voting Technology Workshop (EVT),* 2008.

6. Eliza Newlin Carney. Voting Without a Net in South Carolina. *National Journal,* June 21, 2010. http://www.nationaljournal.com/njonline/rg_20100621_7815.php.

7. Stephen Checkoway, Ariel J. Feldman, Brian Kantor, J. Alex Halderman, Edward W. Felten, and Hovav Shacham. Can DREs Provide Long-Lasting Security? The Case of Return-Oriented Programming and the AVC Advantage. *Proc. USENIX/ ACCURATE Electronic Voting Technology Workshop (EVT),* David Jefferson, Joseph Lorenzo Hall, and Tal Moran, eds., August 2009.

8. Thomas Claburn. Google Seeks Redefinition of 'Responsible Disclosure.' *InformationWeek,* July 2010. http://www.informationweek.com/news/smb/security/showArticle.jhtml?articleID=226100117.

9. Jeremy Epstein, et al. In D.C's Web Voting Test, the Hackers Were the Good Guys. *Washington Post,* Oct 2010. http://voices.washingtonpost.com/local-opinions/2010/10/in_dcs_web_voting_test_the_hac.html.

10. Brandon Fail. HAVA's Unintended Consequences: A Lesson for Next Time. *Yale Law Journal,* 116, 2006. http://www.yalelawjournal.org/pdf/116-2/Fail.pdf.

11. Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten. Security Analysis of the Diebold AccuVote-TS Voting Machine. *Proc. USENIX/ACCURATE Electronic Voting Technology Workshop (EVT),* Boston, MA, Aug 2007.

12. Rop Gonggrijp and Willem-Jan Hengeveld. Studying the Nedap/Groenendaal ES3B Voting Computer: A Computer Security Perspective. *Proc. USENIX/ACCURATE Electronic Voting Technology Workshop (EVT),* 2007.

13. J. Alex Halderman and Ariel J. Feldman. Pac-Man on the Sequoia AVC-Edge DRE Voting Machine. Aug 2010. http://www.cse.umich.edu/~jhalderm/pacman/.

14. Bev Harris. System Integrity Flaw Discovered At Diebold Elections System. *Scoop,* Feb 10, 2003. http://www.scoop.co.nz/stories/HL0302/S00052.htm.

15. Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, and Dan S. Wallach. Analysis of an Electronic Voting System. *Proc. IEEE Symposium on Security and Privacy,* Oakland, CA, pages 27–40, May 2004.

16. Rebecca Mercuri. Trust the Vote? Not in DC! *OpEdNews,* Nov 8, 2010. http://www.opednews.com/articles/Trust-the-vote-Not-in-DC-by-Rebecca-Mercuri-101108-990.html.

17. Ryan Paul. E-vothing Bendor Blocks Security Audit with Legal Threats. *ars technica,* 2008. http://arstechnica.com/tech-policy/news/2008/03/e-voting-blocks-e-voting-security-audit-with-legal-threat.ars.

18. Elliot Proebstel, Sean Riddle, Francis Hsu, Justin Cummins, Freddie Oakley, Tom Stanionis, and Matt Bishop. An Analysis of the Hart Intercivic DAU eSlate. *Proc. USENIX/ACCURATE Electronic Voting Technology Workshop (EVT),* 2007.

19. Vic Rawl for U.S. Senate. Statement of Judge Vic Rawl. June 14, 2010. http://www.vicrawl.com/vicrawl/post/1023-statement-of-judge-vic-rawl.

12

20. Jim Tyre. 2010 Pioneer Award Winner Hari Prasad Defends India's Democracy. *EFF Deeplinks Blog,* Nov 1, 2010. https://www.eff.org/deeplinks/2010/11/2010-pioneer-award-winner-hari-prasad-defends.

21. David A. Wagner, et al. California Secretary of State's Top-to-Bottom Review (TTBR) of Electronic Voting Systems. July 2007.

22. Scott Wolchok, Eric Wustrow, J. Alex Halderman, Hari K. Prasad, Arun Kankipati, Sai Krishna Sakhamuri, Vasavya Yagati, and Rop Gonggrijp. Security Analysis of India's Electronic Voting Machines. *Proc. 17th ACM Conference on Computer and Communications Security (CCS),* Chicago, IL, Oct 2010.

23. Eric Wustrow, Scott Wolchok, Dawn Isabel, and J. Alex Halderman. Security Analysis of the Washington, D.C. Internet Voting System. In preparation, 2010.