In Defense of Pseudorandom Sample Selection

Joseph A. Calandrino[*], J. Alex Halderman[*], and Edward W. Felten[*,†]

[*]Center for Information Technology Policy and Dept. of Computer Science, Princeton University
[†]Woodrow Wilson School of Public and International Affairs, Princeton University
{jcalandr,jhalderm,felten}@cs.princeton.edu

Abstract

Generation of random numbers is a critical component of existing post-election auditing techniques. Recent work has largely discouraged the use of all pseudorandom number generators, including cryptographically secure pseudorandom number generators (CSPRNGs), for this purpose, instead recommending the sole use of observable physical techniques. In particular, simple dice rolling has received a great deal of positive attention [4, 6, 9]. The typical justification for this recommendation is that those less comfortable with mathematics prefer a simple, observable technique. This paper takes a contrary view. Simple, observable techniques like dice rolling are not necessarily robust against sleight of hand and other forms of fraud, and attempts to harden them against fraud can dramatically increase their complexity. With simple dice rolling, we know of no techniques that provide citizens with a reasonable means of verifying that fraud did not occur during the roll process. CSPRNGs, used properly, can be simple, robust, and verifiable, and they allow for the use of auditing techniques that might otherwise be impractical. While we understand initial skepticism towards this option, we argue that appropriate use of CSPRNGs would strengthen audit security.

1 Introduction

A number of well-publicized reports have highlighted flaws in existing electronic voting machines (for example, [7, 5], among others), prompting calls for the use of voter-verifiable paper ballots. Because software cannot change a paper ballot already in a ballot box, auditors can examine those ballots following an election to confirm that software errors and fraud did not change the outcome. To avoid the time and expense of counting all paper ballots, auditors can manually review a subset of the ballots to gain high statistical confidence that the electronically tabulated preliminary outcome is correct. Existing law and literature describe numerous auditing techniques. For example, auditors may sample full precincts, individual voting machines, or individual ballots for manual verification.

If an adversary knows or can influence which subset of ballots will be considered during the audit, that individual can commit fraud in a manner that the audit will not detect. For example, if officials will not examine any ballots in a certain precinct, the adversary may cause that precinct to report an altered tally without fear of discovery. Even partial knowledge, such as knowledge of biases in the selection algorithm, may be sufficient to undermine the effectiveness of the audit. Therefore, subset selection must be unpredictable. To achieve unpredictability, existing election auditing proposals rely on random sampling.[1]

Like auditing in general, selection of a random sample may seem straightforward, but it is non-trivial in practice, particularly if certain participants want to cheat. Numerous attacks, from weighted dice to malicious software, can bias a selection process. In addition, randomness is not the only desirable property of sample selection (see Section 3).

A number of recent papers encourage the use of dice rolling techniques for sample selection and discourage any processes involving computers. One even proposes a legal ban on all computer-based pseudorandom number generators, including cryptographically secure pseudorandom number generators (CSPRNGs), in the audit process [6]. In this paper we take a contrary view. Our three primary contributions are:

- We detail limits to the security and practicality of

[1] An auditing process need not involve randomness: for example, a full manual count would require no randomness. Alternatively, if the set of ballots to be verified depends on the preliminary electronic records via a complicated function, we could make it infeasible to modify a sufficiently large set of the ballots to change a race's outcome without detection. Here we assume that randomness is necessary.

dice rolling.

- We address several misconceptions regarding CSPRNGs.

- We propose means of using CSPRNGs for secure, simple, and transparent sample selection.

Overall, we argue that debate on this topic has failed to adequately weigh the benefits that CSPRNGs can offer to election auditing. Used appropriately, CSPRNGs can offer the desired properties as or more effectively than other proposed options, including dice rolling, and they can be invoked without requiring voters to trust that any particular computer is operating correctly.

We focus on fraud rather than general error or reliability issues, but this should not be interpreted as an indication of the relative likelihood or importance of these problems. Rather, someone looking to commit fraud would wish to produce the hardest-to-detect set of errors yielding the desired outcome. Therefore, in detecting worst-case fraud, we implicitly seek all other forms of outcome-altering error.

The remainder of this paper is organized as follows. Section 2 discusses randomness and pseudorandomness. Section 3 describes necessary properties of sample selection. Section 4 describes example selection processes involving dice rolling alone and utilizing CSPRNGs, and Sections 5 compares these processes in the context of our necessary properties. Finally, Section 6 summarizes the discussion.

Note Before proceeding, we wish to note that, if the use of CSPRNGs is banned, simple dice rolling is one option that election officials should consider. Cordero et al. [4] assume that officials would be unwilling or unable to use CSPRNGs properly and, given that assumption, offer a careful analysis of the use of dice rolling and other techniques in sample selection. We choose to critique dice rolling largely due to its popularity, not because it is necessarily worse than other observable options.

2 Randomness and Pseudorandomness

The term "pseudorandom," as used by cryptographers, may be misleading to a non-technical reader. Cryptographers refer to certain processes that generate a sequence of numbers as cryptographically secure pseudorandom number generators (CSPRNGs). Intuitively, a CSPRNG expands a short random value, known as a seed, into a longer sequence of numbers. That sequence is reproducible given the seed, but given a good CSPRNG, a minor change in the seed yields a very different sequence. Therefore, if some minimum portion of the seed will be chosen via an adequately random process, an adversary would be unable to predict the resulting sequence prior to seed generation—even if the adversary can influence the remainder of the seed. Numerous important systems, ranging from military communications systems to the encryption that protects virtually all online transactions, rely on the functionally equivalent security between "cryptographically secure pseudorandom" and "random."

The procedures used to generate the seed are important. Officials should wait to generate a seed until all results have been irrevocably committed. If an adversary can change results after seed generation, he can cheat knowing exactly which audit units will be selected for verification. In general, officials should typically wait to generate a seed until a set of random numbers is necessary. Also, officials should generate a new seed whenever a new set of random numbers is needed. These choices allow us to generate and disclose any seeds publicly, yielding a more transparent and trustworthy process than one relying on the seed's secrecy. We elaborate on seed generation in Section 4.2.

Our discussion in this paper is limited to CSPRNGs and should not be interpreted as an endorsement of all pseudorandom number generators. Certain pseudorandom number generators do not provide the necessary security properties for auditing, and any method of generation—pseudorandom or not—can be detrimental if used inappropriately.[2] We argue only that proper use specifically of CSPRNGs can strengthen the sample selection process.

3 Necessary Properties of Sample Selection

Cordero et al. [4] identify four desirable properties of sample selection for auditing: simplicity, verifiability, robustness, and efficiency. We build on this list and modify several of its definitions, resulting in a list of five properties that we consider necessary:

Unpredictability. An adversary should be unable to acquire any advance knowledge regarding sample selection that allows for a practical advantage in committing undetected fraud.

[2] Rivest [11] proposes the use of a special-purpose, non-cryptographic PRNG in the selection process and independently draws several conclusions similar to those of this paper ([11] also seeks to achieve a slightly different set of properties than this paper and focuses more heavily on a specific PRNG than the general selection process). Indeed, a CSPRNG is not strictly necessary: other special-purpose PRNGs can potentially achieve all necessary properties for election auditing. We cautiously focus on CSPRNGs primarily because, in addition to meeting all necessary properties by definition, several CSPRNGs are widely used, well-evaluated, and well-trusted.

Verifiability. Citizens must have some reasonable means of gaining acceptable confidence that sample selection was not biased or otherwise affected by "untrusted participants." An untrusted participant is not necessarily dishonest but simply someone that a citizen may not fully trust to protect her interests. For example, you would probably not trust a stranger with your life savings—even though that stranger may actually be extremely trustworthy. Like [4], we distinguish between verifiability and observability. Observability does not necessarily make a process verifiable. For example, sleight of hand dice tricks may occur in plain sight.

Robustness. A single honest participant in the number generation process should be sufficient to ensure a fair sample selection even if all other participants are dishonest. Therefore, if a citizen trusts even one participant, that citizen should be able to trust that the selection process is, at minimum, not biased against that citizen's preferred candidate.

Simplicity. The sample selection process must be straightforward and understandable to all participants. More complex processes are more difficult to follow and tend to result in more mistakes. Note the difference between understanding the steps of a process (for example, "flip a coin, record the result, repeat nine additional times") and understanding why the process meets all requirements (for example, "the binomial distribution dictates that the audit will uncover outcome-altering fraud with 95% probability"), the latter of which would rule out virtually all statistical audit processes. In addition, any processes necessary for verifying appropriate sample selection must be reasonably straightforward.

Efficiency. The process must not place unreasonable expectations on the time and labor of participants or observers.

We remove the requirement from [4] that most citizens must fully understand why the process works. While this property would be beneficial, many people do not understand even the statistics underlying the audit process—subtle issues can be tricky even for experts. If the algorithms are publicly available and widely supported by trusted, impartial experts in the field of cryptography, this endorsement would be similar to allowing the statistics community to supply its blessing for a statistical audit process.

4 Two Sample Selection Processes

Assume that voting is complete, and officials are ready to begin the audit process. Typically, all preliminary results must be reported before sample selection can begin. Otherwise, an adversary could wait until sample selection is complete and report incorrect tallies only for audit units that will not be examined. In [3], we propose techniques for beginning the audit process before all preliminary results are reported, but we ignore those techniques here for simplicity.

Prior to beginning sample selection, officials need to sequentially number all audit units, whether precincts, voting machines, or ballots. This may be performed explicitly or implicitly: for example, precinct number three may be the third precinct appearing in an alphabetically ordered list of their names. The critical point here is that this numbering be agreed upon and unchangeable after the sampling process begins to prevent disagreement over which audit unit to examine when a given number is generated.

In this section, we describe two processes for generating random (or pseudorandom) samples. The first relies on dice rolling alone to select a sample. The second method uses dice in combination with a CSPRNG. In the sections that follow, we compare both methods.

4.1 Simple Dice Rolling

We consider the method suggested in Cordero et al. [4] but simplify or omit minor complications for brevity and clarity. Assume that n audit units exist. For each digit of n, officials require a single ten-sided die. If 100-999 audit units exist, officials require three dice; if 1,000-9,999 audit units exist, officials require four dice. Each die corresponds to a single digit of the numbers to be generated and could be numbered as such. For example, the die used to generate the ones digit could be numbered 0, 1, $\ldots$, 9, and the die used to generate the hundreds digit could be numbered 000, 100, $\ldots$, 900. Cordero et al. propose instead using die color to determine the digit a die generates [4]; that option may be preferable in many cases.

We assume that the number of participating officials equals the necessary number of dice and that each official receives a single die. For each number to be generated, officials simultaneously roll their dice and select the corresponding audit unit. If the numbers 7000-300-20-3 arise, audit unit 7,323 will be selected. If the numbers 4000-000-50-2 arise, audit unit 4,052 will be selected. For each additional sample to be drawn, officials must simultaneously roll their dice an additional time. Therefore, if fifty audit units are to be drawn, officials must roll their dice and record the results fifty times.

Officials need to re-roll if they generate the same number multiple times or if the number generated is greater than the number of audit units. For example, officials would re-roll if they generate the number 1,743 twice or if 5,000 audit units exist and the rolls yield the number 6,449. Cordero et al. propose a more efficient method that reduces the number of necessary re-rolls [4]. We omit these details and, to allow for a fair comparison, assume that re-rolls are never necessary.

4.2 Dice Plus CSPRNGs

Our proposed process has three steps. First, certain interested parties—potentially including election officials, party officials, candidates, and concerned citizens—roll dice in private to generate sequences of numbers. Second, officials enter those sequences into a computer that combines them to produce a seed for a CSPRNG. Finally, the computer uses a CSPRNG to select audit units for review.

The first step is entirely manual and occurs before sample selection is to begin. The state gives each participant in the seed generation process a set of five ten-sided dice, with each die having a different color than the four others, and a recording sheet. The recording sheet contains five columns, each labeled with a different die's color, and ten rows. Officials may distribute these items during training sessions or via mail. If desired, a participant may replace any of the dice with ones from any source that the participant prefers, including trusted party officials. Once a set of dice is chosen, each participant must depart to a private location, such as her home. The participant should roll her dice in unison and record each die's value in the record sheet column corresponding to its color.[3] This roll-and-record process should be repeated nine additional times—ten in all—for a total of fifty generated values. Once complete, the participant should sign and fold her record sheet, place it in an envelope, and sign the sealed envelope. Upon arrival at the sample selection site, she should deposit her envelope in a transparent, publicly observable locked box at the site.[4] She should not share her generated numbers with others until her sheet is revealed.

For the remainder of the process, a single computing device—a general-purpose computer or a simple special-purpose device—loaded with sample selection software will be necessary. The sample selection software will create the seed from user-input data and use a pre-programmed CSPRNG with the seed to select a sample. The selection process does not require that citizens place any trust in this computer or the software loaded to it: the process allows anyone to verify that the computer's output is correct.

The second step of the process occurs when sample selection is to begin and results in the production of a seed. Officials unlock the box and remove the envelopes in alphabetical order by participant name. When an envelope is removed, officials open it and unfold the record sheet in the presence of observers, who may photograph or videotape the sheets, envelopes, and general process. As soon as is practically feasible and before election results are confirmed, the contents of the sheets should be made easily accessible to the public. Officials enter the sequence of numbers contained on a sheet row-by-row into the computing device immediately after the sheet's envelope is opened. Once sequences from all sheets have been entered, the computer combines the values into a seed for a CSPRNG.[5]

The final step of the process occurs immediately after the seed becomes available. The computer uses the CSPRNG, the number of audit units, and the seed to select a sample of audit units based on the number assigned to each unit.[6] As with the record sheets, observers may view, copy, or photograph the generated numbers from the medium on which they are displayed. This set of numbers should be made public before election results are confirmed.

States may wish to allow (or encourage) participants to generate a limited number of additional record sheets using a method of the participant's choice, such as coin flipping. Given appropriate methods for combining numbers into a seed, adding additional record sheets cannot decrease the randomness of the resulting seed. Therefore, even if someone were to place non-random values on her record sheet, the resulting seed would be no less random than if the manipulated values were excluded.

If a new sample is needed at another point during the audit process, officials may need to generate a new seed. If anticipated, it may be beneficial for participants to generate multiple record sheets prior to sample selection. In practice, however, we anticipate that usually only a single set of numbers will be necessary.

Clear alternatives exist to the described proposal, including using fewer or more dice. As with any generation

[3] If concerns exist that participants may write their numbers ambiguously (whether purposely or not), fill-in-the-bubble sheets or other designs may mitigate the issue. To help prevent such problems, usability experts should review record sheets prior to their use.

[4] We assume that the period of time between depositing the envelope in the box and the sample selection process itself is short enough that the participant can reasonably monitor the box. If the envelope is to remain in the box for an extended period of time, careful consideration must be given to means of ensuring the physical security of the record sheets.

[5] For more technical readers: a standard cryptographic hash function such as SHA-256 could combine sequences into a seed. Note that, if the CSPRNG used accepts sufficiently large seeds, the contents of the record sheets could be used directly as the seed, avoiding the combination process. Technically, this could produce a slightly more "random" process, but this difference is of little practical significance.

[6] The CSPRNG can be designed not to output duplicate numbers, avoiding the need for "re-rolling." Use of nearly any popular CSPRNG would suffice, including the generation methods suggested in FIPS 186-2 [8].

method, usability testing would be critical in determining the optimal process.

5 Comparison of Processes

We now compare both processes described in the previous section based on the properties set forth in Section 3. We consider the merits and drawbacks of both processes based on details of the processes alone, but we note that pilot tests may be the most meaningful method for comparing simplicity and efficiency.

5.1 Unpredictability

Encyclopedia Britannica [1] describes numerous methods for cheating with dice. Some tricks would be obvious even to a careless observer, but some are quite subtle. A dishonest official skilled with sleight of hand may be able to substitute crooked dice in and out or to use other tricks to bias the sample selection. Even the ability to make certain sides of a die sticky can bias the roll, and an official could "accidentally" touch glue and transfer that glue to sides of a die. Mitigation techniques exist, such as storing and rolling dice inside a clear canister, but we have difficulty imagining a process that is provably not subject to clever, novel forms of fraud, particularly when the potential for collusion between some subset of officials exists. If a dishonest individual could potentially introduce bias in the process, it is not necessarily unpredictable.

The CSPRNG selection process described in Section 4.2, however, allows each participant sufficient influence on the seed to ensure that the CSPRNG's output is unpredictable.[7] The system is unpredictable assuming that at least one individual generates a record sheet honestly based on the specified process and that the contents of the sheet remain unknown to others until its envelope is opened. As a safeguard, the process uses far more rolls than is strictly necessary. This choice provides sufficient randomness even if the participant unknowingly uses extremely biased dice.[8] Although the process already mitigates bias, a participant suspicious of the provided dice can have them examined or simply replace them with dice distributed by political parties or others. Finally, concerned participants can generate additional sheets using alternative methods. Therefore, the resulting set of selected audit units will be unpredictable to all other parties regardless of their behavior. Each participant's numbers can only increase the randomness of the seed produced by the combination process, and even a single properly generated sheet can ensure a sufficiently random seed. Therefore, each official has the ability to make the sample selection process unpredictable.

5.2 Verifiability

Cordero-style dice rolling is also not necessarily verifiable. Suppose that every citizen—or a trusted party acting on the citizen's behalf—can observe the roll process, study security measures, and rigorously test the dice. In this case, no proof exists that the dice examined are the same dice in the same state as those rolled. Even with cameras focused on her, a dishonest participant may be able to swap altered dice in and out during the roll process. Alternatively, an official could imperceptibly rub an adhesive from the side of a die before returning it. Because parties cannot simultaneously examine the dice, the first to examine the dice could potentially swap or modify them.[9] Finally, even fair dice may be subtly rolled in a manner that biases the results—even casinos use special tables and impose additional requirements to counter this possibility [1]. Mitigation techniques may increase the level of sophistication necessary to perform fraud, but these techniques also may increase cost and complexity or make verification more difficult for an average citizen.

CSPRNGs offer greater verifiability than pure dice rolling. The process of opening the envelopes is more transparent to observers than a roll process, in which observers must trust that all dice are fair and that all officials roll them fairly. The contents of record sheets can immediately be made public, copied by observers, posted to the Internet, etc. The results of the seed construction and CSPRNG algorithms are determined solely by the number of audit units and the contents of the record sheets, and these algorithms should be made public. Therefore, any individual could run the algorithms in her preferred format with these numbers to verify that the sample is correct. Each major political party in addition to concerned citizens could provide an implementation of the algorithms in a user-friendly format. For example, users might be able to enter data into a web page and get back the appropriate sample, and computers at public libraries could allow users to access such web pages. The only portion of this process that is not entirely verifiable is the seed generation. As we argue in the preceding and following subsections, a single honest participant can prevent any number of dishonest participants from biasing seed selection.

[7] For more technical readers: each record sheet contains up to 166 bits of entropy.

[8] Even if each die can only generate three of the ten numbers at random, a cheater would be unable to gain any practical advantage without advances in the state of the art. Depending on the level of bias deemed imperceptible, fewer or more rolls might be useful. In addition, because record sheets are public, anyone can analyze their contents for evidence of heavily biased dice.

[9] This may be accidental: a test that I conduct may destroy evidence of value to a future tester.

5.3 Robustness

Cordero-style dice rolling is not robust against someone that can somehow predict or influence a die. Assume that the attacker can influence the roll of a single die such that it will never generate a number under five (or 50, 500, etc.). Even if he does not know which digit that die will generate, he could freely commit fraud in all audit units numbered only with digits under five—for example, 4,444 or 1,420, but not 1,429. Given any simple correspondence between number rolled and audit unit selected, we can imagine similar tricks.

The CSPRNG process does not suffer from this shortcoming either. Recall that each participant has sufficient influence on the seed to ensure an unpredictable sample selection process. Therefore, if a citizen trusts at least one participant to protect his interests, that citizen can trust the participant to ensure that the process is, at minimum, not biased against his interests. This characteristic makes the scheme robust. Even if all other participants besides a single individual actively seek to undermine the process, that individual can thwart any malicious goals. This is a substantial improvement over Cordero-style dice rolling, in which one official alone can significantly bias sample selection.

5.4 Simplicity

Cordero-style dice rolling is extremely straightforward. Additional safeguards to ensure other properties might decrease simplicity, but the original process is fairly clear. This is an attractive property, since no procedure will be adopted unless officials can reasonably be expected to follow it.

While the example CSPRNG process is not as straightforward as Cordero-style dice rolling, it is not particularly complicated and could be easily demonstrated. Essentially, participants must roll and record five dice ten times, place record sheets in an envelope and locked box, and enter the rolled numbers into a computing device. While maintenance of the computing device may require additional steps, these steps seem reasonable for officials trusted to set up electronic voting machines. Adding CSPRNGs to the audit process yields only a mild increase in the complexity of that process.

5.5 Efficiency

Because each official rolls a die once for each sample item drawn, Cordero-style dice rolling rapidly becomes inefficient for larger samples. Even if an official could roll and record a die once every five seconds, selection of less than 400 audit units would require more than half an hour of continuous rolling. If a process requires excessive time, its tedium might lead to mistakes and shortcuts by participants as well as oversights by observers, both of which are particularly dangerous for unverifiable processes.[10] Unfortunately, this means that closer contests requiring larger sample sizes are precisely those for which Cordero-style dice rolling is most inefficient and dangerous. Additional security measures could exacerbate this inefficiency.

Inefficiency also poses a problem for some promising auditing proposals and suggestions. Several recently proposed schemes require generation of a random number for each audit unit [10, 3]. For a statewide race in New Jersey, given the characteristics described in [2] (approximately 6,300 precincts, 700 voters per precinct, and 50% turnout), this would require hours of continuous rolling if audit units were precincts and months of continuous rolling if units were individual ballots. Thus, requiring Cordero-style dice rolling would effectively amount to a ban on these promising techniques. In addition, Norden et al. [9] suggest sampling at the state level and sampling individual machines rather than precincts, both of which could increase the sample sizes that a process must handle. While Cordero-style dice rolling may be acceptable in certain cases for which a small or even medium quantity of random numbers is necessary, it seems unacceptable in a number of reasonable cases and may prevent other beneficial changes in election auditing.

The proposed CSPRNG process typically requires exactly fifty rolls from each participant regardless of sample size. The various proposals discussed in the previous paragraph would not necessitate an increase in manual effort for sample selection. In addition to fifty rolls, officials must enter the record sheets' contents into a computer, and participants that generate extra record sheets will apply additional effort. Nevertheless, these are upfront costs that do not increase with the sample size. We expect that, when the required sample size reaches several hundred or more, the number of rolls required by Cordero-style dice rolling would make that process less efficient than the example CSPRNG process.

6 Conclusion

Recent work has advised against the use of all pseudorandom number generators, including CSPRNGs, in the audit process. Some work has gone so far as to call for a ban that would include CSPRNGs [6]. We appreciate these concerns and understand (and even support) a healthy reluctance on the part of election officials and citizens, particularly given recent issues with the use of

[10] Each official could roll multiple dice, but this only cuts the length by a small factor and still requires a tedious process.

technology in the voting process. In this paper, however, we have argued that the use of CSPRNGs can result in a process that is more secure and transparent than Cordero-style dice rolling. CSPRNGs can increase the unpredictability and robustness of the sample selection process, and unlike Cordero-style dice rolling, such a process is verifiable by voters. Larger sample sizes, as required to audit closer races, could be impractical for Cordero-style dice rolling, and a ban on CSPRNGs could also equate to an effective ban on several promising audit methods. Given the advantages that CSPRNGs can provide, careful consideration of their appropriate use and further testing would be of far greater benefit than a ban on their use.

Acknowledgments

We thank Eric Rescorla and the anonymous reviewers for their helpful comments and suggestions.

Calandrino performed this research under an appointment to the Department of Homeland Security (DHS) Scholarship and Fellowship Program, administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and DHS. ORISE is managed by Oak Ridge Associated Universities (ORAU) under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORAU/ORISE.

References

[1] Dice: Cheating with dice. Encyclopedia Britannica: Encyclopedia Britannica Online, 2008. http://www.britannica.com/eb/article-1813/dice.

[2] Appel, A. W. Effective audit policy for voter-verified paper ballots in New Jersey, February 2007. http://www.cs.princeton.edu/~appel/papers/appel-nj-audits.pdf.

[3] Calandrino, J. A., Halderman, J. A., and Felten, E. W. Machine-assisted election auditing. In *Proc. 2007 USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)*.

[4] Cordero, A., Wagner, D., and Dill, D. The role of dice in election audits—extended abstract. In *IAVoSS Workshop on Trustworthy Elections 2006*.

[5] Feldman, A., Halderman, J. A., and Felten, E. W. Security analysis of the Diebold Accuvote-TS voting machine. In *Proc. 2007 USENIX/ACCURATE Electronic Voting Technology Workshop (EVT '07)*.

[6] Jefferson, D., Ginnold, E., Midstokke, K., Alexander, K., Stark, P., and Lehmkuhl, A. Post-election audit standards working group report: Evaluation of audit sampling models and options for strengthening California's manual count, July 2007. http://www.sos.ca.gov/elections/peas/final_peaswg_report.pdf.

[7] Kohno, T., Stubblefield, A., Rubin, A., and Wallach, D. Analysis of an electronic voting system. In *Proc. 2004 IEEE Symposium on Security and Privacy*, pp. 27–42.

[8] National Institute of Standards and Technology (U.S. Department of Commerce). Digital signature standard (DSS). FIPS PUB 186-2, January 2000. http://csrc.nist.gov/publications/fips/fips186-2/fips186-2-change1.pdf.

[9] Norden, L., Burstein, A., Hall, J. L., and Chen, M. Post-election audits: Restoring trust in elections, August 2007. http://www.brennancenter.org/dynamic/subpages/download_file_50227.pdf.

[10] Rivest, R. L. On auditing elections when precincts have different sizes, April 2007. http://people.csail.mit.edu/rivest/Rivest-OnAuditingElectionsWhenPrecinctsHaveDifferentSizes.pdf.

[11] Rivest, R. L. A "sum of square roots" (SSR) pseudorandom sampling method for election audits, April 2008. http://people.csail.mit.edu/rivest/Rivest-ASumOfSquareRootsSSRPseudorandomSamplingMethodForElectionAudits.pdf.