SIMULATING THE 2011 REFERENDUM IN NEW ZEALAND

ABSTRACT. On 26 November 2011, an indicative referendum was held in New Zealand with the aim of gauging public support for a change from the current parliamentary electoral system (Mixed Member Proportional) to one of four alternatives. In order to understand the consequences (in terms of the seat distribution of parties in Parliament) of a change in electoral system, we created an online simulator several months before the referendum date. Several interesting research issues arose in this work, which in our opinion deserve greater analysis. We describe the assumptions made in order to create such a simulator, and their consequences.

1. INTRODUCTION

This paper reports on an online simulator [6] developed by the present authors in the year leading up to the Electoral Referendum [10] held in New Zealand in 2011. The aim of the simulator was to give the user information on the seat distribution in Parliament of each party under the various systems listed in the referendum, under various scenarios.

We first describe the specific details of the referendum, and then a more general description of the problem.

1.1. **The referendum rules.** The Electoral Referendum Act 2010 [5] set up an indicative (non-binding) referendum to coincide with the 2011 General Election in New Zealand. The two questions were:

(1) Should New Zealand keep the MMP voting system?
(2) If NZ were to change to another voting system, which voting system would you choose?
  - First Past the Post (FPP)
  - Preferential Vote (PV)
  - Single Transferable Vote (STV)
  - Supplementary Member (SM)

We first describe the alternatives on offer. Terminology is mostly standard.

- MMP (Mixed Member Proportional) is the current system, with the following features. There are 70 (geographically defined) electoral districts, and each voter makes a plurality vote ("electorate vote") for a single candidate in one district. In addition, each voter makes a separate vote ("party vote") for a nationally registered political party, each of which has an ordered list of candidates. Parliamentary seats are allocated to parties via the Sainte-Laguë (Webster) method, based on each party's national total of party votes. The method is modified by a threshold rule, which specifies that a party with no candidate who has won the electorate vote in a district, and less than 5% of the national party vote, will not be allocated any seats. A party's seats are filled, in the first instance, by those of its candidates who are electorate-vote winners in their districts, and then by candidates from the party's list. If necessary, a special "overhang" seat, additional to the usual complement, is created for any electorate-vote winner who finds that his party has been allocated insufficiently many seats to accomodate him.
- FPP has geographical districts, and each voter makes a plurality vote for a single candidate in one district. Each plurality winner is elected to Parliament.
- PV is as FPP, but in each district voters submit a complete preference order and the winner is computed using Instant Runoff (this is called Alternative Vote in some places, for example in the 2011 UK referendum).
- SM allocates some seats exactly as FPP, and others proportionally via a separate party vote as with MMP (possibly with a different threshold). The legislation specified that there would be 90 districts, and 30 seats allocated proportionally.

Date: 28 February 2013.

- STV is an entire family of voting systems, with many parameters not specified by the legislation. The legislation specified that STV should be considered to have between 24 and 30 districts, with between 3 and 7 members each. We made a concrete choice of 30 districts, determined as described in Section 2.2.

In all electoral systems, the total number of seats in the Parliament is fixed at 120; the only exception is the overhang-seat rule of MMP. Note that each system is anonymous — only the number of voters (in each district) with each preference order is important, not their identities. Thus for each district we need only specify the *voting situation* – that is, the number of voters holding each preference order.

1.2. **The general problem.** For the purposes of this paper, the key features to note from the above list are that we must consider both majoritarian and proportional electoral systems, and we must consider both plurality rule and other voting rules that use more detailed preference information and may elect more than one member per electoral district. The general question is how to compare the consequences, in terms of parties' seat distributions, of each of several electoral systems being used. Our viewpoint here is "short-term". Of course, if a change to an electoral system is in fact implemented, we would expect eventual changes in the number and composition of parties as they react to the new environment. Thus we are chiefly concerned with predicting the result of an election under several systems, simultaneously.

The referendum legislation specified several parameters that a research article should leave free, such as the total number of seats. It also failed to specify some parameters that must be fixed in order to produce a simulator of the type described here, such as the precise form of STV to be implemented. We describe in the next section the particular choices we made for various parameters. Unless otherwise stated, we see no difficulty in performing the analysis for general values of these parameters.

2. Assumptions made

An important goal in producing the simulator was to use as few assumptions as possible, to document them and to keep them as simple as possible. The main reason was to ensure usability for non-academic users (that is, the general public), and to ensure that users are confident in the accuracy of results (transparency and reproducibility). These design goals are not always adopted in scientific publications. The modular approach taken as described below allows for systematic improvement of the model by focusing on one component at a time. Note that the source code for our implementation is publicly available, because we used a Javascript implementation [6] readable by most web browsers.

2.1. **Assumptions made for informational/complexity reasons.** In order for such a simulator to be practically usable, it is necessary to limit the amount of information entered by the user. For example, it is not feasible to assume that a user can input all information on voter preferences over parties in each electoral district. Given m parties and n voters, the number of possible voting situations is $\binom{n+m!-1}{n}$. Each voting situation can be specified by a $m!$-tuple of nonnegative integers with sum of entries equal to n. If n is much larger than $m!$, an efficient way to encode this is to list the $m!$ entries. If $m!$ is much larger than n, then most entries will be zero, and it is better to list only the nonzero entries, along with the corresponding preference orders (each a list of length m). But in any case, except for very small values of $n+m$, a full specification of the voting situation is too much information to expect.

Thus we decided early on that no district-specific or candidate-specific information would be allowed to be entered, and restricted the user (with minor exceptions noted below) to entering the national support level for each party. More specifically, for each party we allow only the percentage of voters for which that party is the first choice. This is consistent with opinion poll results in most electoral jurisdictions — it is costly to elicit further preferences even in jurisdictions in which these preferences are considered by the preference aggregation method used.

This solves the usability issue, but the lack of information immediately leads to several challenging problems, discussed in the following sections.

New district	Comprises (existing districts)
0	$\frac{7}{12}$(Whangarei)
1	$\frac{5}{12}$(Whangarei), $\frac{2}{12}$(Northland)
2	$\frac{7}{12}$(Northland)
3	$\frac{4}{12}$(Rodney), $\frac{3}{12}$(Northland)
$\vdots$	

TABLE 1. Some virtual districts formed by combination of existing districts

The party support is translated into votes in different ways, depending on the electoral system:

MMP:: cast party vote for the favoured party;

FPP:: cast vote for the favoured party's candidate;

PV:: cast first-preference vote for the favoured party's candidate, with (possibly) subsequent preferences for other parties' candidates (see below);

STV:: place the favoured party's candidates, in the party's chosen order, at the top of the preference list, followed (possibly) by the candidates of other parties (see below);

SM:: cast party vote for the favoured party, and electorate vote for the favoured party's candidate.

2.2. **Redistricting.** All of the alternatives given involve independent electoral districts, and the legislation setting up the referendum stipulated that the systems should be compared on the assumption of a 120-seat parliament for each. Since there are currently only 70 electoral districts used, this required us to perform redistricting. This is an interesting problem in general, with many known algorithms [4], but somewhat incidental to our main purpose. Thus we adopted a rather ad hoc approach which seems sufficient for our immediate purposes, but which raises some questions for further research.

Our approach was to build on the current electoral districts rather than constructing new ones from scratch. We first considered the existing graph of electoral districts, where two districts are joined by an edge if and only if they are adjacent on the map. (In a few cases, such adjacencies were disregarded: New Zealand is a mountainous country, and two districts with a common border may not be particularly close in human-geographical terms.) We then formed new virtual districts (only the vote percentages in each district are important) by taking convex combinations of old adjacent districts. The table below illustrates the process for the first few virtual districts in a division into 120 districts. The population of each such district must, of course, be 7/12 of the population of an existing district (which are themselves of equal population).

This virtual redistricting process can be described as follows. Let x_{ij} denote the fraction of population of new district i coming from old district j. There are constraints on total population: $\sum_j x_{ij} = 1$ and $\sum_i x_{ij} = 7/12$. The other constraints come from the adjacency matrix: if $a_{j,j'} = 0$ then $x_{ij}x_{ij'} = 0$ for all i. Given these constraints, any desired objective function can be used. For example, we could try to minimize $\sum_{i,j} x_{ij}^2$, in order to obtain districts involving manageable convex combinations.

For the specific problem at hand, a simple greedy algorithm sufficed for our purposes. Starting from a node of degree 1 in the original graph, we can form the districts in a straightforward manner. We feel that the general problem deserves further study.

2.3. **Disaggregating the national vote into districts.** For all systems except MMP, in order to compute the seat allocation in Parliament we must infer district-level information about votes for each party from the input data (the national vote total). This problem is often faced by election forecasters, for example in FPP jurisdictions such as UK and Canada. Of course, variation in preferences among the districts is essential for any interesting results — otherwise, the plurality winner at the national level would win all the seats in Parliament.

There is a growing theory of disaggregation used by election forecasters (see for example [3, 2, 9]). Some translate the "popular vote" (the national support level for each party) directly into numbers of seats, perhaps by computing the expectation with respect to a postulated distribution. These methods typically use very detailed information about past elections to fine-tune parameters in the model. Others compute the result in each seat (perhaps probabilistically) and then aggregate to the national level. Our design goals mentioned above preclude the use of these more complex methods. In addition to issues of transparency to end users, there is always the unpleasant possibility of overfitting. We did consider probabilistic allocation of seats to parties, which should be more accurate (consider a country where one party has probability more than 1/2 of winning each seat). However in order to use probabilistic allocation, a model is needed and we found no compelling simple such model. The accuracy of our results shows that in the New Zealand context, at least, our deterministic allocation worked well, perhaps because relatively few seats involve very close elections.

The two most simplest methods for inferring seat distribution data are the so-called Uniform Swing Hypothesis (USH) and Proportional Loss Hypothesis (PLH), which belong to the second category above. These simple models work as follows. Each starts with a reference point (usually the actual votes cast at the previous election) and uses the change in popular vote to infer changes in each district. USH assumes that each party's vote fraction in each electorate changes additively in the same way that its popular vote does. PLH assumes that the distribution across districts of each party's vote remains the same.

For example, consider a situation with 2 districts (of equal size D, for simplicity) and 2 parties, in which the vote fractions are initially 0.6 for party A and 0.4 for party B in district 1, and 0.4 for party A and 0.6 for party B in district 2. The popular vote fractions are 0.5 for each party. If the popular vote changes to 0.6 for party A and 0.4 for party B, then USH predicts that in district 1, the respective vote fractions will become 0.7 and 0.3, and in district 2 they will change to 0.5 and 0.5. By contrast, PLH predicts that the vote fractions will be proportional to $(0.72D, 0.32D)$ and $(0.48D, 0.48D)$ in districts 1 and 2 respectively. In other words, they are $(0.692, 0.308)$ and $(0.5, 0.5)$.

Note that PLH requires that the number of voters be allowed to vary, while USH can lead to predictions of negative number of votes for a party in a district. Each method determines a definite winner for each district, and these are aggregated to obtain the seat totals in Parliament. An alternative approach is to assign each district to each party with some probability, and then compute the expected number of districts won.

In our simulator we used PLH and the reference point was the national total MMP party vote at the 2008 NZ general election. We used the party vote rather than the electorate vote because we believe it more likely to reflect voters' sincere preferences and to be immune to district- and candidate-specific influence. For example, under FPP it seems likely that votes for local candidates correlate highly with overall party support, and parties are much more stable over time than candidates.

2.4. **Inferring preference orders.** Limiting the user input to just the percentage of first preferences given to each party means that in order to simulate the preference-based methods (in our case, PV and STV), we must infer the distribution of other preferences. We did this by estimating the distribution of preferences conditional on the first preference, for each possible first preference. The estimator was the empirical conditional distribution from survey data [8]. The survey was held just after the 2008 election and included a question allowing users to score each political party out of 10. We converted this to preference orders, breaking ties in scores by assigning the probability mass corresponding to a tie equally to each tied party. Note that this means that a respondent rating all parties as equally good will contribute a more than average weight to the "minor" parties.

When the number of parties is not small, the rapid growth of the number of profiles means that even storing the inferred preference information is nontrivial, as described in Section 2.1.

3. Numerical results

We ran the simulator on some historical New Zealand general elections. Results (using data from [11]) are shown in Table 2 and Table 3. For elections since 1996 under the MMP system, the simulator worked exactly, provided we enter the number of electorate seats for each "minor" party (the only exception to our assumption of no district or candidate-specific information). This is as expected, since there is no need to infer missing information in this case. We present our simulated predictions for the result of the 2011 election under each of the electoral systems under consideration, assuming no change in voting behaviour from the actual election.

Prior to 1996 New Zealand used the FPP system, with a number of districts that was fixed at 80 in 1902 and gradually increased over time from 1969, reaching 99 in 1993. Because of lack of information on the historical district adjacency graph, we simply used the same 120 districts as above and linearly scaled the seat allocations for each party. We used the popular vote as input. We go back only to 1935, because before that the familiar party structure (with two main “left" and “right" parties, and a few much smaller parties) did not exist. The two main parties, which still exist today, are the National Party (NAT) and the Labour Party (LAB). We refer to all other parties and independents as OTHER. The composition of OTHER has varied considerably in the last century. Since 1996 the parties actually represented in Parliament have been the Green Party (GREEN), New Zealand First Party (NZF), ACT New Zealand (ACT), the Maori Party (MAORI), the Alliance Party (ALL), the Mana Party (MANA). The new Conservative Party (CON) would have won a seat under SM according to Table 2.

Method	NAT	LAB	GREEN	NZF	ACT	MAORI	MANA	UNI	CON
MMP	59	34	14	8	1	3	1	1	0
FPP	99	21	0	0	0	0	0	0	0
PV	86	34	0	0	0	0	0	0	0
STV	67	46	4	0	0	3	0	0	0
SM	90	24	3	2	0	0	0	0	1

TABLE 2. Simulated seat distributions in NZ Parliament for 2011 election using MMP party vote as input

Year	NAT simulated	LAB simulated	others simulated	NAT actual	LAB actual	others actual
1993	54	45	0	50	45	4
1990	71	26	0	67	29	1
1987	50	47	0	57	40	0
1984	37	58	0	37	56	2
1981	48	44	0	47	43	2
1978	48	44	0	51	40	1
1975	55	32	0	55	32	0
1972	33	54	0	32	55	0
1969	47	37	0	45	39	0
1966	47	33	0	44	35	1
1963	47	33	0	45	35	0
1960	48	32	0	46	34	0
1957	41	39	0	41	39	0
1954	43	37	0	45	35	0
1951	51	29	0	50	30	0
1949	48	32	0	46	34	0
1946	39	41	0	38	42	0
1943	37	43	0	34	45	0
1938	21	59	0	25	53	2
1935	20	60	0	19	53	8

TABLE 3. Real and simulated seat distributions in NZ Parliament – historical FPP elections using FPP popular vote as input

As suggested by Table 2, the typical outcomes produced by the various electoral systems fall in a roughly linear order. At one end of the spectrum lies the most majoritarian system, FPP, which produces two-party Parliaments in which the more popular party is disproportionately favoured. It is possible that our simulator produces even more lopsided results than might occur in reality, due to its neglect of the characteristics of individual candidates: a relatively popular candidate who has the misfortune to represent the less-popular party may be able to hold on to a seat that would otherwise change hands.

It is interesting that the PV system gives generally similar results to FPP. Considering the lower-order preferences clearly makes a difference – in particular, it prevents the most popular party from winning some districts it could win under FPP, by allowing opposing votes to be united behind a single candidate, and thus produces less disproportionate Parliaments. But, at least for the New Zealand party structure we are considering, the effect of the lower-order preferences almost never propels a third-party candidate into a winning position.

The SM system also gives results rather like those of FPP; this is no surprise as 75% of the seats are being allocated exactly as FPP. Proportionality in the allocation of the remaining seats allows a few seats to be held by minor parties.

At the other end of the spectrum lie the (near-) proportional outcomes of MMP. This system produces essentially proportional results for the two main parties, and is the most favourable to minor parties. It is perhaps a little surprising that the simulated version of STV does not match these outcomes more closely. To some extent, this may reflect the districting choices made: STV with 3-seat districts (which we used in rural parts of New Zealand, to keep the districts to a reasonable geographic size) is more majoritarian than proportional. Only in the 5, 6, or 7-seat districts (used in urban areas) is something like proportionality achievable. It may also reflect the party structure: support for New Zealand's minor parties (except for the Maori Party) is not concentrated in particular districts, but rather uniformly spread over all districts.

3.1. **Validating our approach.** The general procedure we have used rests on a rather simple model (compared to some used by election forecasters, for example). All that is required is an estimate of the overall preference distribution, and an estimate of the distribution of party support by district.

In order to have more confidence in the results, we used a student project to reimplement the simulator more generally, allowing for a general number of districts and parties, variable district magnitude, variable thresholds, more voting rules, and user-supplied reference data as listed in the last paragraph. The same type of numerical results as above were computed for general elections in Canada and UK (difficulties with availability of data for other FPP jurisdictions have so far limited us to these two countries). The results are shown in Tables 4 and 5 (CON, LAB, LIB stand for (Progressive) Conservative, Labour and Liberal parties, and OTHER for the total of all other parties). Note that no redistricting was needed for these results — we used scaling as discussed in the previous subsection for the few years in which the total number of seats differed from the current value. The reference points for data on preferences and spatial distribution of parties were from surveys and election results in 2001/2 and 2005 in UK and 2008 and 2011 in Canada, respectively.

The 2010 UK general election generated substantial interest in the election forecasting research community. Using as input opinion poll data from several months before the election, our simulator results, our simulator gave results that were less accurate than those presented in [3]. Using polling data from much closer to the election gave much more accurate results, similar to those in Table 4 – in fact these results compare favourably to election eve predictions by pollsters. Since our simulator was not intended to forecast changes in party support over an election campaign, this is not surprising. We would need to add a model of preference change to our current model in order to make more realistic forecasts, and this is a possible topic for future work.

The results presented here show that the simple assumptions underlying our simulator appear to have led to a surprisingly accurate simulator of election results in the FPP context, particularly in the New Zealand context. It is not so surprising that it has value in predicting the results of an election given the results of the last election and actual popular vote counts (or opinion poll data from close to the election). However (to us at least) it is surprising how accurate the results are on historical data, when one would expect the model to be stretched to breaking point. The seat numbers, seat boundaries, preference correlations and spatial distribution of party support (not to mention parties themselves) have presumably all changed, yet results are rather accurate.

Time constraints have not permitted us to test our model's predictions for data from a country such as Australia, which uses forms of STV for many public elections. This is an obvious next step in the refinement of the simulator.

Year	CON simulated	LAB simulated	OTHER simulated	CON actual	LAB actual	OTHER actual
2010	291	258	96	306	258	81
2005	196	331	118	198	355	93
2001	186	391	68	166	413	62
1997	186	407	66	165	418	76

TABLE 4. Real and simulated seat distributions in UK House of Commons using FPP popular vote as input

Year	CON simulated	LIB simulated	OTHER simulated	CON actual	LIB actual	OTHER actual
2011	154	95	59	166	103	39
2008	144	97	67	143	77	88
2006	144	97	67	124	103	81
2004	111	128	69	107	135	66

TABLE 5. Real and simulated seat distributions in Canadian House of Commons using FPP popular vote as input

One of the motivations behind this project was to raise the level of public debate on the 2011 New Zealand referendum. The authors believed that the debate generated by the 2011 UK referendum on the Alternative Vote (in our terminology, PV) engendered a very low level of public understanding. It is natural to look at simulated results on this topic. Our methodology predicts that the 2010 UK general election under PV would have led to Labour winning a plurality of seats, but not an outright majority (assuming as usual the same voter behaviour). By contrast, a recent analysis of the same question [7], using survey data and more district-specific analysis, yields results of 284, 248, 89 seats for the Conservative, Labour and Liberal Democrat parties, respectively. This discrepancy seems a good candidate for future work, which we intend to pursue. We also computed the results for Canada. The simulator predicts that the 2011 election, if results were computed under PV with no change in voting behaviour, would have led to the Liberal party remaining as the second biggest party, instead of the dramatic fall it had in reality.

4. CONCLUSION

The simulation procedure adopted here is appealing simple and yet appears to yield rather reliable results. Future refinements are clearly possible, without destroying the basic design simplicity. For example, all known indices of representation, governability, fragmentation, etc can be computed easily from the input and output data.

The only other simulator of this type to our knowledge is the one developed over several years by Bissey and coauthors, known as ALEX [1], which found only after completing the work described above. Our brief perusal of the documentation for lits latest version (4.1) leads us to believe that our assumptions on preferences are more realistic (ALEX uses a single-peaked model) as is the spatial distribution of party support. We have not been able yet to make a direct comparison of the two systems in order to compare their predictions.

Acknowledgement: the authors thank Michael Fowlie and Siyuan Zhang for their assistance with numerical results described in this section.

REFERENCES

[1] Marie-Edith Bissey, Mauro Carini and Guido Ortona. "ALEX3: a Simulation Program to Compare Electoral Systems". In: *Journal of Artificial Societies and Social Simulation* 7 (2004). URL: `http://jasss.soc.surrey.ac.uk/7/3/3.html` (cit. on p. 7).

[2] Stephen D Fisher et al. “From polls to votes to seats: Forecasting the 2010 British general election”. In: *Electoral Studies* 30.2 (June 2011), pp. 250–257. URL: http://linkinghub.elsevier.com/retrieve/pii/S0261379410000946 (cit. on p. 3).

[3] Rachel Gibson and Michael S Lewis-Beck. “Methodologies of election forecasting: Calling the 2010 UK “hung parliament””. In: *Electoral Studies* 30.2 (June 2011), pp. 247–249. URL: http://linkinghub.elsevier.com/retrieve/pii/S0261379410000922 (cit. on pp. 3, 6).

[4] Pietro Grill di Cortona et al. *Evaluation and Optimization of Electoral Systems*. SIAM, 1999 (cit. on p. 3).

[5] New Zealand Government Parliamentary Counsel Office. *Electoral Referendum Act 2010*. Accessed 2012-02-13. URL: http://www.legislation.govt.nz/act/public/2010/0139/latest/DLM2833513.html (cit. on p. 1).

[6] Geoffrey Pritchard and Mark C. Wilson. *2011 referendum simulator*. Accessed 2012-02-13. URL: http://www.stat.auckland.ac.nz/\~geoff/voting/ (cit. on pp. 1, 2).

[7] D Sanders et al. “Simulating the effects of the alternative vote in the 2010 UK general election”. In: *Parliamentary Affairs* 64.1 (2011), pp. 5–23. URL: http://pa.oxfordjournals.org/content/64/1/5.short (cit. on p. 7).

[8] Jack Vowles et al. *2008 New Zealand Election Study*. Accessed 2012-02-13. URL: http://www.nzes.org/exec/show/2008/ (cit. on p. 4).

[9] Paul Whiteley et al. “Aggregate level forecasting of the 2010 general election in Britain: The Seats-Votes model”. In: *Electoral Studies* 30.2 (June 2011), pp. 278–283. URL: http://linkinghub.elsevier.com/retrieve/pii/S0261379410000995 (cit. on p. 3).

[10] Elections New Zealand. *2011 Referendum website*. Accessed 2012-02-13. URL: http://www.referendum.org.nz/ (cit. on p. 1).

[11] Elections New Zealand. *General elections 1890–1993 - seats won by party*. Accessed 2012-02-13. URL: http://www.elections.org.nz/elections/resultsdata/fpp-seats-won.html (cit. on p. 4).