Public Evidence from Secret Ballots[⋆]

Matthew Bernhard[◁], Josh Benaloh[□], J. Alex Halderman[◁], Ronald L. Rivest[◇],
Peter Y. A. Ryan[○], Philip B. Stark[‡], Vanessa Teague[▷], Poorvi L. Vora[†], Dan S. Wallach[⋆]

[□]Microsoft Research [◁]University of Michigan [○]University of Luxembourg
[‡]University of California at Berkeley [⋆]Rice University [▷]University of Melbourne
[†]The George Washington University [◇]Massachusetts Institute of Technology

Abstract. Elections seem simple—*aren't they just about counting?* But they have a unique, challenging combination of security and privacy requirements. The stakes are high; the context is adversarial; the electorate needs to be convinced that the results are correct; and the secrecy of the ballot must be ensured. They also have practical constraints: time is of the essence, and voting systems need to be affordable and maintainable, as well as usable by voters, election officials, and pollworkers. It is thus not surprising that voting is a rich research area spanning theory, applied cryptography, practical systems analysis, usable security, and statistics. Election integrity involves two key concepts: *convincing evidence that outcomes are correct* and *privacy*, which amounts to *convincing assurance that there is no evidence* about how any given person voted. These are obviously in tension. We examine how current systems walk this tightrope.

1 Introduction: What is the evidence?

It is not enough for an election to produce the correct outcome. The electorate must also be convinced that the announced result reflects the will of the people. And for a rational person to be convinced requires evidence.

Modern technology—computer and communications systems—is fragile and vulnerable to programming errors and undetectable manipulation. No current system that relies on electronic technology alone to capture and tally votes can provide convincing evidence that election results are accurate without endangering or sacrificing the anonymity of votes.[1]

Paper ballots, on the other hand, have some very helpful security properties: they are readable (and countable, and re-countable) by humans; they are relatively durable; and they are tamper-evident. Votes cast on paper can be counted using electronic technology; then the accuracy of the count can be checked manually to ensure that the technology functioned adequately well. Statistical methods allow the accuracy of the count to be assessed by examining only a fraction of the ballots manually, often a very small fraction. If there is also convincing evidence that the collection of ballots has been conserved (no ballots added, lost, or modified) then this combination—voter-verifiable paper ballots, a mechanized count, and a manual check of the accuracy of that count—can provide convincing evidence that announced electoral outcomes are correct.

[⋆] A more in-depth version of this paper can be found at https://arxiv.org/abs/1707.08619

[1] Moreover, the systems that come closest are not readily usable by a typical voter.

Conversely, absent convincing evidence that the paper trail has been conserved, a manual double-check of electronic results against the paper trail will not be convincing. If the paper trail has been conserved adequately, then a full manual tally of the ballots can correct the electronic count if the electronic count is incorrect.

These considerations have led many election integrity advocates to push for a voter-verifiable paper audit trail (VVPAT[2]) in the absence of paper ballots.

In the 2016 US presidential election, about three quarters of Americans voted using systems that generated voter-verifiable paper records. The aftermath of the election proved that even if 100% of voters had used such systems, it would not have sufficed to provide convincing evidence that the reported results are accurate.

- No state has laws or regulations to ensure that the paper trail is conserved adequately, and that evidence to that effect is provided.
- No state had laws or regulations that ensured adequate manual scrutiny of the paper to determine that the electronically-generated results were correct.
- Many states that have a paper trail also have laws that make it hard for anyone to check the results using the paper trail—even candidates with war chests for litigation. Not only can other candidates fight attempts to check the results, the states themselves can fight such attempts. This treats the paper as a nuisance, rather than a safeguard.

The bottom line is that the paper trail is not worth the paper it's printed on. Clearly this must change.

Other approaches like *software independence* and *end-to-end verifiability* offer tools to improve electronic voting systems, but these methods have not been broadly applied.

1.1 Why so hard?

Several factors make it difficult to generate convincing evidence that reported results are correct. The first is the trust model.

No one is trusted. In any significant election, voters, election officials, and equipment used to vote cannot necessarily be trusted by anyone with a stake in the outcome. Voters, operators, system designers, and external parties are also all potential adversaries.

The need for evidence Because officials and equipment may not be trustworthy, elections should be *evidence-based*. Any observer should be able to verify the reported results based on trustworthy evidence from the voting system. Many in-person voting systems fail to provide sufficient evidence; and as we shall see Internet systems scarcely provide any at all.

The secret ballot Perhaps the most distinctive element of elections is the *secret ballot*, a critical safeguard that defends against vote selling and voter coercion. In practical terms, voters should not be able to prove how they voted to anyone, *even if they wish to do so*. This restricts the types of evidence that can be produced by the voting system. For example, the voting system may not provide votes encrypted by the voter as evidence,

---

[2] The VVPAT consists of a cash-register style printout of each vote that the voter can check but cannot touch. Voter-marked paper ballots or ballots marked using a ballot-marking device are preferable to VVPAT because voters may not check the VVPAT record.

because the voters may choose to reveal their selections and the randomness used during encryption in response to bribery or coercion.

The challenge of voting is thus to use fragile technology to produce trustworthy, convincing *evidence* of the correctness of the outcome while protecting voter *privacy* in a world *where no person or machine may be trusted*. The resulting voting system and its security features must also be *usable* by regular voters.

The aim of this paper is to explain the important requirements of secure elections and the solutions already available from e-voting research, then to identify the most important directions for research.

Prior to delving into our discussion, we need to make a distinction in terminology. *Pollsite* voting systems are those in which voters record and cast ballots at predetermined locations, often in public areas with strict monitoring. *Remote* voting refers to a system where voters fill out ballots anywhere, and then send them to a central location to cast them, either physically mailing them in the case of vote-by-mail, or sending them over the Internet in the case of Internet voting.

Section 2 defines the requirements related to notions of election evidence, section 3 on privacy and voter authentication, and section 4 on more general usability, availability and local regulatory requirements. Section 5 describes the cryptographic, statistical, and engineering tools that have been developed for designing voting systems with verifiably correct election outcomes. Section 6 concludes with the promise and problems associated with Internet voting.

2 Secure Voting Requirements: Trust, Verifiability, and Evidence

For an election to be accepted as legitimate, the outcome should be convincing to all—and in particular to the losers—leaving no valid grounds to challenge the outcome. Whether elections are conducted by counting paper ballots by hand or using computer technology, the possibility of error or fraud necessitates assurances of the accuracy of the outcome.

It is clear that a naive introduction of computers into voting introduces the possibility of wholesale and largely undetectable fraud. If we can't detect it, how can we prevent it?

2.1 Risk Limiting Audits

Statistical post-election audits provide assurance that a reported outcome is correct, by examining some or all of an *audit trail* consisting of durable, tamper-evident, voter-verifiable records. Typically the audit trail consists of paper ballots.

The *outcome* of an election is the set of winners. An outcome is incorrect if it differs from the set of winners output by a perfectly accurate manual tabulation of the audit trail.

Definition 1. *An audit of an election contest is a* **risk-limiting audit (RLA)** with risk limit α *if it has the following two properties:*

1. *If the reported contest outcome under audit is incorrect, the probability that the audit enables outcome correction is at least* $1-\alpha$.
2. *The audit will not indicate a need to alter a reported outcome that is correct.*

Together, these two properties imply that post-RLA, either the reported set of winners is the set that a perfectly accurate hand count of the audit trail would show, or an event which altered the election's outcome has occurred and was not detected by the audit (this event has probability no larger than bound α). RLAs amount to a limited form of probabilistic error correction: by relying on the audit trail, they have a known minimum probability of correcting the reported outcome if it is incorrect. They are not designed to detect (or correct) an incorrectly-reported tally, only an incorrectly-reported outcome.

The following procedure is a trivial RLA: with probability $1 - \alpha$, perform a full manual tally of the audit trail. Amend the outcome to match the set of winners the full hand count shows if that set is different.

The art in constructing RLAs consists of maintaining the risk limit while performing *less work* than a full hand count when the outcome is correct. Typically, this involves framing the audit as a sequential test of the statistical hypothesis that the outcome is incorrect. To reject that hypothesis is to conclude that the outcome is correct. RLAs have been developed for majority contests, plurality contests, and vote-for-k contests and complex social choice functions including D'Hondt—see below. RLAs have also been devised to check more than one election contest simultaneously [84].

2.2 Software Independence

Rivest and Wack introduced a definition targeted specifically at detecting misbehavior in computer-based elections:

Definition 2. *[68] A voting system is* **software independent** *if an undetected change or error in its software cannot cause an undetectable change or error in an election outcome.*

Software independence clearly expresses that it should not be necessary to trust software to determine election outcomes, but it does not say what procedures or types of evidence should be trusted instead. A system that is not software independent *cannot* produce a convincing evidence trail, but neither can a paper-based system that does not ensure that the paper trail is complete and intact, a cryptographic voting system that relies on an invalid cryptographic assumption, or a system that relies on audit procedures but lacks a means of assuring that those procedures are properly followed. We could likewise demand independence of many other kinds of trust assumptions: hardware, paper chain-of-custody, cryptographic setup, computational hardness, procedures, good randomness generation *etc.*

Rivest and Wack also define a stronger form of the property that includes error recovery:

Definition 3. *[68] A voting system is* **strongly software independent** *if it is software independent and a detected change or error in an election outcome (due to the software) can be corrected without rerunning the election.*

A strongly software-independent system can recover from software errors or bugs, but that recovery in turn is generally based on some other trail of evidence.

Software independence can be viewed as a form of tamper-evident system: a material software problem leaves a trace. Strongly software independent systems are resilient: not only do material software problems leave a trace, the overall election system can recover from those problems.

One mechanism to provide software independence is to record votes on a paper record that provides physical evidence of voter's intent, can be inspected by the voter prior to casting the vote, and—if preserved intact—can later be manually audited to verify the election outcome. Risk-limiting audits (see Section 5.2) can then achieve a prespecified level of assurance that results are correct; machine assisted risk-limiting audits [20], can help minimize the amount of labor required.

Open problems:

- How can systems handle errors in the event that elections don't verify? Can they recover?

2.3 End-to-end verifiability

The concern regarding fraud and desire for transparency has motivated the security and crypto communities to develop another approach to voting system assurance: *end-to-end verifiability* (E2E-V). An election that is end-to-end verifiable achieves software independence together with the analagous notion of hardware independence as well as independence from actions of election personnel and vendors. Rather than attempting to verify thousands of lines of code or closely monitor all of the many processes in an election, E2E-V focuses on providing a means to detect errors or fraud in the process of voting and counting. The idea behind E2E-V is to enable voters themselves to monitor the integrity of the election. This is challenging because total transparency is not possible without undermining the secret ballot, hence the mechanisms to generate such evidence have to be carefully designed.

Definition 4. *(adapted from [14])*

A voting system is **end-to-end verifiable** *if it has the following three kinds of verifiability:*

- **Cast as intended:** *Voters can independently verify that their selections are correctly recorded.*
- **Collected as cast:** *Voters can independently verify that the representation of their vote is correctly collected in the tally.*
- **Tallied as collected:** *Anyone can verify that every well-formed, collected vote is correctly included in the tally.*

If verification relies on trusting entities, software, or hardware, the voter and/or auditor should be able to choose them freely. Trusted procedures, if there are any, must be open to meaningful observation by every *voter.*

Note that the above definition allows each voter to check that her vote is correctly collected, thus ensuring that attempts to change or delete cast votes are detected. In addition, it should also be possible to check the list of voters who cast ballots, to ensure that votes are not added to the collection (*i.e.,* to prevent ballot-box stuffing). This is called *eligibility verifiability* [53,81].

2.4 Collection Accountability

In an E2E-V election protocol, voters can check whether their votes have been properly counted, but if they discover a problem, there may not be adequate evidence to correct it. An election system that is *collection-accountable* provides voters with evidence of any failure to collect their votes.

Definition 5. *An election system is* **collection accountable** *if any voter who detects that her vote has not been collected has, as part of the vote-casting protocol, convincing evidence that can be presented to an independent party to demonstrate that the vote has not been collected.*

Another form of evidence involves providing each voter with a code representing her votes, such that knowledge of a correct code is evidence of casting a particular vote [27]. Yet another mechanism is a suitable paper receipt. Forensic analysis may provide evidence that this receipt was not forged by a voter [11,7].

Open problems:

- Can independently verifiable evidence be provided by the voting system for incorrect ballot casting?

2.5 Dispute Resolution

While accountability helps secure the election process, it is not very useful if there is no way to handle disputes. If a voter claims, on the basis of accountability checks provided by a system, that something has gone wrong, there needs to be a mechanism to address this. This is known as *dispute resolution*:

Definition 6. *[46] A voting system is said to have* **dispute resolution** *if, when there is a dispute between two participants regarding honest participation, a third party can correctly resolve the dispute.*

An alternative to dispute resolution is dispute freeness:

Definition 7. *[50] A* **dispute free** *voting system has built-in prevention mechanisms that eliminate disputes between the active participants; any third party can check whether an active participant has cheated.*

Open problems:

- Can effective dispute resolution for all classes of possible errors exist in a given system?
- Are there other reasonable definitions and mechanisms for dispute resolution?
- Can a system offer complete dispute resolution capabilities in which every dispute can be adjudicated using evidence produced by the election system?

2.6 From Verifiable to Verified

Constructing a voting system that creates sufficient evidence to reveal problems is not enough on its own. That evidence must actually be used—and used appropriately—to ensure the accuracy of election outcomes.

An election result may not be verified, even if it is generated by an end-to-end verifiable voting system. For verification of the result, we need several further conditions to be satisfied:

- Enough voters and observers must be sufficiently diligent in performing the appropriate checks.
- Random audits (including those initiated by voters) must be sufficiently extensive and unpredictable that attempts at election manipulation have a high chance of being detected.

– If checks fail, this must be reported to the authorities who, in turn, must take appropriate action.

These issues involve complex human factors, including voters' incentives to participate in verification. Little work has been done on this aspect of the problem.

An E2E-V system might give an individual voter assurance that her vote has not been tampered with *if* that voter performs certain checks. However, sufficiently many voters must do this in order to provide evidence that the election outcome as a whole is correct. Combining risk-limiting audits with E2E-V systems can provide a valuable layer of protection in the case that an insufficient number of voters participate in verification.

Finally, another critical verification problem that has received little attention to date is how to make schemes that are recoverable in the face of errors. We do not want to have to abort and rerun an election every time a check a fails. Certain levels of detected errors can be shown to be highly unlikely to affect the outcome, and hence can be tolerated. Other types and patterns of error can be handled and corrected for, either post hoc or dynamically.

Both Küsters *et al.* [55] and Kiayias *et al.* [52] model voter-initiated auditing [10] and its implications for detection of an incorrect election result. Both definitions turn uncertainty about voter initiated auditing into a bound on the probability of detecting deviations of the announced election result from the truth.

Open problems:

– Can systems be designed so that the extent and diligence of checks performed can be measured?

– Can verification checks be abstracted from the voter, either by embedding them in other election processes or automating them?

3 Secure Voting Requirements: Voter Authentication, Privacy, Receipt-Freeness and Coercion-Resistance

This section focuses on secure voting system requirements related to authenticating the voter and ensuring that the evidence provided for verifiability cannot be used to coerce or bribe the voter to vote in a certain manner.

3.1 Voter Authentication

A significant challenge for election systems is the credentialing of voters to ensure that all eligible voters, and no one else, can cast votes. This presents numerous questions: what kinds of credentials should be used? How should they be issued? Can they be revoked or de-activated? Are credentials good for a single election or for an extended period? How difficult are they to share, transfer, steal, or forge? Can the ability to create genuine-looking forgeries help prevent coercion? These questions must be answered carefully, and until they are satisfied for remote voting, pollsite voting is the only robust way to address these questions—and even then, in-person credentialing is subject to forgery, distribution, and revocation concerns (for instance, the Dominican Republic recently held a pollsite election where voters openly sold their credentials [35]). In the U.S., there is concern that requiring in-person credentialing, in the form of voter ID, disenfranchises legitimate voters.

Open problems:
- Is there a sufficiently secure way to distribute credentials for Internet voting?
- Is a traditional PKI sufficient to ensure eligibility for remote voting?
- How does use of a PKI change coercion assumptions?

3.2 Privacy, Receipt Freeness, and Coercion Resistance

In most security applications, privacy and confidentiality are synonymous. In elections, however, privacy has numerous components that go well beyond typical confidentiality. Individual privacy can be compromised by "normal" election processes such as a unanimous result. Voters may be coerced if they can produce a proof of how they voted, even if they have to work to do so.

Privacy for votes is a means to an end: if voters don't express their true preferences then the election may not produce the right outcome. This section gives an overview of increasingly strong definitions of what it means for voters to be free of coercion.

Basic Confidentiality We will take *ballot privacy* to mean that the election does not leak any information about how any voter voted beyond what can be deduced from the announced results. Confidentiality is not the only privacy requirement in elections, but even simple confidentiality poses significant challenges. It is remarkable how many deployed e-voting systems have been shown to lack even the most basic confidentiality properties (e.g., [42,34,24,21,59]).

Perhaps more discouraging to basic privacy is the fact that remote voting systems (both paper and electronic) inherently allow voters to eschew confidentiality. Because remote systems enable voters to fill out their ballots outside a controlled environment, anyone can watch over the voter's shoulder while she fills out her ballot.

In an election—unlike, say, in a financial transaction—even the candidate receiving an encrypted vote should not be able to decrypt it. Instead, an encrypted (or otherwise shrouded) vote must remain confidential to keep votes from being directly visible to election authorities.

Some systems, such as code voting [26] and the Norwegian and Swiss Internet voting schemes, defend privacy against an attacker who controls the computer used for voting; however, this relies on assumptions about the privacy and integrity of the code sheet. Some schemes, such as JCJ/Civitas [45], obscure who has voted while providing a proof that only eligible votes were included in the tally.

Several works [33] [55], following Benaloh [16] formalize the notion of privacy as preventing an attacker from noticing when two parties swap their votes.

Open problems:
- Can we develop more effective, verifiable forms of assurance that vote privacy is preserved?
- Can we build means of privacy for remote voting through computer-based systems?

Everlasting Privacy Moran and Naor expressed concern over what might happen to encrypted votes that can still be linked to their voter's name some decades into the future, and hence decrypted by superior technology. They define a requirement to prevent this:

Definition 8. *[60] A voting scheme has* **everlasting privacy** *if its privacy does not depend on assumptions of cryptographic hardness.*

Their solution uses perfectly hiding commitments to the votes, which are aggregated homomorphically. Instead of privacy depending upon a cryptographic hardness assumption, it is the integrity of an election that depends upon a hardness assumption; and only a real-time compromise of the assumption can have an impact.

Systemic Privacy Loss We generally accept that without further information, a voter is more likely to have voted for a candidate who has received more votes, but additional data is commonly released which can further erode voter privacy. Even if we exclude privacy compromises, there are other privacy risks which must be managed. If voters achieve privacy by encrypting their selections, the holders of decryption keys can view their votes. If voters make their selections on devices out of their immediate control (*e.g.* official election equipment), then it is difficult to assure them that these devices are not retaining information that could later compromise their privacy. If voters make their selections on their own devices, then there is an even greater risk that these devices could be infected with malware that records (and perhaps even alters) their selections (see, for instance, the Estonian system [82]).

Open problems:

- Are there ways to quantify systemic privacy loss?
- Can elections minimize privacy loss?
- Can elections provide verifiable integrity while minimizing privacy loss?

Receipt-freeness The problem of preventing coercion and vote-selling was considered solved with the introduction of the *Australian* ballot. The process of voting privately within a public environment where privacy can be monitored and enforced prevents improper influence. Recent systems have complicated this notion, however. If a voting protocol provides a receipt but is not carefully designed, the receipt can be a channel for information to the coercive adversary.

Benaloh and Tuinstra [15] pointed out that passive privacy is insufficient for resisting coercion in elections:

Definition 9. *A voting system is* **receipt free** *if a voter is unable to prove how she voted* even if she actively colludes with a coercer and deviates from the protocol in order to try to produce a proof.

Traditional elections may fail receipt-freeness too. In general, if a vote consists of a long list of choices, the number of possible votes may be much larger than the number of likely voters. This is sometimes called (a failure of) the *short ballot assumption* [71]. Prior to each election, coercers assign a particular voting pattern to each voter. When the individual votes are made public, any voter who did not cast their pattern can then be found out. This is sometimes called the *Italian attack*, after a once prevalent practice in Sicily. It can be easily mitigated when a vote can be broken up, but is difficult to mitigate in systems like IRV in which the vote is complex but must be kept together. Mitigations are discussed in Sections 5.2 and 5.3.

Incoercibility has been defined and examined in the universally composable framework in the context of general multiparty computation [22,90]. These definitions examine whether the protocol introduces additional opportunities for coercion that are not present when the computation is performed by a trusted party. With some exceptions (such as [5]), they usually focus on a passive notion of receipt-freeness, which is not strong enough for voting.

Coercion Resistance Schemes can be receipt-free, but not entirely resistant to coercion. Schemes like Prêt à Voter [74] that rely on randomization for receipt-freeness can be susceptible to *forced randomization*, where a coercer forces a voter to always choose the first choice on the ballot. Due to randomized candidate order, the resulting vote will be randomly distributed. If a specific group of voters are coerced in this way, it can have a disproportionate impact on the election outcome.

If voting rolls are public and voting is not mandatory, this has an effect equivalent to *forced abstention*, wherein a coercer prevents the voter from voting. Schemes that rely on credentialing are also susceptible to coercion by *forced surrender of credentials*.

One way to fully resist forced abstention is to obscure who voted. However, this is difficult to reconcile with the opportunity to verify that only eligible voters have voted (eligibility verifiability), though some schemes achieve both [41].

Moran and Naor [60] provide a strong definition of receipt freeness in which a voter may deviate actively from the protocol in order to convince a coercer that she obeyed. Their model accommodates forced randomization. A scheme is resistant to coercion if the voter can always pretend to have obeyed while actually voting as she likes.

Definition 10. *A voting scheme S is* **coercion resistant** *if the following holds:*

There exists a strategy for a coerced voter V such that, for any strategy adopted by the Coercer C, V is able to cast her intended vote in a manner that is indistinguishable to C from her having followed C's instructions.

Coercion resistance is defined in [45] to include receipt freeness and defence against forced-randomization, forced abstention and the forced surrender of credentials. More general definitions include [56], which incorporates all these attacks along with Moran and Naor's notion of a coercion resistance strategy.

Note that if the coercer can monitor the voter throughout the vote casting period, then resistance is futile. For in-person voting, we assume that the voter is isolated from any coercer while she is in the booth (although this is questionable in the era of mobile phones). For remote voting, we need to assume that voters will have some time when they can interact with the voting system (or the credential-granting system) unobserved.

More Coercion Considerations Some authors have tried to provide some protection against coercion without achieving full coercion resistance. *Caveat coercitor* [39] proposes the notion of *coercion evidence* and allows voters to cast multiple votes using the same credential.

Open problems:

- Can we design usable, verifiable, coercion-resistant voting for a remote setting?

4 Other Secure Voting Requirements

In this section we briefly review more general secure voting system requirements such as usability, availability and those resulting from local election regulations.

4.1 Availability

Denial-of-Service (DoS) is an ever-present threat to elections which can be mitigated but never fully eliminated. A simple service outage can disenfranchise voters, and the threat of attack from foreign state-level adversaries is a pressing concern. Indeed, one of the countries that regularly uses Internet voting, Estonia, has been subject to malicious outages [89].

A variant of DoS specific to the context of elections is *selective DoS*, which presents a fundamentally different threat than general DoS. Voting populations are rarely homogeneous, and disruption of service, for instance, in urban (or rural) areas can skew results and potentially change election outcomes. If DoS cannot be entirely eliminated, can service standards be prescribed so that if an outcome falls below the standards it is vacated? Should these standards be dependent on the reported margin of victory? What, if any, recovery methods are possible? Because elections are more vulnerable to minor perturbations than most other settings, selective DoS is a concern which cannot be ignored.

4.2 Usability

A voting system must be *usable* by voters, poll-workers, election officials, observers, and so on. Voters who may not be computer literate—and sometimes not literate at all—should be able to vote with very low error rates. Although some error is regarded as inevitable, it is also critical that the interface not drive errors in a particular direction. For instance, a list of candidates that crosses a page boundary could cause the candidates on the second page to be missed. Whatever security mechanisms we add to the voting process should operate without degrading usability, otherwise the resulting system will likely be unacceptable. A full treatment of usability in voting is beyond the scope of this paper. However, we note that E2E-V systems (and I-voting systems, even when not E2E-V) add additional processes for voters and poll workers to follow. If verification processes can't be used properly by real voters, the outcome will not be properly verified. One great advantage of statistical audits is to shift complexity from voters to auditors.

Open problems:

- How effectively can usability be integrated into the design process of a voting system?
- How can we ensure full E2E-V, coercion resistance, etc., in a usable fashion?

4.3 Local Regulatory Requirements

A variety of other mechanical requirements are often imposed by legal requirements that vary among jurisdictions. For example:

- Allowing voters to "write-in" vote for a candidate not listed on the ballot.
- Mandating the use of paper ballots (in some states without unique identifying marks or serial numbers; in other states *requiring* such marks)
- Mandating the use of certain social choice functions (see section on Complex Election Methods below).
- Supporting absentee voting.
- Requiring or forbidding that "ballot rotation" be used (listing the candidates in different orders in different jurisdictions).

- Requiring that voting equipment be certified under government guidelines.

Newer electronic and I-voting systems raise important policy challenges for real-world adoption. For example, in STAR-Vote [7], there will be multiple copies of every vote record: mostly electronic records, but also paper records. There may be instances where one is damaged or destroyed and the other is all that remains. When laws speak to retention of “the ballot”, that term is no longer well-defined. Such requirements may need to be adapted to newer voting systems.

Complex Election Methods Many countries allow voters to *select, score, or rank* candidates or parties. Votes can then be tallied in a variety of complex ways [19,76]. None of the requirements for privacy, coercion-resistance, or the provision of verifiable evidence change. However, many tools that achieve these properties for traditional ”first-past-the-post” elections need to be redesigned.

An election method might be complex at the voting or the tallying end. For example, party-list methods such as D’Hondt and Sainte-Laguë have simple voting, in which voters select their candidate or party, but complex proportional seat allocation. Borda, Range Voting, and Approval Voting allow votes to be quite expressive but are simple to tally by addition. Condorcet’s method and related functions [80,88] can be arbitrarily complex, as they can combine with any social choice function. Instant Runoff Voting (IRV) and the Single Transferable Vote (STV) are both expressive and complicated to tally. This makes for several challenges.

Open problems:

- Which methods for cast-as-intended verification (e.g. code voting [26]) work for complex voting schemes?
- How do Risk-limiting audits apply to complex schemes? (See Section 5.2)
- How can complex ballots mitigate failures of the *short ballot assumption* [71]?
- Can we achieve everlasting privacy for complex elections?

5 How can we secure voting?

The goal of this section is to provide a state-of-the-art picture of current solutions to voting problems and ongoing voting research, to motivate further work on open problems, and to define clear directions both in research and election policy.

5.1 The Role of Paper and Ceremonies

Following security problems with direct-recording electronic voting systems (DREs) noted in [59,21,34,92] and others, many parts of the USA returned to the use of paper ballots. If secure custody of the paper ballots is assumed, paper provides durable *evidence* required to determine the correctness of the election outcome. For this reason, when humans vote from untrusted computers, cryptographic voting system specifications often use paper for security, included in the notions of dispute-freeness, dispute resolution, collection accountability and accountability [54] (all as defined in Section 2).

Note that the standard approach to dispute resolution, based on non-repudiation, cannot be applied to the voting problem in the standard fashion, because the human voter does not have the ability to check digital signatures or digitally sign the vote (or other messages that may be part of the protocol) unassisted.

Dispute-freeness or accountability are often achieved in a polling place through the use of cast paper ballots, and the evidence of their chain of custody (e.g., wet-ink signatures). Paper provides an interface for data entry for the voter—not simply to enter the vote, but also to enter other messages that the protocol might require—and data on unforgeable paper serves many of the purposes of digitally signed data. Thus, for example, when a voter marks a *Prêt à Voter* [74] or *Scantegrity* [27] ballot, she is providing an instruction that the voting system cannot pretend was something else. The resulting vote encryption has been physically committed to by the voting system—by the mere act of printing the ballot—before the voter "casts" her vote.

Physical ceremony, such as can be witnessed while the election is ongoing, also supports verifiable cryptographic election protocols (see Section 5.3). Such ceremonies include the verification of voter credentials, any generation of randomness if required for the choice between cast and audit, any vote-encryption-verification performed by election officials, etc.

The key aspect of these ceremonies is the chance for observers to see that they are properly conducted.

Open problems:

- Can we achieve dispute-resolution or -freeness without the use of paper and physical ceremony?

5.2 Statistics and Auditing

Two types of Risk Limiting Audits have been devised: *ballot polling* and *comparison* [57,12,83]. Both types continue to examine random samples of ballots until either there is strong statistical evidence that the outcome is correct, or until there has been a complete manual tally. "Strong statistical evidence" means that the p-value of the hypothesis that the outcome is incorrect is at most α, within tolerable risk.

Both methods rely on the existence of a *ballot manifest* that describes how the audit trail is stored. Selecting the random sample can include a public ceremony in which observers contribute by rolling dice to seed a PRNG [31].

Ballot-polling audits examine random samples of individual ballots. They demand almost nothing of the voting technology other than the reported outcome. When the reported outcome is correct, the expected number of ballots a ballot-polling audit inspects is approximately quadratic in the reciprocal of the (true) margin of victory, resulting in large expected sample sizes for small margins.

Comparison audits compare reported results for randomly selected subsets of ballots to manual tallies of those ballots. Comparison audits require the voting system to commit to tallies of subsets of ballots ("clusters") corresponding to identifiable physical subsets of the audit trail. Comparison audits have two parts: confirm that the outcome computed from the commitment matches the reported outcome, and check the accuracy of randomly selected clusters by manually inspecting the corresponding subsets of the audit trail. When the reported cluster tallies are correct, the number of clusters a comparison audit inspects is approximately linear in the reciprocal of the reported margin. The efficiency of comparison audits also depends approximately linearly on the size of the clusters. Efficiency is highest for clusters consisting of individual ballots: individual cast vote records. To audit at the level of individual ballots requires the voting system to commit to the interpretation of each ballot in a way that is linked to the corresponding element of the audit trail.

In addition to RLAs, auditing methods have been proposed with Bayesian [70] or heuristic [69] justifications.

All post-election audits implicitly assume that the audit trail is adequately complete and accurate that a full manual count would reflect the correct contest outcome. *Compliance audits* are designed to determine whether there is convincing evidence that the audit trail was curated well, by checking ballot accounting, registration records, pollbooks, election procedures, physical security of the audit trail, chain of custody logs, and so on. *Evidence-based elections* [86] combine compliance audits and risk-limiting audits to determine whether the audit trail is adequately accurate, and if so, whether the reported outcome is correct. If there is not convincing evidence that the audit trail is adequately accurate and complete, there cannot be convincing evidence that the outcome is correct.

Audits in Complex Elections Generally, in traditional and complex elections, whenever an election margin is known and the infrastructure for a comparison audit is available, it is possible to conduct a rigorous risk-limiting comparison audit. This motivates many works on practical margin computation for IRV [58,25,79,18].

However, such an audit for a complex election may not be efficient, which motivates the extension of Stark's *sharper discrepancy measure* to D'Hondt and related schemes [85]. For Schulze and some related schemes, neither efficient margin computation nor any other form of RLA is known (see [43]); a Bayesian audit [70,28] may nonetheless be used when one is able to specify suitable priors.

Open problems:

- Can comparison audits for complex ballots be performed without exposing voters to "Italian" attacks?
- Can risk-limiting or other sound statistical audits be developed for systems too complex to compute margins efficiently?
- Can the notion of RLAs be extended to situations where physical evidence is not available (i.e. Internet voting)?

5.3 Cryptographic Tools and Designs

Major Approaches to Voting Cryptography Typically E2E-V involves providing each voter with a *protected receipt*—an encrypted or encoded version of their vote—at the time the vote is cast. The voter can later use her receipt to check whether her vote is included correctly in the tabulation process. Furthermore, given the set of encrypted votes (as well as other relevant information, like the public keys), the tabulation is *universally verifiable*: anyone can check whether it is correct. To achieve this, most E2E-V systems rely on a public bulletin board, where the set of encrypted ballots is published in an append-only fashion.

The votes can then be turned into a tally in one of two main ways. *Homomorphic encryption* schemes [30,16] allow the tally to be produced on encrypted votes. *Verifiable shuffling* transforms a list of encrypted votes into a shuffled list that can be decrypted without the input votes being linked to the (decrypted) output. There are efficient ways to prove that the input list exactly matches the output [77,63,87,6,40].

Techniques for Cast-as-Intended Verification How can a voter verify that her cast vote is the one she wanted? *Code Voting*, first introduced by Chaum [26], gives each voter a sheet of codes for each candidate. Assuming the code sheet is valid, the voter can cast a vote on an untrusted machine by entering the code corresponding to her chosen candidate and waiting to receive the correct confirmation code. Modern interpretations of code voting include [93,44,73].

The alternative is to ask the machine to encrypt a vote directly, but verify that it does so correctly. Benaloh [9] developed a simple protocol to enable vote encryption on an untrusted voting machine. A voter uses a voting machine to encrypt any number of votes, and casts only one of these encrypted votes. All the other votes may be "audited" by the voter. If the encryption is audited, the voting system provides a proof that it encrypted the vote correctly, and the proof is public. The corresponding ballot cannot be cast as the correspondence between the encryption and the ballot is now public, and the vote is no longer secret. Voters take home receipts corresponding to the encryptions of their cast ballots as well as any ballots that are to be audited. They may check the presence of these on a bulletin board, and the correctness proofs of the audited encryptions using software obtained from any of several sources. However, even the most diligent voters need only check that their receipts match the public record and that any ballots selected for audit display correct candidate selections. The correctness proofs are part of the public record that can be verified by any individual or observer that is verifying correct tallying.

A rigorous understanding of E2E-V protocols In addition to the work of Adida on assisted-human interactive proofs (AHIPs, see [1]), there has been some work on a rigorous understanding of one or more properties of single protocols, including the work of Moran and Naor [62,61] and Küsters et al. [54].

There have also been formalizations of voting protocols with human participants, such as by Moran and Naor [61] (for a polling protocol using tamper-evident seals on envelopes) and Kiayias *et al.* [51]. However, there is no one model that is sufficient for the rigorous understanding of the prominent protocols used/proposed for use in real elections. The absence of proofs has led to the overlooking of vulnerabilities in the protocols in the past, see [47,49,48,38].

Many systems use a combination of paper, cryptography, and auditing to achieve E2E-V in the polling place, including Markpledge [64,3], Moran and Naor's scheme [60], Prêt à Voter [74], Scantegrity II [23], Wombat [72,8], STAR-Vote [7] and Demos [52]. Their properties are summarised more thoroughly in the following section.

The cryptographic literature has numerous constructions of end-to-end verifiable election schemes (e.g., [36,65,74,23,71,64,72,78,7,44]). There are also various detailed descriptions of what it means to verify the correctness of the output of E2E-V systems (e.g., [52,15,60]). Others have attempted to define alternative forms of the E2E-V properties [66,32,54]. There are also less technical explanations of E2E-V intended for voters and election officials [14,91].

Open problems:

- Can we develop a rigorous model for human participants and the use of paper and ceremonies in cryptographic voting protocols?
- Can we examine rigorously the combination of statistical and cryptographic methods for election verification?

Techniques for Coercion Resistance Some simple approaches to coercion resistance have been suggested in the literature. These include allowing multiple votes with only the last counting and allowing in-person voting to override remotely cast votes (both used in Estonian, Norwegian, and Utah elections [82,37,17]). It is not clear that this mitigates coercion at all. Alarm codes can also be provided to voters: seemingly real but actually fake election credentials, along with the ability for voters to create their own fake credentials. Any such approach can be considered a partial solution at best, particularly given the usability challenges.

One voting system, *Civitas* [29], based on a protocol by Juels, Catalano and Jakobsson [45], allows voters to vote with fake credentials to lead the coercive adversary into believing the desired vote was cast. Note that the protocol must enable universal verification of the tally from a list of votes cast with both genuine and fake credentials, proving to the verifier that only the ones with genuine credentials were tallied, without identifying which ones they were.

Open problems:

- Can we develop cryptographic techniques that provide fully coercion resistant remote voting?

Cryptographic Solutions in Complex Elections Cast-as-intended verification based on creating and then challenging a vote works regardless of the scheme (*e.g.* Benaloh challenges). Cut-and-choose based schemes such as Prêt à Voter and Scantegrity II need to be modified to work.

Both uses of end-to-end verifiable voting schemes in government elections, the Takoma Park run of Scantegrity II and the Victorian run of Prêt à Voter, used IRV (and one used STV). Verifiable IRV/STV counting that doesn't expose individual votes to the Italian attack has been considered [13], but may not be efficient enough for use in large elections in practice, and was not employed in either practical implementation.

Open problems:

- Is usable cast-as-intended verification for complex voting methods possible?

Table 1 summarizes how the various election systems satisfy the defintions given in Sections 2, 3, and 4.

6 A Look Ahead

Voting has always used available technology, whether pebbles dropped in an urn or marked paper put in a ballot box; it now uses computers, networks, and cryptography. The core requirement, to provide public evidence of the right result from secret ballots, hasn't changed in 2500 years.

Computers can improve convenience and accessibility over plain paper and manual counting. In the polling place there are good solutions, including Risk Limiting Audits and end-to-end verifiable systems. These must be more widely deployed and their options for verifying the election result must actually be used.

Many of the open problems described in this paper—usable and accessible voting systems, dispute resolution, incoercibility—come together in the challenge of a remote

	fielded	coercion resistance	everlasting privacy	software independence	take-home evidence	ballot cast assurance	collection accountable	verifiably cast-as-intended	verifiably collected-as-cast	verifiable counted-as-collected	paper/electronic/hybrid	write-ins supported	preferential ballots supported
Poll-site techniques in widespread use													
Hand-counted in-person paper	●	●	●	●	○	○	○	●	○	●[7]	p	●	●
Optical-scan in-person paper	●	●	●	●	○	○	○	●	○	●[7]	h	●	●
DRE (with paper audit trail)	●	●	○	●	○	○	○	◐[7]	●	●[7]	h	●	●
Paperless DRE	●	●	○	○	○	○	○	○	○	○	e	●	●
Poll-site systems from research													
Prêt-à-voter [74]	●[1]	●	○	●	●	◐[12]	●	●	●	●	h	○	●
Scantegrity [27]	●[1]	●	○	●	●	◐[12]	●	●	●	●	h	◐	◐
STAR-Vote [7]	◐[2]	●	◐[6]	●	●	◐[13]	◐[14]	●	●	●	h	○	○
Wombat [72]	◐[3]	●	◐[6]	●	●	◐[13]	◐[14]	●	●	●	h	○	○
VeriScan [11]	○	●	◐[6]	●	●	◐[13]	◐[14]	●	●	●	h	○	◐
Scratch and Vote [4]	○	●	○	●	●	◐[12]	●	●	●	●	h	◐	◐
MarkPledge [64]	○	●	◐[6]	●	●	●	●	●	●	●	e	○	○
ThreeBallot [67]	○	●	●	●	●	◐[12]	●	●	◐	●	h	○	○
Remote voting systems and techniques													
Helios [2]	◐[3]	○[5]	◐[6]	●	●[8]	◐	◐[7]	●	●	●	e	○	○
Remotegrity [93]	◐[1]	○	◐[6]	●	●	○	●	●	●	●	h	○	◐
Civitas [29]	◐[3]	●	◐[6]	○	●[8]	○	◐	●	●	●	e	◐	●
Selene [75]	○	●	◐[6]	●	○	◐[7]	●	●	●	●	e	○	●
Norway [37]	●	○[5]	○	○	○	○	◐[11]	○	○	◐[7]	e	●	●
Estonia [82]	●	○[5]	○	○	◐[9]	○	○	○	○	○	e	○	○
iVote [42]	●[4]	○[5]	○	○	○	○	○	◐[10]	◐[10]	○	e	○	●
Paper ballots returned by postal mail	●	○	●	●	○	○	○	○	○	●[7]	p	●	●

● = provides property ○ = does not provide property ◐ = provides property with provisions

[1] Used in small trial elections
[2] Pending deployment
[3] Used in private sector elections
[4] Absentee voting only
[5] Allows multiple voting
[6] Possible with PPAT
[7] With sufficient auditing
[8] Receipts sent by email
[9] Temporary email receipt
[10] Queryable (phone system)
[11] Queryable (code sheets)
[12] Enhanced with pre- and post-election auditing
[13] Enhanced with auditing during elections
[14] To the extent the paper resists forgery

Table 1: **Applying our threat model to fielded and proposed voting schemes** — Note that certain features like credentialing and availability are excluded, as these factors impact all systems in roughly equivalent ways. The Utah system has not been made available for rigorous security analysis, and is excluded.

voting system that is verifiable and usable without supervision. The open problem of a system specification that (a) does not use any paper at all and (b) is based on a simple procedure for voters and poll workers, will motivate researchers for a long time. Perhaps a better goal is a hybrid system combining paper evidence with some auditing or cryptographic verification.

Research in voting brings together knowledge in many fields—cryptography, systems security, statistics, usability and accessibility, software verification, elections, law and policy to name a few—to address a critical real-world problem.

The peaceful transfer of power depends on confidence in the electoral process. That confidence should not automatically be given to any outcome that seems plausible—it must be earned by producing evidence that the election result is what the people chose. Insisting on evidence reduces the opportunities for fraud, hence bringing greater security to citizens the world over.

Acknowledgments

This work was supported in part by the U.S. National Science Foundation awards CNS-1345254, CNS-1409505, CNS-1518888, CNS-1409401, CNS-1314492, and CNS-1421373, the Center for Science of Information STC (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370, the Maryland Procurement Office under contract H98230-14-C-0127, and FNR Luxembourg under the PETRVS Mobility grant.

References

1. B. Adida. *Advin Cryptographic Voting Systems*. PhD thesis, MIT, July 2006.
2. B. Adida. Helios: Web-based open-audit voting. 17th USENIX Security Symp., Aug. 2008. vote.heliosvoting.org.
3. B. Adida and C. A. Neff. Efficient receipt-free ballot casting resistant to covert channels. *IACR Cryptology ePrint Archive*, 2008, 2008.
4. B. Adida and R. L. Rivest. Scratch and Vote: Self-contained paper-based cryptographic voting. In *ACM Workshop on Privacy in the Electronic Society*, 2006.
5. J. Alwen, R. Ostrovsky, H.-S. Zhou, and V. Zikas. Incoercible Multi-Party Computation and Universally Composable Receipt-Free Voting. In *Advin Cryptology*. 2015.
6. S. Bayer and J. Groth. Efficient zero-knowledge argument for correctness of a shuffle. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2012.
7. S. Bell, J. Benaloh, M. D. Byrne, D. DeBeauvoir, B. Eakin, G. Fisher, P. Kortum, N. McBurnett, J. Montoya, M. Parker, O. Pereira, P. B. Stark, D. S. Wallach, and M. Winn. STAR-vote: A secure, transparent, auditable, and reliable voting system. *USENIX Journal of Election Technology and Systems*, 1(1), Aug. 2013.
8. J. Ben-Nun, N. Fahri, M. Llewellyn, B. Riva, A. Rosen, A. Ta-Shma, and D. Wikström. A new implementation of a dual (paper and cryptographic) voting system. In *5th International Conference on Electronic Voting*, 2012.
9. J. Benaloh. Simple verifiable elections. In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2006.
10. J. Benaloh. Ballot casting assurance via voter-initiated poll station auditing. In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2007.
11. J. Benaloh. Administrative and public verifiability: Can we have both? In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2008.
12. J. Benaloh, D. Jones, E. Lazarus, M. Lindeman, and P. B. Stark. Soba: Secrecy-preserving observable ballot-level audit. In *proc. Proc. USENIXAccurate Electronic Voting Technology Workshop*, 2011.
13. J. Benaloh, T. Moran, L. Naish, K. Ramchen, and V. Teague. Shuffle-sum: Coercion-resistant verifiable tallying for STV voting. *IEEE Transactions on Information Forensics and Security*, 4(4), 2009.
14. J. Benaloh, R. Rivest, P. Y. Ryan, P. Stark, V. Teague, and P. Vora. End-to-end verifiability, 2015. arXiv:1504.03778.
15. J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections. In *26th ACM Symp. on Theory of Computing*, 1994.
16. J. D. C. Benaloh. *Verifiable Secret-ballot Elections*. PhD thesis, Yale, 1987. AAI8809191.
17. M. Bernhard. What happened in the Utah GOP caucus. https://mbernhard.com/Utahvoting.pdf.
18. M. Blom, P. J. Stuckey, V. J. Teague, and R. Tidhar. Efficient computation of exact IRV margins, 2015. arXiv:1508.04885.
19. S. Brams. *Mathematics and democracy*. Princeton Univ. Press, 2008.
20. J. A. Calandrino, J. A. Halderman, and E. W. Felten. Machine-assisted election auditing. In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2007.

21. California Secretary of State's Office. Top-to-bottom review of electronic voting systems, 2007. http://wwws.os.ca.gov/elections/voting-systems/oversight/top-bottom-review/.
22. R. Canetti and R. Gennaro. Incoercible multiparty computation. In *37th IEEE Symp. on Foundations of Computer Science*, 1996.
23. R. Carback, D. Chaum, J. Clark, J. Conway, A. Essex, P. S. Herrnson, T. Mayberry, S. Popoveniuc, R. L. Rivest, E. Shen, A. T. Sherman, and P. L. Vora. Scantegrity II municipal election at Takoma Park: The first E2E binding governmental election with ballot privacy. In *18th USENIX Security Symp.*, Aug. 2010.
24. Carter Center. Expert study mission report—Internet voting pilot: Norway's 2013 parliamentary elections, Mar. 2014. http://www.regjeringen.no/upload/KRD/Kampanjer/valgportal/valgobservatorer/2013/Rapport_Cartersenteret2013.pdf.
25. D. Cary. Estimating the margin of victory for instant-runoff voting. In *USENIX/ACCURATE Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2011.
26. D. Chaum. SureVote: Technical overview. In *IAVoSS Workshop on Trustworthy Elections*, 2001.
27. D. Chaum, R. Carback, J. Clark, A. Essex, S. Popoveniuc, R. L. Rivest, P. Y. A. Ryan, E. Shen, and A. T. Sherman. Scantegrity II: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes. In *USENIX/ACCURATE Electronic Voting Workshop*, Aug. 2008.
28. B. Chilingirian, Z. Perumal, R. L. Rivest, G. Bowland, A. Conway, P. B. Stark, M. Blom, C. Culnane, and V. Teague. Auditing australian senate ballots. *arXiv preprint arXiv:1610.00127*, 2016.
29. M. Clarkson, S. Chong, and A. C. Myers. Civitas: A secure remote voting system. Technical report, Cornell University Computing and Information Science Technology Report, May 2007. http://www.truststc.org/pubs/545.html.
30. J. B. (Cohen) and M. J. Fischer. A robust and verifiable cryptographically secure election scheme. In *26th Annual Symp. on Foundations of Computer Science*, 1985.
31. A. Cordero, D. Wagner, and D. Dill. The role of dice in election audits – extended abstract. In *IAVoSS Workshop On Trustworthy Elections*, 2006.
32. V. Cortier, D. Galindo, R. Küsters, J. Müller, and T. Truderung. Verifiability Notions for E-Voting Protocols. Technical report, Technical Report 2016/287, Cryptology ePrint Archive, 2016. Available at http://eprint. iacr. org/2016/287.
33. S. Delaune, S. Kremer, and M. Ryan. Verifying privacy-type properties of electronic voting protocols: A taster. In *Towards Trustworthy Elections*. 2010.
34. A. J. Feldman, J. A. Halderman, and E. W. Felten. Security analysis of the Diebold AccuVote-TS voting machine. In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2007.
35. E. Fieser. People Openly Sell Votes for $20 in the Dominican Republic. http://www.bloomberg.com/news/articles/2016-05-16/people-openly-sell-their-votes-for-20-in-the-dominican-republic, May 2016.
36. K. Fisher, R. Carback, and A. T. Sherman. Punchscan: Introduction and system definition of a high-integrity election system. In *IAVoSS Workshop on Trustworthy Elections*, 2006.
37. K. Gjøsteen. The Norwegian Internet voting protocol. In *3rd International Conference on E-Voting and Identity*, 2011.
38. M. Gogolewski, M. Klonowski, P. Kubiak, M. Kutylowski, A. Lauks, and F. Zagórski. Kleptographic attacks on e-voting schemes. In *International Conference on Emerging trends in Information and Communication Security*, 2006.
39. G. S. Grewal, M. D. Ryan, S. Bursuc, and P. Y. Ryan. Caveat coercitor: Coercion-evidence in electronic voting. In *34th IEEE Symp. on Security and Privacy*, 2013.
40. J. Groth. A verifiable secret shuffle of homomorphic encryptions. *Journal of Cryptology*, 23(4), 2010.
41. R. Haenni and O. Spycher. Secure Internet Voting on Limited Devices with Anonymized DSA Public Keys. Aug. 2011.
42. J. A. Halderman and V. Teague. The New South Wales iVote system: Security failures and verification flaws in a live online election. In *5th International Conference on E-Voting and Identity*, Aug. 2015.
43. L. A. Hemaspaandra, R. Lavaee, and C. Menton. Schulze and ranked-pairs voting are fixed-parameter tractable to bribe, manipulate, and control. In *International Conference on Autonomous Agents and Multiagent Systems*, 2013.
44. R. Joaquim, C. Ribeiro, and P. Ferreira. Veryvote: A voter verifiable code voting system. In *International Conference on E-Voting and Identity*, 2009.
45. A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant Electronic Elections. In *ACM Workshop on Privacy in the Electronic Society*, Nov. 2005.
46. T. Kaczmarek, J. Wittrock, R. Carback, A. Florescu, J. Rubio, N. Runyan, P. L. Vora, and F. Zagórski. Dispute Resolution in Accessible Voting Systems: The Design and Use of Audiotegrity. In *E-Voting and Identify - 4th International Conference*, 2013.
47. C. Karlof, N. Sastry, and D. Wagner. Cryptographic voting protocols: A systems perspective. In *14th USENIX Security Symp.*, Aug. 2005.
48. J. Kelsey, A. Regenscheid, T. Moran, and D. Chaum. Attacking paper-based E2E voting systems. In *Towards Trustworthy Elections, New Directions in Electronic Voting*, 2010.
49. S. Khazaei and D. Wikström. Randomized partial checking revisited. In *Topics in Cryptology, CT-RSA 2013*. 2013.
50. A. Kiayias and M. Yung. Self-tallying elections and perfect ballot secrecy. In *5th International Workshop on Practice and Theory in Public Key Cryptosystems*, 2002.
51. A. Kiayias, T. Zacharias, and B. Zhang. Ceremonies for end-to-end verifiable elections. *IACR Cryptology ePrint Archive*, 2015, 2015.
52. A. Kiayias, T. Zacharias, and B. Zhang. End-to-end verifiable elections in the standard model. In *Advin Cryptology—EUROCRYPT 2015*. 2015.
53. S. Kremer, M. Ryan, and B. Smyth. Election verifiability in electronic voting protocols. In *4th Benelux Workshop on Information and System Security*, Nov. 2009.
54. R. Küsters, T. Truderung, and A. Vogt. Accountability: Definition and relationship to verifiability. In *17th ACM Conference on Computer and Communications Security*, 2010.
55. R. Küsters, T. Truderung, and A. Vogt. Verifiability, privacy, and coercion-resistance: New insights from a case study. In *32nd IEEE Symp. on Security and Privacy*, 2011.

56. R. Küsters, T. Truderung, and A. Vogt. A game-based definition of coercion resistance and its applications. *Journal of Computer Security*, 20(6), 2012.
57. M. Lindeman, P. B. Stark, and V. S. Yates. BRAVO: Ballot-polling risk-limiting audits to verify outcomes. In *USENIX Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2012.
58. T. R. Magrino, R. L. Rivest, E. Shen, and D. Wagner. Computing the margin of victory in IRV elections. In *USENIX Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2011.
59. P. McDaniel et al. EVEREST: Evaluation and validation of election-related equipment, standards and testing, Dec. 2007. http://www.patrickmcdaniel.org/pubs/everest.pdf.
60. T. Moran and M. Naor. Receipt-free universally-verifiable voting with everlasting privacy. In *Advin Cryptology*, 2006.
61. T. Moran and M. Naor. Basing cryptographic protocols on tamper-evident seals. *Theoretical Computer Science*, 2010.
62. T. Moran and M. Naor. Split-ballot voting: Everlasting privacy with distributed trust. *ACM Transactions on Information and System Security*, 13(2), 2010.
63. C. A. Neff. A verifiable secret shuffle and its application to e-voting. In *ACM Conference on Computer and Communications Security*. ACM, 2001.
64. C. A. Neff. Practical high certainty intent verification for encrypted votes, 2004. http://www.votehere.net/vhti/documentation.
65. S. Popoveniuc and B. Hosp. An Introduction to Punchscan. In *IAVoSS Workshop on Trustworthy Elections*, Aug. 2006.
66. S. Popoveniuc, J. Kelsey, A. Regenscheid, and P. L. Vora. Performance requirements for end-to-end verifiable elections. In *USENIX Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2010.
67. R. L. Rivest. The Three Ballot voting system, 2006. https://people.csail.mit.edu/rivest/Rivest-TheThreeBallotVotingSystem.pdf.
68. R. L. Rivest. On the notion of "software independence" in voting systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2008.
69. R. L. Rivest. DiffSum: A simple post-election risk-limiting audit. *CoRR abs/1509.00127*, 2015.
70. R. L. Rivest and E. Shen. A Bayesian method for auditing elections. In *USENIX Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2012.
71. R. L. Rivest and W. D. Smith. Three voting protocols: ThreeBallot, VAV, and Twin. In *USENIX/ACCURATE Electronic Voting Technology Workshop*, Aug. 2007.
72. A. Rosen, A. Ta-shma, B. Riva, and J. Y. Ben-Nun. Wombat voting system. http://www.wombat-voting.com.
73. P. Y. A. Ryan. Prêt à voter with confirmation codes. In *2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Aug. 2011.
74. P. Y. A. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia. Prêt à Voter: A voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security*, 4(4), 2009.
75. P. Y. A. Ryan, P. B. Roenne, and V. Iovino. Selene: Voting with Transparent Verifiability and Coercion-Mitigation. Cryptology ePrint Archive, Report 2015/1105, 2015. http://eprint.iacr.org/.
76. D. G. Saari. *Geometry of voting*. 2012.
77. K. Sako and J. Killian. Receipt-free mix-type voting scheme: A practical solution to the implementation of a voting booth. In *Advanced in Cryptology—EUROCRYPT '95*, 1995.
78. D. R. Sandler, K. Derr, and D. S. Wallach. VoteBox: A tamper-evident, verifiable electronic voting system. In *17th USENIX Security Symp.*, July 2008.
79. A. D. Sarwate, S. Checkoway, and H. Shacham. Risk-limiting audits and the margin of victory in nonplurality elections. *Statistics, Politics and Policy*, 4(1), 2013.
80. M. Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2), 2011.
81. B. Smyth, M. Ryan, S. Kremer, and M. Kourjieh. Towards automatic analysis of election verifiability properties. In *Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*. 2010.
82. D. Springall, T. Finkenauer, Z. Durumeric, J. Kitcat, H. Hursti, M. MacAlpine, and J. A. Halderman. Security analysis of the Estonian Internet voting system. In *21st ACM Conference on Computer and Communications Security*, 2014.
83. P. Stark. Conservative statistical post-election audits. *Annals of Applied Statistics*, 2008.
84. P. B. Stark. Super-simple Simultaneous Single-ballot Risk-limiting Audits. In *Proceedings of the 2010 International Conference on Electronic Voting Technology/Workshop on Trustworthy Elections*, Aug. 2010.
85. P. B. Stark and V. Teague. Verifiable European elections: Risk-limiting audits for d'hondt and its relatives. *USENIX Journal of Election Technology and Systems*, 3(1), 2014.
86. P. B. Stark and D. A. Wagner. Evidence-Based Elections. *IEEE Security and Privacy Magazine*, 10(05), 2012.
87. B. Terelius and D. Wikström. Proofs of restricted shuffles. In *Progress in Cryptology–AFRICACRYPT 2010*. 2010.
88. T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3), 1987.
89. I. Traynor. Russia accused of unleashing cyberwar to disable Estonia, May 2007. http://www.theguardian.com/world/2007/may/17/topstories3.russia.
90. D. Unruh and J. Müller-Quade. Universally composable incoercibility. In *Advin Cryptology*. 2010.
91. U.S. Vote Foundation and Galois. The future of voting: End-to-end verifiable Internet voting specification and feasibility assessment study, 2015. https://www.usvotefoundation.org/sites/default/files/E2EVIV_full_report.pdf.
92. D. Wallach. Security and Reliability of Webb County's ES&S Voting System and the March 06 Primary Election. *Expert Report in Flores v. Lopez*, 2006.
93. F. Zagórski, R. Carback, D. Chaum, J. Clark, A. Essex, and P. L. Vora. Remotegrity: Design and use of an end-to-end verifiable remote voting system. In *11th International Conference on Applied Cryptography and Network Security*, 2013.