

## ON THE SECURITY OF BALLOT MARKING DEVICES

DAN S. WALLACH\*

### CONTENTS

I. INTRODUCTION .....	559
II. WHAT, EXACTLY, IS A BMD? .....	563
III. THREAT AND AUDIT MODELS .....	565
IV. LIVE AUDITING OF BMDS .....	570
V. NON-UNIFORM MALICIOUS BEHAVIOR .....	573
VI. REALIZING THE AUDITING SCHEME.....	575
VII. ADDITIONAL ARGUMENTS.....	582
VIII. DISCUSSION AND CONCLUSIONS .....	585

---

\* Department of Computer Science, Rice University. I would like to thank my colleagues for reviewing drafts of this paper and offering a variety of constructive suggestions: Claudia Acemyan, Ben Adida, Mike Byrne, Joseph Hall, Philip Kortum, and Whitney Quesenbery. I would also like to thank the editors of the Ohio State Technology Law Journal for their detailed work to improve this manuscript.

## I. Introduction

Every voting system must protect against a variety of security threats. It is the essential purpose of any voting system to provide evidence that its stated outcomes are correct, even in the face of adversaries who may wish to tamper with it. Every voting system must also provide usability and accessibility features, because errors in human voters' operation of the voting system can lead to changes in the outcome, particularly if the margin of victory is smaller than the margin of human error.

In the early 2000's, paperless electronic voting systems gained prominence for their ability to offer important accessibility features (e.g., optionally large text, button boxes, multiple languages, headphones), but these systems also created unacceptable security vulnerabilities.<sup>1</sup> Tampered or even buggy software could corrupt or destroy all evidence of voters' original intent.<sup>2</sup>

Electronic ballot marking devices (BMDs) would seem to bridge the gap between the fundamental security properties of paper, which cannot be overwritten or tampered by any computer and thus create the potential for elections to be software independent, and the variety of usability features available with computers, which cannot be provided in an equivalent manner by paper-and-pen. BMDs thus have the potential to provide the best of both worlds.

Recently, Andrew Appel, Richard DeMillo, and Phillip Stark (hereafter, "ADS") staked out some important security claims, arguing against BMDs.<sup>3</sup> They argue that voters can easily be fooled and will

---

<sup>1</sup> Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, & Dan S. Wallach, *Analysis of an Electronic Voting System*, in PROC. OF IEEE SYMPOSIUM ON SECURITY & PRIVACY (2004).

<sup>2</sup> *Id.*

<sup>3</sup> Andrew W. Appel, Richard DeMillo, & Philip B. Stark, *Ballot-Marking Devices (BMDs) Cannot Assure the Will of the Voters* (April 12, 2019) (unpublished manuscript), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3375755](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3375755).

neither notice deliberate errors, nor even if they do notice will they have any meaningful proof of the BMD's misbehavior.<sup>4</sup>

We need to consider exactly how often a voter might notice an error, what common electoral processes will do next, and how they might be enhanced. We'll also need to discuss the properties of hand-marked paper ballots, considered by ADS and many others to be the "gold standard" for election security.

#### a. Why not just mark ballots by hand?

A central question, posited by many election integrity activists, is why we don't just stick with hand-marked paper ballots. This question is important to address directly.

- Not every voter has the ability to do all the tasks necessary to read, mark, and cast a paper ballot. Some voters have low vision or zero vision. Some voters have limited motor control. Some voters are illiterate or dyslexic. Some voters have multiple such issues. BMDs have the potential to make voting far more accessible to these populations. BMDs can also offer a variety of different languages, both in text and voice, offering greater assistance to non-native English speakers. Furthermore, Federal and state laws generally make these features mandatory.
- Ballot marking devices also have the advantage of eliminating complete classes of voting errors that can occur with hand-marked paper ballots. For example, with a BMD it is impossible to "overvote";<sup>5</sup> the BMD can enforce common rules like "only one vote per contest."<sup>6</sup> Enforcing such rules is even more important with voting methods that allow multiple

---

<sup>4</sup> *Id.*

<sup>5</sup> *See, e.g.*, UNITED STATES ELECTION ASSISTANCE COMMISSION, 2015 VOLUNTARY VOTING SYSTEM GUIDELINES 45-46 (2015) [hereinafter VVSG], [https://www.eac.gov/sites/default/files/eac\\_assets/1/28/VVSG.1.1.VOL.1.FINAL1.pdf](https://www.eac.gov/sites/default/files/eac_assets/1/28/VVSG.1.1.VOL.1.FINAL1.pdf) (using the term "EBM" (electronically-assisted ballot marker) where we use "BMD").

<sup>6</sup> *Id.*

selections in a contest, such as rank-choice or instant-runoff voting. BMDs additionally do not allow voters to make stray or ambiguous marks, since everything on the paper was printed by the computer. If a voter needs to change their mind after the ballot is printed, they can “spoil” it and start over again.<sup>7</sup> For contrast, consider the 2008 recount of the very close Minnesota Senate race between Coleman and Franken.<sup>8</sup> Ambiguous hand-marked ballots were individually considered in litigation after the election.<sup>9</sup>

- A well-designed BMD can also help every voter accurately convey their intent. For example, a BMD will commonly have a confirmation screen at the end of the process that can highlight contests that a voter might have accidentally skipped.<sup>10</sup> Features like this become even more important as ballots grow longer and more complicated. Likewise, a BMD does not face the space constraints of a hand-marked paper ballot, allowing each question to appear on a separate screen, and thus help prevent voters from accidentally skipping over a contest. For contrast, consider the paper ballot in Broward County, Florida in 2018, where the contests for U.S. Senate and Congressional Representative were placed under the long ballot instructions in the left column, leading a potentially significant number of voters to miss them entirely<sup>11</sup>.

---

<sup>7</sup> Procedures for voters to spoil their ballots vary from state to state. In Texas, for example, voters are entitled to no more than three attempts at filling out a ballot. *See* Tex. Elec. Code Ann. § 64.007(b) (West 1986).

<sup>8</sup> D’Angelo Gore, *Mining the Minnesota Recount*, FACTCHECK.ORG (Jan. 28, 2009), <https://www.factcheck.org/2009/01/mining-the-minnesota-recount/> [<https://perma.cc/NN2U-4GGM>].

<sup>9</sup> *Id.*

<sup>10</sup> VVSG, *supra* note 5, at 46-47.

<sup>11</sup> Appel has written a summary of this issue in Florida 2018, alongside other famous ballot layout failures. Andrew Appel, *Florida is the Florida of Ballot-Design Mistakes*, FREEDOM TO TINKER (Nov. 14, 2018), <https://freedom-to-tinker.com/2018/11/14/florida-is-the-florida-of-ballot-design-mistakes/> [<https://perma.cc/6E7G-ND64>]. Bad ballot layout can induce high undervote rates in any voting technology, but at least BMDs can operate without the constraint of compressing a lengthy ballot to fit onto a sheet of paper.

- An important nationwide trend is the consolidation of polling places, both for early voting and on election days. Such “vote centers” allow any voter to cast any one of potentially thousands of distinct ballot styles. To run a vote center with hand-marked paper ballots, this requires having laser printers for “ballot on demand” printing.<sup>12</sup> Unfortunately, laser printers have their own issues. Most notably, they require a significant power draw (i.e., several kilowatts) to warm up the toner drum, which might be unavailable in some locations. Similarly, consumer-grade UPS (uninterrupted power supply) devices may have inadequate power for a laser printer.<sup>13</sup> If the power goes out, and a UPS isn’t sized to handle the laser printer, the election is dead in the water. Conversely, BMDs generally use thermal printers, which are low power and have no consumables like ink or toner cartridges<sup>14</sup>. Commercial BMDs have (or should have) enough battery to run for hours without power. BMD-based elections will be more robust in the face of power failures.

Consequently, a fundamental challenge we face in any BMD implementation is trying to combine the security properties of hand-marked paper ballots with the usability and operational benefits of a BMD. In Section II, we try to define BMDs. In Section III, we consider how we might model the threat and what we might hope to accomplish with an audit. In Section IV, we describe a “live auditing”

---

<sup>12</sup> See, e.g., *Ballot on Demand Solution*, ELECTION SYS. & SOFTWARE, <https://www.essvote.com/products/ballot-on-demand/> [<https://perma.cc/VSG8-TSEQ>].

<sup>13</sup> See, e.g., Tony, *APC Battery Backup for Laser Printer*, EXCESSUPS (Feb. 6, 2012), <https://excessups.com/blog/guides/apc-battery-backup-to-protect-laser-printer/> [<https://perma.cc/C88B-NSME>].

<sup>14</sup> Direct thermal printers are commonly used to print restaurant and other “point of sale” receipts. They typically use an array of fixed heating elements. As the paper moves by, if it touches a hot element, the paper turns black. The paper, itself, has a temperature-sensitive coating to make this work. Consequently, there are no consumable inks or toners. See, e.g., *Difference between Direct Thermal and Thermal Transfer Printing*, ONLINE LABELS (June 2019), <https://www.onlinelabels.com/articles/direct-thermal-vs-thermal-transfer.htm> [<https://perma.cc/V4PX-54RX>].

process and analyze how likely it might be to detect misbehaving BMDs. In Section V, we consider attackers' ability to hide malicious behavior and the ability for an election official to detect it, such as by observing anomalously high rates of spoiled ballots. In Section VI, we consider how an election official might be able to tactically improve their chances of catching malware. In Section VII, we consider a variety of other arguments that have been made, including the question of why BMDs might or might not be preferable to hand-marked paper ballots. We conclude in Section VIII.

## II. What, exactly, is a BMD?

Fundamentally, a BMD is a device that knows about all the different ballot styles that a voter might see. By inserting an unfilled ballot, or perhaps a blank sheet of paper with only a barcode indicating the specific ballot style, the BMD can then present a touch-screen interface to the voter to select their choices.<sup>15</sup> BMDs typically include a variety of accessibility features (button boxes, headphones with audio output, font size settings, and other features as a supplement to the touch-screen), allowing a larger number of voters to operate these devices without assistance.<sup>16</sup> When the voter is finished, the BMD does what its name says: it prints a marked ballot.<sup>17</sup>

After that, BMDs come in two varieties: stateful and stateless.<sup>18</sup> The former retains an electronic copy of every ballot, while the latter promptly forgets what it saw and starts over again. Stateful BMDs might allow for faster tallies, and provide redundancy against catastrophic failures (e.g., lost ballot boxes<sup>19</sup>). Stateless BMDs might

---

<sup>15</sup> VVSG, *supra* note 5, at 34-35.

<sup>16</sup> *Id.* at 63-76.

<sup>17</sup> *Id.* at 46-47.

<sup>18</sup> *Id.* at 144-151. What we refer to as a “stateful BMD” would be classified as a “VVPAT” (voter-verified paper audit trail system) by the VVSG standards, subjecting the system to additional engineering requirements.

<sup>19</sup> See, e.g., Stacy Fernández, *A Ballot Box Found Weeks After Election Day has Flipped a \$569 Million School Bond Vote in Midland*, TEX. TRIBUNE (Jan. 17, 2020, 6:00 PM), <https://www.texastribune.org/2020/01/17/lost-midland-ballot-box-flips-569-million-school-bond-election/> [<https://perma.cc/J5N2-9PPG>].

be simpler to construct and provide stronger guarantees against the impact of malware within the machine.

What happens with the paper ballots after they are printed varies from vendor to vendor. Typically, the voter will carry it by hand to a ballot box, which then has a scanner on top.<sup>20</sup> For hand-marked ballots, these scanners can flag common error modes, including when a voter has indicated more than one vote in a contest which allows at most one vote (i.e., “overvoting,” which can never happen with a BMD, but is possible with hand-marked paper).<sup>21</sup> Some vendors offer “privacy sleeves” to allow poll workers to do this operation on behalf of voters who do not have the necessary manual dexterity, while still preserving the voter’s privacy.<sup>22</sup> One system, Los Angeles County’s VSAP, has the printer and ballot box integrated together, so that the entire process can be completed independently and privately.<sup>23</sup>

Integrating the printing and casting would seem to have desirable usability properties, particularly for voters with limited manual dexterity. However, such BMDs create additional security concerns, where a BMD might print something contrary to the voter’s stated desires and automatically cast it without any opportunity for the voter to intervene.<sup>24</sup>

The rest of this paper will focus primarily on stateless BMDs, wherein the paper ballot is the only way to know the voter’s intent, and where vote casting is a manual process, where the voter moves the ballot to a

---

<sup>20</sup> *VVSG Vol 1, Version 1.1*, section 3.2.2.2, pp. 48-49. These scanners are commonly referred to as “precinct count optical scanners” (PCOS).

<sup>21</sup> *Id.* at 45.

<sup>22</sup> *Id.* at 49-50.

<sup>23</sup> *Review L.A. County’s 2020 Election Administration Plan*, VSAP, <https://vsap.lavote.net/> [<https://perma.cc/B5M5-EQPQ>].

<sup>24</sup> Andrew W. Appel, *Serious Design Flaw in ESS ExpressVote Touchscreen: “Permission to Cheat”*, FREEDOM TO TINKER (Sept. 14, 2018), <https://freedom-to-tinker.com/2018/09/14/serious-design-flaw-in-ess-expressvote-touchscreen-permission-to-cheat/> [<https://perma.cc/U7LN-Q7DX>].

physically distinct ballot box.<sup>25</sup> This simplifies the discussion, and makes it clear exactly what “the ballot” actually is. In particular, this makes it clear what happens during a recount, where “recounting the ballots” means looking at paper ballots, not electronic records. This also simplifies our discussion of what it means to “cast” a ballot.

### III. Threat and audit models

Every BMD is just a computer. Like any computer, it might have bugs in its software that don’t turn up in testing and might then impact the voter’s experience. Also, like any computer, its software could include malware, not intended by the manufacturer or election official to be present, but perhaps surreptitiously inserted when nobody was looking. Plenty of opportunities for this exist in modern elections, where voting machines may be delivered days or weeks in advance of an election.<sup>26</sup>

In the mid-2000’s, researchers were concerned with how this sort of attack might play out with paperless electronic voting systems (which typically went by the unwieldy acronym “direct recording electronic”—DRE), since a voting machine might appear to be operating correctly, displaying exactly what the voter intended, but secretly record the ballot internally in a very different fashion.<sup>27</sup> A related issue is that a paperless electronic system can also retain the ballots in the order cast, or randomize them in a reversible fashion, allowing ballot secrecy to be compromised by anybody who observes the order in which voters arrive at the polls.<sup>28</sup>

---

<sup>25</sup> VVSG, *supra* note 5, at 48-49. These scanners are commonly referred to as “precinct-count optional scanners” (PCOS).

<sup>26</sup> This is colloquially referred to as the “sleepover voting machine problem”. A variety of physical security protocols have been deployed to mitigate these threats, but this is beyond the scope of this paper. *See, e.g.*, Anahad, O’Connor, *Mom, Can My Voting Machine Spend the Night?*, N.Y. TIMES: THE LEDE BLOG (Aug. 19, 2008, 1:59 PM) <https://thelede.blogs.nytimes.com/2008/08/19/mom-can-my-voting-machine-spend-the-night/> [<https://perma.cc/6MMG-M7U7>],

<sup>27</sup> Kohno, Stubblefield, Rubin, & Wallach, *supra* note 1, at 2 (2004).

<sup>28</sup> *Id.* at 16.



The mitigations that were used against these attacks, at the time, were not particularly impressive. Logic and accuracy testing (commonly shortened to “L&A”), conducted prior to the start of the election, would run a small and predetermined set of tests votes through the machine, verifying that the proper tally appeared at the end.<sup>29</sup> Of course, if the machines watched their internal clocks, they could behave correctly while under test and then be malicious only on Election Day. Similarly, the number of votes used in L&A is typically much smaller than will appear in a real election, providing additional opportunities for a voting machine to distinguish between test conditions and a real election, and thus behave properly during L&A.<sup>30</sup>

A more aggressive mitigation, only used in a handful of jurisdictions, was to conduct a parallel test.<sup>31</sup> With this, some fraction of the voting machine population is randomly selected and then, rather than being deployed to the field, is instead set up in the elections warehouse where an operator enters a full day’s worth of votes according to a script.<sup>32</sup> As with L&A, the post-election tally from these machines under test has a known correct outcome and any deviation from this would indicate a serious problem.<sup>33</sup> Because no real votes are being cast, video can also be captured to ensure that operator data-entry errors can be differentiated from malicious vote flips.<sup>34</sup>

A thought experiment on how to defeat this, which can possibly be attributed to Avi Rubin, is to have a secret knock.<sup>35</sup> This is an input that no rational testing process would ever contain, but which malware

---

<sup>29</sup> See, e.g., CENT. FOR ELECTION SYS., LOGIC AND ACCURACY TESTING (2015 ed. 2010),

[https://www.eac.gov/sites/default/files/eac\\_assets/1/28/Logic\\_and\\_Accuracy\\_Testing\\_Manual\\_Final\\_v1.4.pdf](https://www.eac.gov/sites/default/files/eac_assets/1/28/Logic_and_Accuracy_Testing_Manual_Final_v1.4.pdf).

<sup>30</sup> Ariel J. Feldman, Alex J. Halderman, & Edward Felten, *Security Analysis of the Diebold AccuVote-TS Voting Machine*, CTR. INFOR. TECH. POL’Y (2006).

<sup>31</sup> BRENNAN CTR. FOR JUSTICE, THE MACHINERY OF DEMOCRACY: PROTECTING ELECTIONS IN AN ELECTRONIC WORLD 18 (2006), <https://www.brennancenter.org/sites/default/files/publications/Machinery%20of%20Democracy.pdf>.

<sup>32</sup> *Id.* at 72.

<sup>33</sup> *Id.* at 122

<sup>34</sup> *Id.* at 52.

<sup>35</sup> See AVI RUBIN, BRAVE NEW BALLOT: THE BATTLE TO SAFEGUARD DEMOCRACY IN THE AGE OF ELECTRONIC VOTING (Morgan Road Books, New York, 2006).

would look for.<sup>36</sup> The canonical example is a write-in vote for Mickey Mouse, although that might well happen in practice, so an attacker would need to select something more obscure. The general idea, then, is that the malware will act identically to the legitimate software until it sees the secret knock, and only then start misbehaving.<sup>37</sup> This would be effectively undetectable without potentially destructive forensic testing, although it requires co-conspirators to perform the secret knock during the real election, creating significant risks for the conspiracy.

**a. Would all these attacks and defenses still work in a BMD?**

This is the crux of ADS's argument. Malicious software in a BMD can certainly show one thing on the screen and print something else on the paper. This is particularly troublesome with some vendors who print two different encodings of the ballot: barcodes for machine-readable data and printed text for humans. While voters can verify the printed text, no voter will be able to detect errors in the barcodes. Let's call this an "inconsistent barcode" attack. Alternatively, the machine might produce a completely consistent paper ballot (i.e., the barcode and the human-readable text are in total agreement), but the paper ballot differs from what the voter entered on the touchscreen. Let's call that a "switched-intent" attack.

There's a very simple solution to the inconsistent barcode attack: get rid of the barcodes and have the only record of the voter's intent be human-readable text. Any computer-printed text that's readable by a voter will also be readable by a computer scanner with exceptionally high accuracy.<sup>38</sup> While many current BMDs do not operate this way, this can be addressed through regulatory mandates and software updates from the vendors.

---

<sup>36</sup> *Id.*

<sup>37</sup> *Id.*

<sup>38</sup> To achieve the necessary accuracy, the OCR software likely needs to use state-of-the-art "deep learning" techniques. Until such an OCR technology can be proven out in practice, barcodes will be necessary as a bridging technology.

### **b. Auditing and inconsistent ballot detection**

So long as paper ballots have both human-readable text and a barcode, we need to consider how existing audit processes can be adapted to detect these attacks. A number of audit processes, including recounts and risk limiting audits (RLAs), provide these opportunities.

If even a single ballot is inconsistent, that's evidence of a serious problem—either a major software bug or an inconsistent-barcode security attack—that would require emergency procedures.<sup>39</sup> Because of this, inconsistent barcode attacks are unlikely to be mounted by attackers in the real world, because even a single inconsistent ballot represents incontrovertible evidence that something has gone wrong. Attackers who wish to quietly manipulate an election outcome would not want to leave this kind of evidence so easily available for discovery.

What about a switched intent attack then? This seems preferable for an attacker, since it's not immediately obvious when it occurs. How can we discover one? ADS base their argument against such discoveries on three factors: that voters are unlikely to notice these attacks, that even if a voter does discover such an attack there is no good process to respond to such discoveries, and that there's no alternative process in place that might reliably detect the attacks.

### **c. Will voters notice a switched intent attack?**

A number of studies were conducted at Rice,<sup>40</sup> where researchers created a paperless voting machine that deliberately lied on its

---

<sup>39</sup> See *infra* Section VI.A for discussion of such emergency procedures.

<sup>40</sup> See Bryan A. Campbell & Michael D. Byrne, Now do Voters Notice Review Screen Anomalies? A Look at Voting System Usability, Presented at Electronic Voting Technology/Workshop on Trustworthy Elections 2009, (Aug. 10 2009), [http://chil.rice.edu/research/pdf/CampbellByrne\\_EVT\\_\(2009\).pdf](http://chil.rice.edu/research/pdf/CampbellByrne_EVT_(2009).pdf); see also Sarah P. Everett, The Usability of Electronic Voting Machines and How Votes can be Changed Without Detection (May 2007) (unpublished Ph.D dissertation, Rice University) (on file with Rice University).

summary screen. Their goal was to detect how many research participants, drawn from the local population, would notice that the machine changed their inputs and would then go back and fix their “mistake.”<sup>41</sup> Depending on exactly how they set up the experiment, between 1/3 and 1/2 of the voters noticed the introduced errors.<sup>42</sup>

In most of these experiments, the participants were given made-up names on a printed sheet and asked to vote for them.<sup>43</sup> It’s entirely possible that with real candidate names, in a real election, particularly at the top of the ticket where name recognition will be higher, voters might be more likely to notice discrepancies. This suggests that, if there were systemic vote flipping malware, something that tried to move thousands or tens-of-thousands of votes, we would have large numbers of regular voters who recognize the error when it happens.

A more recent manuscript by Richard DeMillo et al. describes two studies.<sup>44</sup> The first considers timing data from observed live voters.<sup>45</sup> The second presents results from an exit survey of voters, presented with a blank ballot and asked if it was equal to the actual ballot they voted.<sup>46</sup> Neither of these were controlled studies, so their observations are unreliable for predicting verification rates.

What might happen if a voter notices an error on a printed BMD? Most voters will likely head back to the poll workers’ table, perhaps sheepishly admit to having made a mistake, and request a chance to repeat the process with a fresh ballot. This process, commonly called ballot spoiling, is a completely standard part of any election process involving paper ballots. In Texas, for example, a voter is entitled to three attempts.<sup>47</sup>

---

<sup>41</sup> Campbell & Byrne, *supra* note 40, at 1; Everett, *supra* note 40 at 4-5.

<sup>42</sup> Campbell & Byrne, *supra* note 40 at 1; Everett, *supra* note 40 at 4.

<sup>43</sup> Campbell & Byrne, *supra* note 40 at 5; Everett, *supra* note 40 at 42.

<sup>44</sup> Richard DeMillo et al., What Voters are Asked to Verify Affects Ballot Verification: A Quantitative Analysis of Voters’ Memories of Their Ballots 1 (April 13, 2019) (unpublished manuscript), <https://ssrn.com/abstract=3292208>.

<sup>45</sup> *Id.* at 7.

<sup>46</sup> *Id.* at 8.

<sup>47</sup> Tex. Elec. Code Ann. § 64.007 (West 1987).

We can expect there to be a certain background rate of spoiled ballots, no matter the accuracy of the ballot marking devices, so it's only when the spoilage rate gets reliably above the background rate that we'll have a useful signal. Clearly, poll workers need to track every time a voter spoils a ballot and election administrators need this data available as the election is ongoing, giving them real-time situational awareness of problems as they manifest.

Given all this information, what should an election official do when the spoiled ballot rate is higher than expected? Preferably, they would have a variety of different responses available, from deploying additional auditors to more serious emergency procedures.<sup>48</sup>

#### **IV. Live auditing of BMDs**

Election officials need procedures for conducting audits on BMDs, in the field, while the election is ongoing. Because BMDs retain no internal memory of cast votes, the only hard requirement for conducting any sort of live audit is that any ballots printed during the audit must be kept out of the ballot box, which is basically the same process used for a spoiled ballot.<sup>49</sup> Such a process will be naturally transparent to voters or election observers, who would be free to witness the process.

Who should conduct the audits? Audits might be conducted by poll workers, as part of their regular duties, or they might be conducted by dedicated auditors, working for the election administration, driving from one polling location to another during the election period. The essential attributes of a good auditing process are that "enough" tests are conducted to observe rare events, and that these tests are sufficiently random that a malicious BMD has no way to reliably determine whether it is operating with a real voter or with an auditor.

---

<sup>48</sup> See *infra* Section V.D for discussion of the actual likelihood of such detection.

<sup>49</sup> Tex. Elec. Code Ann. § 64.007. Similar rules should apply in every state.

If a BMD is going to misbehave, the auditor will have a chance to catch it. And if any auditor, anywhere in the county, catches even one malicious machine in the act, the game is over. Call the police; we've got evidence of a serious crime.<sup>50</sup>

This idea of live auditing has been around since at least Benaloh's challenge mechanism,<sup>51</sup> and was quickly adopted by research voting systems like Helios<sup>52</sup> and VoteBox.<sup>53</sup> Even without cryptography, the concept is the same. We wish to test a machine to prove that it is generating correct output. The machine does not know that it is being tested. The machine must commit to its output, and then we can verify the correctness of that output, or alternatively arrive at concrete proof of the machine's misbehavior.

---

<sup>50</sup> See *infra* Section VI.A for a discussion of emergency procedures.

<sup>51</sup> Josh Benaloh, Simple Verifiable Elections, presented at the USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '06) (June 14, 2006), [https://www.usenix.org/legacy/event/evt06/tech/full\\_papers/benaloh/benaloh.pdf](https://www.usenix.org/legacy/event/evt06/tech/full_papers/benaloh/benaloh.pdf); Josh Benaloh, Ballot Casting Assurance via Voter-Initiated Poll Station Auditing, presented at the 2nd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)(Jun. 28 2007), <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.446.6091&rep=rep1&type=pdf>.

<sup>52</sup> See Ben Adida, *Helios: Web-Based Open-Audit Voting*, 17TH USENIX SECURITY SYMPOSIUM 335 (2008), [https://www.usenix.org/legacy/events/sec08/tech/full\\_papers/adida/adida.pdf](https://www.usenix.org/legacy/events/sec08/tech/full_papers/adida/adida.pdf); see also Ben Adida et al., Electing a University President using Open-Audit Voting: Analysis of Real-World use of Helios, presented at Electronic Voting Technology Workshop / Workshop on Trustworthy Elections (EVT/WOTE '09) (June 25, 2009), [https://www.usenix.org/legacy/events/evtwote09/tech/full\\_papers/adida-helios.pdf](https://www.usenix.org/legacy/events/evtwote09/tech/full_papers/adida-helios.pdf).

<sup>53</sup> Daniel Sandler & Dan S. Wallach, Casting Votes in the Auditorium, Presented at the 2nd USENIX/ACCURATE Electronic Voting Tech. Workshop (Aug. 6, 2007), [https://www.usenix.org/legacy/events/evt07/tech/full\\_papers/sandler/sandler.pdf](https://www.usenix.org/legacy/events/evt07/tech/full_papers/sandler/sandler.pdf); Daniel R. Sandler, Kyle Derr & Dan S. Wallach, *VoteBox: A Tamper-Evident, Verifiable Electronic Voting System*, 17TH USENIX SECURITY SYMP. 349, 349 (2008), [https://www.usenix.org/legacy/event/sec08/tech/full\\_papers/sandler/sandler.pdf](https://www.usenix.org/legacy/event/sec08/tech/full_papers/sandler/sandler.pdf).

Prob. Cheating (p)	Audits (n)	Prob. Detection
1%	40	33.10%
	80	55.25%
	120	70.06%
	160	79.97%
	200	86.60%
	240	91.04%
	280	94.00%
	320	95.99%
5%	500	99.34%
	10	40.13%
	20	64.15%
	30	78.54%
	40	87.15%
	50	92.31%
10%	60	95.39%
	10	65.13%
	20	87.84%
	30	95.76%
	40	98.52%
15%	50	99.48%
	10	80.31%
	20	96.12%
	30	99.24%
	40	99.85%
	50	99.97%

Table 1: Probabilities of discovering a cheating voting machine as a function of the odds of cheating occurring ( $p$ ) and the number of audits ( $n$ ) conducted.

#### a. Baseline audits

An election director must conduct some amount of auditing no matter what, and in the event suspiciously high spoiled ballot rates are reported, the election director might adaptively deploy more auditors.

What is the probability of catching at least one malicious machine in the act? The math is straightforward. Let us say that a malicious BMD does a switched intent attack with probability  $p$ . A randomly audited machine would then be caught cheating, again with probability  $p$ . Equivalently, the BMD gets away with its malice with probability  $1 - p$ . The probability of the BMD getting away with it after  $n$  audits is then  $(1 - p)^n$ . The probability of detecting the malware is  $1 - (1 - p)^n$ .

Table 1, above, shows some real numbers for  $p$  and  $n$ . Gilbert has also suggested an auditing process like this.<sup>54</sup>

From this, we can see that if the attacker wishes to modify only 1% of the ballots, an election official wishing to detect such an attack will need to conduct somewhere above 300 audits. To achieve a 99% confidence of detecting the attack, 468 audits would be necessary.

Of particular note, the probability of detecting a switched-intent attack has no dependency on the number of votes cast in the election. This means that the proportional cost of reaching a given detection probability shrinks as the voting population grows.

## V. Non-uniform malicious behavior

After a Twitter discussion of this idea, Stark wrote two drafts of an essay in response.<sup>55</sup> This section responds to some of his arguments.

Stark's strongest claim is that the adversary can be far more selective about which voters to attack.<sup>56</sup> For example, the attacker might only tamper with ballots for voters who operate very slowly, far more slowly than any auditor. Similarly, an attacker might only tamper with ballots for voters that use button boxes, large fonts, or other accessibility features.

Stark argues that the complexity of the variations that specify a given cast ballot, including the votes themselves, the time of day, the amount of time spent, and so forth, create a highly dimensional space that cannot be efficiently audited. In his words:

---

<sup>54</sup> Juan E. Gilbert, Ballot Marking Verification Protocol (unpublished manuscript), <http://www.juangilbert.com/BallotMarkingVerificationProtocol.pdf> (last visited Mar. 27, 2020).

<sup>55</sup> See Philip B. Stark, There is no Reliable Way to Detect Hacked Ballot-Marking Devices (Aug. 21, 2019) (unpublished manuscript), <https://www.stat.berkeley.edu/~stark/Preprints/bmd-p19.pdf>.

<sup>56</sup> *Id.*



Live tests need to probe *every* subset of voter preferences, BMD settings, and voter interactions with the BMD that could alter *any* contest outcome, and they need to probe every such subset enough to have a high probability of detecting any changes to selections in that subset.<sup>57</sup>

**a. Thought experiment: The shoulder-surfing auditor**

Consider the following thought experiment: we assign auditors to voters selected at random from the general voting population. Those voters will then do their normal voting process, but an auditor will watch them and will double-check the veracity of the printed ballot. This means that we are selecting from the distribution of real voters rather than from the distribution of all possible voter attributes.

This thought experiment is equivalent to Stark’s “oracle bound.”<sup>58</sup> He offers as an example an election with 20 BMDs, each of which prints 140 ballots during an election, for a total of 2800 ballots, of which 14 ballots have been altered by malware in the BMDs.<sup>59</sup> He concludes that the shoulder-surfing auditor would need to observe  $n = 539$  voters to achieve a 95% chance of detecting the malware.<sup>60</sup> Stark uses the following equation to find  $n$ :

$$\frac{2800-14}{2800} \cdot \frac{2799-14}{2799} \dots \frac{2800-(n-1)-14}{2800-(n-1)} \leq 0.05 \quad (61)$$

This equation models a sequence of shoulder-surfing audits. The fraction on the left is the probability that the malware survives the first audit, which is to say, the probability that the first selected ballot was clean, which is slightly less than 1.0. As we move to the right, we’re computing the probability that the first audit was clean and the second

---

<sup>57</sup> *Id.* at 7.

<sup>58</sup> *Id.* at 9.

<sup>59</sup> *Id.* at 10.

<sup>60</sup> *Id.*

<sup>61</sup> *Id.*

audit was clean, which we're thus multiplying together. We're interested in how far to the right we need to get before the malware wins with probably less than 5%.

Is Stark's math correct? How else might we model this? Using the previous equation, with  $p = \frac{14}{2800} = 0.5\%$  and  $n = 539$ , we derive a detection probability of 93.29%. Stark's version is more precise in its counting, representing an error in the math of Section 4 of 1.71%.<sup>62</sup> As the number of ballots cast in the election grows, the results of the two equations will converge.

Stark's argument hinges on his selection of a very small election, with only 2800 ballots cast while also selecting a very small fraction of malware activity, 0.5%.<sup>63</sup> This essay's live auditing strategy demands a fixed number of audits regardless of the total number of ballots cast, much like risk limiting audits select a number of ballots as a function of the margin of victory, not of the number of ballots cast. Stark's math is correct, but his example is cherry-picked to represent the very worst possible case for a live audit on a BMD election.

BMDs seem to be of great interest to large counties and states that will use them to collect millions of votes, with smaller counties often selecting paper ballots because they do not need BMD-only features like support for multiple languages, accessibility features, or thousands of distinct ballot styles. Stark's numbers do not reflect the relative effort of live audits as they might be conducted by the kinds of election jurisdictions that are favoring BMDs.

## VI. Realizing the auditing scheme

Before Stark leaves behind his "oracle bound" model, he dismisses it as being "impossibly optimistic."<sup>64</sup> We next consider how realistic this model might be in practice.

---

<sup>62</sup> *Id.*

<sup>63</sup> *Id.*

<sup>64</sup> *Id.*

While we cannot assign auditors at random to specific voters, we can create a probabilistic model of how real voters will behave. There are two parts to a model like this: voting preferences (i.e., what the voter ultimately selects in each contest), and machine-observable behaviors of the voter (e.g., how slow or fast the voter enters their preferences).

With historical ballots in hand, we can easily construct a probabilistic model that reflects how often voters will vote straight-ticket for one party or split their ticket; we can differentiate this by precinct, since we can examine prior cast ballots on a precinct-by-precinct basis. While we cannot connect prior cast ballots with the exact time they were cast—the cast ballots should have been randomized—we can still look at data from the event logs to determine how long voters took to create and cast their ballots. For features that we cannot observe from event logs, such as how many times a voter back-tracked and changed a preference, we can estimate these features from usability experiments.

As such, the challenge for the auditor is to create a “random voter” model that reflects all the voters across each voting precinct, including their preferences and behaviors. This then produces an auditing script, which will express the votes to be selected and the manner in which to do it. The script must also specify the time and location of each voting machine to be audited.

Political consultants regularly produce detailed models of voter behavior, typically used by “get out the vote” campaigns and other efforts to influence voter behavior. Similar techniques could be used to create the audit scripts. Making this model realistic, for example, recognizing that slow input behaviors may correlate with age, and age may correlate with party preferences, is a significant part of the challenge, but this challenge is still a tractable engineering problem.

#### **a. Down-ballot tampering**

Stark also suggests that an attacker might target a down-ballot race, where the number of audits that include that race might be quite small.

For example, in the Harris County (Houston, Texas) general election in November 2018, roughly 1.2 million ballots were cast for statewide contests.<sup>65</sup> For contrast, voters in the City of Baytown, on the east side of Harris County, cast roughly 9,000 votes in their propositions on the same ballot.<sup>66</sup>

The actual election results from Baytown were not close at all. The closest contest was decided by a nearly 2:1 margin. A malicious attacker would have needed to tamper with roughly 1,500 ballots to change its outcome. To be really sure, since the attacker could never have had such a precise prediction in advance, the attacker would have probably chosen at least 2,000 ballots to tamper. If 300 audits were performed, the odds of a county-wide “shoulder-surfing” auditor observing one of those tampered 2,000 voters would be roughly 40%. In a hypothetical variant of Baytown with a tighter election and only 400 votes tampered, the odds of observing a tampered vote would be only 10%.

These numbers suggest that an election official will need other techniques, besides live audits, to detect focused down-ballot tampering. We discuss other processes next.

#### **b. Spoiled ballot rates as a signal of problems**

Stark analyzes the question of how many spoiled ballots would be necessary to represent a statistically significant signal over a base rate of “expected” spoiled ballots. He notes that the answer is a function of the number of BMDs used; it is actually a function of the number of ballots cast, but the math is the same. With an assumed 1% spoilage rate, and 200,000 cast ballots, we might normally expect 2000 spoiled

---

<sup>65</sup> HARRIS COUNTY CLERK’S OFF. ELECTION DIVISION, HARRIS COUNTY, TEXAS — GENERAL AND SPECIAL ELECTIONS — NOVEMBER 06, 2018, at 1, 5 (2018), <https://harrisvotes.com/HISTORY/20181106/cumulative/cumulative.pdf>.

<sup>66</sup> *Id.*

ballots,<sup>67</sup> with 74 additional spoiled ballots being a signal that the spoilage rate was outside of the 95% confidence interval.

Stark uses an assumption that only 10% of voters would notice BMD-introduced errors.<sup>68</sup> Several researchers have examined this question and have found very different results. Matt Bernhard et al. presented its participants with varying signage or spoken prompts, yielding detection rates as low as 6.5% (with no interventions) and as high as 71.4% (with the most successful intervention).<sup>69</sup> Philip Kortum et al. separately considered whether voters attempt to check for errors and, if they do, whether they succeed, ultimately finding that 23.1% of their voters attempted to do the check and 76% of those successfully found an anomaly.<sup>70</sup> These studies suggest that, with the right voter prompts and instructions, we can achieve much higher than Stark's 10% assumption.

Nonetheless, with Stark's 10% assumption, he concludes that the malware can get away with roughly 730 attacks, corresponding to changing the margin of victory by 0.73%; this includes the computation that one tampered ballot causes one candidate to lose a vote and the other to gain a vote.<sup>71</sup> The table below generalizes Stark's math to three different election sizes and three different likelihoods that voters will detect tampered ballots and spoil them. The reported Margin  $\Delta\%$  indicates the maximum change to a margin of victory that malware might hope to accomplish without being detected.

---

<sup>67</sup> Until we see large deployments of BMD voting systems, we won't know the actual rate of spoiled ballots on these systems and how it varies with other effects, such as the size of the ballot. This data is certainly important to collect and study.

<sup>68</sup> Stark, *supra* note 55, at 13.

<sup>69</sup> Matt Bernhard, Allison McDonald, Henry Meng, Jensen Hwa, Nakul Bajaj, Kevin Chang, and J. Alex Halderman, Can Voters Detect Malicious Manipulation of Ballot Marking Devices?, Proceedings of the 41st IEEE Symp. on Security and Privacy, at 7 (2020).

<sup>70</sup> Philip Kortum, Michael D. Byrne, & Julie Whitmore, Voter Verification of BMD Ballots is a Two-Part Question: Can They? Mostly, They Can. Do They? Mostly They Don't, arXiv e-print, (March 10, 2020) (unpublished manuscript), <https://arxiv.org/abs/2003.04997>.

<sup>71</sup> Stark, *supra* note 55, at 13.

Detection%	Num ballots	Margin $\Delta\%$
10%	9,000	3.556%
	200,000	0.740%
	1,200,000	0.300%
30%	9,000	1.185%
	200,000	0.247%
	1,200,000	0.100%
50%	9,000	0.711%
	200,000	0.148%
	1,200,000	0.060%

Table 2: Possible malware changes to the margin of victory as a function of the detection rate and number of ballots cast. Code to generate this table is presented in Appendix A.

Table 2 shows the impact of increasing the size of the electorate as well as increasing the likelihood that voters will notice and spoil tampered ballots. For a county-wide contest in Harris County, any malware attack that moved the margin of victory more than 0.3%, even with Stark's most pessimistic assumption about voters noticing erroneous ballots, would pass the 95% confidence interval on the expected ballot spoilage rate. With a more optimistic assumption on voters spoiling these ballots, the malware could not hope to move the margin of victory by more than six one-hundredths of a percent without exceeding the 95% confidence interval.

Consequently, for elections in large jurisdictions, real-time tracking of ballot spoilage events should provide an effective mechanism for election officials to detect switched-intent attacks. For much smaller elections, however, such as Baytown's local election, a switched-intent attack would seem to have more room to operate.

### c. Election procedures and emergencies

As discussed in Section III.C and quantified in Section V.D, election officials gain power from having situational awareness of the rate of ballots being spoiled. With this, election officials can attempt to intuit strategies being taken by malware and adaptively create auditing

strategies that might catch the malware in the act. If, for example, there were an unusually high spoiled ballot rate in Baytown, or if there were a hotly contested race there with large political stakes,<sup>72</sup> then the election official could choose to deploy additional auditors to those areas, on top of the baseline of countywide live audits and spoiled ballot tracking.

As such, the “game” becomes less like flipping coins and more like playing poker. Statistics still play a role, but the players must spend a significant amount of energy trying to intuit each other's strategies, including bluffing and other forms of subterfuge. Furthermore, this is a game where the attacker must move first, committing to a malware attack that will have a specific impact on the outcome (whether a switched-intent attack or something else). The election official gets to make responsive moves up to the day of election, after the voting machines are potentially beyond the reach of the attacker.

In the wake of the 2016 election, the Department of Homeland Security declared elections to be “critical infrastructure.”<sup>73</sup> One of the consequences of this is the existence of the Elections Infrastructure Information Sharing and Analysis Center (EI-ISAC),<sup>74</sup> creating a structured program for the federal government to share threat intelligence information with election officials, as well as for election officials to aggregate and share such intelligence amongst themselves. Certainly, the compressed timetable of elections means that this sort of sharing needs to happen quickly, but if, for example, the various

---

<sup>72</sup> Baytown is adjacent to a number of petrochemical refineries, including one of the largest in the U.S., owned by Exxon-Mobil, so it's easy to imagine complicated politics between the citizens of the city and the industry around them. *Baytown Refinery*, WIKIPEDIA, [https://en.wikipedia.org/wiki/Baytown\\_Refinery](https://en.wikipedia.org/wiki/Baytown_Refinery) (last modified Oct. 7, 2019, 8:56 AM) [<https://perma.cc/CAE4-JXW3>]; BAYTOWN AREA OPERATIONS, EXXONMOBIL (Aug. 5, 2019), <https://corporate.exxonmobil.com/Locations/United-States/Baytown-area-operations-overview#aboutUs> [<https://perma.cc/9PCK-QLQ2>].

<sup>73</sup> ELECTION SECURITY, U.S. DEP'T OF HOMELAND SECURITY, <https://www.dhs.gov/topic/election-security> (last modified Feb. 7, 2020) [<https://perma.cc/G57Z-QGJM>].

<sup>74</sup> EI-ISAC CHARTER, CTR. FOR INTERNET SECURITY, <https://www.cisecurity.org/ei-isac/ei-isac-charter/> [<https://perma.cc/6CUA-QSB8>].

Federal intelligence agencies reach a conclusion that attacks are likely in specific states or localities, then EI-ISAC provides the necessary infrastructure to disseminate this information, allowing election officials to perhaps get ahead of the malware before the election even begins.

#### **d. Responding to an emergency**

Stark raises this point in his responses, and it is an important issue to discuss. While the legal process for managing elections varies from state to state, consider what happens when a natural disaster strikes right before or during an election. Exactly this happened with Hurricane Sandy, which struck the northeastern seaboard on October 29, 2012, causing notably large damage in New Jersey and New York, right before a presidential election. In the wake of this storm, many politicians recognized the need for emergency procedures<sup>75</sup> and the National Association of Secretaries of State began a push to get states to adopt laws and procedures to take disasters into account.<sup>76</sup> Ultimately, a cyberattack on an election can and should be treated much the same as a hurricane or other natural disaster. If the scope and reach of a cyberattack is large enough that the outcome of the election is in doubt, then suitable disaster procedures would allow a governor to declare an emergency and re-run an election, perhaps with a different voting technology.

Note that modern elections are not actually finalized on the night of the election, even though losing candidates will customarily concede to the victors at the time. Instead, all elections have a canvass period after the election.<sup>77</sup> During this period, a wide variety of activities

---

<sup>75</sup> See, e.g., Thomas Kaplan, *Using Hurricane Sandy as a Lesson for Future Elections*, N.Y. TIMES (Nov. 12, 2013), <https://www.nytimes.com/2013/11/13/nyregion/lessons-from-hurricane-sandy-being-applied-to-election-planning.html> [<https://perma.cc/JZ7X-6VUV>].

<sup>76</sup> NAT'L ASS'N OF SECRETARIES OF STATE, STATE LAWS & PRACTICES FOR THE EMERGENCY MANAGEMENT OF ELECTIONS 4 (2017).

<sup>77</sup> *State Election Canvassing Timeframes and Recount Thresholds*, NAT'L ASS'N OF SECRETARIES OF STATE (2018), <https://www.nass.org/resources/2018-election-information/Canvassing-Timeframes-and-Recount-Thresholds>.



occur, which includes processes like tabulating vote-by-mail ballots and resolving provisionally cast ballots. (See, e.g., California's canvass information page<sup>78</sup>) The canvass period is typically when a risk-limiting audit will be conducted, and is also a suitable time for cyber-forensics to be conducted on BMDs that were discovered during audits or simply flagged by voters spoiling their ballots.

Still, once a vote has been tampered, you cannot determine the intent of the voter. So what do you do? Procedurally, this should be no different than a ballot box, or potentially a warehouse of every ballot box, being lost or destroyed in a flood. It is an emergency, and you need emergency procedures to resolve the problem. While it would be politically sensitive to declare that a cyberattack damaged an election and as such it had to be re-run, the likelihood of an emergency response mitigates against the risks of cyberattacks. In other words, if the attacker does not think they will get away with it, they are less likely to bother with the attack.

## VII. Additional arguments

A wide variety of other arguments have been made by Stark and others against ballot marking devices. This section tries to respond to these arguments.

The adversary will know the election official's auditing strategy, giving them an advantage, Stark notes.<sup>79</sup> As discussed in Section 5.2, the election official must create an auditing strategy that mirrors the real-world distribution of voter's preferences and behaviors. This will inevitably require a sophisticated software tool that constructs an auditing script based on real-world data taken from prior elections. If this tool and the data it uses are open-source, then the only input that the adversary does not know is the random seed that drives the production of the audit script. If the model has weaknesses that do not capture real-world voter behavior, the adversary could further take

---

<sup>78</sup> *The Official Canvass of the Vote*, CAL. SEC'Y OF STATE, <https://www.sos.ca.gov/elections/official-canvass/> [https://perma.cc/H58Q-JAYR].

<sup>79</sup> Stark, *supra* note 55, at 12.

advantage of these weaknesses. On the flip side, the human auditors executing the auditing script will inevitably add randomness on their own (e.g., making errors with respect to the scripted vote inputs and needing to correct those).

A more serious threat is that the computer generating the auditing script is, itself, controlled by the adversary. This can be partly mitigated by having multiple, independent computers generating the audit script, with the random seed produced by rolling physical dice. The resulting scripts should be identical, otherwise we have evidence of malware, which again leads us to emergency processes.

These auditing procedures would require significant additional staffing and training. Let us again use the November 2018 election in Harris County, Texas, to get some realistic numbers. In this election, 1.2 million votes were cast across a fleet of roughly 10,000 voting machines.<sup>80</sup> In this election, roughly 29% of ballots were cast on Election Day, with 63% of ballots cast during the early voting period.<sup>81</sup> The remaining 8% were absentee postal mail ballots. In the 2019 general election, there were 52 early voting centers and on Election Day there were 750 local voting locations.<sup>82</sup>

Let us assume that an auditing script for 2020 would follow the same distribution of votes that we saw in 2018, with the bulk of ballots cast in early voting locations. If we had one audit team (perhaps two trained auditors) per early voting location, we would then need 104 trained auditors. If we want 500 live audits across the full election, then this will average out to roughly six audits per early voting location, spread across the two-week early voting period, i.e., less than one audit per day. As such, the auditors would most likely be regular poll workers given extra training for the auditing function.

---

<sup>80</sup> HARRIS COUNTY CLERK'S OFF. ELECTION DIVISION, *supra* note 65.

<sup>81</sup> *Id.*

<sup>82</sup> Voting Information, HARRIS VOTES, <https://www.harrisvotes.com/VotingInfo> (last visited May 25, 2020).

On Election Day, the same 52 teams of auditors could then be charged with executing roughly three audits per team, representing a very reasonable workload to achieve around 500 audits across the election. Even in the case where the probability of a BMD cheating was a low 1%, this audit would have a 99.34% chance of detecting it (see Table 1). It is also straightforward to see these auditors given additional auditing tasks, adaptively, as described in Section 6.

These auditing procedures depend on the margin of victory, which is not known while the audits are under way. If the final margin of victory is larger than the margin used for the audit script, then we will have confidence that any possible attack would be smaller than the margin of victory, and therefore we have increased confidence in the outcome of the election. Conversely, if the final margin of victory is smaller than the margin used for the audit script, then the audit process might well have missed an attack that could well have changed the outcome.

This paper's analysis shows that live audits plus live tracking of spoiled ballots can drive down the available margin for malware to tamper with an election result without being detected. This raises a related question of whether an adversary might even attempt this sort of malware-based attack when other attacks (e.g., tampering with the voter registration database, or bombarding social media with propaganda) might be more likely to succeed and have a larger impact.

Perhaps more importantly, once the margin of victory gets within 1% of the number of votes cast, a variety of other factors come into play, particularly around the usability of the voting system.<sup>83</sup> BMDs have the potential to perform much better than hand-marked paper in these circumstances, as we discuss next.

---

<sup>83</sup> See, e.g., Sarah P. Everett et al., *Electronic Voting Machines versus Traditional Methods: Improved Preference, Similar Performance*, 2008 CHI PROC. 883.

### VIII. Discussion and Conclusions

BMDs give us the opportunity to build more sophisticated voting systems with “end to end” (e2e) security guarantees.<sup>84</sup> While none of today’s BMDs have features like this, the research literature has a variety of designs that give voters a “receipt” that allows them to prove that their vote was correctly included in the final tally (i.e., “counted as cast”).<sup>85</sup> There are also many clever techniques that can be used to audit voting devices to catch them if they are cheating (i.e., verifying that voters are “cast as intended”).<sup>86</sup> If we ever want to have e2e elections, then we will likely require BMD-like devices which produce regular paper ballots as well as computing the necessary cryptography.

How realistic is it that e2e will make the jump from the research literature to commercial production? To pick one example, Microsoft is investing in an open-source toolkit called ElectionGuard,<sup>87</sup> and they have announced partnerships with many of the vendors of election equipment.<sup>88</sup> It is quite likely that the next generation of BMDs, and perhaps even current-generation BMD hardware with new software, will adopt these techniques.

The risks of malware in current-generation BMDs are non-trivial, but they can be mitigated through human-centered ballot design, careful auditing procedures, and suitable election emergency laws. They also keep the door open to new cryptographic techniques, such as used in ElectionGuard, that have the potential to protect against a variety of other election threats.

---

<sup>84</sup> See, e.g., Susan Bell et al., *STAR-Vote: A Secure, Transparent, Auditable, and Reliable Voting System*, 1 USENIX J. ELECTION TECH. & SYS. 1 (2013).

<sup>85</sup> *Id.*

<sup>86</sup> *Id.*

<sup>87</sup> Tom Burt, *Protecting Democratic Elections Through Secure, Verifiable Voting*, MICROSOFT (May 6, 2019), <https://blogs.microsoft.com/on-the-issues/2019/05/06/protecting-democratic-elections-through-secure-verifiable-voting/> [<https://perma.cc/2PY7-6RVC>].

<sup>88</sup> Tom Burt, *Another Step in Testing ElectionGuard*, MICROSOFT (Feb. 17, 2020), <https://blogs.microsoft.com/on-the-issues/2020/02/17/wisconsin-electionguard-polls/> [<https://perma.cc/WLD9-MKKT>].

Unlike the paperless electronic voting systems that BMDs are being purchased to replace, the paper ballots that come out of BMDs give us the ability to consider the security procedures contemplated here. BMDs are the best technology available today that combines the security benefits of paper with the accessibility benefits of computers.

### Appendix - A Computing detectable spoiled ballot rates

This short Python program shows how to generate Table 2, here in a simple comma-separated-value format. In my attempts to reproduce Stark's numbers, he suggested I use NumPy's `poisson.ppf`— the Poisson cumulative distribution function— which is more accurate than the more commonly used normal approximation to the Poisson distribution:  $\mu + 1.96\sqrt{\mu}$ . (Wikipedia explains where the number 1.96 comes from in this equation.<sup>89</sup>)

```
from scipy.stats import poisson

print("Detection%,Num ballots,Margin Delta%")

expectedSpoilage = 0.01
for detectionFraction in [0.1, 0.3, 0.5]:
    for electionSize in [9000, 200000, 1200000]:
        mu = electionSize * expectedSpoilage
        cd = poisson.ppf(0.95, mu)
        print("%d%%,%d,%.3f%%" %
              (detectionFraction * 100, electionSize,
               200.0 * (cd -
                mu) / (electionSize * detectionFraction)))
```

---

<sup>89</sup> 1.96, WIKIPEDIA, <https://en.wikipedia.org/wiki/1.96> [<https://perma.cc/TS5E-F2AG>] (last modified Mar. 15, 2020).