

**Submission  
No 6**

# **PREFERENCE COUNTING IN LOCAL GOVERNMENT ELECTIONS IN NSW**

**Name:** Dr Vanessa Teague  
**Date Received:** 27 September 2017

# Submission to the Inquiry into preference counting in local government elections in NSW

Dr Andrew Conway  
[REDACTED]

Dr Michelle Blom, Dr Chris Culnane, Prof Peter J. Stuckey, Dr Vanessa Teague  
School of Computing and Information Systems  
University of Melbourne  
[REDACTED]

Prof Rajeev Goré  
Research School of Computer Science  
Australian National University  
[REDACTED]

This submission was prepared in the authors' personal capacity. It addresses the rules and processes for NSW local government ballot counting.

None of the authors have any financial interest in electronic voting or counting. Some have Australian Research Council grants for research in the area. Vanessa Teague is an advisory board member of Verified Voting, a non-governmental US organization working toward accuracy, integrity and verifiability of elections. Andrew Conway is a member of the Secular Party. Michelle Blom is a member of the Victorian Greens.

## 1 Introduction

This submission describes strengths, weaknesses, errors and suggested improvements in the process used by NSWEC to count votes in local government elections.

Some NSW local government elections are counted by the councils themselves. We do not have any information about their processes and have not included them in our report. The absence of demonstrated errors should not be taken to indicate that their processes are more accurate than those of NSWEC. On the contrary, NSWEC's complete preference files and distribution of preferences files have allowed us to double-check their results and to identify a number of errors and weaknesses, some of which have been corrected. This is a good thing for NSW democracy. The councils that run their own elections may be just as error prone, but are much harder for us to check. Our main recommendations are ways that NSWEC can open more of their process to independent scrutiny, which will increase the likelihood that other mistakes and weaknesses can be found and fixed.

NSW counting legislation is excessively complex and very difficult to understand or implement accurately. It is hardly surprising that NSWEC's implementation has had a series of errors, because the legislation itself is so hard to interpret. Rather than make minor amendments, we suggest removing the current counting specification completely and starting again, perhaps by looking at other jurisdictions such as Victoria's Legislative Council or the Senate, though they have unnecessarily complex details too.

We suggest two obvious simplifications of the NSW counting rules: removing the last parcel and removing randomness. We also recommend improvements to the transparency of the electronic counting process.

## 1.1 Main recommendations

### Method of Ballot Counting:

- Remove randomness and do deterministic weighted transfer of all votes,
- Remove last parcel and transfer all votes.

This would align NSW counting with most other Australian STV rules, including those of the Senate and Victoria's Legislative Council.

There are also various ways of improving the use of transfer values when a candidate is elected. The current Senate and Victorian systems simply assign the same transfer value to every vote regardless of its starting value—the JSCEM should consider fairer alternatives, which we assume will be described in submissions from other experts.

### Electronic Process:

- Open the source code to public scrutiny,
- Do random audits of the paper ballots against the published data files.

This would increase the likelihood of software errors being identified and corrected, and would also allow NSWEC to demonstrate that random choices were made fairly if randomness could not be removed from the counting rules.

## 2 An overview of the NSW local government counting process

In NSW local government elections administered by NSWEC, paper ballots are digitised. The electronic records are then counted by computer and also posted online, along with a full distribution of preferences.

So there are two stages in which an error or security problem might cause the election result to be wrong: the digitization could go wrong, and produce electronic records that do not faithfully reflect the paper votes, or the counting could go wrong. We can only check the count. The only way to check the digitization is by a rigorous statistical audit of the paper ballots against the published preferences, which we also suggested for the Senate [1]. However, the small sizes and close margins in NSW make it hard to derive high confidence without auditing a very large fraction of votes. The rest of this submission concentrates on checking the count, because that is what we have been able to do, but auditing the input is probably even more important.

We have found two issues in the electronic count:

- *randomness in the counting rules* can sometimes mean that different runs of the same election, with different randomly-selected ballots, get different results.
- *errors in the software* have caused incorrect distributions of preferences and, in at least one local council in 2012, an election outcome that was probably incorrect.

In most local government elections, neither the software bugs nor the randomness made any difference. But each year, a handful of councils have close enough results that the last one or two seats to be assigned can vary according to the randomness.

The software errors were not related to randomness. We have found three errors, relating to the “last parcel” computation, the rounding of partial votes, and the breaking of three-way ties. The technical details are described in [2].

Although the software errors are a completely separate issue from the effect of randomness, the combination is particularly unfortunate. When a software error affects the count in an election close enough for randomness to matter, it is very hard to define who should have won.

### 3 Making random choices transparently

Although it would be better to eliminate randomness, it is quite possible to conduct a transparent randomised process that is verifiably fair. Consider the Tattslotto draw: a machine gives a public display of the random selection of numbers. If Tattslotto numbers were not drawn in the open, people would not play.

When computers implement randomised algorithms, they do not generally require new random input each time they need to make a random choice. Instead they are initialised with a “seed” value which is supposed to be random, then they run an algorithm that appears to be random but is actually completely determined by the seed value. This is called “pseudorandom,” because it seems random, but in fact is precisely repeatable by anyone who knows the algorithm and the seed.

This would allow NSWEC to do its random selection in a fair and transparent way, even if randomness could not be removed from the counting rules. They would need to make all the source code for the count openly available.

The process would be:

1. publish full source code for the count,
2. publish full preference data files (exactly as they already do),
3. generate the random seed in a public way, such as by throwing dice or borrowing the Tattslotto machine.

Then any interested scrutineers or citizens could download the code and rerun the randomized count with the same seed value—they should get exactly the same outcome. This method would show everyone that the result was fair, and could not be manipulated or biased.

### 4 A brief history of our investigations into NSW counting

In 2015 we implemented the NSW counting rules and re-ran each of the 2012 local government elections, for which data was available, 1 million times.<sup>1</sup> In 33 of the 217 contests, there was more than one possible outcome due to randomness. In 16 of those contests, the most likely winner did not win. For example, in Mosman, Anne Connon wins 97% of runs but did not win a seat. This does not prove that there was a mistake, but it suggests that the randomised process should change.

We also found a software bug in the computation of the “last parcel,” which made the official distribution of preferences in the council of Griffith incorrect. Candidate Rina Mercuri wins 91% of runs when the algorithm is correct, but only 10% using the incorrect code used by NSWEC for that

<sup>1</sup>Source code is on GitHub at <https://github.com/SiliconEconometrics/PublicService>  
The complete 2012 results are at <https://siliconeconometrics.github.io/PublicService/CountVotes/NSWLGE2012MillionRuns/>

count. She did not win a seat. Although she might possibly have lost even if the software had been correct, there is a compelling argument that she was disadvantaged by the software error.

We recommend removing the notion of “last parcel” altogether, because it arbitrarily selects some votes (but not others) for distribution, and because it introduces complexity into the code that makes it difficult to implement correctly. Even when correctly implemented, it can distort the result [3].

We notified NSWEC of the error, though it was too late to overturn the 2012 election results. They acknowledged and corrected the software error.<sup>2</sup> We recommended opening the source code and choosing the randomness transparently, which they promised to consider but did not implement.

When the results of the 2016 local government elections were published we immediately re-ran each count for which data was available, 1 million times.<sup>3</sup> We found two more software errors and one candidate who wins 56% of runs but did not win a seat. We informed NSWEC about these issues on Friday 23rd September 2016, with full details of all runs available on Monday 26th—the letter we sent is attached here. They did not make any public acknowledgement of the issues. We do not know whether the software errors were corrected.

In 2017, we ran each election 1 million times<sup>4</sup> and found four cases where randomness mattered.

- In the Graystones ward of Cumberland, Eddy Sarkis missed out by 15 votes, but actually won 77% of runs. The official winner, Glen Richardson, won the other 23%. This result was overturned after a recount.
- In the Wollstonecraft ward of North Sydney, Tim Kelly missed out by 11 votes but won 21% of runs. The official winner, Samuel Gunning, won 79% of the time.
- In Oberon the reported margin was 11. Two candidates other than the official winners could win with tiny probabilities: 4% for Ian Tucker and 0.25% for Neil Francis.
- In Armidale, Jim Maher won 0.025% of runs.

The current process does not give the losing candidates or their scrutineers any evidence whatsoever that the random choices were made fairly.

## 5 The effect of randomness on Legislative Council elections

Legislative Council counting should also be reconsidered. At present, NSWEC uses the same code for both, since the counting algorithm is almost the same.<sup>5</sup> Since the Legislative Council rules cannot easily be changed, we see two options, neither ideal:

1. change the Legislative Council rules (by referendum if necessary) to match new non-random local government counting rules,
2. leave the Legislative Council rules as they are; mandate open source counting code and a public generation of the random seed, as described in Section 3.

---

<sup>2</sup>[http://www.office.elections.nsw.gov.au/\\_data/assets/pdf\\_file/0008/218681/PRCC\\_statement.pdf](http://www.office.elections.nsw.gov.au/_data/assets/pdf_file/0008/218681/PRCC_statement.pdf)

<sup>3</sup>The complete 2016 results are at <https://siliconeconometrics.github.io/PublicService/CountVotes/NSWLGE2016MillionRuns/>

<sup>4</sup>Complete 2017 results are at <https://siliconeconometrics.github.io/PublicService/CountVotes/NSWLGE2017MillionRuns/>

<sup>5</sup>It is not exactly identical. For example, the tiebreaking rules are different.

Legislative Council elections are less likely to be affected by randomness than local government elections for two reasons: there are fewer of them, and they have more votes. However, it is certainly possible that a future Legislative Council election might have a close enough margin that randomness can change the result, so the conduct of those elections must demonstrate the random values were chosen fairly.

We re-ran the 2015 Legislative Council elections one thousand times to assess the effects of randomness. The same winners were elected every time. The final margin varied from 2884 to 3335.<sup>6</sup>

Although random selection made no significant difference to the 2015 Legislative Council elections, it is certainly possible that some future election might have a close enough margin. Based on our testing, a margin of a few hundred could be enough to make the election outcome vary randomly. It is not very likely in any particular election, but unlikely things can happen—in the 2013 Western Australian Senate race, one lost ballot box was sufficient to change the outcome. It was a small difference in an early stage of the counting that made the change—the difference between the last two candidates is not always the true electoral margin.

It would not be possible to redesign the NSW electoral process after a close result was discovered—the electoral process must be open and transparent just in case randomness matters.

## References

- [1] Berj Chilingirian et al. “Auditing australian senate ballots”. In: *arXiv preprint arXiv:1610.00127* (2016). <https://arxiv.org/pdf/1610.00127.pdf>.
- [2] Andrew Conway et al. “An analysis of New South Wales electronic vote counting”. In: *arXiv preprint arXiv:1611.02015* (2016). <https://arxiv.org/pdf/1611.02015.pdf>.
- [3] Rajeev Goré and Ekaterina Lebedeva. “Simulating STV Hand-Counting by Computers Considered Harmful: ACT”. In: *International Joint Conference on Electronic Voting*. Springer. 2016, pp. 144–163.

---

<sup>6</sup>The mean was 3108, which was close to the official result. The standard deviation was 65.6.

## Errors in the 2016 NSW LG elections

We found two different types of error in the count for the 2016 NSW LG elections, that resulted in incorrect distribution of preferences in four observed instances. Our guess is that they did not change who was elected in any contest, but we cannot be sure, and there is some reason for doubt.

We pointed out these two errors and three of the instances (we found the fourth later) to the NSW Electoral Commission (NSWEC) a couple of days after the election results were posted, but have not yet had a response from them. In the prior LG election we pointed out a different error to them that they acknowledged it was an error and probably – but not certainly – meant the wrong person was elected.

### Error 1: 3 Way ties for elimination

When a candidate must be chosen for exclusion, and there is a tie for the lowest candidate, then the tie is broken via a countback. From the published spec (Functional Requirements Specification for the Vote Count v 3.2), section 1.4.26:

If the Election is a Local Government Election and there are 2 or more Candidates with equal lowest current Progressive Totals, then the Count process must be reviewed to go back and determine the previous Count at which the Progressive Totals for these 2 or more Candidates were last unequal and Exclude the Candidate with the lowest Progressive Total at that point. If the 2 or more Candidates have had equal Progressive Totals at all preceding Counts (including Count 1), then a draw must be conducted to determine the Candidate to be Excluded.

This is somewhat ambiguous in the case of 3 or more way ties. Consider the case of Hawkesbury. There is a 3 way tie at count 37, on 21 votes. Countbacks show:

Count	VIGOUROUX Tara	SMITH Fiona	MERCER Sidonie
36	21	21	21
35	20	21	21
34	20	21	21
33	20	21	21
32	20	21	21
31	20	21	21
30	20	21	21
29	20	21	21
28	20	21	17

...

On count 36, they are all equal. On count 35 they are unequal, and indeed there is a unique lowest, VIGOUROUX Tara, and it would seem very reasonable to exclude VIGOUROUX Tara. One could however interpret as having to go back until all three are unequal, that is step 28. Then the lowest is MERCER Sidonie. In the official results, it is VIGOUROUX Tara who is excluded (as seems reasonable), so this is evidence that the official interpretation is to not require all unequal, just a clear loser. The same choice is made in other equivalent places (e.g. count 12 of Albury, 4 way tie).

Similarly, at count 12 there is a 3 way tie. This time there is never a count when all are unequal. However, there is still an obvious loser in the countback, GINNINGS Margaret, who is indeed eliminated:

3 way tie at count 12 on 4 votes. GINNINGS Margaret was excluded.

Count	LUGTON Jen	GINNINGS Margaret	CRAWFORD Edward
-------	------------	-------------------	-----------------

11	4	4	4
10	4	4	4
9	4	4	4
8	4	4	4
7	4	3	4
6	4	3	4
5	4	3	4
4	4	3	4
3	4	3	4
2	4	3	4
1	3	3	4

This interpretation is confirmed in an unpublished (as far as we know) addendum to the specification version 3.2 that we have been sent (personal communication).

**However, this is contradicted in counts 9 and 20 Hawkesbury.**

In count 9, MIDDLEBROOK Shaun is excluded rather than the more reasonable CARTER Michelle.

Count	BLACK Peter	CARTER Michelle	MIDDLEBROOK Shaun
8	3	3	3
7	3	3	3
6	3	3	3
5	3	3	3
4	3	3	3
3	3	3	3
2	3	2	3
1	3	2	3

Similarly, in count 20, JAY Graeme was excluded instead of THOMAS John.

Count	REARDON Roger	THOMAS John	JAY Graeme
19	9	9	9
18	9	9	9
17	9	9	9
16	9	9	9
15	9	9	9
14	9	9	9
13	9	9	9
12	9	9	9
11	9	8	9
10	9	8	9
9	9	8	9
8	9	8	9
7	9	8	9
6	9	8	9
5	9	8	9
4	9	8	9
3	9	8	9
2	9	6	9
1	9	5	9

Also in Cambelltown, count 17, RIAD Youssef was excluded instead of SIMMONS Carla.

Count	RIAD Youssef	FANOUS Anthony	SIMMONS Carla
16	3	3	3
15	3	3	3
14	3	3	2
13	3	3	2



12	3	3	2
11	3	3	2
10	3	3	2
9	3	3	2
8	3	3	2
7	3	3	2
6	3	3	2
5	3	3	2
4	3	3	2
3	3	3	2
2	3	3	2
1	3	3	2

These three exclusions directly contradict both the expected behaviour demonstrated in other rounds, and the (unpublished as far as we know) specification we were sent.

The thing these three implausible results have in common is that while there is a clear loser (which is needed), it is not possible to distinguish amongst the others (not that you need to).

### Does it matter?

The wrong order of exclusion can clearly directly cause the election to change. It can also indirectly cause the results to change in not-obvious ways, as if the effect it has is to cause a random draw, then this will take numbers from the pseudo random number generator, changing the results of future draws.

This probable bug happened to not make a difference in this case, as Cambelltown's results are the same when run a million times without this putative bug. Hawkesbury is similar, although it may change the order of election. Not having access to the source code we cannot confirm that the bug is restricted to these two contests.

## Error 2: Rounding of votes to be transferred during excess distribution

When distributing the excess from an elected candidate, only a portion (defined by a transfer value) of the votes get transferred. This is done by distributing the appropriate papers to each candidate based on the next continuing preference, and then taking a portion of these. Usually, this portion will not be an integer. Rounding is done by a set of rounding rules defined in section 1.4.17.1 of the spec. A fixed number will be rounded up, ordered by largest fractional value (as makes sense). In case of ties, preference goes to the largest integer value, then highest values in a countback, then a random draw.

In Bland Shire Council, count 2, the transfer value is exactly 0.2. Four candidates end up with a fractional value of 0.6; three of those 4 need to be chosen. The integers are:

MONAGHAN Brian 9

GRELLMAN Peter 7

BAKER Bruce 5

THOMAS Muzz 3

It is unambiguous that the first three should be rounded up, and the last rounded down. But instead, BAKER Bruce is the candidate who was rounded down, and THOMAS Muzz was rounded up.

There is some rounding that the spec requires to fixed number of decimal digits. However, this does not affect the outcome as decimal rounding does not affect multiples of 0.2. Of course careless use of a combination of decimal rounding and binary arithmetic could cause this problem, but it is also likely to cause many other problems which don't seem to exist. We do not understand the cause of this bug (not having access to the source code), but will point out that this is the only example of a rounding determined by the integers where the transfer value is perfectly representable in base ten but not in base two.

### **Does it matter?**

In principle a rounding error could make a difference in a close fight – many of the contests are decided by a couple of votes. However in the specific case of Bland Shire, the same candidates are elected in a million runs (although it may affect order of election).

The more serious effect is what this could have on random number generation (in case the bug causes a random draw, which is somewhat unlikely to cause the observed outcome, but other possibilities are also somewhat problematic). This could have a large knock on effect. In particular, a very similar situation arises in count 3 of Yass Valley Council. In this particular case the NSWEC rounding appears correct, but an undetectable (from the outside) effect on random number generation is likely to change the result of the election, as our million runs show BUTLER Greg should be elected roughly 56% of the time (but wasn't in the official count) while TURNER Kim should be elected roughly 36% of the time (and was in the official count). Not having access to the source code we cannot confirm that this result was not affected by this bug.