Submission No 22

PROPOSALS TO INCREASE VOTER ENGAGEMENT, PARTICIPATION AND CONFIDENCE

Name: A/Prof Vanessa Teague

Date Received: 16 August 2024

Submission to the NSW JSCEM inquiry into proposals to increase voter engagement, participation and confidence

A/Prof. Vanessa Teague vanessa.teague@anu.edu.au

August 16, 2024

I am a cryptographer with a particular interest in election security. Over the last decade, my colleagues and I have found numerous security and privacy problems in electronic voting systems, including the iVote system.

Apart from the obvious privacy and reliability problems, the main concern is *undetected* or *undetectable* failures. Although this can happen in paper-based elections too, there is a greater risk that a poorly-designed electronic electoral process may seem to complete successfully, but diverge from the true wishes of voters due to an undetected software error or security problem.

This inquiry's terms of reference focus on engagement, participation, and confidence, all of which are very important. However, integrity should come first—trustworthiness before trust. We should not seek to make people highly confident in electoral processes that can actually be easily manipulated—we should design processes that give scrutineers and the public solid evidence that the results are accurate. In the computer science literature, this is called *verifiability*: the opportunity for people to check, using something independent of the system, that the system accurately recorded, retained, and counted the votes.

Recommendation 1. Focus on integrity, transparency and verifiability first, with public trust as a consequence, rather than focusing on persuading people to trust something without evidence.

The rest of this submission addresses the technical background for the various proposals to involve computers in NSW elections, particularly emphasising the protection of the secret ballot and the opportunity to verify the results. I have focused on the technologies outlined in the NSW Electoral Commission's "Technology Assisted Voting - Paper 3 - Final Review Report" [Com23] ("The TAV report"), but would be happy to extend the discussion to any other proposals the committee is interested in.

1 Specific points in response to the TAV Report

Election failures undermine public trust. However, the TAV Report reads as if this reduction in trust is a result of the election re-run rather than the initial failure. The following quotation illustrates this.

10. The cost and adverse impact on public trust in democratic processes of re-running an invalid election due to a material irregularity arising from TAV requires mitigating legislation. It would be appropriate and proportionate to those risks if legislation protected, in specified circumstances, the validity of an election result despite technical performance issues with a TAV channel; for example where it is not available for all or some eligible electors to use or where votes that are already cast cannot be verified or counted.

[Com23], p.4.

Both the assumption and the conclusion are wrong. It is the failed election, not the decision to re-run it, that (quite rationally) undermines public trust. No legislation could "protect the validity of an election result" that is in fact invalid. The obvious example is the Local Government Election 2021, in which thousands of voters were disenfranchised due to an iVote technical failure. The NSW Supreme Court voided the results on the basis that this failure had a material impact on the election.¹ The way to avoid the problem is to avoid catastrophic electoral failures, not to refuse to re-run elections that have failed. The NSWEC should not run a TAV channel if it reintroduces a serious risk of this kind of event.

The quotation above conflates two very different situations: a long-planned decision not to run TAV, and a sudden unplanned failure of TAV. A sudden unplanned failure can disenfranchise a significant number of voters, as we saw in the 2021 LGE. This should be treated like any other arbitrary disenfranchisement in a democracy. If the courts decide that the number of disenfranchised individuals was not material, then that is their decision, but there is absolutely no justification for a pre-emptive exclusion of voters from Supreme Court protections merely because their disenfranchisement was due to a technical failure.

Recommendation 2. Do not enact any provision that treats arbitrary voter disenfranchisement through technical failures any less seriously than any other arbitrary voter disenfranchisement.

2 Internet voting, including phone voting

"Internet voting," for this submission, means a form of voting in which a person votes from outside a controlled location (such as a polling place or embassy) and conveys

¹This does not necessarily mean that the results were wrong, but that enough voters were disenfranchised that the results could be wrong.

their vote without paper. This includes iVote, email voting, and phone voting. These systems have the same main advantages and the same main risks: they are convenient and relatively cheap, but they leave the voter with no evidence that their vote was accurately recorded. They also have significant privacy risks, which depend on the specifics of the system.

There are generally more scientific papers identifying security problems in (non-phone) Internet voting systems than phone voting systems, but that is not because phone voting is inherently any more secure than any other form of Internet voting. The security and privacy implications tend to depend on existing infrastructure rather than the system itself. For example, a vote cast via a properly-authenticated end-to-end encrypted phone call (which is rare) is hard to eavesdrop in transit, whereas an ordinary VoIP or cell phone call may be readable by the service provider.

Some Internet voting systems provide mechanisms that allow voters to verify—under certain assumptions—that their vote was accurately recorded. The Helios system [Adi08] allows voters to challenge the encrypted vote that is going to be sent, then allows the public to verify the tallying cryptographically. However, this process also provides an avenue for people to prove how they voted, hence exposing them to coercion. The Estonian Internet voting system [HMVW16] is based on similar principles, but attempts to address the coercion problem by allowing repeat voting and by hiding the transcript from public view. Unfortunately, these measures break the verifiability guarantee [Per22]. The Swiss Internet voting system ² uses randomly-generated codes to allow a voter to check that the vote they sent matched their intentions, but this is dependent on significant assumptions, including the secure and private delivery of a paper code sheet. This method does not incorporate preference-based voting and is difficult for voters with accessibility challenges.

I can discuss any of these examples in greater detail, but the summary is that usable verification, with reasonable assumptions and without exposing voters to coercion, is not currently feasible, and is not likely to be feasible in a reasonable time. Without usable verification, there is a significant risk of undetectable fraud.

Recommendation 3. Discontinue Internet voting.

If Internet voting (including phone voting) must be retained, make it as transparent as possible and ensure that all voters are fully and frankly informed of the substantial risks to both privacy and integrity.

3 Computerised (kiosk) voting in a polling place

There are many different methods of using computers to assist with voting in a polling place, some of which are well described in the TAV Report. There are some cryptographic schemes that give voters some extra verifiability properties in addition to plain paper (see below), but again there is really no usable, genuine verification method without paper.

²https://www.bk.admin.ch/bk/en/home/politische-rechte/e-voting.html

The following lists some relevant projects, and discusses their privacy and verifiability.

3.1 Internet voting in a polling place

In past elections, NSWEC has made iVote-enabled terminals, with no paper evidence trail, available in polling places. This combines the convenience of pollsite voting with the security and privacy of Internet voting.

3.2 Direct-recording electronic (DRE) voting machines

DREs are computers that accept votes electronically, via a touchscreen or some other interface, and produce an aggregated tally (either paper or electronic) at the end of voting. They have been used extensively in both Brazil, where they are universal, and the USA, where their use has declined as a result of public distrust and legal challenges. Paperless DREs do not provide meaningful verifiability, even if a voter can see their vote on a confirmation screen, because the voter is not able to check that the electronic record matches what they see. If the tally results are electronic, scrutineers cannot verify that voting data is accurately maintained and transferred. There is a substantial history of scientific demonstrations of serious security and privacy problems in paperless DREs, though this has slowed somewhat in recent years as their use in the USA declines.³

Paperless DREs are not used in most Australian states or territories. The only exception I know of is the ACT's EVACS system.⁴ EVACS votes are communicated electronically to a server inside the polling place, and then manually transferred on digital media to a central counting location. Independent analyses have shown serious problems in EVACS, affecting both privacy⁵ and accuracy.⁶

Paperless DREs suffer from essentially the same main risk as Internet voting: the opportunity for undetected errors or security breaches to undetectably alter election results.

Recommendation 4. Do not adopt paperless Direct Recording Electronic (DRE) voting machines.

3.3 Voting machines with Voter-Verifiable paper audit trail, or "Ballot marking devices"

Some electronic voting machines print a human-readable record of the vote, which is shown to the voter for verification, then retained in the polling place and either counted or audited. Some authors distinguish "ballot marking devices", in which the paper record is counted along with other (manually marked) ballot papers, from "voting machines

³https://verifiedvoting.org/verifier/#mode/navigate/map/ppEquip/mapType/normal/year/ 2024

⁴The TAV Report also mentions a Western Australian trial of paperless DREs.

 $^{{}^{5} \}texttt{https://github.com/teor2345/Elections2018/blob/master/ElectionsACTDisclosure.md}$

⁶See our submission at https://www.parliament.act.gov.au/__data/assets/pdf_file/0005/ 1751693/Submission-003-Vanessa-Teague,-Andrew-Conway,-Thomas-Haines,-T-Wilson-Brown.

pdf

with Voter-Verifiable paper audit trail", in which the primary count is electronic and the paper record is only for auditing. However, this distinction is not very consistently applied, and many authors use the two terms interchangeably.

Although there is controversy in the USA about the use of ballot marking devices for general voters (see [ADS20] and [KBW21] for examples of opposing views), this seems to be a good solution for voters who are not physically able to fill in their own ballot. It gives voters an opportunity to see and check that their printout matches their intention, though of course this is only directly useful to sighted voters. An implementation of this idea would need to be carefully designed to give voters the opportunity to verify their printout, and send it to a ballot box automatically (without assuming they were able to do so manually), or to reject it if the printout did not reflect their intentions.

The paper printout could then be included in the normal scrutineering process, along with the other (manually-marked) paper ballots, thus completing an evidence trail from the voter verifying that it is accurately recorded to the scrutineer verifying it is accurately counted. Alternatively, if the primary record is electronic, an audit of the paper could be conducted in the presence of scrutineers, perhaps at the same time as the audit of any digitised ballots (see subsection 4.2).

Voting machines always entail some privacy risk, because if the machine is compromised then the attacker can learn individual votes. This risk needs to be clearly communicated to voters, so they can compare it with the more obvious risk involved in asking another person to complete their ballot.

Although not perfect, this solution probably offers the best tradeoff among the various goals of integrity, privacy, access, and convenience. It is more expensive than Internet voting because it requires a computer with a printer in the polling place, but the software design need not be excessively expensive because the system could be very simple.

The contests available on a voting computer could be much broader than the electorate in which the computer was situated—all the computers in the state (and overseas or interestate voting centres, if used) could be loaded with all the candidate names for all contests.

Recommendation 5. Consider providing computers in a polling place, equipped with a method of printing a voter-verifiable paper record that can be automatically submitted into a physical ballot box after the voter verifies its contents.

There are alternative, cryptographic, methods of voter verification designed to run in a polling place, for example Victoria's now-discontinued vVote project [CRST15]. These systems have many good properties, notably the opportunity for individuals to verify the proper processing of their vote even after they have left the polling place. However, their complexity is substantial, and although I would gladly support anyone who wanted to revive that work, I now believe that the simpler solutions described above are probably a better use of resources.

4 Other recommendations

There are numerous other opportunities for reform, allowing NSW to benefit from the improved speed and convenience of computers without precipitating another electoral failure by unwisely depending upon them. This section makes some other general notes and recommendations about technology in voting, relevant to NSW and the rest of Australia.

4.1 Electronic candidate information and paper returns

A simple way to halve the time required for remote postal voting is to allow voters to download candidate information (which is public) from the Internet, then print the ballot, fill it in, and either mail it back or deliver it to an embassy or interstate voting centre. The software for supporting this could reuse a lot of the code for the kiosk-based voting system described above.

Recommendation 6. Consider making candidate information available online, so that remote voters could print a blank ballot and either mail it back or deliver it to an interstate or overseas voting centre.

This still suffers from the security and privacy problems of (perhaps international) mail, but it would be harder to manipulate a large number of ballots undetectably than it is via Internet voting.

4.2 Auditing the ballot digitization process

Like many other Australian jurisdictions, NSW uses computers to digitise and count upper house (Legislative Council) ballots. The trustworthiness of this process would be greatly improved by a rigorous audit, in the presence of scrutineers, taking a random selection of ballot papers and comparing them to their corresponding digital record. I believe the NSWEC conducts some auditing on its own initiative, but there is no legislation specifying a level of rigour or transparency. Recent Commonwealth legislation now requires the Australian Electoral Commission to conduct and audit for Senate votes.⁷ If implemented properly, this greatly reduces the risk that an undetected error or security problem in the scanning process could undetectably alter the outcome. The audit should be available to scrutineers, who can check that it is being done properly, and should become part of a standard transparent electoral process.

Recommendation 7. Legislate for a rigorous, public, audit of the ballot digitisation process, in which randomly selected paper ballots are compared with their digitised preferences to estimate the rate of error.

⁷See the Commonwealth Electoral Act, 273AC "Ballot paper sampling assurance throughout computerised scrutiny of votes in Senate election"

4.3 Source code transparency

Some electoral commissions, including in recent cycles the NSWEC, have made some election-related code available for independent examination. This process has huge benefits because it allows the public an opportunity to identify and correct errors before the election. My colleagues and I have found (and mostly had corrected) numerous errors related to privacy, integrity, and counting accuracy, in Australia and overseas, often through source code access.

Although practice has improved, the law lags far behind. NSW legislation still includes the ridiculous imposition of criminal penalties for sharing the code without the Commissioner's permission.

(159(2)) A person must not disclose to any other person any source code or other computer software that relates to technology assisted voting under the approved procedures, except in accordance with the approved procedures or in accordance with any arrangement entered into by the person with the Electoral Commissioner.

Maximum penalty—200 penalty units or imprisonment for 2 years, or both.

[NSW Electoral Act 2017]

It is only because Switzerland had the opposite law (a requirement for source code openness) that the cryptographic errors in iVote in 2019 came to light.

Recommendation 8. Repeal 159(2) and replace it with a requirement that source code be openly available for public scrutiny.

4.4 Open source software and federalism

The New South Wales electoral commission has for many years advocated a unified national approach to electronic voting, a recommendation repeated in the TAV Report. As far as I know, apart from one small iVote run in Western Australia, no other electoral commission has expressed much enthusiasm. Different Australian jurisdictions, despite some similarities, often have quite different processes, regulations and risk appetites. For example, the ACT *Electoral Act 1992* Section 118A(2) states "The commissioner may approve a program under subsection (1) (a) only if the program will— ... (d) not allow a person to find out how a particular elector cast their vote."⁸ NSW has no such requirement, so a system that was perfectly legal in NSW might be disallowed in the ACT.

Requirements for code secrecy (above) are another good example. It seems highly unlikely that anyone would want to share a source code base with NSW if there was a risk that disclosing it for public scrutiny could result in criminal penalties.

In short, a single, unified, proprietary codebase is probably unworkable.

⁸https://www.legislation.act.gov.au/View/a/1992-71/current/html/1992-71.html

A much better solution would be a genuinely open-source model, in which different electoral commissions could work co-operatively on a publicly-owned codebase under an open arrangement such as the Gnu Public License. Specific customisations, to comply with local regulations or add features relevant to only one state, could be implemented by a single commission for their own needs. This would have many of the advantages of a unified system—different commissions could cooperate on those parts that were common—without most of the disadvantages of forced uniformity. It is shameful that so much public money has been spent, in NSW and in other Australian jurisdictions, on software that should be owned by the public but is unavailable for either scrutiny or improvement.

I have recently worked on several related US election auditing projects, one in San Francisco, one in Boulder, and one for the statewide election audits in Colorado.⁹ For election auditing software in the USA, this open source model is uncontroversial and works well—each of these projects effectively used pre-existing open source software, customised it to their needs, and put back a reusable version that others could build on.

I do not understand why Australian electoral commissions are so averse to the open source software model, which can be highly effective in reducing duplication and inefficiencies, with none of the downsides of forced standardisation. This model would be perfect for a polling-place voting system, in which different electoral commissions might want to customise the user experience, while sharing the background code for serving or printing the ballot information. Audit software could also be developed in this way indeed, some of the US auditing software mentioned above could be openly reused in Australia.

Recommendation 9. Consider a genuinely open-source software model, with optional customisations by each state or territory, rather than the coordinated nationwide approach envisaged in the TAV Report.

The two options are not strictly inconsistent, of course.

4.5 The NSW Electoral Commission's analysis of "material impact"

The TAV Report describes how the NSW Supreme Court decided to assess the vital question of whether the exclusion of a certain number of voters in the 2021 LGE constituted a "material impact" on the election result. I submitted an affidavit to the court, describing algorithmic methods for calculating whether a certain number of excluded voters was sufficient to change the outcome.¹⁰ The TAV report states

The Electoral Commissioner employed a Monte Carlo simulation, which calculated the probability of a different outcome by considering a thousand

⁹The state of Colorado contracted my not-for-profit, Democracy Developers Ltd, to write the code for the system, which will be publicly available and publicly owned. Our code is at https://github. com/DemocracyDevelopers

¹⁰Analysis at https://github.com/AndrewConway/ConcreteSTV/blob/main/reports/ NSWLGE2021Report.pdf, joint work with Andrew Conway.

simulations of the missing iVotes based on random selections of actual vote preferences in impacted contests.

This characterisation is false—the simulations provided by NSWEC to the Supreme Court did not "calculate the probability of a different outcome," and were explicitly rejected in the Supreme Court judgement because they are based on invalid assumptions. It is disappointing that this invalid method is still being used—it tends to understate the implications of election failures.

A similar invalid approach seems to have been applied in TAV Section 305 to examine what "changes to votes are sufficient to alter election outcomes." That section describes the implications of removing a *random* 5% of votes, which has a much lower probability of changing outcomes than removing a *specific*, *targeted* 5% of votes.

I would, as always, be happy to discuss any of these issues with the committee.

References

- [Adi08] Ben Adida. Helios: Web-based open-audit voting. In In Proceedings of the 17th USENIX Security Symposium (Security '08, 2008.
- [ADS20] Andrew W Appel, Richard A DeMillo, and Philip B Stark. Ballot-marking devices cannot ensure the will of the voters. *Election Law Journal: Rules, Politics, and Policy*, 19(3):432–450, 2020.
- [Com23] "New South Wales Electoral Commission". Technology assisted voting paper 3 - final review report, Nov 2023.
- [CRST15] Chris Culnane, Peter YA Ryan, Steve Schneider, and Vanessa Teague. vvote: a verifiable voting system. ACM Transactions on Information and System Security (TISSEC), 18(1):1–30, 2015.
- [HMVW16] Sven Heiberg, Tarvi Martens, Priit Vinkel, and Jan Willemson. Improving the verifiability of the estonian internet voting scheme. In International Joint Conference on Electronic Voting, pages 92–107. Springer, 2016.
- [KBW21] Philip Kortum, Michael D Byrne, and Julie Whitmore. Voter verification of ballot marking device ballots is a two-part question: Can they? mostly, they can. do they? mostly, they don't. *Election Law Journal: Rules, Politics,* and Policy, 20(3):243–253, 2021.
- [Per22] Olivier Pereira. Individual verifiability and revoting in the estonian internet voting system. In International Conference on Financial Cryptography and Data Security, pages 315–324. Springer, 2022. https://eprint.iacr.org/ 2021/1098.pdf.