

The Overhang model

Polls, voter flows, seats and election night

How the forecast behind Ten Thousand Beehives works¹, 21 October 2026

Abstract. The model turns published polls into a forecast of the New Zealand party vote, the seats that follow from it, and, on election night, a running estimate from voting places as they report. Party shares are treated as a composition and modelled in isometric log-ratio coordinates. Opinion is a latent state that drifts week by week, with the covariance of its movements built from voter flows fitted to the NZ Election Study. Pollsters have house effects and polls are noisier than sampling error alone. The latent paths are drawn exactly by Kalman filtering, a damped drift lets trends persist, and the dynamics are chosen by how well they predict later polls. A systematic error that grows with distance from the election, calibrated on four past elections, is added to the forecast, which is then turned into seats by the threshold, Sainte-Laguë and overhang rules, with the electorate winners from a model calibrated on past elections. A spatial model of voting-place swings, still being built, will update the forecast as places report.

Status: sections 3 to 8 are implemented and tested. Section 9 describes the election-night model: the data layer is built, the spatial fit is in progress, and the live update is not yet written.

Contents

1	The pipeline	2
2	Compositions	3
3	The latent opinion model	3
4	Exact latent paths and momentum	11
5	From the latent state to the election	14
6	Seats and government	16
7	Electorate seats	20
8	Scenarios	31
9	Election night	34
10	Assumptions and limitations	37
11	Appendix: notation	39
12	Appendix: derivations	44

¹In 1999, 81% of voters asked for a smaller Parliament. We built ten thousand of them instead. Sorry.

1 The pipeline




Figure 1: The components and where each is described. The polling model runs before the election and supplies the prior for the election-night model; both forecasts feed the seat model.

The code is organised in layers that match this picture. Table 1 maps each section to its modules. Each symbol is defined where it first appears, and the notation appendix (Section 11) lists them all with their types.

Section	Content	Modules (beehives.)
2	compositions, ILR, multinomial noise	core.compositional, core.parties
3	latent opinion model, flows, gravity model	inference.model, flows.flows, flows.survey_flows, sources.nzes
4	exact paths, momentum, calibration	inference.smoothing, inference.backtest
5	forecast and systematic error	inference.model, inference.retrospective
6	seats and government	electoral.seats, electoral.forecast, core.parties
7	electorate seats	electoral.electorate_model, electoral.swing_meter, sources.electorate_panel, sources.electorate_baseline
8	scenarios	electoral.scenarios, electoral.scenario_catalogue, electoral.scenario_runs, sources.ministers
9	election night	sources.booths, sources.clean, inference.booth_panel, inference.swing_field

Table 1: Where each part of the model lives.

2 Compositions

A party-vote result is a composition: $K = 8$ non-negative parts that sum to one, in the order National, Labour, Green, ACT, NZ First, Te Pāti Māori, Opportunity, Other. Write $\mathbf{x} \in \Delta^{K-1} = \{\mathbf{x} \in \mathbb{R}^K : x_i \geq 0, \sum_i x_i = 1\}$. Other is a residual bucket for every smaller party. The party vote is a composition because only ratios between parties carry information, and because a gain for one party is a loss for the others. We therefore avoid modelling shares directly and use log-ratios.

2.1 The isometric log-ratio transform

Let $V \in \mathbb{R}^{K \times (K-1)}$ be the Helmert-type basis whose column j ($j = 1, \dots, K-1$) is

$$V_{ij} = \begin{cases} 1/\sqrt{j(j+1)} & \text{if } i < j \\ -j/\sqrt{j(j+1)} & \text{if } i = j \\ 0 & \text{if } i > j \end{cases}$$

(parts indexed from 0). The columns are orthonormal and each sums to zero. The ILR coordinates of a composition are

$$\mathbf{z} = \text{ilr}(\mathbf{x}) = V^\top \log \mathbf{x} \in \mathbb{R}^{K-1},$$

with x_i floored at 10^{-3} and renormalised before the logarithm so that a party reported at 0% does not give an infinite value. Coordinate j is the balance of parts $0, \dots, j-1$ against part j : it equals $\sqrt{j/(j+1)} \log(\text{gm}(x_0, \dots, x_{j-1})/x_j)$ with gm the geometric mean. The inverse is the softmax

$$\begin{aligned} \mathbf{x} &= \text{softmax}(V\mathbf{z}) \\ \text{softmax}(\mathbf{q})_i &= \frac{\exp(q_i)}{\sum_k \exp(q_k)}. \end{aligned}$$

Because V is orthonormal, Euclidean distance between ILR vectors is the Aitchison distance between compositions, and priors on the coordinates do not depend on which party is treated as a reference. This is why ILR is used throughout in place of an additive log-ratio, which would single one party out.

2.2 Sampling noise of a poll

A poll of effective size n with observed shares $\mathbf{x}$ has ILR covariance, by the delta method,

$$\Sigma(\mathbf{x}, n) = \frac{1}{n} V^\top \text{diag}(1/\mathbf{x}) V, \quad (1)$$

(the usual rank-one term vanishes because the columns of V sum to zero). Small shares have large log-ratio variance, which is correct: a party at 3% is measured far less precisely, in relative terms, than one at 30%. The sections below write Σ_p for its value for poll p (Section 3.4). The notation appendix (Table 10) lists every symbol.

3 The latent opinion model

3.1 State and dynamics

Opinion is a latent ILR vector $\mathbf{z}_t$ on a weekly grid $t = 0, 1, \dots, T$ that starts at the previous election. The start is the previous election's result $\mathbf{z}^{\text{prev}} = \text{ilr}(\mathbf{x}^{\text{prev}})$, held softly:

$$\mathbf{z}_0 \sim \mathcal{N}(\mathbf{z}^{\text{prev}}, \sigma_0^2 I),$$

with $\sigma_0 = 0.15$. A daily Gaussian random walk observed weekly is a weekly walk with variance scaled by the grid step $h = 7$ days, so the model is sampled on the weekly grid:

$$\begin{aligned} \mathbf{z}_{t+1} &= \mathbf{z}_t + \boldsymbol{\varepsilon}_t \\ \boldsymbol{\varepsilon}_t &\sim \mathcal{N}(\mathbf{0}, hQ). \end{aligned}$$

In the sampler the walk is written in non-centred form,

$$\mathbf{z}_t = \mathbf{z}^{\text{prev}} + \sigma_0 \boldsymbol{\varepsilon}^{\text{init}} + \sum_{t' < t} \sqrt{h} L \boldsymbol{\varepsilon}_{t'}^{\text{std}}$$

with $LL^\top = Q$ and $\boldsymbol{\varepsilon}_{t'}^{\text{init}}, \boldsymbol{\varepsilon}_{t'}^{\text{std}} \sim \mathcal{N}(\mathbf{0}, I)$, which makes the posterior geometry easy for the optimiser. Section 4 adds a drift to this state.

3.2 Where the covariance Q comes from: voter flows

A random walk with independent coordinates would let National and ACT move independently of Labour and Green, but voters move between particular parties. Q is therefore built from a set of directed flows, one for each ordered pair of distinct parties ($n_{\text{flow}} = K(K - 1) = 56$). Suppose a fraction χ of party j 's voters move to party i , the flow $f = (j \rightarrow i)$. To first order the log-shares change by $\Delta \log x_j = -\chi$ and $\Delta \log x_i = \chi x_j / x_i$, and the others do not change. A composition is defined up to scale, so the change is centred. Let $\mathbf{a}_f$ be the vector with $(\mathbf{a}_f)_j = -1$, $(\mathbf{a}_f)_i = x_j^{\text{ref}} / x_i^{\text{ref}}$ and zeros elsewhere, evaluated at reference shares $\mathbf{x}^{\text{ref}}$ (the average of the polls' reported shares). The direction of the flow in log-share space is $\mathbf{a}_f - \bar{a}_f \mathbf{1}$, with $\bar{a}_f$ the mean of its entries, and in ILR coordinates it is the row

$$D_f = (\mathbf{a}_f - \bar{a}_f \mathbf{1})^\top V.$$

If flow f moves a fraction with standard deviation r_f per day, the covariance of movements is

$$\begin{aligned} Q &= \sum_f r_f^2 D_f^\top D_f + (\sigma_{\text{iso}}^2 + 10^{-10}) I \\ &= D^\top \text{diag}(\mathbf{r}^2) D + (\sigma_{\text{iso}}^2 + 10^{-10}) I, \end{aligned} \quad (2)$$

where $D \in \mathbb{R}^{n_{\text{flow}} \times 7}$ stacks the rows D_f , $\mathbf{r}$ collects the rates, $\sigma_{\text{iso}} \sim \text{HalfNormal}(0.002)$ is a small isotropic term for movement not explained by any listed flow, and the jitter 10^{-10} keeps Q positive definite. Parties that trade voters therefore move together (one's gain is the other's loss) and parties with no flow between them, such as ACT and Te Pāti Māori, barely move together.

The rates are priors with a learned multiplier:

$$\begin{aligned} r_f &= r_f^0 \exp(u_f / 2) \\ u_f &\sim \mathcal{N}(0, 1). \end{aligned} \quad (3)$$

3.3 Where the base rates come from: a gravity model of switching

The base rates r_f^0 are fitted to the NZ Election Study (NZES) 2023, using respondents' recalled 2020 vote and 2023 vote. A voter who voted for party j at the previous election chooses destination i by a conditional logit (the same direction as the flow $f = (j \rightarrow i)$),

$$\begin{aligned} P(i | j) &= \frac{\exp(v_{ji})}{\sum_k \exp(v_{jk})} \\ v_{ji} &= \lambda_j \mathbf{1}[j = i] + \pi_i + \beta \psi_{ji}, \end{aligned}$$

with a loyalty λ_j for each source party (the utility of staying), a pull π_i for each destination, and one covariate ψ_{ji} : how much party j 's voters like party i relative to their own party (from the NZES like-scores), with weight β . It was the only covariate retained, because it predicted held-out parties best. Parameters are fitted by penalised maximum likelihood (a ridge on loyalty and pull) to the observed switching counts, with 15 parameters for 64 cells, so sparse cells borrow strength from the rest. A fitted switching probability over one election cycle of n_{cyc} days is turned into a daily rate by

$$r_f^0 = \frac{P(i | j)}{\sqrt{n_{\text{cyc}}}}, \quad (4)$$

for the flow $f = (j \rightarrow i)$, so that the displacement of the random walk over one cycle, with standard deviation $\sqrt{n_{\text{cyc}}}r_f^0 = P(i \mid j)$, equals the observed probability of switching. The electorate model fits a gravity model of the same form to how voters split their party and candidate votes (Section 7.2).




Figure 2: Probability that a voter of the source party moves to the target party over one election cycle, from the 2020 to 2023 NZES. Edge width and opacity grow with the probability, edges under 3% are omitted, and the percentage in each node is the share who stay.

Utilities in the gravity model. The affinity coefficient on the like-score is 0.49.




Figure 3: Fitted parameters of the NZES gravity model. Loyalty is the utility of staying with the party and pull is the utility of moving to it. The coefficient on the affinity covariate is 0.49.

Caveat. One cycle (2020 to 2023, the unwinding of Labour’s landslide) is not a stable measure of affinity. The multiplier u_f in Equation 3 lets the polls correct each rate. More NZES waves (2014 to 2020) would give more cycles, and allow a test of whether flows are stable between cycles.

Figure 4 and Figure 5 show what the polls do to these rates. Every flow is shrunk or left alone and none is substantially enlarged: the largest ratios of learned to prior rate are about 1.1 (Labour to Opportunity), and the biggest cuts are ACT to National (0.24), National to ACT (0.27), Labour to NZ First (0.30) and Other to National (0.35). The survey cycle was unusually fluid, so a prior taken from it alone overstates how much voters move in a typical three years.




Figure 4: The same probabilities after fitting to the polls, from the posterior median of each flow rate, on the same scale as Figure 2.




Figure 5: Ratio of the learned to the prior daily flow rate for each pair of parties. Orange cells are flows the polls shrank relative to the NZES prior, and no flow is much larger than its prior.

3.4 The observation model

Poll p is taken by pollster $a(p)$ with fieldwork midpoint at day t_p after the cycle start, so its grid position is t_p/h . Write $t_{l_0} = \lfloor t_p/h \rfloor$ and $w_p = t_p/h - t_{l_0}$. The latent state at the poll is the linear interpolation

$$\mathbf{z}(t_p) = (1 - w_p)\mathbf{z}_{t_{l_0}} + w_p\mathbf{z}_{t_{l_0}+1}.$$

Polls report different parties. A party a poll does not report (a dash in the source table) is not a zero: its votes sit in Other for that poll. Polls are therefore grouped into blocks m by which parties they report. Block m reports K_m categories (always including Other), and $M_m \in \{0, 1\}^{K \times K_m}$ sends each party to its category, with every unreported party sent to Other. With V_m the ILR basis for K_m parts, the model prediction for a poll in block m is the ILR of the merged composition, given by the observation function

$$\text{obs}_m(\mathbf{z}, \boldsymbol{\eta}) = V_m^\top \log(M_m^\top \text{softmax}(V(\mathbf{z} + \boldsymbol{\eta}))) \in \mathbb{R}^{K_m-1}. \quad (5)$$

The observation $\mathbf{y}_p$ is the ILR of the poll's reported shares $\hat{\mathbf{x}}_p$ in that basis, and

$$\mathbf{y}_p \sim \mathcal{N}(\text{obs}_m(\mathbf{z}(t_p), \boldsymbol{\eta}_{a(p)}), \delta \Sigma_p)$$

where $\Sigma_p = \Sigma_m(\hat{\mathbf{x}}_p, n_p)$. Σ_m is Equation 1 in the basis V_m (with $n_p = 1000$ if a poll reports no sample size) and $\delta \geq 1$ is a design effect that inflates polls' noise beyond sampling error: $\delta = 1 + \text{HalfNormal}(1)$.

Figure 6 shows the observations the model is given: the 128 polls in the data, with fieldwork ending between 6 November 2023 and 19 October 2026, one point per poll and party, with the shape marking the polling firm. These are the reported shares before any house effect is removed, so the firms sit at different heights within a panel. Opportunity, for example, sits at a median of 2.0% in the polls ending before 2026 and 5.7% in the 20 polls ending since 1 July 2026.




Figure 6: Every poll in the data by party, as reported. A party a poll did not report is left out of its panel, not drawn at zero. Each panel has its own vertical scale.

3.5 House effects

Each pollster a has a bias $\eta_a \in \mathbb{R}^7$ in ILR coordinates, relative to the consensus of the others. To pool series that share a firm $g(a)$ (every Curia series, for instance),

$$\eta_a = \sigma_{\text{firm}} (\xi_{g(a)} - \bar{\xi}) + \sigma_{\text{ser}} \zeta_a, \quad (6)$$

with standard-normal draws $\xi_g, \zeta_a \sim \mathcal{N}(0, I)$, the centring $\bar{\xi}$ over firms (so the effects are relative and not absolute), $\sigma_{\text{firm}} \sim \text{HalfNormal}(0.08)$ and a smaller series-level deviation $\sigma_{\text{ser}} \sim \text{HalfNormal}(0.04)$. House

effects are constant in time. Because they are centred across firms, a bias shared by every pollster is not identifiable from polls alone and is the job of the systematic error in Section 5.2.

Figure 7 shows the fitted effects. Roy Morgan is the clear outlier: at the latest state it runs about 3.8 points below the consensus for Labour and about 1.8 above for Green.




Figure 7: Effect of each pollster's house effect on each party's share at the latest latent state, in percentage points, as the median and 80% interval over posterior draws. Each panel has its own scale.

Quantity	Prior	Role
z_0	$\mathcal{N}(z^{\text{prev}}, 0.15^2 I)$	start at the previous result
u_f	$\mathcal{N}(0, 1)$	log-scale multiplier on each flow rate
σ_{iso}	HalfNormal(0.002)	movement not due to any listed flow
$\sigma_{\text{firm}}, \sigma_{\text{ser}}$	HalfNormal(0.08), HalfNormal(0.04)	firm and series house effects
δ	$1 + \text{HalfNormal}(1)$	design effect

Table 2: Priors of the polling model (SVI stage).

3.6 Inference

The joint posterior over the path, flow multipliers, house effects, design effect and σ_{iso} is approximated by stochastic variational inference (SVI): maximise the evidence lower bound over a low-rank multivariate normal guide (rank 10), with the ELBO compiled by tracing for speed (about 2,000 steps, about two minutes). The low-rank covariance captures the strongest correlations along the random walk that a mean-field guide would miss.

NUTS is available for a slow, final fit. SVI is good at the low-dimensional hyperparameters, but not at how tightly the polls pin a long path, so the paths are replaced in the next section.

4 Exact latent paths and momentum

4.1 Exact paths given the hyperparameters

SVI's path uncertainty is too wide and grows with time (the guide cannot express how tightly weekly polls pin a long walk). But given one draw of the hyperparameters $(Q, \boldsymbol{\eta}, \delta)$ the model is linear and Gaussian apart from the merging of unreported parties, so the path can be drawn exactly.

Each poll is assigned to its nearest grid week. A poll that reported every party observes the state directly: $\tilde{\mathbf{y}}_p = \mathbf{y}_p - \boldsymbol{\eta}_a = \mathbf{z}_t + \boldsymbol{\nu}_p$ with $\boldsymbol{\nu}_p \sim \mathcal{N}(\mathbf{0}, \delta \Sigma_p)$. A poll with merged parties is linearised about the current state estimate $\mathbf{z}^*$: with $\mathbf{x} = \text{softmax}(V(\mathbf{z}^* + \boldsymbol{\eta}))$ and $\mathbf{x}^{\text{mrg}} = M_m^\top \mathbf{x}$,

$$\begin{aligned} \mathbf{y}_p &\approx \text{obs}_m(\mathbf{z}^*, \boldsymbol{\eta}) + J(\mathbf{z} - \mathbf{z}^*) \\ J &= V_m^\top \text{diag}(1/\mathbf{x}^{\text{mrg}}) M_m^\top (\text{diag}(\mathbf{x}) - \mathbf{x} \mathbf{x}^\top) V, \end{aligned} \quad (7)$$

where obs_m is Equation 5. Write $\tilde{\mathbf{y}}_p$ for the poll's observation reduced to the linear form $\tilde{\mathbf{y}}_p = H \mathbf{z}_t + \boldsymbol{\nu}_p$, with $H = I$ for a fully reporting poll and $H = J$ for a merged one. Then the Kalman filter folds in each week's polls. With filtered state mean $\hat{\mathbf{z}}_{t|t}$ and covariance P_t , and for the random walk $\hat{\mathbf{z}}_{t|t-1} = \hat{\mathbf{z}}_{t-1|t-1}$:

$$\begin{aligned} P_{t|t-1} &= P_{t-1} + hQ \\ S &= H P_{t|t-1} H^\top + \delta \Sigma_p \\ G &= P_{t|t-1} H^\top S^{-1} \\ \hat{\mathbf{z}}_{t|t} &= \hat{\mathbf{z}}_{t|t-1} + G(\tilde{\mathbf{y}}_p - H \hat{\mathbf{z}}_{t|t-1}) \\ P_t &= P_{t|t-1} - G H P_{t|t-1}. \end{aligned} \quad (8)$$

A backward pass draws a path: $\mathbf{z}_T \sim \mathcal{N}(\hat{\mathbf{z}}_{T|T}, P_T)$ and then, for $t = T - 1, \dots, 0$ (forward filtering, backward sampling),

$$\begin{aligned} \mathbf{z}_t \mid \mathbf{z}_{t+1} &\sim \mathcal{N}(\hat{\mathbf{z}}_{t|t} + B_t(\mathbf{z}_{t+1} - \hat{\mathbf{z}}_{t|t}), P_t - B_t P_t) \\ B_t &= P_t(P_t + hQ)^{-1}. \end{aligned}$$

One path is drawn for each draw of the hyperparameters. The result is a posterior of the latent path whose width reflects how tightly the polls pin each week, not the variational approximation. (In one fit the raw SVI standard deviation grew from 0.36 to 2.96 points while the exact paths stayed at 0.7 to 1.3.)

4.2 A damped drift: opinion has momentum

The weekly random walk has no memory, but opinion trends. Over six months a party that has been rising tends to keep rising for a while, so the uncertainty of a forecast made months ahead grows faster than a random walk allows. The state is extended with a weekly drift $\mathbf{d}_t \in \mathbb{R}^7$:

$$\begin{aligned} \mathbf{z}_{t+1} &= \mathbf{z}_t + \mathbf{d}_t + \boldsymbol{\varepsilon}_t \\ \mathbf{d}_{t+1} &= \varphi \mathbf{d}_t + \boldsymbol{\varepsilon}_t^{\text{drift}} \\ \boldsymbol{\varepsilon}_t &\sim \mathcal{N}(\mathbf{0}, c^2 h Q) \\ \boldsymbol{\varepsilon}_t^{\text{drift}} &\sim \mathcal{N}(\mathbf{0}, \kappa^2 c^2 h Q) \\ \mathbf{d}_0 &\sim \mathcal{N}(\mathbf{0}, (\kappa^2 c^2 h Q)/(1 - \varphi^2)), \end{aligned} \quad (9)$$

with the drift's persistence $\varphi \in [0, 1)$ (its weekly decay), a level-shock scale $c > 0$ (a multiple of the fitted step size), and a trend scale $\kappa \geq 0$ (the size of drift shocks relative to level shocks). Both kinds of shock use the flow

covariance Q , so momentum moves along the directions in which voters trade. With $\kappa = 0$ and $c = 1$ this is the random walk. In the augmented state $\mathbf{z}_t^+ = (\mathbf{z}_t, \mathbf{d}_t) \in \mathbb{R}^{14}$ the dynamics are linear,

$$\begin{aligned}\mathbf{z}_{t+1}^+ &= F\mathbf{z}_t^+ + \boldsymbol{\varepsilon}_t^+ \\ F &= \begin{pmatrix} I & I \\ 0 & \varphi I \end{pmatrix} \\ \text{Cov}(\boldsymbol{\varepsilon}_t^+) &= \begin{pmatrix} c^2 h Q & 0 \\ 0 & \kappa^2 c^2 h Q \end{pmatrix},\end{aligned}$$

and polls observe only the level, $H = [H_z, 0]$, so Equation 8 applies to $\mathbf{z}_t^+$ with $\hat{\mathbf{z}}_{t|t-1}^+ = F\hat{\mathbf{z}}_{t-1|t-1}^+$ and $P_{t|t-1} = FP_{t-1}F^\top + \text{Cov}(\boldsymbol{\varepsilon}^+)$.

Figure 8 shows the fitted latent path for the larger and the smaller parties, with the 80% band from the exact paths of this section and the polls behind it. The band is narrow where polls are frequent and the path follows the sustained moves, such as Labour’s rise to a median of 34.4% on 28 February 2026 and its fall to 28.2% in the last grid week (24 October 2026). Dotted rules mark dated events.




Figure 8: Latent party shares by week, as the median and 80% band over posterior draws, over the polls (points), for the fit of 20 October 2026 to polls ending 19 October 2026. Top: the larger parties. Bottom: the smaller parties, with the 5% threshold dashed. Dotted verticals mark dated events, which the notebook names on hover.

4.3 Choosing the dynamics by how well they predict later polls

The dynamics (c, φ, κ) and a scale ρ on the design effect (the random walk's fit tends to blame poll noise for trends it cannot follow, so with momentum the noise may be lower, $\delta' = \max(\rho\delta, 1)$) are not learned by SVI. Fitting κ inside the variational model collapses it to zero, because the in-sample bound prefers smooth paths to good forecasts. Instead they are chosen by predictive performance, a rolling-origin backtest.

For each origin week t_0 (from week 52, leaving earlier weeks as warm-up) filter using only polls up to t_0 , predict ahead $s = 1, \dots, 26$ weeks: $z_{t_0+s}^+ \mid \text{polls up to } t_0 \sim \mathcal{N}(F^s \hat{z}_{t_0|t_0}^+, P_{t_0+s|t_0}^{zz})$ and for each fully reporting poll p in week $t_0 + s$ the predictive distribution is

$$y_p \sim \mathcal{N}(\hat{z}_{t_0+s|t_0} + \eta_a, P_{t_0+s|t_0}^{zz} + \delta' \Sigma_p),$$

where $\hat{z}_{t_0+s|t_0}$ and $P_{t_0+s|t_0}^{zz}$ are the level block of the s -step prediction. The score is the mean log density over all origins and polls,

$$\text{LS}(c, \varphi, \kappa, \rho) = \frac{1}{n_{\text{pred}}} \sum_{t_0, p} \log \mathcal{N}(y_p; \hat{z}_{t_0+s|t_0} + \eta_a, P_{t_0+s|t_0}^{zz} + \delta' \Sigma_p),$$

a proper scoring rule, so intervals that are too wide or too narrow are both penalised. LS is maximised over a grid ($c \in \{1, 0.5, 0.35, 0.25, 0.18, 0.12, 0.06\}$, $\varphi \in \{0.8, 0.9, 0.95, 0.98\}$, $\kappa \in \{0.25, 0.5, 1, 2, 4\}$, $\rho \in \{1, 0.75, 0.5\}$), which always includes the plain random walk, so the result cannot score worse than it. After choosing, the exact paths of the previous section are redrawn under the chosen dynamics, and the posterior carries the drift so forecasts can continue a trend.

Result on the 19 October 2026 fit. The selected dynamics improve the mean held-out log score from -1.04 (random walk) to -0.54 . The share of later polls inside the model's 80% predictive region, by how far ahead they were predicted, is below. The score surface is flat along a ridge where $c\kappa$ is roughly constant, so c and κ are not separately identified: the data prefer a smooth, persistent trend over a jittery level.

Weeks ahead	1–2	3–4	5–8	9–13	14–26
Random walk	91%	85%	81%	73%	61%
With momentum	89%	85%	79%	80%	71%

Beyond a couple of months the random walk is overconfident, which is what momentum fixes in part. A residual under-coverage at the longest horizon remains.

5 From the latent state to the election

5.1 Carrying the state to polling day

Let (z_T, d_T) be a draw of the state at the latest grid week and τ the number of days to the election. Take $n_{\text{wk}} = \lfloor \tau/h \rfloor$ whole weekly steps of Equation 9, $z \leftarrow z + d + \varepsilon$ and $d \leftarrow \varphi d + \varepsilon^{\text{drift}}$, then a final part-week of fraction $w_\tau = \tau/h - n_{\text{wk}}$: $z \leftarrow z + w_\tau d + \sqrt{w_\tau} \varepsilon$. Each draw of the posterior (hyperparameters, path, drift) carries its own shocks, so the forecast's width includes parameter, path and future-movement uncertainty. Call the result z_E . The election-day shares are

$$\begin{aligned} x_E &= \text{softmax}(V(z_E + \varepsilon_{\text{sys}})) \\ \varepsilon_{\text{sys}} &\sim \mathcal{N}(\mathbf{0}, \sigma_{\text{sys}}(\tau)^2 I), \end{aligned} \tag{10}$$

where the systematic error $\sigma_{\text{sys}}(\tau)$ is described next.

5.2 The systematic poll-to-election error

Polls as a whole can miss the result by more than their noise allows (shared methodological bias, late swings, differential turnout). This does not average away, and because house effects are centred across firms (Equation 6) it is not absorbed elsewhere. It is estimated from past elections. For election e and a forecast made τ days before polling day from only the polls then available, let $\omega_{e,\tau}$ be the difference between the actual result's ILR and the mean ILR of the forecast draws, and $\sigma_{\text{fc},e,\tau,j}^2$ the forecast's own variance in each coordinate j . A moment estimate of the extra variance, over the n_{elec} elections used, is

$$\hat{\sigma}_{\text{sys}}^2(\tau) = \max\left(0, \frac{1}{n_{\text{elec}}(K-1)} \sum_{e,j} \left((\omega_{e,\tau,j})^2 - \sigma_{\text{fc},e,\tau,j}^2\right)\right).$$

Using the 2014, 2017, 2020 and 2023 elections, with the dynamics calibrated afresh from only the polls available each time, gives

Days before polling day τ	7	21	60	120
Needed $\hat{\sigma}_{\text{sys}}(\tau)$ per coordinate	0.053	0.142	0.226	0.258

and a curve is fitted to these variances,

$$\sigma_{\text{sys}}^2(\tau) = v_{\text{max}}(1 - e^{-\tau/\tau_s})$$

$$v_{\text{max}} = 0.0816$$

$$\tau_s = 68 \text{ days},$$

so the systematic standard deviation is zero on polling day, rises over about two months and levels off at $\sqrt{v_{\text{max}}} \approx 0.29$. At fourteen days out $\sigma_{\text{sys}} \approx 0.12$ per coordinate, about 3.5 points of standard deviation for a party at 28% Mean absolute error of the central forecast on party votes was 1.2 points at 7 and 21 days, 1.9 points at 60 days and 2.3 at 120 days. 2017's late swing dominates the estimate (National 6.3 points above and Labour 4.9 below the 7-day forecast), so with four elections the amount is uncertain, and it is exposed as a multiple (systematic_scale) in the forecast.

Figure 9 shows Equation 10 applied from each week's smoothed state, for National and Labour, with the systematic error at its calibrated size. For the week of 14 October 2023 National's 80% interval for election day runs from 19.1% to 50.8% around a median of 32.5%. From the last grid week, 24 October 2026, it runs from 24.0% to 31.2% around a median of 27.6%, and Labour's from 24.0% to 31.8% around 28.0%. The two intervals overlap almost entirely, so the model does not separate the two parties' election-day shares.




Figure 9: Election-day share of National (top) and Labour (bottom) implied by each week's latent state, as the median and 80% interval, carried to 7 November 2026 with the calibrated systematic error. The smoothed state is used throughout, so earlier weeks show how uncertainty depends on distance from the election, not what a forecaster would have said at the time. Fit of 20 October 2026.

6 Seats and government

Each posterior draw of $\mathbf{x}_E$ is turned into a Parliament by the MMP rules. Write x_i for the share of party i in $\mathbf{x}_E$.

Eligibility. A party takes part in the list allocation if its party-vote share is at least 5% or it wins at least one electorate: $\mathcal{E} = \{i : x_i \geq 0.05 \text{ or } n_i^{\text{win}} \geq 1\}$, where n_i^{win} is its electorate wins; Other never takes a seat.

Sainte-Laguë. The 120 list-allocation seats go one at a time to the party with the largest quotient,

$$i^* = \arg \max_{i \in \mathcal{E}} \frac{x_i}{2n_i^{\text{list}} + 1},$$

where n_i^{list} is the seats party i holds so far (divisors 1, 3, 5, ...). The final n_i^{list} is its list entitlement.

Overhang. A party keeps $n_i^{\text{seat}} = \max(n_i^{\text{list}}, n_i^{\text{win}})$ seats. Any excess $n_i^{\text{over}} = n_i^{\text{seat}} - n_i^{\text{list}}$ is overhang, which enlarges the House without reducing anyone else's entitlement, so the House has $n_{\text{house}} = 120 + \sum_i n_i^{\text{over}}$ seats.

Electorate wins. The number of electorates each party wins, n_i^{win} , comes from the electorate model of Section 7, which decides the winner of each of the 71 electorates l in every posterior draw, so that a party whose vote rises wins more seats everywhere and nearby seats move together. A seat won by a party outside the seven is recorded as untracked and gives no tracked party a seat. The code keeps an older rule as a fallback and a comparison: the holder i_l of an electorate keeps it with probability $\min(1, x_{i_l}/x_{i_l}^{2023})^\gamma$ with $\gamma = 0.5$, and a few contested seats have stated probabilities $\varpi_{l,i}$ that party i wins.

Government. Group the parties by position: Left (Labour, Green, Te Pāti Māori), Centre (Opportunity), Right (National, ACT, NZ First). A group $\mathcal{G}$ holds a majority if its seats exceed half of the House, $2 \sum_{i \in \mathcal{G}} n_i^{\text{seat}} > n_{\text{house}}$. Over draws, the probabilities of four exclusive outcomes are reported: the right governs alone; the left governs alone; the centre decides (neither alone has a majority but one does with the centre); no majority even with the centre.

Figure 10 shows the vote share of each side over time. In the last grid week the left is at a median of 44.3% (80% band 43.1% to 45.6%), the right at 47.8% (46.6% to 48.9%) and the centre at 6.2% (5.6% to 6.9%). Figure 11 shows the seats from 500 simulated draws at the forecast’s default settings. National has a median of 40 seats (80% interval 35 to 45) and Labour 34 (29 to 39); the House has a median of 127 seats (122 to 132). The right alone has a majority in 62.6% of draws, the left alone in 3.0%, the centre decides in 33.8% and neither side has a majority even with the centre in 0.6%. Opportunity has a median of 8 seats with an 80% interval of 6 to 9, and enters Parliament in 92% of draws. The electorate seats come from the electorate model of Section 7.




Figure 10: Party-vote share of the left (Labour, Green, Te Pāti Māori), centre (Opportunity) and right (National, ACT, NZ First) by week, as the median and 80% band of the summed shares in each posterior draw. Fit of 20 October 2026.




Figure 11: Seats by party from 500 simulated draws of election-day shares, as the median and 80% interval; the dashed line is 61 seats. Electorate seats come from the electorate model of Section 7.

Figure 12 shows the whole distribution behind each of those intervals: the share of the 500 draws in which a party wins each number of seats, electorate and list seats together and overhang included. Each panel has its own scales, so a small party's few seats are as easy to read as a large party's many. The two large parties have by far the widest distributions, with standard deviations of about four seats: National's most likely number is 41 seats, in 14.6% of draws, and Labour's is 36, in 11.0%. The smaller parties are tighter, from Green (standard deviation 2.3 seats, most likely 18) and NZ First (1.9, most likely 13) to ACT (1.6, most likely 11). Te Pāti Māori is the tightest: it wins exactly three seats in 43.4% of draws and two to four in 89.6%, and none in 0.8%. Opportunity is the only party whose distribution has two parts. It wins no seat in 8.4% of draws, when it falls short of the threshold and wins no electorate, and otherwise it almost always wins between six and ten (90.0% of draws), with 7 the most likely (30.4%). Only 0.2% of draws fall between, with one to five seats, because a party that clears 5% of the party vote has about six. That gap is why Opportunity decides who governs in a third of draws and is out of Parliament in one draw in twelve.

Figure 13 draws one of those simulations as a House, the draw closest to the median seats of every party. It has 129 seats, 120 plus nine overhang seats, arranged from the left end of the chamber to the right in the order Te Pāti Māori, Green, Labour, Opportunity, National, NZ First, ACT. A ringed seat is overhang, and the dashed line marks the middle of the House: a majority is 65 seats, so a bloc must reach past it from its own end.




Figure 12: The chance of each party winning each number of seats, from 500 simulated draws, with electorate and list seats together and overhang included. The parties are in order from the left end of the chamber to the right, and each panel has its own scales. Fit of 20 October 2026.





Figure 13: A typical Parliament: one simulated draw, chosen as the one nearest the median seats of every party, with the parties from left to right. Ringed seats are overhang, and the dashed line marks the middle of the House.

A single Parliament hides how much the draws differ, and the scenarios of Section 8 are all questions asked of the whole set. Figure 14 therefore puts every draw on one page. Each horizontal line is one simulated Parliament, built from the same blocks as Figure 13, one seat to a unit of width, with the parties in the same order from the left end of the chamber to the right. Each line is centred on the middle of its House, so the solid line down the middle is the majority: a bloc holds one when its blocks reach past that line. The dashed lines mark the ends of a House of 120 seats, so a line that runs beyond them has overhang, as most of the lines do. The lines are sorted by the four outcomes of government above, with the left governing alone at the top, then the draws in which the centre decides, the few with no majority even with the centre, and the right governing alone at the bottom. Within each outcome they run from the most left-leaning to the most right-leaning, measured by the seats of the left less those of the right. A strip down the left edge shows the outcome of each line. The left governs alone in the 3.0% of draws at the top. The 33.8% in which Opportunity decides are the lines where its teal block lies across the middle line, so that neither side reaches a majority without it. The right governs alone in the 62.6% below, where National, NZ First and ACT together pass the middle line.

■ Te Pāti Māori ■ Green ■ Labour ■ Opportunity ■ NZ First ■ National ■ ACT ■ Other



Who can govern

Left governs alone (3%) Opportunity decides (29%)
Stalemate (0%) Right governs alone (68%)

Figure 14: All 500 simulated Parliaments, one horizontal line each. A line is the seats of each party from Te Pāti Māori at the left to ACT at the right, centred on the middle of its House: the solid line is the majority and the dashed lines are the ends of a House of 120 seats, so a longer line has overhang. Lines are sorted by who can govern (the strip on the left, with the share of draws in its legend), and within each group from the most left-leaning to the most right-leaning. Fit of 20 October 2026.

7 Electorate seats

The party vote fixes how many seats each party is entitled to, and the electorate results decide how many of those seats come from the list and whether there is overhang (Section 6). The 71 electorates are separate contests that share one national swing, and the model decides each of them in every posterior draw.

7.1 The model

For electorate l and party i (the seven named parties and Other), the model predicts the log of the candidate-vote share from the result of the previous election in the same electorate and four terms,

$$\log c_{li} = \log c_{li}^{\text{prev}} + \theta_{\text{swing}} \mu_i + \theta_{\text{split}} \text{gap}_{li} + \theta_{\text{stand}} \text{stand}_{li} + \theta_{\text{retire}} \text{retire}_{li} + \varepsilon_{li}^{\text{elec}}. \quad (11)$$

Here c_{li}^{prev} is the notional candidate-vote share of party i at the previous election, in the boundaries of the coming one, and $\mu_i = \log x_i - \log x_i^{2023}$ is the log change in the party's national party-vote share in a posterior draw (shares are floored at 0.5% first). A draw in which a party's vote falls therefore lowers its candidates in every seat together. The indicators stand_{li} and retire_{li} are 1 for the party of the sitting member when that member stands again and when they do not (one half each where the status is unknown), and 0 for every other party. A member who changes party takes a share of their old party's candidate vote with them, 0.53, estimated from four cases.

The term gap_{li} compares a party's candidate vote with the candidate vote its electorate's party vote would give if the voters there split their two votes as voters usually do. Write x_{lj}^{prev} for party j 's party-vote share in the electorate at the previous election, and $T_{ji}^{(l)}$ for the probability that a voter who gives their party vote to j gives their candidate vote to the candidate of i (Section 7.2). Then

$$\text{gap}_{li} = \log \sum_j x_{lj}^{\text{prev}} T_{ji}^{(l)} - \log c_{li}^{\text{prev}}, \quad (12)$$

with 0.005 added to each share before the logs. A candidate vote far above what the party vote and the usual splits explain, such as a sitting member's personal vote, has a negative gap and is pulled back toward them by a positive θ_{split} . The gap used before was the log of the party's own party-vote share less the log of its candidate share, which counts every party vote as a vote for the party's own candidate. Under the transfers about half of ACT's party voters and a fifth of NZ First's give their candidate vote to National, so a National candidate in a seat with many ACT party voters is not expected to fall short of the National party vote alone.

Only the parties that contest a seat take part in it; call them $\mathcal{S}_l$. For 2026 they are the parties with a candidate in the seat in the 2026 candidate tables, and Other when anyone stands as an independent or for a party outside the seven (Section 7.5). For past elections, which the fit uses and for which no candidate lists are kept, they are the parties with at least 3% of the previous candidate vote. The transfers are a softmax over the parties standing,

$$T_{ji}^{(l)} = \frac{\exp(\lambda_j^s \mathbf{1}[i = j] + \pi_i^s + \alpha_{ji}^s)}{\sum_{k \in \mathcal{S}_l} \exp(\lambda_j^s \mathbf{1}[k = j] + \pi_k^s + \alpha_{jk}^s)} \quad \text{for } i \in \mathcal{S}_l, \quad (13)$$

and $T_{ji}^{(l)} = 0$ for a party i not standing, so each row sums to one over $\mathcal{S}_l$ and the voters of a party with no candidate pass to the candidates who do stand. This has the form of the gravity model of Section 3.3: a loyalty λ_j^s of party j 's voters to its own candidate, a pull π_i^s of party i 's candidates, and an affinity α_{ji}^s for each pair, which a ridge shrinks toward the log of the survey's switching probability $P(i | j)$.

The noise is drawn in ILR coordinates, as in the national model, so it moves shares within the simplex. For electorate l it is a vector $\mathbf{z}_l^{\text{elec}} \in \mathbb{R}^{K-1}$, and it enters Equation 11 as its centred log-ratio, $\varepsilon_{li}^{\text{elec}} = (V \mathbf{z}_l^{\text{elec}})_i$. Within an electorate its covariance is built from the split-vote flows, as Q is built from the switching flows in Equation 2:

$$\Sigma_{\text{elec}} = \sigma_{\text{flow}}^2 D_s^\top \text{diag}(\mathbf{r}_s^2) D_s + (\sigma_{\text{iso}}^e)^2 I + \sigma_{\text{oth}}^2 \mathbf{v} \mathbf{v}^\top. \quad (14)$$

The rows of D_s are the directions D_f of Section 3.2 for the 42 flows between the named parties, evaluated at the 2023 national party vote, and the rates $\mathbf{r}_s$ are the off-diagonal entries of the transfers with every party standing. The noise therefore moves candidate vote along the edges that split voting uses, such as between National and ACT and between Labour and the Greens, and σ_{flow} scales it. The isotropic term, with scale σ_{iso}^e , covers movement along no flow. Other has no flows. Its candidates are mostly independents and small parties that come and go and do not move with the party vote, so it has a term of its own along its centred log-ratio,

$\mathbf{v} = V^\top \mathbf{e}_{\text{Other}}$ with $\mathbf{e}_{\text{Other}}$ the unit vector of Other, which changes only the log-ratios between Other and each named party.

Across electorates the noise has three parts, each with covariance Σ_{elec} : a draw shared by every electorate, a Gaussian field over the electorate centroids with an exponential kernel of length ℓ_{elec} , and an independent draw for each electorate,

$$\text{Cov}(\mathbf{z}_l^{\text{elec}}, \mathbf{z}_{l'}^{\text{elec}}) = m_{\text{elec}}^2 b_l^{\text{elec}} b_{l'}^{\text{elec}} (w_g + w_s \exp(-d_{ll'}/\ell_{\text{elec}}) + (1 - w_g - w_s) \mathbf{1}[l = l']) \Sigma_{\text{elec}},$$

where $d_{ll'}$ is the distance between the centroids of electorates l and l' , w_g and w_s are the weights of the shared and spatial parts, b_l^{elec} widens the noise of a seat whose boundaries changed, and the multiplier m_{elec} is chosen so that the predicted probabilities are calibrated on the elections left out of the fit. Nearby seats therefore move together beyond the common swing, so a poor result in one seat says something about its neighbours.

In each draw the candidate-vote shares are the softmax over $\mathcal{S}_l$ of the right-hand side of Equation 11, so they sum to one over the parties standing and a party with no candidate gets none. The winner is the party with the largest share, and the margin of the seat in that draw is the winner's share less the runner-up's. A seat won by Other has no margin, because Other's runner-up is not tracked.

7.2 The split-vote gravity model

The held-out check below uses the elections from 2002 to 2023, but the electorate model's own fit uses those of 2005 to 2023: the transfers are fitted on the voting places of the elections whose contests are being fitted, so a fit that leaves an election out does not see its booths. The transfers of Equation 13 are fitted on the voting places of these elections, whose results give both the party vote and the candidate vote of each place. The candidate votes of place b are taken as multinomial with expected shares $\sum_j \hat{x}_{bj} T_{ji}^{(l)}$, from the place's party-vote shares $\hat{x}_{bj}$ and the parties with a candidate there, so each place counts in proportion to its votes. This is an ecological fit: it learns from places with more ACT party voters giving more candidate votes to National, not from what any one voter did. Special votes, which the Electoral Commission reports only for the electorate as a whole, are left out. In 2002 to 2014 the files also have rows for the ordinary votes cast before polling day, each an aggregate of many places; these are flagged and left out of the fit, and including them changes the held-out log likelihood below by 0.0002 per vote.

The prior is the table of NZES 2023 switching probabilities $P(i | j)$ of Section 3.3. The fit maximises the multinomial log likelihood less a ridge penalty on the loyalties, the pulls and the departures of the affinities from $\log P(i | j)$, so the transfers shrink toward the survey and a fit with a very large ridge returns it exactly. The ridge also stops a party the data barely see from running off to complete loyalty, and a party with no party votes in the years fitted keeps the prior's row.

The Commission's split-vote reports, which cross-tabulate the two votes directly, could not be obtained, so the transfers are checked on the voting places themselves. Each election was left out in turn, the transfers were fitted on the others, and the candidate shares of the held-out places were predicted from their party shares. Table 3 compares the predictions with the no-split rule, which gives each candidate their party's party-vote share renormalised over the parties standing. The root-mean-square error is weighted by votes and averaged over the parties standing, and the log likelihood is per vote.

Held out	Places	RMSE, gravity	RMSE, no split	Log lik., gravity	Log lik., no split
2002	4917	0.083	0.089	−1.241	−1.270
2005	5287	0.074	0.078	−1.146	−1.163
2008	5519	0.082	0.089	−1.080	−1.097
2011	5508	0.074	0.082	−1.080	−1.113
2014	5458	0.080	0.088	−1.125	−1.168
2017	6988	0.064	0.075	−1.110	−1.171
2020	9426	0.070	0.084	−1.130	−1.191

2023	7988	0.072	0.078	−1.211	−1.247
all	51091	0.074	0.083	−1.141	−1.180

Table 3: Leave-one-election-out check of the split-vote transfers on the voting places, against the no-split rule, with the chosen ridge weight of 300. Aggregate rows are not scored. The pooled row weights each election by its votes.

The transfers beat the no-split rule in all eight held-out elections, with a pooled log likelihood of -1.1411 per vote against -1.1795 . The gain is largest in 2017 and 2020 (0.06 per vote) and smallest in 2005 and 2008 (under 0.02). The error falls by only about a tenth, from 0.083 to 0.074, because most of the error at a single place is noise that no rule for splitting removes. The ridge weight, 300, was chosen from a grid of 10 to 10^5 by the pooled held-out log likelihood. The score is flat from 10 to 3000 (a spread of 0.001 per vote) and falls only toward the prior (-1.1519 at 10^5), so the choice matters little, but it was made on the same years it is reported on. Every fold also uses the 2023 survey as its prior, including the fold that holds out 2023.

The electorate model’s fit uses the panels of its own seven elections, 2005 to 2023, with 46,174 places. With every party standing, National and Labour party voters give 94% and 95% of their candidate votes to their own party’s candidate, and Te Pāti Māori’s give 98%. ACT party voters give 45% to ACT and 49% to National, NZ First party voters give 74% to NZ First and 20% to National, and Green party voters give 54% to the Green candidate and 40% to Labour’s.

7.3 Data and fit

The model is fitted to the electorate results of 2005 to 2023: 472 contests with a known result and at least two parties in the contest, from the Electoral Commission’s candidate-vote and party-vote files. Electorates are matched between elections by name. The boundaries were redrawn for 2008, 2014 and 2020, and the previous election’s voting places were not re-aggregated into those new boundaries; instead the noise of a redrawn seat is widened by the factor b_i^{elec} , which the fit puts at 1.0, so no extra is measurable. A sitting member is matched to the previous winner by name. Previous winners who stood again held their party’s seat in 84% of 376 contests, and when the previous winner retired their party held it in 70% of 83 contests.

The transfers are fitted first, on the voting places of the same elections (Section 7.2). The four coefficients of Equation 11 are then fitted by least squares to the change in each party’s log candidate share relative to the other parties in the contest. The noise terms are fitted to the residuals, taken as log-ratios between pairs of parties that stood in the same contest: the mean square of each pair’s residual log-ratio is matched to its variance under Equation 14 by non-negative least squares, reweighted so that each pair’s relative error counts alike and the large residuals of Other do not dominate. A party with no votes in a result is treated as not standing there. The shared and spatial weights and the length scale are then fitted to the products of the residuals of pairs of seats in the same election, binned by distance. The noise is fitted on the contests whose boundaries did not change, and every contest of 2008, 2014 and 2020 had new boundaries, so only 2005, 2011, 2017 and 2023 inform it.

Term	Symbol	Estimate	Reading
national swing	θ_{swing}	0.66 ± 0.11	a party’s candidates move by about two thirds of its national log swing
split	θ_{split}	0.30 ± 0.11	about a third of the way back toward the candidate vote the transfers imply
member stands	θ_{stand}	0.31 ± 0.06	bonus in log share for the sitting member’s party
member retires	θ_{retire}	0.15 ± 0.10	a party keeps part of its advantage when its member leaves
flow scale	σ_{flow}	0.69	scale of the noise along the split-vote flows
isotropic scale	σ_{iso}^e	0.25	noise along no flow
Other scale	σ_{oth}	0.75	extra noise in the log-ratios between Other and each named party
shared weight	w_g	0.19	share of the noise variance common to every electorate

spatial weight	w_s	0.81	share in the spatial field, which leaves none for the independent part
spatial length	ℓ_{elec}	5.1 km	of the electorate centroids
boundary factor	b_l^{elec}	1.0	no extra noise is measurable where boundaries changed
multiplier	m_{elec}	1.0	chosen on the held-out elections

Table 4: Fitted terms of the electorate model. The intervals are the standard deviation over 200 resamples of whole elections. The transfers are fitted once on all the elections and not refitted in the resamples, so the intervals leave out their uncertainty. The noise terms are not given an interval.

A standing member is worth 0.31 in log share and a retiring member’s party keeps 0.15, so standing is worth about 0.16 more than retiring. The retire term is positive because a party that held a seat keeps some of its advantage when its member leaves. With the plain gap the same fit gives 0.69, 0.34, 0.33 and 0.17, so the transfers change the coefficients little; the Māori seats keep the plain gap and shrink toward those values (Section 7.6). Other’s log-ratios carry the largest term of the noise.

The spatial part is the least certain. Its fit has two minima of similar depth: one with a weight of about 0.8 on a field of length about 5 km, so that seats are close to independent beyond their immediate neighbours and the rest of the variance is shared by all seats, and one with a weight of about 0.2 at the upper bound of 1000 km, a field that is nearly the same over the whole country. The fit starts from nine lengths between 3 and 1000 km and keeps the minimum with the lower objective, here the short one (27.0 against 29.7). The four elections it rests on do not settle the choice: leaving out 2011, 2014, 2017 or 2023 gives the short minimum (4.8 to 7.7 km), and leaving out 2005, 2008 or 2020 gives the long one. The length is much shorter than the 9 to 60 km found for voting places in Section 9.3, and the multiplier, chosen afterwards, absorbs part of the difference between the two minima.

7.4 Checking the model

Each election was left out in turn: the model, transfers included, was refitted without it, and then forecast the electorates of that election from the previous result and the election’s actual national party-vote swing. These tests measure the electorate step, not the polling error that feeds it. Table 5 compares the model with the earlier version of the electorate model, which had the plain gap, independent noise for each party with a scale of its own, and the 3% rule for who contests a seat, and with three rules that use no model. The Brier score is summed over the parties, so a forecast that is always wrong scores 2.

Method	Multiplier	Seats called right	Brier score	Log score
Both at multiplier 1.0, 2000 draws				
earlier model	1.0	86.0%	0.200	−0.386
this model	1.0	86.7%	0.192	−0.369
this model, no stand or retire terms	1.0	87.3%	0.187	−0.374
Each at its own calibrated multiplier, 4000 draws				
earlier model	0.8	86.0%	0.198	−0.367
this model	1.0	87.5%	0.192	−0.368
previous winner holds		81.1%	0.377	−1.30
retention rule ($\gamma = 0.5$)		81.1%	0.323	−0.87
constant retention		81.1%	0.344	−0.87

Table 5: Leave-one-election-out scores over the 472 contests of 2005 to 2023. The upper block is the comparison the change was required to pass; the lower block scores each model at the multiplier its own calibration chose. The earlier model’s lower-block scores were recorded before this version replaced it and cannot be recomputed from the current code.

The change was required to score no worse than the earlier model on the held-out log score with both at a multiplier of 1.0, and it does: -0.369 against -0.386 . That comparison favours it, though. The earlier model’s own calibration chose a multiplier of 0.8, so at 1.0 its noise was too wide. Scored at each model’s own multiplier, the earlier model has -0.367 and this one -0.368 , which is level within the noise of the simulation. Election by election this model is ahead in 2005, 2008, 2011, 2014 and 2017 and behind in 2020 (-0.471 against -0.420) and 2023 (-0.572 against -0.495). The gains of this version are therefore in its structure rather than in a better score: the parties standing in 2026 come from the candidate tables, the split gap follows how voters actually split, the noise is correlated along the split-vote flows, and each draw gives a margin for every seat.

The stand and retire terms add only 0.005 to the log score, and without them the Brier score and the share of seats called right are slightly better. They help in 2005, 2011, 2014 and 2017 and hurt in 2008, 2020 and 2023, most in 2023 (-0.555 with them against -0.449 without). They are kept, since the log score is the measure the model is chosen by, but the evidence for them is thin.

The multiplier was chosen by the pooled held-out log score over a grid from 0.5 to 1.3 (1000 draws, with its own resamples and seed, so these scores come from a different setup from the table above and the two values at a multiplier of 1.0 are not comparable with each other, only within the grid). The score peaks at 1.0 (-0.3602) and falls on both sides (-0.3611 at 0.9, -0.3613 at 1.1, -0.3679 at 1.3 and -0.4319 at 0.5). The grid first stopped at 1.0; it was extended to 1.3 after the first calibration landed on that edge, and the extension changed neither the multiplier nor the fit. The multiplier is chosen on the same held-out elections it is then scored on, so the lower block of Table 5 is mildly optimistic for both models. Other choices also saw every election. The 2023 survey prior and the 2023 reference shares of Equation 14 are used in every fold, including the one that holds out 2023, the ridge of the transfers was chosen by the held-out check of Table 3 over all the years, and the starting lengths of the spatial fit were set by looking at the fit on all the contests.

Predicted chance	Party-contests	Mean predicted	Observed
0% to 10%	230	3.1%	2.6%
10% to 20%	61	14.3%	19.7%
20% to 30%	48	24.3%	25.0%
30% to 40%	40	35.0%	35.0%
40% to 50%	29	45.4%	41.4%
50% to 60%	28	55.0%	53.6%
60% to 70%	35	64.4%	68.6%
70% to 80%	45	75.6%	71.1%
80% to 90%	56	85.8%	82.1%
90% to 100%	303	97.9%	96.4%

Table 6: Reliability of the held-out probabilities of this model at its own multiplier: each party’s chance of winning each contest, grouped by the chance predicted, against how often the party won.

Table 6 shows that the probabilities are close to calibrated but slightly overconfident at the top. Of the 303 chances above 90%, 96.4% were won against 97.9% predicted, and in the 80% to 90% band 82.1% against 85.8%. The earlier model was closer above 90% (97.2% against 97.5%, from 281) and similar in the 80% to 90% band (81.7% against 85.3%, from 60).

7.5 The 2026 baseline

For 2026 the baseline c_{ii}^{prev} is the 2023 result in the boundaries of the 2025 review. Each 2023 voting place goes to the 2026 electorate, on the same roll (general or Māori), that contains it. Votes that cannot be placed, namely special votes, votes at places without coordinates and votes cast outside the voter’s own electorate, are shared among the pieces of their own 2023 electorate in proportion to its placed votes. Of the 22 electorates whose boundary did not move, 19 reproduce their actual 2023 totals exactly. Three differ, by up to 12%, because a small boundary change moves a large booth (Banks Peninsula -11.9% , Ikaroa-Rāwhiti $+4.5\%$, Te Tai Tonga -4.3%),

so those are treated as changed. Nineteen electorates are unchanged, 42 are changed and 10 are new. Only one notional winner differs from the same-named 2023 winner: Mt Albert becomes National by 153 notional votes, where Labour won by 18 in 2023. Port Waikato’s 2023 candidate poll was postponed, so its baseline uses the by-election count, which overstates National’s share because Labour, the Greens and ACT did not stand.

The status of each sitting member comes from dated news reports, listed in docs/electorate-news-2026.md: 60 are standing, eight are retiring (Duncan Webb, Paulo Garcia, Maureen Pugh, Judith Collins, Shane Reti, Brooke van Velden, Andrew Bayly and Chris Penk), two changed party (Mariameno Kapa-Kingi, who founded the Te Tai Tokerau Party, and Tākuta Ferris, now an independent) and one, Megan Woods, is standing on the list only.

The parties that contest each seat come from the 2026 candidate tables, since nominations have closed. In the 64 general seats National and Labour stand everywhere, the Greens in 50, ACT and NZ First in 48 each and Opportunity in 41, and an independent or a candidate of a smaller party stands in 56. In the seven Māori seats Labour and Te Pāti Māori stand in all, the Greens in five and National and Opportunity in one each. A party with no candidate in a seat never wins it in any draw. The tables change the contesting parties of 64 of the 71 seats from the 3% rule, mostly by adding small parties with little of the 2023 vote and by removing parties that do not stand again. In Ilam, for example, Opportunity’s candidate took 25.4% of the 2023 candidate vote and Opportunity has no candidate now, and National’s chance there rises from 70.7% under the earlier model to 95.7%.

A party standing for the first time in a seat, or with less than 3% of its 2023 candidate vote, has almost no baseline to carry forward. Its log share would start near the 0.005 added before the logs, and the split term would move it only part of the way toward what its party vote implies, leaving it near 1 or 2% whatever its party vote. Such a party’s baseline is set instead to the candidate vote its party vote implies under the transfers, $\sum_j x_{lj}^{\text{prev}} T_{ji}^{(l)}$, so its gap in Equation 12 is zero, and the softmax over the parties standing normalises it with the others. This applies to 112 party-seat pairs, mostly Opportunity (37), Other (27) and NZ First (26), then the Greens (11), ACT (10) and Labour in Port Waikato; the four in the Māori seats are handled by their own model (Section 7.6). Past elections use the 3% rule, so the fit never sees this case and the leave-one-election-out check does not test it. At the mean national swing of the election-day draws, Labour’s central share in Port Waikato, whose baseline is the by-election in which it did not stand, is 18.0% (1.8% without the seed) and National’s 54.1% (74.9%), and the median margin over the draws is 34 points rather than 54. The Greens in Tāmaki go from 2.0% to 9.2% and in Waimakariri from 1.3% to 8.4%, and NZ First in Taranaki-King Country from 3.0% to 10.3%. No seat’s favourite changes, and no party’s chance of winning a seat moves by more than 0.3 points.

Tāmaki shows what the transfers change. It is reported here as a check and was not used to choose anything. Its notional baseline gives ACT 43.1% of the candidate vote but 12.3% of the party vote, the personal vote of its member, Brooke van Velden, who is retiring, and gives National 33.3% of the candidate vote but 52.2% of the party vote. The plain gap compared ACT’s candidate vote only with ACT’s own party vote, and pulled it back toward 12.3% by about a third. The transfers say instead that ACT’s party voters give about half their candidate vote to National, so the candidate vote the party vote implies is 7.2% for ACT and 58.1% for National. ACT’s gap in Equation 12 becomes -1.74 against -1.23 before, and National’s $+0.55$ against $+0.44$, so the pull from ACT toward National is larger even with the smaller θ_{split} . With election-day shares from the posterior (10,000 draws), National now wins Tāmaki in 47.2% of draws, ACT in 44.6%, Labour in 8.1% and the Greens in under 0.1%, where the earlier model gave ACT 55.5%, National 37.5% and Labour 7.1%. It is the only seat whose favourite changes. The other large moves keep the favourite and make it more likely: besides Ilam, Northland (63.9% to 84.0%), Dunedin, Whangārei, Mt Albert, Christchurch Central, East Cape, West Coast-Tasman and Kaikōura each move by 7 to 14 points. Part of this comes from removing parties that have no candidate.

7.6 The Māori seats

The seven Māori electorates are forecast from their own history and from polls. The general-seat fit rests on 64 seats, while the Māori seats have only seven elections of history and a different set of contestants, so its terms cannot be assumed to carry over.

Whakaata Māori commissioned Curia polls of all seven seats in September and October 2026, each of 500 respondents (margin of error 4.5%), and the 2020 and 2023 elections have one Curia poll per seat. Comparing

those 14 polls with the results gives a root-mean-square error of 0.34 in log share. Sampling accounts for only 0.09 of it, so about 0.33 is error that more respondents would not remove. The error was 0.22 in 2020 and 0.40 in 2023, Labour was overstated by about 29% and Te Pāti Māori understated by about 16%, and the estimate is weak because three of the 2020 polls could only be read from summaries.

The history of the seven seats, 48 contests over seven elections, gives a prior for each seat. A Māori-seat fit of Equation 11 finds a swing coefficient of 0.67 (standard error 0.27), close to the 0.69 of the plain-gap general fit, so the Māori seats follow the national party-vote swing about as general seats do, and the history gives no ground for a smaller swing. It also finds a persistence term of 0.78 (0.14), meaning the Māori seats polarise and do not regress toward a mean, and incumbency terms that are negative and poorly identified. The coefficients are shrunk toward the general-seat values in proportion to their precision, using the general fit with the plain gap so that the features mean the same in both, with an allowed difference of 0.3 that is a judgement, since seven elections cannot identify it. In these seats the noise scale is about 0.64 for Te Pāti Māori and 0.43 for Labour, and about 40% of the noise variance is shared by the seven seats. Leaving each election out, the Māori-seat prior has a root-mean-square error of 0.80 in the log change of a share, against 0.89 for the general coefficients and 0.97 for no change.

The polls and the prior are combined by precision, with the polls' sampling error and the historical poll error added, and the total weight of the polls in a seat is capped at one half. History therefore supplies about half of each forecast for Labour and Te Pāti Māori. The Greens and the other candidates, who have little history in these seats, take more of their forecast from the polls. The polls rearrange Labour, the Greens and Te Pāti Māori inside Te Tai Tonga, Te Tai Tokerau and Waiariki, but barely move Te Pāti Māori's total: its expected electorate wins are 2.94 with and without them, and the probability that the right governs alone changes from 61.4% to 62.6%. Doubling or halving the weight of the polls changes a seat's probability by a few points, for example Labour's chance in Hauraki-Waikato runs from 0.37 to 0.48.

Two checks remain open. The Māori-roll party-vote polls are stored but not used in the model, and they disagree with the seat polls: Whakaata Māori–Curia puts Te Pāti Māori at 29% of the Māori roll, while the seat polls put its party vote at 10% to 23%. And Te Tai Tokerau, where the Te Tai Tokerau Party of Mariameno Kapa-Kingi and a Green candidate both stand, is the least settled seat.

7.7 The swing meter

The numbers and figures of this section and the next, are from the fit of 20 October 2026, made with the earlier version of the electorate model (Table 5), so they still show ACT ahead in Tāmaki. The seat-distribution, Parliament and scenario figures come from the same fit and are regenerated after the merge.

The central forecast sets the noise in Equation 11 to zero and the national swing to its mean over the draws. The lead of the first party over the second, lead_l , in points of candidate vote, is then the margin of seat l , and half of it, $\text{flip}_l = \text{lead}_l / 2$, is the swing from the leader to the runner-up that would tie the seat. The meter ranks seats by flip_l and bands them as knife-edge (under 2 points), marginal (under 5), fairly safe (under 10) and safe. Of the 71 seats, 2 are knife-edge, 5 marginal, 17 fairly safe and 47 safe.




Figure 15: The 20 electorates closest to changing hands. A bar is the swing to the runner-up that would tie the seat, in the leader's colour and fainter where the simulation is less sure the leader holds it, and the dot is the runner-up's colour. Dashed lines mark the band edges at 2, 5 and 10 points. Fit of 20 October 2026.

Figure 15 shows the closest seats. The two knife-edge seats are West Coast-Tasman (1.3 points, a retiring National member) and Waitākere (1.4, Labour ahead of National in a seat with a retiring National member). Banks Peninsula (2.6), Tāmaki (2.7, ACT ahead of National) and Hauraki-Waikato (2.9, Te Pāti Māori ahead of Labour) follow. Hutt South needs 3.1 points and National holds it in 64% of draws. Its baseline lead over Labour is 0.07 in log share, the national swing since 2023 takes 0.25 off, and Bishop standing again adds 0.33, leaving a central lead of 0.14 against a noise standard deviation of 0.41 for the difference. Botany needs 25 points and National holds it in more than 99%. The one reported poll of Hutt South, a Victor Consulting poll with Labour at 42 and Bishop at 30 reported by The Spinoff, would put Labour well ahead of this, but it is a single report.

National holds 44 seats in the baseline and wins 39.9 on average, with 27 more than 90% likely and 20 less than 10% likely. Labour wins 20.7 on average, with 7 more than 90% likely. Te Pāti Māori wins 3.0, Green 4.5 and ACT 1.4. The old retention rule gave National 37.7 and Labour 17.5.

7.8 The outlook map

Each electorate is called for the party that wins it in the largest share of draws, and the call is faded by that share, so a solid hexagon is a safe call and a pale one a toss-up. National is called in 43 electorates, Labour in 20, Te Pāti Māori in 3, Green in 3 and ACT in 2, and a call is at least 90% likely in 27 of National's seats and 7 of Labour's. Four seats are expected to change hands. Te Tai Tokerau, Te Tai Tonga and Tāmaki Makaurau go from Te Pāti Māori to Labour, with Labour's chance 61%, 54% and 56%, and Waitākere goes from National to Labour at 49%, which is close to a toss-up. Figure 16 shows the calls as one hexagon per electorate, grouped by city, and Figure 17 draws them on the electorate boundaries, with insets for Auckland, Wellington and Christchurch.




Figure 16: The 2026 outlook by electorate, one hexagon per electorate. The party is the one that wins the seat in most of 500 draws, and the shade shows how often. Fit of 20 October 2026.

■ National ■ Labour ■ Green ■ ACT **Confidence**
■ Te Pāti Māori ■ 50% ■ 75% ■ 100%




Figure 17: The same calls on the 2025 electorate boundaries, with the shade showing the chance of the call and insets for the three largest cities.

7.9 What the model does not do

- **The Māori seats rest on thin data:** seven elections of history, 24 polls and a poll error estimated from 14 of them. They keep their own poll-informed model with the plain gap. The shrinkage allowance, the cap on the polls' weight and the noise prior are judgements, and the Māori-roll party-vote polls are not used.
- **Other has no split of its own.** Independents and every party outside the seven are one category, with its own noise term and no flows, because minor parties and Māori parties have come and gone and no single set of transfers fits them.
- **The transfers are ecological and partly checked.** They are fitted to voting places, not to voters, and checked on held-out places rather than against the Commission's split-vote reports. They are fitted once on all the elections and not refitted in the resamples, and their prior comes from one survey, 2023.
- **A retiring member's personal vote has no term of its own.** The gap pulls it back toward what the party vote and the usual splits imply, and the retire coefficient is one value for every retirement.
- **The spatial field is one exponential kernel (a Gaussian field) over the electorate centroids,** and its fit has two minima that four elections do not separate (Table 4). A model over the electorate areas is being explored.
- **A party standing for the first time in a seat starts at its implied split share** (Section 7.5), a rule the leave-one-election-out check cannot test because past contests use the 3% rule.
- **The candidate tables are taken as complete.** Nominations have closed, but a candidate missing from the tables would leave that party unable to win the seat.
- **A first-term member and a long-serving one are worth the same,** because tenure was not tested.
- **Boundary changes before 2020 are treated as extra noise,** and the 2026 baseline was not checked against a second source.
- **The noise terms are fixed,** with no uncertainty of their own.
- **Statuses are assigned by rule:** a member who moves to another seat counts as standing, and an unknown status counts half as standing and half as retiring.
- **The results depend on the polls.** National's party vote of 27.6% in the posterior is what puts so many of its seats in play.

8 Scenarios

Who governs, whether a particular coalition has a majority, how large the House is and whether a minister keeps their seat are events that depend on the whole simulation at once. The scenario generator names such events and reads their probabilities off the posterior draws. Each is a question asked of the Parliaments of Figure 14, and its probability is the share of those lines on which it holds.

8.1 Events on the draws

A scenario is a yes or no test on one simulated draw. It sees the party-vote shares of the draw, the seats of each party, the House size, the overhang, the winner of every electorate and, for named people, whether each is in Parliament. The probability $P(A)$ of a scenario A is the share of the n_{draw} draws in which it holds, with standard error $\sqrt{P(A)(1 - P(A))/n_{\text{draw}}}$, which is about two points at $n_{\text{draw}} = 500$. Scenarios combine with and, or and not, and the probability of A given B is $P(A \mid B) = P(A \wedge B)/P(B)$ over the draws in which B holds. The generator reports the number of those draws, and warns when there are fewer than 50, since a conditional probability on a handful of draws is not worth quoting. A bloc has a majority when its seats exceed half of the House, which is $\lfloor n_{\text{house}}/2 \rfloor + 1$ seats.

The catalogue has five groups: who governs, coalitions, overhang and thresholds, seats and candidates, and ministers and leaders. Beyond the catalogue the notebook has a builder that combines up to three conditions, with an optional given condition, and two kinds of re-simulation. A what-if moves one party's election-day vote share by a number of points, rescaling the others in proportion, and runs the seat simulation again. A stress test adds the polling error of a past election, taken from the backtest of Section 5.2 at the horizon nearest the days left, to the election-day draws in log-ratio coordinates, and runs it again.

8.2 Who governs and coalitions

Figure 18 shows the main governing outcomes, which are the four groups of lines in Figure 14. The right bloc has a majority alone in 62.6% of draws and the left bloc in 3.0%. In a further 33.8% neither has a majority alone but one has it with Opportunity, and in 0.6% no bloc reaches a majority even with Opportunity. National and ACT never have a majority without NZ First in these draws, so the chance that NZ First holds the balance is the same 62.6% as the chance that the right governs alone. National is the largest party in 79.4% of draws and Labour in 17.4%, and Labour is ahead of National on the party vote in 54.8%.




Figure 18: The main governing scenarios as probabilities over 500 posterior draws. The bar for each scenario is its probability; the standard error is about two points. Fit of 20 October 2026.

The generator works out, for every combination of the seven parties, how often their seats together make a majority, and lists the minimal winning coalitions, those that lose their majority if any one party leaves. By default it leaves out combinations that cross the line between left and right unless they are named in `config/scenarios.toml`. Table 7 gives the main ones.

Coalition	Has a majority	Minimal and winning
National, ACT, NZ First	62.6%	62.6%
National, ACT, NZ First, Opportunity	96.2%	33.6%
Labour, Green, Te Pāti Māori, Opportunity	32.6%	22.4%
National, NZ First, Opportunity		21.4%
National, ACT, Opportunity	9.6%	
Labour, Green, Opportunity	8.4%	
Labour, Green, Te Pāti Māori	3.0%	

Table 7: Chance that a coalition has a majority, and that it is a minimal winning coalition. A blank means the coalition is not in that list.

8.3 Overhang, seats and candidates

The House has at least 121 seats in 96.8% of draws, at least 123 in 87.6%, at least 125 in 71.8% and at least 130 in 26.6%. National has an overhang of at least one seat in 92.6% of draws, because it wins more electorates than its party vote entitles it to, and Te Pāti Māori has one in 54.8% and two or more in 24.8%. At least 5% of the party vote is wasted on parties that do not enter Parliament in 8.4% of draws, almost all of them draws in which Opportunity falls under the threshold. Opportunity enters Parliament in 91.6% and enters only through an electorate win in 2.0%.

National wins at least 40 electorates in 59.6% of draws and Labour at least 20 in 60.0%. Te Pāti Māori wins at least four Māori seats in 28.6% and Labour wins at least two in 87.6%. The Te Tai Tokerau Party wins Te Tai

Tokerau in 6.8%. At least 10 electorates change hands against the notional 2023 winners in 91.2% of draws, and at least 20 in 6.4%.

8.4 Ministers and leaders

The table of ministers and leaders, built from Wikipedia, is joined to the candidate lists, so each person has a probability of being in Parliament through an electorate or a list seat. A person who is not standing is never in Parliament, and the scenarios count only people who are standing and matched to a candidate. The Prime Minister is out of Parliament in 0.4% of draws. At least two Cabinet ministers are out in 86.8% and at least three in 42.4%. The Finance Minister is out in 95.6%, because Nicola Willis is a list-only candidate and National's electorate wins use up its list seats in nearly every draw. The same arithmetic puts a National list-only minister or former minister, such as Melissa Lee, nearly as likely to miss out. Table 8 lists those most likely to miss out.

Person	Party	Standing	Not in Parliament
Nicola Willis	National	list	95.6%
Melissa Lee	National	Mt Albert and list	90.4%
Ron Mark	NZ First	Wairarapa and list	89.6%
Michael Wood	Labour	Mt Roskill	83.0%
Deborah Russell	Labour	list	78.8%
Paul Goldsmith	National	Epsom and list	78.0%
Chris Bishop	National	Hutt South and list	34.2%
Mike Butterick	National	Wairarapa and list	26.8%

Table 8: Ministers and former ministers most likely to miss out, from 500 draws. Ministers of 2023 to 2026 and former ministers of the 2017 to 2023 governments are included.

8.5 Conditional events, what-ifs and stress tests

Among the ready-made conditional events, the right governs alone in 65.4% of the draws in which Opportunity falls under 5%, though there are only 52 such draws. The left governs alone in 12.7% of the 118 draws in which Labour has more than 30% of the party vote, and NZ First holds the balance in 45.8% of the 275 draws in which National has less than 28%.

Table 9 shows how the main probabilities move when National loses two more points and when the polls miss as they did in 2017, 2020 and 2023. A 2017-style miss halves the chance that the right governs alone, a 2020-style miss almost removes it, and a 2023-style miss costs about eight points. Each re-simulation uses the same random draws as the base run, so the differences between rows are not sampling noise in the vote.

Scenario	Base	National —2	Miss like 2017	Miss like 2020	Miss like 2023
Right bloc alone governs	62.6%	48.4%	30.2%	9.4%	55.0%
Opportunity decide	33.8%	46.8%	63.0%	84.2%	43.0%
National largest party	79.4%	69.0%	34.4%	70.6%	80.0%
Labour largest party	17.4%	27.0%	59.4%	25.4%	15.8%
House of at least 130	26.6%	36.8%	7.6%	50.8%	18.0%

Table 9: Probabilities after a what-if shift of National's vote and after adding the polling error of past elections, at the 21-day horizon nearest the 17 days left.

8.6 Reading the scenarios

The probabilities carry every assumption of the forecast: the polls, the systematic error, the electorate model and the placeholder seat probabilities of its fallback. The standard error of about two points applies to the draws and not to the model. The minister scenarios depend on a list of ministers taken from Wikipedia and on the

match of names to the candidate tables. And the stress tests add the error of a single past election, so they describe how the forecast would change if the polls were wrong in that particular way, not how likely that is.

9 Election night

Status: the data layer (Section 9.1) is built and validated; the fit of the spatial kernel (Section 9.3) is under development; the update and the replay tests are planned.

9.1 Data

For each of eight elections (2002 to 2023) the Electoral Commission publishes, for every electorate, the party votes at each voting place and the special votes by category (overseas, before polling day, on polling day, party only, postal, hospital), with a stated total that the rows must sum to (checked on load). From 2008 there are coordinates for the voting places. From 2017 places are split into advance and election-day places. A voting place serves voters of several electorates, so places are keyed by address, not by electorate. Only about 44 to 48% of votes lie in places that match by address between consecutive elections (25 to 29% from 2017 to 2023): venues change. The model therefore does not rely on booths persisting.

Figure 19 shows the 2023 electorate winners as one hexagon per electorate, and Figure 20 the swing to National since 2020 at each voting place in Auckland. In the 2023 returns National leads 43 of the 71 electorates with results, Labour 17, Te Pāti Māori 6, Green 3 and ACT 2. Port Waikato has no candidate result in the files and is drawn grey. Each dot is a polling place, and the two kinds of voting are drawn apart, on one colour range, because they behaved differently. Weighted by their own votes, the 801 advance places swung 15.2 points to National nationally and 18.6 in the Auckland window, where all 200 swung towards National. The 2,330 election-day places swung 6.5 points nationally and 8.6 in Auckland, where 95% of the 447 places swung towards National.





Figure 19: Leading candidate's party in each electorate at the 2023 election, one hexagon per electorate in a tile layout. Port Waikato is grey because its result is missing from the candidate files.




Figure 20: Swing in National's share of the party vote at each polling place in the Auckland area, 2020 to 2023, in percentage points, for advance voting (a) and election-day voting (b) on one colour range. Dot area is the number of votes at the place. A place with no polling place within 300 m in 2020 takes its electorate's swing.

9.2 A baseline from the previous election, smoothed over space

Let place b be at location $\mathbf{q}_b$ (NZTM metres) with ILR composition $\mathbf{y}_b$ computed from its counts with +0.5 added to each party. Its baseline is the previous election's composition smoothed over space,

$$\mathbf{m}_b = \sum_{c \in \mathcal{N}_{\text{nb}}(b)} \alpha_{bc} \mathbf{y}_c^{\text{prev}}$$

$$\alpha_{bc} \propto n_c \exp(-\|\mathbf{q}_b - \mathbf{q}_c\|/\ell_{\text{base}}),$$

over its $n_{\text{nb}} = 10$ nearest previous-election places $\mathcal{N}_{\text{nb}}(b)$, with n_c valid votes, $\sum_c \alpha_{bc} = 1$ and $\ell_{\text{base}} = 2$ km, computed separately for advance and election-day places where the kind is known. No place needs to exist in both elections. The baseline's own sampling variance is carried forward.

9.3 A spatial swing field

The ILR composition of place b in the new election is modelled as

$$\mathbf{y}_b = \mathbf{m}_b + \mathbf{s} + \mathbf{f}(\mathbf{q}_b) + \mathbf{u}_b + \boldsymbol{\nu}_b, \quad (15)$$

with a national swing $\mathbf{s} \in \mathbb{R}^7$, a spatial field $\mathbf{f}$, a place-level nugget $\mathbf{u}_b \sim \mathcal{N}(\mathbf{0}, \sigma_{\text{nug}}^2 \Sigma_{\text{fld}})$ and sampling noise $\boldsymbol{\nu}_b \sim \mathcal{N}(\mathbf{0}, \iota \Sigma(\hat{\mathbf{x}}_b, n_b) + \text{baseline noise})$ with an overdispersion ι . The field is a zero-mean Gaussian process, with cross-coordinate covariance given by a matrix Σ_{fld} (taken from the voter-flow covariance of Equation 2 or estimated) and a spatial kernel ker ,

$$\text{Cov}(\mathbf{f}(\mathbf{q}), \mathbf{f}(\mathbf{q}')) = \text{ker}(\mathbf{q}, \mathbf{q}') \Sigma_{\text{fld}}$$

$$\text{ker}(\mathbf{q}, \mathbf{q}') = \sigma_{\text{fld}}^2 \exp(-\|\mathbf{q} - \mathbf{q}'\|/\ell_{\text{fld}}) \quad (16)$$

(the exponential family, or Matérn 3/2: $\text{ker}(\mathbf{q}, \mathbf{q}') = \sigma_{\text{fld}}^2 (1 + \sqrt{3}\|\mathbf{q} - \mathbf{q}'\|/\ell_{\text{fld}}) \exp(-\sqrt{3}\|\mathbf{q} - \mathbf{q}'\|/\ell_{\text{fld}})$), with a sill σ_{fld}^2 and length scale ℓ_{fld} fitted per transition and kind of place. The +0.5 pseudo-count, the minimum votes per place and places with no nearby previous-election booth are handled explicitly.

Evidence from the data. After removing the national mean and subtracting sampling noise, the semivariance of the swing rises from about 0.03 at 125 m to about 0.12 at 100 km and beyond. Exponential fits on election-day places give a length scale of 48 km (2017 to 2020) and 26 km (2020 to 2023), with a place-level nugget of about 0.05 to 0.06 per coordinate. Electorate means explain only 11 to 17% of the leftover swing variance, so a smooth field over location carries more information than an electorate-level swing does.

9.4 Updating as places report

Let $\mathcal{R}$ be the places that have reported and $\mathcal{U}$ those that have not. Under Equation 15, $(\mathbf{y}_b)_{b \in \mathcal{R} \cup \mathcal{U}}$ is jointly Gaussian given the hyperparameters, so the swing at unreported places is the exact Gaussian conditional,

$$\begin{aligned} \mathbf{f}_{\mathcal{U}} \mid \tilde{\mathbf{y}}_{\mathcal{R}} &\sim \mathcal{N}(\hat{\mathbf{f}}_{\mathcal{U}}, \mathbf{C}_{\mathcal{U} \mid \mathcal{R}}) \\ \hat{\mathbf{f}}_{\mathcal{U}} &= \mathbf{C}_{\mathcal{U}\mathcal{R}}(\mathbf{C}_{\mathcal{R}\mathcal{R}} + \Lambda_{\mathcal{R}})^{-1} \tilde{\mathbf{y}}_{\mathcal{R}} \\ \mathbf{C}_{\mathcal{U} \mid \mathcal{R}} &= \mathbf{C}_{\mathcal{U}\mathcal{U}} - \mathbf{C}_{\mathcal{U}\mathcal{R}}(\mathbf{C}_{\mathcal{R}\mathcal{R}} + \Lambda_{\mathcal{R}})^{-1} \mathbf{C}_{\mathcal{R}\mathcal{U}}, \end{aligned}$$

where $\tilde{\mathbf{y}}_{\mathcal{R}}$ are the reported swings with the baseline and the national mean removed, $\mathbf{C}$ are blocks of the covariance Equation 16 (a kernel value times Σ_{fld} for each pair of places and coordinates) and $\Lambda_{\mathcal{R}}$ the noise covariance (nugget plus sampling plus baseline noise). The national swing $\mathbf{s}$ has as its prior the difference between the polling model’s election-day forecast (with its systematic error, Equation 10) and the previous result; as places report, its posterior narrows, and the forecast moves from the poll prior toward the observed returns. With about 2,600 places a dense solve is feasible (a 2600×2600 kernel per coordinate); a nearest-neighbour approximation is the fallback. The update is repeated each time a batch of reports arrives.

9.5 Aggregating, and special votes

Expected turnout at each unreported place is its baseline place’s turnout, spatially smoothed because venues change, scaled to the electorate’s enrolment. Draws of place compositions are weighted by turnout and summed to electorate and national party shares, so intervals carry the correlation between places. Electorate boundaries for 2026 are applied by locating places inside the new boundary polygons, which also gives notional 2023 results for redrawn seats.

Special votes, about a fifth of the total (20.5% in 2023), are counted after election night and differ from ordinary votes. For election e and electorate l let $\mathbf{o}_{el} = \text{ilr}(\text{special shares}) - \text{ilr}(\text{ordinary shares})$. The model for it is

$$\begin{aligned} \mathbf{o}_{el} &= \mathbf{o}_e^{\text{nat}} + \mathbf{v}_{el} \\ \mathbf{v}_{el} &\sim \mathcal{N}(\mathbf{0}, \Omega), \end{aligned}$$

a national mean offset plus an electorate effect, learned across the eight elections, with the total special-vote count forecast from enrolment and its share in past elections. The forecast of the final result is the ordinary-vote forecast plus the special-vote forecast, so the central estimate moves again when the specials land.

9.6 Validation

The election-night model will be judged by replay. Take 2020 and 2023, leave the election out of the fit of the kernel and the special-vote model, reveal places in simulated orders (uniform random; largest first; by region; advance places first), and at each fraction counted record the error and the 80% interval coverage of national and electorate party shares. It is compared with three baselines: uniform national swing applied to the baseline; electorate-level swing; and the polling prior alone. A model that cannot beat uniform swing after 20% counted should be dropped. The same kriging cross-validation (hide 5% to 50% of places, predict from the rest) tests the swing field on its own.

10 Assumptions and limitations

- **Gaussian in log-ratio space:** opinion movements, house effects and the systematic error are Gaussian in ILR coordinates, so shares are log-normal-like with the right behaviour near zero. Sudden shocks (a leadership change) show up as large steps, not jumps.

- **Flows from one cycle:** the base rates come from the NZES 2020 to 2023 transition, and the multiplier u_f has only a modest prior range ($\exp(\pm 1/2)$ at one standard deviation). Stability of flows between cycles is untested.
- **Constant house effects and a shared design effect:** a pollster's bias does not change over time, and every pollster shares one design effect.
- **Tuned dynamics:** the drift settings are chosen on this cycle's history, dominated by the last year of sustained moves (NZ First and Opportunity up, Te Pāti Māori down). They are not given parameter uncertainty and may over-extrapolate if those trends reverse. The scale parameters c and κ are not separately identified.
- **Systematic error from four elections:** the amount is uncertain; 2017 dominates.
- **Seats:** list seats follow the MMP rules exactly. Electorate seats come from the electorate model, whose limits are listed in Section 7.9; the retention exponent is used only by the fallback rule.
- **Spatial model:** stationary kernel, one length scale per kind of place; mobile voting teams have no location and are excluded; special votes are handled at electorate level only.

11 Appendix: notation

The same conventions are used in every section. A plain italic letter is a scalar (h, δ, K). A bold lower-case letter is a vector ($\mathbf{z}, \boldsymbol{\eta}$), and an upper-case italic letter is a matrix (Q, V, P); the integers K and T are scalars. A calligraphic letter is a set or a distribution family: $\mathcal{N}(\mathbf{m}, \Sigma)$ is a Gaussian with mean $\mathbf{m}$ and covariance Σ , $\mathcal{E}$ and $\mathcal{G}$ are sets of parties, $\mathcal{R}$ and $\mathcal{U}$ are sets of voting places, and $\mathcal{N}_{\text{nb}}(b)$ is the set of neighbours of place b . Functions and named operators are upright, set with a text name and a subscript if they depend on a choice: ilr , softmax , diag , Cov , HalfNormal , and the observation function $\text{obs}_m(\mathbf{z}, \boldsymbol{\eta})$. A random variable is written with its distribution, $z \sim \mathcal{N}(\mathbf{m}, \Sigma)$, and the parameters of a distribution sit inside the brackets, or after a bar for a conditional ($z_t \mid z_{t+1}$). A hat marks an estimate, so $\hat{z}_{t|t}$ is the filtered mean of the state; a bar marks an average over an index ($\bar{\xi}$ is the mean of the firm vectors); a tilde marks an observation with its known offsets removed ($\tilde{\mathbf{y}}$) and a prime a dummy copy of an index (t'). A transpose is $A^\top$ and an inverse is A^{-1} . A subscript is an index or a name. A text subscript or superscript ($\sigma_{\text{iso}}, \mathbf{x}^{\text{prev}}$) names a quantity and is never an index. Upright subscripts such as “E” (polling day) are labels. A prefix Δ means “change in”, and Δ^{K-1} is the simplex; no other symbol uses Δ . Where a function such as $\Sigma(\cdot)$ returns a matrix, its subscripted forms are its values (for example Σ_p for poll p) or, when the subscript is a name, a separate named matrix (Σ_{fld}).

Symbol	Type	Meaning	Sec.
Indices (each letter has one meaning; a letter reused in one section only is said so)			
i, j, k	index	party or ILR coordinate ($i, j, k = 0, \dots, K - 1$ for parts, $1, \dots, K - 1$ for coordinates); in Section 3.3 j is the source and i the destination of a flow	2
t, t'	index	grid week, $0, \dots, T$; t' a dummy copy	3
p	index	poll	3.4
a	index	pollster (series); $a(p)$ is the pollster of poll p	3.4
g	index	polling firm; $g(a)$ is the firm of pollster a	3.4
f	index	directed flow $f = (j \rightarrow i)$ from party j to party i , $f = 1, \dots, n_{\text{flow}}$	3.2
m	index	block of polls that report the same parties	3.4
s	index	steps ahead (weeks) in the backtest, $s = 1, \dots, 26$	4.3
e	index	election	5.2
l	index	electorate, $l = 1, \dots, 71$	6
b, c	index	voting place	9
Compositions (Section 2)			
K	scalar	number of parties, 8; the ILR vector has $K - 1 = 7$ coordinates	2
Δ^{K-1}	set	the simplex of compositions	2
$\mathbf{x}$	vector	a composition (party shares), $x_i \geq 0, \sum_i x_i = 1$. $\mathbf{x}^{\text{prev}}$ the previous election, $\mathbf{x}^{\text{ref}}$ the reference shares for flows (Section 3.2), $\mathbf{x}_E$ the election-day shares (Section 5); x_i^{2023} is a 2023 share (Section 6)	2
$\hat{\mathbf{x}}_p$	vector	shares reported by poll p ; $\hat{\mathbf{x}}_b$ the counts of place b as shares	3.4
$\mathbf{z}$	vector	ILR coordinates in $\mathbb{R}^{K-1}$. $\mathbf{z}_t$ the latent state at week t ; $\mathbf{z}(t_p)$ its interpolation at poll p ; $\mathbf{z}^{\text{prev}}$ the ILR of the previous result	2
V	matrix	Helmert basis, $K \times (K - 1)$; V_m the same for the K_m categories of block m	2
$\text{ilr}, \text{softmax}, \text{gm}$	function	ILR transform, its inverse, geometric mean	2
$\text{diag}(\mathbf{a})$	function	diagonal matrix with the entries of a vector	2
$\mathbf{1}, \mathbf{0}, I$	vector, matrix	ones, zeros, identity	2

Symbol	Type	Meaning	Sec.
$\Sigma(\mathbf{x}, n)$	function	sampling covariance of the ILR of a poll of size n with shares $\mathbf{x}$ (Equation 1); Σ_m is the same in the basis V_m , and $\Sigma_p = \Sigma_m(\hat{\mathbf{x}}_p, n_p)$ its value for poll p	2
n_p, n_b	scalar	effective sample size of poll p ; valid votes at place b (or c)	2, 9
Polling model (Section 3)			
t_p	scalar	fieldwork midpoint of poll p , in days after the cycle start; t_p/h is its grid position	3.4
t_{10}, w_p	scalar	$\lfloor t_p/h \rfloor$ and the interpolation weight $t_p/h - t_{10}$	3.4
K_m, M_m	scalar, matrix	categories of block m ; the $K \times K_m$ matrix with entries in $\{0, 1\}$ that sends each party to its category	3.4
obs_m	function	observation function of block m : model prediction of a poll's ILR (Equation 5)	3.4
$\mathbf{y}_p, \tilde{\mathbf{y}}_p$	vector	ILR of the poll's reported shares; the same with known offsets removed (the house effect, or the linearisation point)	3.4, 4
$\boldsymbol{\nu}_p$	vector	sampling noise of poll p (ν_b for place b)	4
δ, ρ, δ'	scalar	design effect; its scale in the backtest; $\delta' = \max(\rho\delta, 1)$	3.4, 4.3
$\boldsymbol{\eta}_a$	vector	house effect of pollster a	3.4
$\boldsymbol{\xi}_g, \boldsymbol{\zeta}_a$	vector	standard-normal firm and series draws behind $\boldsymbol{\eta}_a$	3.4
$\sigma_{\text{firm}}, \sigma_{\text{ser}}$	scalar	firm and series house-effect scales	3.4
σ_0	scalar	prior sd of the start of the path, 0.15	3
$\mathcal{N}, \text{HalfNormal}$	distribution	Gaussian; half-normal with a given scale	3
Voter flows (Section 3.2)			
n_{flow}	scalar	number of directed flows, $K(K-1) = 56$	3.2
χ	scalar	fraction of a party's voters that moves along a flow	3.2
$\mathbf{a}_f, \bar{a}_f$	vector, scalar	raw flow vector: -1 at the source, $x_j^{\text{ref}}/x_i^{\text{ref}}$ at the destination, 0 elsewhere; the mean of its entries	3.2
D_f, D	row, matrix	row of the centred flow in ILR coordinates, $n_{\text{flow}} \times 7$ matrix stacking them	3.2
$r_f, \mathbf{r}, r_f^0$	scalar, vector	daily flow rate; vector of rates; its prior value	3.2
u_f	scalar	log-scale multiplier on flow f	3.2
Q	matrix	per-day covariance of opinion movements	3.2
σ_{iso}	scalar	isotropic movement not due to a listed flow	3.2
$P(i j)$	function	probability that a voter of party j switches to party i over one cycle	3.3
$v_{ji}, \lambda_j, \pi_i, \beta, \psi_{ji}$	scalar	utility; loyalty of source j ; pull of destination i ; weight on, and value of, the like-score covariate	3.3
n_{cyc}	scalar	days in the election cycle	3.3
Dynamics and exact paths (Section 4)			
h, T	scalar	grid step, 7 days; last grid week	3
$\boldsymbol{\varepsilon}_t$	vector	level shock of week t , $\mathcal{N}(\mathbf{0}, c^2 h Q)$ ($c = 1$ for the plain walk)	3
$\boldsymbol{\varepsilon}^{\text{init}}, \boldsymbol{\varepsilon}_t^{\text{std}}$	vector	standard-normal draws of the non-centred walk	3
L	matrix	Cholesky factor, $LL^\top = Q$	3

Symbol	Type	Meaning	Sec.
$\mathbf{d}_t$	vector	weekly drift of the state	4
$\boldsymbol{\varepsilon}_t^{\text{drift}}$	vector	drift shock, $\mathcal{N}(\mathbf{0}, \kappa^2 c^2 h Q)$	4
φ, c, κ	scalar	drift persistence; level-shock scale; trend scale	4
$\mathbf{z}_t^+, F, \boldsymbol{\varepsilon}_t^+$	vector, matrix	augmented state ($\mathbf{z}_t, \mathbf{d}_t$); its transition; its shock ($\boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_t^{\text{drift}}$)	4
H, H_z, J	matrix	observation matrix (J for a merged poll, Equation 7); its level block; Jacobian of obs_m	4
$\hat{\mathbf{z}}_{t t}, P_t$	vector, matrix	filtered mean and covariance of the state at week t (of $\mathbf{z}_t^+$ once drift is included); $\hat{\mathbf{z}}_{t t-1}, P_{t t-1}$ their one-step predictions	4
S, G, B_t	matrix	Kalman innovation covariance and gain; smoother gain	4
$\mathbf{z}^*, \mathbf{x}, \mathbf{x}^{\text{mrg}}$	vector	linearisation point; $\text{softmax}(V(\mathbf{z}^* + \boldsymbol{\eta}))$; its merged form $M_m^\top \mathbf{x}$	4
$t_0, \hat{\mathbf{z}}_{t_0+s t_0}, P_{t_0+s t_0}^{zz}$	scalar, vector, matrix	backtest origin week; level block of the s -step prediction mean and covariance	4.3
LS, n_{pred}	function, scalar	mean log predictive score; number of (origin, poll) pairs	4.3
Forecast and systematic error (Section 5)			
τ	scalar	days from the forecast to polling day	5
n_{wk}, w_τ	scalar	whole weeks to polling day, $\lfloor \tau/h \rfloor$; the part-week $\tau/h - n_{\text{wk}}$	5
$\mathbf{z}_E$	vector	state carried to polling day	5
$\boldsymbol{\varepsilon}_{\text{sys}}, \sigma_{\text{sys}}(\tau)$	vector, function	systematic error and its sd per coordinate	5
$\boldsymbol{\omega}_{e,\tau}, \sigma_{\text{fc},e,\tau,j}^2$	vector, scalar	actual minus mean forecast in ILR; the forecast's own variance in coordinate j	5.2
$\hat{\sigma}_{\text{sys}}^2(\tau), n_{\text{elec}}$	scalar	moment estimate of the extra variance; number of past elections used	5.2
v_{max}, τ_s	scalar	plateau variance and time scale of the curve	5.2
Seats (Section 6)			
$\mathcal{E}, \mathcal{G}$	set	parties eligible for list seats; parties in a government group	6
$n_i^{\text{list}}, n_i^{\text{win}}, n_i^{\text{seat}}, n_i^{\text{over}}$	scalar	list entitlement, electorate wins, total seats $\max(n_i^{\text{list}}, n_i^{\text{win}})$, overhang $n_i^{\text{seat}} - n_i^{\text{list}}$ of party i	6
n_{house}	scalar	size of the House, $120 + \sum_i n_i^{\text{over}}$	6
$i_l, \varpi_{l,i}$	index, scalar	holder of electorate l ; stated probability that party i wins a contested electorate l	6
γ	scalar	retention exponent, 0.5	6
Electorate model (Section 7)			
$c_{li}, c_{li}^{\text{prev}}$	scalar	candidate-vote share of party i in electorate l ; its notional share at the previous election in the current boundaries	7
μ_i	scalar	log change in party i 's national party-vote share since the previous election, $\log x_i - \log x_i^{2023}$	7
$\theta_{\text{swing}}, \theta_{\text{split}}, \theta_{\text{stand}}, \theta_{\text{retire}}$	scalar	coefficients of the national swing, the pull toward the candidate vote the party vote implies, a standing and a retiring member's party	7
$\text{gap}_{li}, \text{stand}_{li}, \text{retire}_{li}$	scalar	log of the candidate share the party vote implies under the transfers minus log candidate share at	7

Symbol	Type	Meaning	Sec.
		the previous election (Equation 12); indicators of a standing and of a retiring sitting member	
$x_{lj}^{\text{prev}}, \hat{x}_{bj}$	scalar	party-vote share of party j in electorate l at the previous election; at voting place b	7, 7.2
$\mathcal{S}_l$	set	parties contesting electorate l	7
$T_{ji}^{(l)}$	scalar	probability that a party-vote voter for j gives the candidate vote to i , over the parties in $\mathcal{S}_l$ (Equation 13)	7.2
$\lambda_j^s, \pi_i^s, \alpha_{ji}^s$	scalar	loyalty, pull and affinity of the split-vote transfers	7.2
$\varepsilon_{li}^{\text{elec}}, \mathbf{z}_l^{\text{elec}}$	scalar, vector	electorate noise in log shares; the same in ILR coordinates	7
$\Sigma_{\text{elec}}, D_s, \mathbf{r}_s, \mathbf{v}$	matrix, vector	ILR covariance of one electorate's noise (Equation 14); directions and rates of the split-vote flows; ILR direction of Other	7
$\sigma_{\text{flow}}, \sigma_{\text{iso}}^e, \sigma_{\text{oth}}$	scalar	scales of the electorate noise along the flows, along no flow, and for Other	7
$w_g, w_s, \ell_{\text{elec}}, b_l^{\text{elec}}, m_{\text{elec}}$	scalar	weights of the shared and spatial noise, the spatial length scale, the boundary factor and the noise multiplier	7
$d_{ll'}$	scalar	distance between the centroids of electorates l and l'	7
$\text{lead}_l, \text{flip}_l$	scalar	lead of the first party over the second in points of candidate vote; half of it, the swing that would tie the seat	7
Scenarios (Section 8)			
A, B	set	events on the draws: a scenario and a condition	8
$P(A), P(A \mid B)$	scalar	share of the draws in which A holds, and in which it holds among the draws where B holds	8
n_{draw}	scalar	number of posterior draws, 500	8
Election night (Section 9)			
$\mathbf{q}_b$	vector	location of place b (NZTM metres)	9
$\mathbf{y}_b, \mathbf{m}_b$	vector	ILR composition of place b ; its baseline from the previous election	9
$\alpha_{bc}, \mathcal{N}_{\text{nb}}(b), n_{\text{nb}}$	scalar, set, scalar	baseline weight of previous place c ; the $n_{\text{nb}} = 10$ nearest previous places; their number	9
$\ell_{\text{base}}, \ell_{\text{fld}}, \sigma_{\text{fld}}^2$	scalar	baseline length scale (2 km); field length scale and sill	9
$\mathbf{s}$	vector	national swing in ILR	9.3
$\mathbf{f}(\mathbf{q})$	function	spatial swing field, a zero-mean Gaussian process	9.3
$\text{ker}(\mathbf{q}, \mathbf{q}')$	function	scalar spatial kernel; the field covariance is $\text{ker} \cdot \Sigma_{\text{fld}}$	9.3
Σ_{fld}	matrix	cross-coordinate covariance of the field and the nugget	9.3
$\mathbf{u}_b, \sigma_{\text{nug}}$	vector, scalar	place-level nugget and its scale	9.3
ι	scalar	overdispersion of sampling noise	9.3
$\mathcal{R}, \mathcal{U}$	set	places that have reported; places that have not	9
$C_{\mathcal{U}\mathcal{R}}, \Lambda_{\mathcal{R}}$	matrix	blocks of the field covariance; noise covariance of the reported places	9
$\hat{\mathbf{f}}_{\mathcal{U}}, C_{\mathcal{U} \mid \mathcal{R}}$	vector, matrix	conditional mean and covariance of the field at unreported places	9

Symbol	Type	Meaning	Sec.
$\boldsymbol{o}_{el}, \boldsymbol{o}_e^{\text{nat}}, \boldsymbol{v}_{el}$	vector	special-vote offset (ILR of special minus ordinary shares), its national mean, its electorate effect	9
Ω	matrix	covariance of the electorate effect $\boldsymbol{v}_{el}$	9

Table 10: Notation. The last column gives the section where the symbol is introduced.

12 Appendix: derivations

12.1 The multinomial covariance in ILR coordinates

For n draws from a multinomial with probabilities $\mathbf{x}$, the observed proportions $\hat{\mathbf{x}}$ have covariance $(\text{diag}(\mathbf{x}) - \mathbf{x}\mathbf{x}^\top)/n$. The map $\hat{\mathbf{x}} \mapsto V^\top \log \hat{\mathbf{x}}$ has Jacobian $V^\top \text{diag}(1/\mathbf{x})$, so by the delta method

$$\begin{aligned} \text{Cov}(V^\top \log \hat{\mathbf{x}}) &= V^\top \text{diag}(1/\mathbf{x})(\text{diag}(\mathbf{x}) - \mathbf{x}\mathbf{x}^\top) \text{diag}(1/\mathbf{x})V/n \\ &= V^\top (\text{diag}(1/\mathbf{x}) - \mathbf{1}\mathbf{1}^\top)V/n \\ &= V^\top \text{diag}(1/\mathbf{x})V/n, \end{aligned}$$

using $V^\top \mathbf{1} = \mathbf{0}$, which is Equation 1.

12.2 The linearisation of a merged poll

For merged poll p in block m , $\text{obs}_m(\mathbf{z}, \boldsymbol{\eta}) = V_m^\top \log(M_m^\top \mathbf{x}(\mathbf{z}))$ with $\mathbf{x}(\mathbf{z}) = \text{softmax}(V(\mathbf{z} + \boldsymbol{\eta}))$. The softmax has Jacobian $\text{diag}(\mathbf{x}) - \mathbf{x}\mathbf{x}^\top$ with respect to its argument, and the argument has Jacobian V with respect to $\mathbf{z}$. The logarithm of the merged vector $\mathbf{x}^{\text{mrg}} = M_m^\top \mathbf{x}$ has Jacobian $\text{diag}(1/\mathbf{x}^{\text{mrg}})M_m^\top$ with respect to $\mathbf{x}$, and the final basis gives $V_m^\top$. Multiplying gives Equation 7.

12.3 Sainte-Laguë as sequential quotients

Awarding seats one at a time to the largest quotient $x_i/(2n_i^{\text{list}} + 1)$ gives the same allocation as dividing each party's vote by 1, 3, 5, ... and taking the 120 largest quotients overall, which is the Sainte-Laguë method.

12.4 Why the flow rates scale with $\sqrt{n_{\text{cyc}}}$

A random walk with daily standard deviation r_f has standard deviation $r_f\sqrt{n}$ after n days (variance adds). Setting that equal to the switching probability over a cycle of n_{cyc} days gives Equation 4: the displacement of each flow over one cycle is, in standard deviations, the fitted probability that a voter moves.