**9938A**

# Measuring Unmeasurability: Evidence of Strategic Vagueness in Singaporean Political Communication

**Abstract:** Do electoral incentives make politicians speak more ambiguously? In a signal model with rational voters anchoring on prior reputation, vagueness shields a favourable reputation but cannot rescue an unfavourable one. I test this on 62,169 hedged Singapore parliamentary claims using *semantic dispersion*, an LLM-derived measure validated at 92.3% agreement with human raters. A redistricting shock produces opposing-signed responses ruling out uniform confounders: MPs who lost favourable voters became vaguer by 0.15 standard deviations; those who gained unfavourable voters became more precise by 0.13 (sign-flip test $F = 7.89$, $p = 0.005$). A pre-election difference-in-differences using Singapore's Nominated MP institution corroborates.



**Word Count:** 7,490

# 1 Introduction

Consider a minister who, asked when a policy will be implemented, replies "we will study the matter in due course." The statement has the grammar of a commitment but not the substance. It is compatible with acting next week and with never acting, with a serious review and a cursory one, with implementation and with shelving. The minister has not lied: the phrase is defensible against any actual outcome. She has also not committed to anything voters could later check. This is a third option between truth-telling and deception: a statement whose range of possible meanings is wide enough that no realisation of the world will contradict it, while a precise reply ("within six months") spans a narrow range. Call this range strategic vagueness. Whether electoral incentives drive politicians to widen it is the question of this paper. I answer it with two Singapore natural experiments: a redistricting shock that produces opposing-signed responses across MPs with different prior reputations, and a within-chamber control of legislators facing zero electoral pressure.

Democratic accountability rests on the supply of verifiable content. Voters hold representatives to account by comparing claims against outcomes; an MP who reduces the verifiable content of her own speech defeats that comparison. The friction is asymmetric information, and electoral pressure should tighten it where accountability most needs to bite. Whether it does, and by how much, is an open empirical question.

Downs (1957), Shepsle (1972), and Glazer (1990) suggested electoral incentives reward this kind of language. No one has tested the prediction in observational data: vagueness resists quantification, and politicians facing different electoral pressures are rarely otherwise comparable.

This paper engages four literatures: strategic ambiguity, cheap talk and information design, political agency, and computational text analysis. None alone delivers the *sign-flip* (vagueness helping a favourable prior, hurting an unfavourable one); the combination does.

Theoretical work on strategic ambiguity treats vagueness as a campaign-platform choice (Alesina and Cukierman, 1990; Aragonès and Neeman, 2000; Callander and Wilson, 2008; Downs, 1957; Glazer, 1990; Shepsle, 1972); I test whether the same incentives shape continuous legislative speech. The model draws on two formal traditions: cheap talk (Crawford and Sobel, 1982; Kartik, 2009), where a self-interested sender chooses messages after observing the state, and information design (Kamenica and Gentzkow, 2011), where the sender commits to a signal structure ex ante. Cheap talk supplies state-dependent sender preferences but no ex-ante stylistic commitment; information design supplies commitment but no state-dependent preferences. The architecture this paper requires is the combination. The political-agency literature (Ashworth, 2012; Besley, 2006; Canes-Wrone et al., 2001; Maskin and Tirole, 2004; Morris, 2001) frames re-election incentives as a binary pander-or-lead choice; the obstacle to extending this to continuous precision has been measurement. Maskin and Tirole's (2004) accountable-versus-independent distinction underwrites the NMP / elected-MP comparison this paper exploits.

Methodologically, Gentzkow et al. (2019b) measure congressional partisanship via word-choice differences across speaker groups; Djourelova (2023) measures slanted language in media coverage. Neither targets interpretive breadth: Gentzkow et al. capture dimensions of difference between known groups, Djourelova captures slant. Semantic dispersion measures something different: the spread of meanings a single hedged claim admits, derived from LLM-generated paraphrase distance. Ornstein et al. (2025) and Asirvatham et al. (2026) establish the LLM-measurement instrument class but do not apply it to strategic communication.

Empirically, a growing computational literature uses text analysis to study political language: Ash et al. (2017) pair a signal model with text analysis to show that electoral incentives induce divisive "posturing"; Di Tella et al. (2023) document rhetorical-complexity convergence between campaign rounds. These computational approaches study what politicians emphasise; the closest experimental anticipations of what they leave undetermined come from Thaler et al. (2024), who document strategic vagueness with naïve receivers, and Tomz and Van Houweling (2009), who show that voter responses to candidate ambiguity depend on prior beliefs. The two establish the mechanism's components separately: prior-dependence on the receiver side and senders exploiting imprecision. The present paper introduces the receiver-regime distinction: rational-receiver behaviour, not naïve reception, is what produces the opposing-signed responses to a single shock that the model predicts and the data display.

I address both obstacles. Vagueness's resistance to quantification falls to the constructed measure (semantic dispersion). Comparability across electoral-pressure regimes falls to two Singapore institutions: the Nominated MP scheme, which seats non-partisan experts who debate fully but face zero electoral pressure, and the Electoral Boundaries Review Committee, which deliberates in secret and publishes without notice, generating prior shocks incumbents cannot anticipate. The pair (zero-incentive control plus exogenous prior shifts) is rare in observational data. Section 2 details both.

The paper makes three contributions. First, I develop a signal model in which a politician commits to communication precision before learning what she will need to say. A rational voter discounts vague messages and falls back on her prior. Vagueness thus *shields* a favourable reputation; only precision (a *lottery ticket*) can move an unfavourable one. Second, I construct semantic dispersion, a transferable measure of interpretive breadth, and validate it through a pairwise forced-choice study giving 92.3% agreement with human judgements ($p < 10^{-6}$). Third, I provide causal evidence. The redistricting result is the headline: a single exogenous shock produces opposing-signed responses for MPs whose priors moved in opposite directions. Reduced-prior MPs (whose constituencies lost reliably favourable voters) became vaguer (+0.026, $p = 0.009$); mixed-prior MPs (whose constituencies absorbed new voters with different priors) became more precise (−0.023, $p = 0.17$); a formal sign-flip test rejects equal responses ($F = 7.89$, $p = 0.005$). A corroborating difference-in-differences finds elected MPs increase vagueness by 15% of a standard deviation relative to NMPs before elections ($p = 0.041$); senior MPs show larger pre-election shifts (+0.041, $p = 0.024$), as predicted by the dynamic extension. The redistricting effect survives topic fixed effects with 90% intact (MPs become vaguer about the same things, not vaguer because they shift to vaguer topics), and within-MP persistence is high ($\hat{\rho} = 0.42$) with level shifts at electoral transitions, confirming vagueness is a sustained style rather than a claim-by-claim reaction to bad news.

Section 2 describes Singapore's institutional setting. Section 3 develops the signal model and derives testable predictions. Section 4 describes the data, measurement pipeline, and human validation. Section 5 presents the empirical analysis. Section 6 addresses robustness and Section 7 concludes.

## 2 Institutional Background

The People's Action Party (PAP) has governed Singapore without interruption since independence in 1965; opposition representation has never exceeded 12% of parliamentary seats. Within the PAP, the Central Executive Committee re-selects candidates through a multi-stage interview process that weighs incumbents' parliamentary contributions (Mauzy and Milne, 2002). Incumbent MPs therefore face two accountability channels: general-election competition (externally) and re-selection by the party (internally). Both rise as an election approaches.

The Electoral Boundaries Review Committee (EBRC) redraws constituency boundaries before each general election. The EBRC deliberates in secret, receives its terms of reference from the Prime Minister's Office, and publishes its report without prior notice, typically two to four months before polling. Individual MPs cannot anticipate or influence boundary changes.

Most Singapore constituencies are Group Representation Constituencies (GRCs), in which slates of three to six MPs are elected as a single ticket; the remainder are Single Member Constituencies (SMCs). The team-ticket structure of GRCs diffuses individual accountability, motivating the GRC-vs-SMC heterogeneity reported in Section 5.1.

Parliament also seats Nominated MPs (NMPs), appointed by the President on the recommendation of a Special Select Committee for two-and-a-half-year terms. Drawn from civil society (academia, the professions, the arts, labour), NMPs sit on Select Committees and deliver substantive policy speeches, but cannot stand for election or contest constituencies. They provide the control group for the difference-in-differences in Section 5.1.

## 3 Theoretical Framework

This section develops a commitment-based signalling model of strategic vagueness. The model has ex-ante commitment to a signal structure (as in information design) with a self-interested sender whose preferences depend on the realised state (as in cheap talk). An MP commits to communication precision before learning her performance; a rational voter discounts vague messages.

### 3.1 Setup

Two players: an elected MP (sender) and a representative voter (receiver). MP performance is $\theta \sim \mathcal{N}(\theta_0, \sigma^2)$; $\theta_0$ is the voter's prior mean and is common knowledge.[1] Normal distributions deliver closed-form comparative statics; directional predictions (Propositions 1–3) extend to a wider class of signal structures where the posterior mean shifts toward the prior as noise rises.

The game proceeds in four stages:

1. The MP chooses a level of vagueness $k$, a communication precision for statements about her performance.
2. Nature draws the MP's performance $\theta \sim \mathcal{N}(\theta_0, \sigma^2)$.
3. The MP observes $\theta$ and sends a message.
4. The voter observes the message, updates her belief via Bayes' rule, and retains the MP if and only if the posterior mean of $\theta$ exceeds threshold $\bar{V}$.

In plain terms: the MP picks a style now, observes her performance later, speaks to that style, and is judged. The critical feature is that **vagueness $k$ is chosen at Stage 1, before the MP learns $\theta$**.

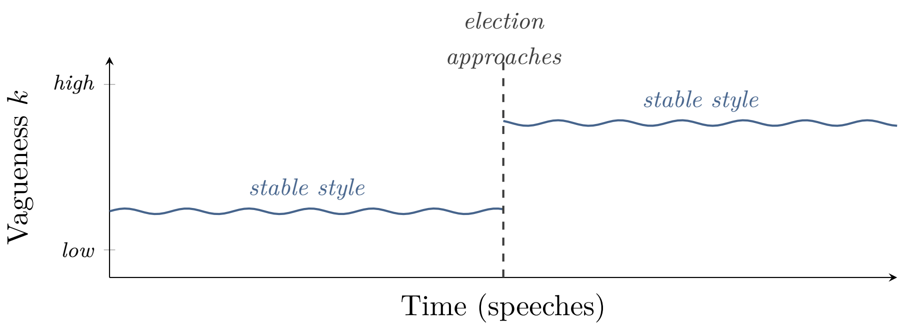


The specific claim content is chosen later at Stage 3 after observing $\theta$. This timing is not a modelling convenience but a requirement for the strategy. A poker player who only raises with

[1] $\theta$ is realised performance, not effort or ability in the career-concerns sense.

strong hands gets read by rational opponents, and the bluff collapses; an MP who only goes vague when hiding bad outcomes would be read the same way.[2] For vagueness to function as a shield, it must be sustained as a consistent *style*, deployed regardless of the realised state, so that the voter cannot extract information about $\theta$ from the choice of $k$. I test this assumption directly in Section 5.3.

Assume that at Stage 3, the voter's observation of a message with vagueness $k$ is equivalent to receiving a signal

$$y \mid \theta \sim \mathcal{N}(\theta, k^2), \tag{1}$$

so higher $k$ means a noisier signal. The empirical counterpart of $k$ is semantic dispersion, defined in Section 4.

A key assumption: voters do not learn the true performance $\theta$ before the election; the message $y$ is all they observe.

At Stage 4, the voter retains the MP if the posterior mean exceeds threshold $\bar{V}$:[3]

$$\mathbb{E}[\theta \mid y] = \underbrace{\frac{\sigma^2}{\sigma^2 + k^2}}_{\equiv \rho(k)} y + \underbrace{\frac{k^2}{\sigma^2 + k^2}}_{1-\rho(k)} \theta_0 \;\geq\; \bar{V}. \tag{2}$$

The weight $\rho(k) \equiv \sigma^2/(\sigma^2 + k^2)$ is the credibility weight: how much the voter trusts the message versus the prior. The voter observes $k$, assesses precision, and discounts rationally.[4] Since $\rho'(k) < 0$, vagueness reduces signal informativeness: high $k$ pushes the voter toward the prior $\theta_0$, low $k$ toward the signal. The MP anticipates this and optimises over $k$.

### 3.2 Retention Probability and the Sign-Flip Property

The unconditional distribution of the signal is $y \sim \mathcal{N}(\theta_0, \sigma^2 + k^2)$. Defining $z(k) \equiv \frac{(\bar{V}-\theta_0)\sqrt{\sigma^2+k^2}}{\sigma^2}$, which measures how far the re-election threshold lies from the prior in units scaled by total uncertainty, retention probability is

$$P(k) = 1 - \Phi(z(k))$$

**Proposition 1 (Sign-Flip).**

$$\partial P/\partial k \gtreqless 0 \iff \theta_0 \gtreqless \bar{V}.$$

[2]Formally, if $k$ were chosen after observing $\theta$, voters would infer $\theta$ from $k(\theta)$, and standard unravelling (Milgrom, 1981) would push $k \to 0$.

[3]Retention approaches a deterministic threshold in the large-electorate limit, where idiosyncratic voter shocks average out; stochasticity in this model enters only through the signal. Singapore's dominant-party margins fit this limit.

[4]Experimental evidence supports partial but incomplete discounting: Thaler et al. (2024) find receivers fail to fully correct for sender imprecision even when incentives are known, consistent with $0 < \rho(k) < 1$ for $k > 0$.

*Proof.* Differentiating via the chain rule:

$$\frac{dP}{dk} = -\phi(z) \cdot z'(k), \quad z'(k) = \frac{(\bar{V} - \theta_0)k}{\sigma^2\sqrt{\sigma^2 + k^2}}.$$

Since $\phi(z) > 0$ and $k > 0$: $\text{sign}(dP/dk) = -\text{sign}(\bar{V} - \theta_0) = \text{sign}(\theta_0 - \bar{V})$. □

Higher $k$ does two things to the voter's posterior: it shifts the mean toward the prior $\theta_0$, and it concentrates the posterior around that shifted mean. When $\theta_0 > \bar{V}$, both effects push probability mass into the safe region above $\bar{V}$. Vagueness shields the MP against a bad draw of $\theta$. When $\theta_0 < \bar{V}$, both effects push mass below $\bar{V}$, and vagueness locks in defeat. Only precision gives the voter's posterior a chance of crossing the threshold, acting as a lottery ticket. Whether vagueness helps or hurts flips with the prior (Figure 1).

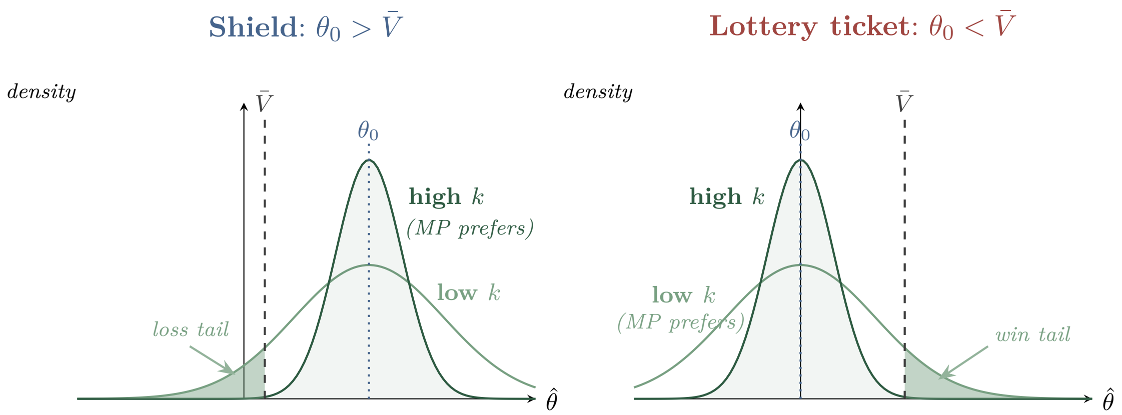


**Figure 1.** The sign-flip property. Each panel plots the density of the voter's posterior mean $\hat{\theta}$ under a vague strategy (high $k$; narrow distribution) and a precise strategy (low $k$; wide distribution). Higher $k$ produces a *tighter* distribution around the prior $\theta_0$ because the voter rationally discounts the noisy signal and anchors on the prior; lower $k$ produces a wider distribution because the voter updates more fully on the signal. Retention requires $\hat{\theta} \geq \bar{V}$. **Left (shield, $\theta_0 > \bar{V}$):** the vague strategy keeps almost all posterior mass safely above $\bar{V}$; the precise strategy risks a loss tail below $\bar{V}$. **Right (lottery ticket, $\theta_0 < \bar{V}$):** the vague strategy locks the voter at a prior below the re-election threshold; only the precise strategy offers a chance of crossing $\bar{V}$ through its upper tail.

### 3.3 The MP's Problem

The MP maximises

$$\Pi(k) = \lambda \cdot P(k) - C(k),$$

where $\lambda$ is the weight the MP places on being re-elected.[5] For NMPs, $\lambda = 0$: they cannot stand for election, so the re-election term drops out. For elected MPs, $\lambda$ rises as the next election approaches. $C(k)$ is the politician's cost of being vague, which is a continuously differentiable function satisfying $C(0) = 0$, $C'(k) > 0$ for $k > 0$, and $C''(k) > 0$. These assumptions say no penalty for speaking clearly, increasing marginal penalty for greater imprecision, and strict convexity (which pins down an interior solution in the shield regime). The cost bundles forces

[5]The career-concerns structure parallels Holmström (1999), with obfuscation cost replacing effort cost.

that disfavour maximal vagueness outside the voter-updating channel: policy commitment, professional reputation within the legislature, and constituent-service demands for specificity.[6] Without this cost, any MP with $\theta_0 > \bar{V}$ would choose maximum vagueness regardless of electoral pressure; the cost creates an interior solution and lets $\lambda$ move $k^*$. The equilibrium concept is Bayesian Nash: $k^*$ is the MP's best response to the voter's updating rule, and the voter's updating rule is rational given $k^*$.

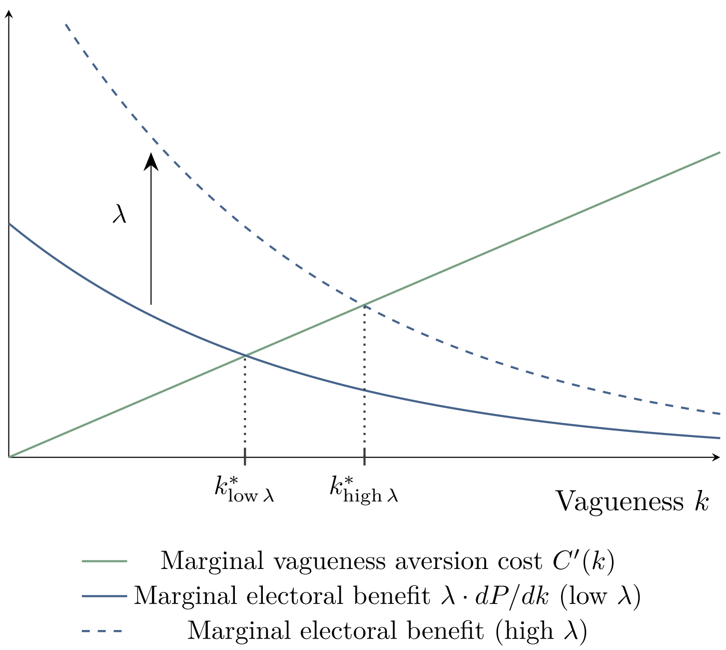


NMPs have $\lambda = 0$, so their objective collapses to $\min_k C(k)$ and $k^* = 0$: with no electoral benefit, the cost dominates and they speak as precisely as possible. This is the empirical control.

### 3.4 Optimal Strategy

The solution requires:
(i) the MP chooses $k^*$ to maximise $\Pi(k) = \lambda \cdot P(k) - C(k)$ anticipating the voter's updating rule;
(ii) the voter updates via Bayes' rule, $\hat{\theta} = \rho(k)\, y + (1 - \rho(k))\, \theta_0$;
(iii) the voter retains iff $\hat{\theta} \geq \bar{V}$. Since the voter's response is uniquely determined for any $k$, the strategic content lies in the MP's anticipation of rational discounting: the MP knows the voter will down-weight vague messages, and optimises $k$ accordingly.

**Proposition 2 (Electoral incentives raise vagueness).**

$$\textit{When } \theta_0 > \bar{V} : dk^*/d\lambda > 0.$$

[6] Page (1976) proposed an alternative mechanism in which politicians de-emphasise divisive specifics rather than offer risky lotteries; the present model nests this as the limiting case where $k$ operates through salience rather than noise.

*Proof.* The first-order condition (FOC) is $\lambda \cdot dP/dk - C'(k) = 0$. By the implicit function theorem (IFT), $dk^*/d\lambda = -\Pi_{k\lambda}/\Pi_{kk}$. At any interior optimum $k^*$ the second-order condition $\Pi_{kk} = \lambda \cdot d^2P/dk^2 - C''(k^*) < 0$ must hold, which requires $C$ to be sufficiently convex relative to any local convexity of $P$; I maintain this assumption throughout the shield regime. The cross-partial is $\Pi_{k\lambda} = dP/dk > 0$ when $\theta_0 > \bar{V}$ (Proposition 1). Hence $dk^*/d\lambda > 0$. □

When electoral stakes $\lambda$ rise, the shielding benefit of vagueness grows relative to the vagueness aversion cost, tipping the margin in favour of less precise communication. The first-order condition equates the marginal electoral benefit of vagueness (the shielding effect weighted by $\lambda$) with the marginal vagueness aversion cost ($C'(k)$). As $\lambda$ increases, the benefit curve shifts up and the optimal $k^*$ increases.

Very safe MPs have less use for the shield than moderately safe ones. An MP who is near-certain to win gains little from vagueness; a moderate decrease in $\theta_0$ moves the MP into the range where shielding has real marginal value. Formally, applying the IFT to the FOC with respect to $\theta_0$ yields $\text{sign}(dk^*/d\theta_0) = \text{sign}(1 - z^2)$, so when $|z| > 1$ (the MP is comfortably safe), $dk^*/d\theta_0 < 0$: a decrease in the prior raises optimal vagueness. This comparative static drives the redistricting prediction (Section 5): losing supportive voters reduces $\theta_0$ while keeping it above $\bar{V}$, increasing the incentive to shield.

**Proposition 3 (Corner solution).**

$$\textit{When } \theta_0 < \bar{V} : k^* = 0 \quad \forall \lambda.$$

*Proof.* The FOC requires $\lambda \cdot dP/dk = C'(k)$. When $\theta_0 < \bar{V}$, Proposition 1 gives $dP/dk < 0$ for all $k > 0$, so LHS $< 0$ and RHS $> 0$. No interior solution exists; $k^* = 0$. □

Proposition 3 is the empirical counterpart of the sign-flip: vagueness is harmful at every margin when $\theta_0 < \bar{V}$, so no amount of electoral pressure can make it attractive. In practice, natural language imposes a floor on interpretive breadth (no claim can have exactly zero dispersion), so the empirical prediction is that MPs in this regime push *toward* that floor; the test is the sign of the change, not the level.

### 3.5 Dynamic Interpretation: Investing and Harvesting

In a multi-period setting, the voter's prior $\theta_0$ evolves: $\theta_{0,t+1} = \rho(k_t)y_t + (1 - \rho(k_t))\theta_{0,t}$. Precise messages update the prior quickly, vague messages leave it roughly unchanged. Early in a term, when the election is distant and $\lambda$ is low, the cost of being vague dominates and the MP speaks precisely. Doing so builds up $\theta_0$: precision lets good performance register. Closer to the election, $\lambda$ rises, the shield becomes worth the cost, and the MP switches to vagueness, cashing in the accumulated reputation. The model therefore predicts precision early in a term and vagueness late, through both the static cost-benefit tilt and the dynamic logic of building reputation before harvesting it.[7]

These dynamics generate two further empirical implications. First, senior MPs should exhibit larger pre-election vagueness increases: more time in the investment phase builds higher $\theta_0$, making the swing to harvesting more pronounced. Second, redistricting that introduces unfamiliar voters functions as a regime reset. New voters have no prior attachment, so $\theta_0$ is low. The prediction is that MPs absorbing such voters become *more precise*, even close to an election.

### 3.6 Predictions

The model produces four testable predictions. First, Proposition 2 implies that elected MPs should become more vague as elections approach (when $\lambda$ rises); NMPs, with $\lambda = 0$, should not. The dynamic interpretation sharpens this: senior MPs, who have spent longer in the investment phase building $\theta_0$, should show larger pre-election increases. Both predictions are tested by a pre-election difference-in-differences with seniority interactions. Second, Proposition 1 implies that MPs whose constituencies lose high-support polling districts should become *more* vague (shield regime); third, Proposition 3 implies that MPs whose constituencies absorb low-support polling districts should become *more precise* (lottery-ticket regime). The opposing signs from a single institutional shock are tested jointly in the redistricting natural experiment. Fourth, pre-commitment implies that within-MP vagueness should be persistent across speeches and shift only at electoral transitions, not claim-by-claim with individual policy outcomes.

### 3.7 Mapping the model to Singapore

The two regimes correspond to the sign of $\theta_0 - \bar{V}$. In Singapore (Section 2), $\theta_0 > \bar{V}$ is overwhelmingly typical for PAP incumbents: the party has governed continuously since 1965 and voter priors are shaped primarily by disposition toward the PAP rather than the individual candidate. Past PAP vote share therefore serves as the cleanest publicly observable proxy for $\theta_0$. The $\theta_0 < \bar{V}$ regime activates only under specific shocks, most cleanly redistricting that introduces voters with no prior attachment. All elected MPs in the analysis sample are PAP; opposition representation has never exceeded 12% of seats and no opposition MP is affected by the 2020 boundary changes.

**Redistricting predictions.** When boundaries change, MP $i$'s effective prior shifts toward the priors of voters who join (and away from those who leave):

$$\Delta\theta_{0i} \;\propto\; \sum_{d \in \mathcal{D}_i} s_d \left(\mathbb{E}[\theta_0 \mid V_d] - \theta_{0i}\right), \tag{3}$$

[7]The pre-commitment assumption at Stage 1 follows from the repeated game: voters condition their updating on realised $k_t$ as a second signal, so the equilibrium keeps $k$ constant within the style set by the current phase. Formally, the infinite-horizon problem is $\max_{\{k_t\}} \sum_{t=0}^{\infty} \beta^t \big(\lambda_t P_t(k_t) - C(k_t)\big)$, with $\beta \in (0, 1)$, $\lambda_t$ rising toward the election, and the prior updating by Bayes' rule. A full numerical characterisation of the optimal switching point is listed as an extension in Section 7.

where $\mathcal{D}_i$ is the set of polling districts entering ($s_d > 0$) or leaving ($s_d < 0$) MP $i$'s constituency, weighted by population share $|s_d|$, and $V_d$ is the source-constituency PAP vote share at the previous election. The maintained assumption is monotonicity of the proxy: $\mathbb{E}[\theta_0 \mid V_d]$ is increasing in $V_d$, so voters from a historically higher-PAP origin yield a higher expected prior. Equation 3 identifies the *sign* of each MP's prior shift, not its magnitude: polling-district-level priors are not directly observable, only their constituency-level proxy. Two cases follow. A *reduced-prior* MP whose constituency cedes polling districts from high-support origins ($\mathbb{E}[\theta_0 \mid V_d] > \theta_{0i}$ with $s_d < 0$) sees $\theta_{0i}$ fall while remaining above $\bar{V}$: the model predicts vagueness rises (shield). A *mixed-prior* MP whose constituency absorbs polling districts from low-support origins ($\mathbb{E}[\theta_0 \mid V_d] < \theta_{0i}$ with $s_d > 0$) sees $\theta_{0i}$ fall, potentially below $\bar{V}$: the model predicts precision rises (lottery ticket). The two cases share a common direction of $\Delta\theta_{0i}$ but straddle $\bar{V}$, so the predicted response flips sign.

# 4 Data and Measurement

Testing these predictions requires a measure of communication precision that varies at the claim level and a setting with clean variation in electoral incentives. Singapore provides both.

## 4.1 Parliamentary Speech Corpus

I scrape the complete Singapore Hansard (the official record of parliamentary debates, 2006–2025) and parse speeches into individual claims using GPT-4o-mini. The extraction prompt forces the model's output into a fixed template with the same fields for every speech. The prompt asks the model to pick out statements that make a factual claim, whether quantitative (with numeric values) or qualitative (with vague quantity language). Procedural language ("Mr Speaker," "I thank the Minister"), questions, and pure expressions of support or opposition are excluded. GPT-4o-mini identifies 355,756 factual claims across the corpus.

Hedged statements are classified via a two-tier architecture. The first tier matches against a curated lexicon of approximately 100 entries across twelve categories, including epistemic modals ("may," "might," "could"), approximators ("about," "approximately," "roughly"), plausibility shields ("I think," "I believe"), and a Singapore-specific political category of stock phrases that signal no commitment ("we will study," "in due course").[8] Unambiguous matches receive high confidence scores (0.85–0.95). The second tier handles ambiguous words ("can," "will," "would," "should") that function as hedges in some contexts but not others. A logistic regression classifier trained on 134 labelled examples from CoNLL-2010, BioScope, and political speech corpora examines a five-word context window around each target to determine hedge status. Inter-method agreement between the lexical and ML components exceeds 90%.

For each hedged claim, GPT-4o-mini generates twenty substitutions: ten numerical interpretations spanning the plausible quantitative range and ten lexical interpretations spanning an intensity spectrum from weak to strong.[9] The interpretation prompt instructs the model to replace each vague term with a specific alternative, with type-specific guidance for quantities, temporal references, evaluative terms, approximators, and comparatives. "Many" yields numerical substitutions [25%, 35%, . . . , 85%] and lexical substitutions [*slightly*, *marginally*, . . . , *predominantly*].

The analysis sample contains 62,169 claims with valid semantic dispersion scores (the measure is defined in Equation 4): 58,620 from 243 elected MPs, 1,798 from 28 NMPs, and 1,751 from 8 NCMPs. The NMP sample is small in absolute terms but sufficient for the DiD, whose identification comes from the *interaction* of MP type with the pre-election period.

[8] The full lexicon is available on request.

[9] All LLM calls use temperature 0 with a fixed seed (42) and strict JSON schema output, ensuring deterministic, reproducible generation. Because outputs are deterministic, the dispersion measure captures semantic properties of the input claim, not randomness in the generation process.

### 4.2 Semantic Dispersion

Semantic dispersion measures the interpretive breadth a hedged claim admits, using LLM-generated paraphrases scored in embedding space. The pipeline (Figure 2) applies to any corpus where hedged claims can be identified, so the measure is portable beyond Singapore.

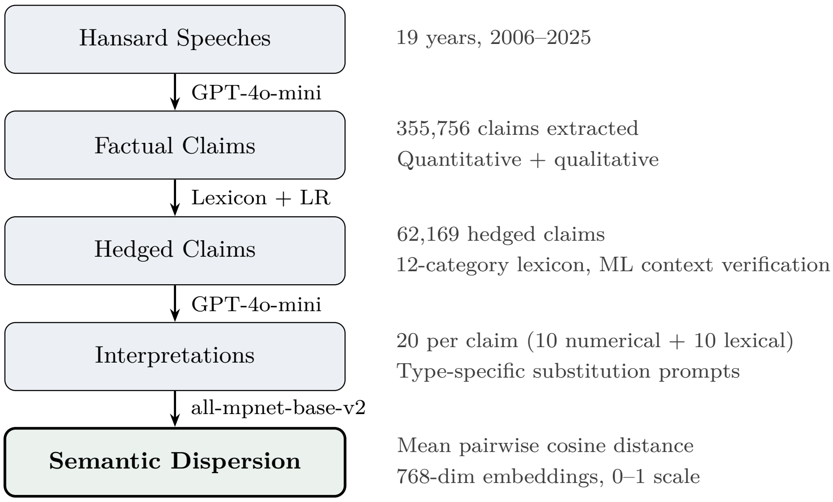


**Figure 2.** Measurement pipeline. Speeches are parsed into factual claims, hedged claims are identified via a hybrid lexical-ML classifier, GPT-4o-mini generates twenty substitutions per claim, and `all-mpnet-base-v2` sentence embeddings yield the semantic dispersion score.

Figure 3 provides geometric intuition: a vague claim like "improve the economy" scatters its interpretations widely in embedding space, while a precise claim like "reduce inflation by 2%" clusters them tightly.

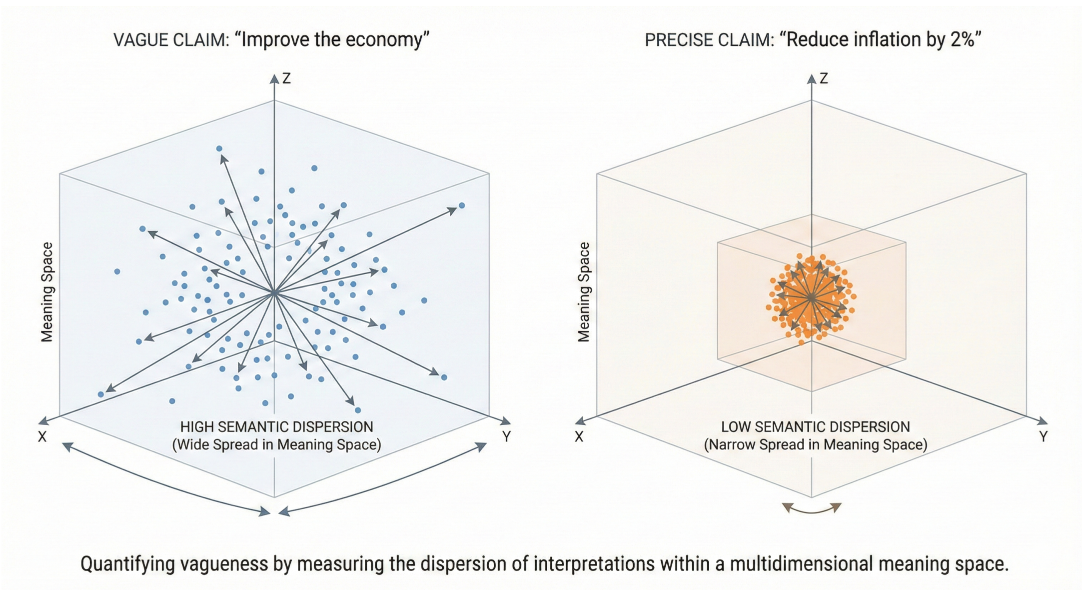


**Figure 3.** Semantic dispersion in meaning space. **Left**: A vague claim ("Improve the economy") generates interpretations that scatter widely across a multidimensional embedding space, yielding high dispersion. **Right**: A precise claim ("Reduce inflation by 2%") generates tightly clustered interpretations, yielding low dispersion.

The outcome variable is the average pairwise cosine distance between 768-dimensional sentence embeddings from `all-mpnet-base-v2` of the generated interpretations:[10]

$$\text{SemanticDispersion}_i = \frac{2}{J(J-1)} \sum_{j<j'} \left(1 - \frac{e_j \cdot e_{j'}}{\|e_j\|\|e_{j'}\|}\right). \tag{4}$$

Table 1 reports sample composition and dispersion descriptives by MP type.

**Table 1.** Sample Composition

| | Claims | MPs | Mean Disp. | SD | Elections |
|---|---|---|---|---|---|
| Elected MPs | 58,620 | 243 | 0.436 | 0.179 | 5 |
| Nominated MPs | 1,798 | 28 | 0.458 | 0.169 | — |
| Non-Constituency MPs | 1,751 | 8 | 0.451 | 0.171 | — |
| Total | 62,169 | 279 | 0.437 | 0.178 | 5 |

*Notes:* Elections: 2006, 2011, 2015, 2020, 2025. Semantic dispersion = mean pairwise cosine distance between `all-mpnet-base-v2` embeddings of LLM-generated interpretations.

Table 2 illustrates the measure on parliamentary claims spanning the dispersion distribution.

**Table 2.** Illustrative Claims and Semantic Dispersion Scores

| Claim | Dispersion | Interpretive range |
|---|---|---|
| "about 230 school fields" | 0.01 | Narrow: 200–260 fields |
| "About 45% of all cases were reported through the National Anti-Violence Helpline" | 0.01 | Narrow: 40–50% |
| "The government has made significant progress" | 0.43 | Wide: 5% improvement to structural reform |
| "We must continue to evolve a more compassionate education system" | 0.66 | Very wide: incremental tweaks to wholesale redesign |

*Notes:* All claims from the Singapore Hansard corpus. Dispersion defined in Equation 4.

The first two claims contain numerical approximators ("about 230," "About 45%") whose interpretations cluster tightly because the plausible range is anchored by the stated figure. The third and fourth use evaluative language ("significant progress," "more compassionate") that admits interpretations spanning orders of magnitude. The measure captures this distinction: it responds to the width of the interpretive window, not to whether a hedge is present.

[10] Re-running 50 claims after a GPT model update reproduces individual interpretations only 42.8% of the time, but the dispersion *scores* (averaged over twenty interpretations per claim) are substantially more stable. All main-text results use cached outputs from a single generation pass.

Just as $k$ governs the spread of the voter's signal around the true state, semantic dispersion computes the spread of plausible interpretations around the claim's core meaning: the empirical counterpart of $k$. The choice to measure *interpretive range* rather than lexical hedging follows Thaler et al. (2024): senders exploit any form of imprecision, including numerical intervals. Semantic dispersion is uncorrelated with hedge *frequency* ($r \approx 0$): the model predicts adjustment on the intensive margin (precision of each hedge), not the extensive margin (frequency of hedging). A potential concern is that the measure conflates strategic vagueness with topical breadth, since broad topics mechanically scatter interpretations. Section 5.1 addresses this by conditioning on the claim's own embedding.

### 4.3 Human Validation

I validate semantic dispersion against human judgement using a pairwise forced-choice study, grounded in Thurstone (1927)'s finding that humans are more reliable at relative than absolute perceptual judgements.[11] The binary outcome (correct/incorrect identification of the higher-dispersion claim) permits exact binomial inference without distributional assumptions.

Twenty claim pairs were constructed from the extreme tails of the dispersion distribution (top vs. bottom quintile), matched to within 30% on word count, plus one gold-standard pair (constructed by the author with an obvious correct answer, used as an attention check). Fifteen raters recruited via Prolific evaluated all 21 pairs, answering: *"Which statement could be interpreted in MORE DIFFERENT WAYS?"* One rater who failed an attention check was excluded ($N = 14$).

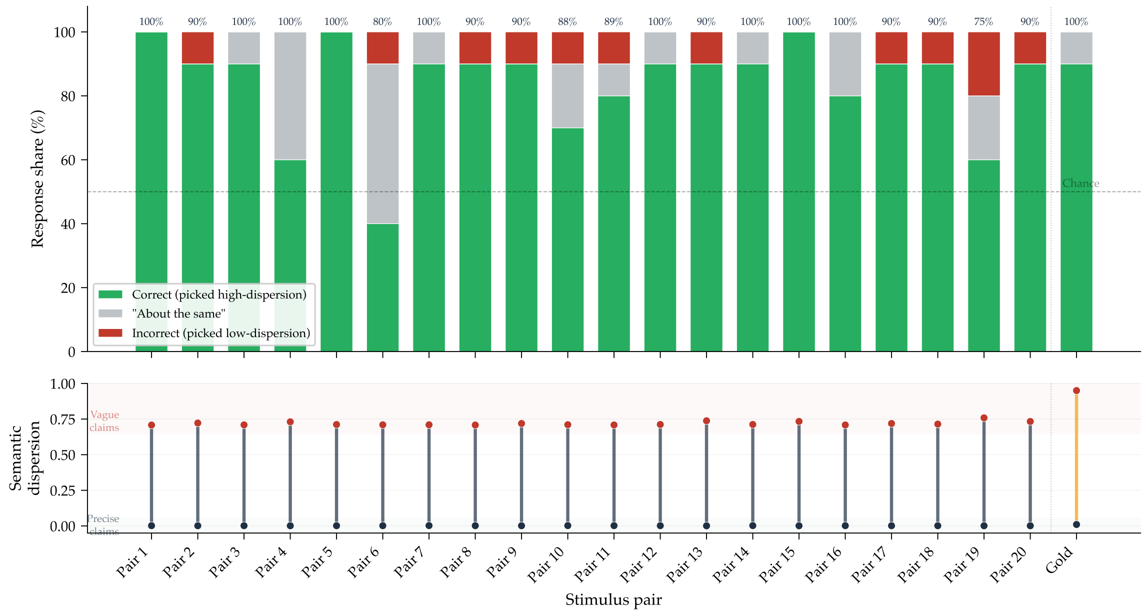


**Figure 4.** Pairwise forced-choice validation ($N = 14$ raters, 20 real pairs + 1 gold standard). Top panel: share selecting the correct (high-dispersion) claim. Dashed line = chance (50%). Bottom panel: dispersion gap between paired claims. All 21 pairs achieve majority-correct classification.

[11] Absolute rating scales conflate *referential specificity* (does this contain numbers?) with *interpretive latitude*; pairwise comparison surfaces interpretive breadth alone.

Raters identified the higher-dispersion claim in 92.3% of 261 non-neutral decisions ($p < 10^{-6}$, binomial test against 50%; the 33 "About the same" responses, 11.2% of 294, are excluded), with all 21 pairs achieving majority-correct classification ($p < 10^{-6}$, sign test). Pair-level accuracy ranged from 75% to 100%; the gold-standard pair achieved 100%. The "About the same" response rate (11.2%) was concentrated in pairs where both claims were short factual statements, suggesting raters default to the neutral option when interpretive breadth is genuinely similar rather than choosing incorrectly. Pairs were drawn from the distribution tails to maximise discrimination power given the sample size; near-centre validation, where the expected per-pair effect size is smaller, would require a substantively larger rater pool to achieve adequate statistical power. The measure therefore tracks human judgement on interpretive breadth.

### 4.4 Electoral and Redistricting Data

Constituency-level results are from Data.gov.sg (95.8% match rate). Redistricting data are hand-coded from EBRC reports (2011–2025), covering all 36 cross-constituency polling district transfers and all 166 MP-year redistricting events in the sample window. Each transfer is classified by the PAP vote-share margin in the source constituency at the preceding election, $V_d$, which proxies the MP's rational expectation of the transferred voters' prior (Section 5.2 describes the classification procedure and the 20-pp margin cut-off).

### 4.5 Descriptive Balance

Table 3 compares elected MPs and NMPs at the MP level. The groups differ substantially on parliamentary activity volume (elected MPs produce roughly four times as many claims), but the outcome variable is well-balanced: mean semantic dispersion is 0.477 for elected MPs and 0.471 for NMPs ($p = 0.440$). Volume differences are absorbed by MP fixed effects in all specifications. Hedge frequency and semantic dispersion are uncorrelated ($r \approx 0$), so MPs who speak more are not thereby vaguer.

**Table 3.** Balance Table: Elected MPs vs. Nominated MPs

| | Elected ($N$=243) | NMP ($N$=28) | Diff. | $p$ |
|---|---|---|---|---|
| | Mean (SD) | Mean (SD) | | |
| *Parliamentary activity* | | | | |
| Claims per MP | 197.3 (246.2) | 56.6 (34.0) | 140.7 | <0.001 |
| Debates participated | 61.3 (91.9) | 18.2 (11.4) | 43.0 | <0.001 |
| Speeches per MP | 94.1 (130.7) | 22.4 (13.8) | 71.8 | <0.001 |
| *Claim characteristics* | | | | |
| Claim length (words) | 17.7 (2.1) | 18.8 (2.0) | −1.0 | 0.012 |
| Hedge density | 0.086 (0.013) | 0.081 (0.011) | 0.005 | 0.039 |
| Hedges per claim | 1.20 (0.13) | 1.22 (0.07) | −0.01 | 0.385 |
| Frac. quantitative | 0.147 (0.094) | 0.112 (0.062) | 0.035 | 0.011 |
| *Outcome variable* | | | | |
| Semantic dispersion | 0.477 (0.032) | 0.471 (0.035) | 0.006 | 0.440 |
| *Topic & role* | | | | |
| Topic concentration (HHI) | 0.148 (0.204) | 0.116 (0.066) | 0.032 | 0.081 |
| Frac. minister | 0.393 | 0.000 | — | — |

*Notes:* Unit of observation is the MP; each cell reports mean (SD) of MP-level averages. $p$-values from Welch's $t$-tests. "Frac. minister" has no $p$-value because NMPs are ineligible by definition.

# 5 Empirical Analysis

The empirical work proceeds in two stages. A pre-election difference-in-differences exploiting Singapore's NMP institution tests the level prediction (Proposition 2). A redistricting natural experiment then tests the sign-flip property (Proposition 1) via opposing sign-responses from a single institutional shock.

## 5.1 Pre-Election Difference-in-Differences

Proposition 2 implies that elected MPs should become vaguer than NMPs as elections approach (when $\lambda$ rises); NMPs, with $\lambda = 0$ throughout, should not.[12] I estimate:

$$\text{Dispersion}_{it} = \beta_0 + \beta_1(\text{Elected}_i \times \text{PreElection}_t) + \alpha_i + \delta_t + \mathbf{X}_{it}'\gamma + \varepsilon_{it}, \tag{5}$$

where $\alpha_i$ are MP fixed effects, $\delta_t$ year fixed effects, and $\mathbf{X}_{it}$ topic controls and session indicators. Standard errors are clustered at the MP level (267 clusters; 12 of the 279 claim-bearing MPs are excluded for insufficient panel length). The coefficient $\beta_1$ captures the differential change in dispersion for elected MPs relative to NMPs in the 12 months before an election.[13]

The corner-solution prediction ($k^* = 0$ for $\lambda = 0$) is directional. Natural language imposes a floor on interpretive breadth, so the empirical test is whether NMPs sit closer to that floor than elected MPs. Three reasons NMP dispersion is not numerically zero: NMPs face a small effective $\lambda$ from possible reappointment to public panels; their topic composition skews technical, mechanically inflating dispersion; and MP fixed effects absorb baseline differences in the DiD. The claim-level variance ordering matches the model: elected-MP SD is 0.179 against 0.169 for NMPs (Table 1), with Levene's test rejecting variance equality at $W = 5.08$, $p = 0.024$.

The identifying assumption is parallel trends in dispersion absent electoral incentives. NMPs share the elected MPs' information environment (same debates, policy shocks, media coverage) while removing the channel the model identifies: $\lambda$. Any residual reappointment-driven incentive among NMPs attenuates the DiD rather than inflating it, so the comparison is conservative. The outcome is balanced at baseline (Table 3, $p = 0.440$), consistent with parallel trends but not dispositive; I test directly below. A positive $\hat{\beta}_1$ is a direct test of Proposition 2.

Table 4 presents the results. Across progressive specifications, elected MPs increase semantic dispersion by 0.024–0.026 relative to NMPs before elections. The preferred specification (column 2, with topic controls) yields $\hat{\beta}_1 = 0.026$ ($p = 0.041$). Standard event-studies are infeasible: synchronised elections make event-time bins collinear with calendar time and the NMP group is too small for bin-level estimation. The magnitude is 15% of a standard deviation, in line

[12] Parliamentary speech carries electoral relevance through two channels: the PAP's internal candidate-selection process evaluates incumbents' parliamentary contributions, and Meet-the-People Sessions plus local media coverage link chamber performance to constituency-level visibility.

[13] Singapore's synchronised elections (all seats contested on the same date) eliminate the bias that arises when different units receive treatment at different times (Goodman-Bacon, 2021).

with effect sizes reported in the slanted-language literature on political text. The result is direct evidence that electoral incentives increase legislative vagueness; given the pre-trend diagnostics below, I treat it as suggestive and rely on the redistricting design in Section 5.2 for the primary causal identification.

**Table 4.** Difference-in-Differences: Elected × Pre-Election on Semantic Dispersion

| | (1) | (2) | (3) |
|---|---|---|---|
| Elected × PreElection | 0.024** | 0.026** | 0.025** |
| | (0.011) | (0.013) | (0.012) |
| MP FE | Yes | Yes | Yes |
| Year FE | Yes | Yes | Yes |
| Topic Controls | No | Yes | Yes |
| Session Indicators | No | No | Yes |
| Observations | 62,169 | 62,169 | 62,169 |
| $R^2$ | 0.087 | 0.104 | 0.106 |
| Clusters | 267 | 267 | 267 |

*Notes:* PreElection = 1 if claim is within 12 months before a general election. MP-clustered SEs in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

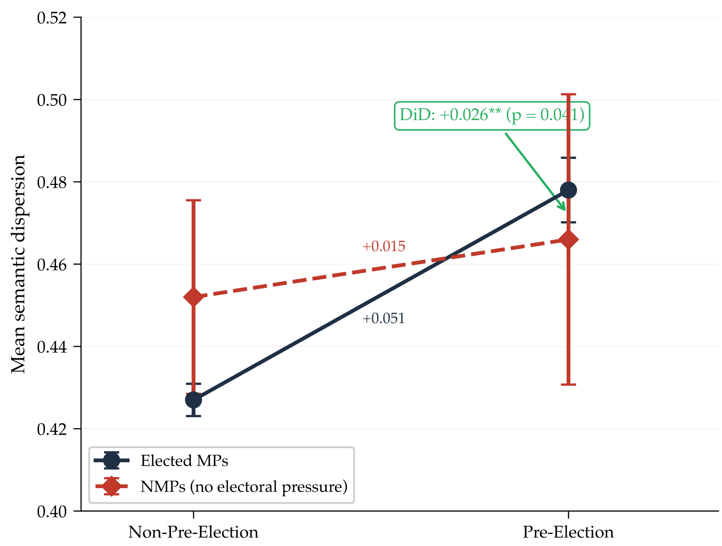


**Figure 5.** Pre-election difference-in-differences. Elected MPs increase semantic dispersion more than NMPs in the 12 months before elections. Raw difference of differences (0.036) exceeds the regression coefficient ($\hat{\beta}_1 = 0.026$, Table 4) because the regression absorbs MP and year fixed effects and topic controls. Error bars show 95% confidence intervals with MP-clustered standard errors.

The dynamic interpretation in Section 3.5 predicts that MPs with more time in the investment phase should have accumulated higher $\theta_0$ and exhibit larger pre-election swings. Senior MPs (4+ terms) show a pre-election DiD of +0.041 ($p = 0.024$), versus smaller and insignificant effects for junior MPs. Seniority affects $\theta_0$ through two channels: cumulative speech (the mechanism of interest) and differential survival across prior elections (a selection artefact). Both push in the predicted direction, so the data cannot separate them.

**Content controls.** A natural concern about the DiD is that elected MPs may have shifted toward inherently vaguer topics before elections rather than hedging more ambiguously on the same topics. Content controls apply only to the DiD because the redistricting design's within-MP variation around a single boundary-change shock makes topic drift implausible as a confounder there, and the theoretical prediction of opposing signs from the same shock already constrains non-mechanistic alternatives. Following standard text-as-data practice (Gentzkow et al., 2019a), I test this by including the first $K$ principal components of each claim's sentence embedding (`all-mpnet-base-v2`) as controls. The principal components capture the most important dimensions of variation in claim content, so the specification conditions on *what* the claim says and isolates variation in *how precisely* it says it. Adding 50 content components shifts the coefficient from 0.026 to 0.024 ($p = 0.068$), within standard error of the original estimate. The point estimate is stable across $K = 5, 10, 20, 50$.[14]

A stronger test interacts the content components with the treatment indicator. If elected MPs simply migrated toward different topics, the content–dispersion mapping would remain stable; only the distribution of claims across content space would shift. Conditional on discussing the same topic, elected MPs generate systematically more dispersed interpretations pre-election ($F_{20} = 2.08$, $p = 0.005$). The effect concentrates along specific semantic dimensions rather than operating as blanket hedging, consistent with the model's prediction that the precision parameter $k$ responds to electoral incentives on the intensive margin.

Decomposing by topic reveals where the shield is most valuable. Budget and fiscal debates show the largest pre-election vagueness increases (+0.072 to +0.096); defence and security debates show significant decreases ($-0.12$). Performance is hardest to verify on budgets and easiest to verify on defence outcomes: the shield pays off where voters have least independent information.[15]

**Pre-trend diagnostics.** A formal pre-trend test reveals a significant differential trend: $\beta_3 = 0.0044$ per year, equivalent to 2.5% of a standard deviation per year ($p = 0.03$).[16] The parallel trends assumption is not fully satisfied. Linear extrapolation of the pre-trend implies the DiD understates the true effect: under the dynamic mechanism, MPs adjust gradually rather than discretely at the 12-month cut-off, so the placebo period itself reflects mechanism leakage and the estimated pre-election shift functions as a conservative lower bound. The redistricting design

[14] Claim and interpretation embeddings share the MPNet encoder but operate on distinct text. Controlling for claim-PCs is conservative: any embedding-linked dispersion is absorbed, biasing the surviving coefficient toward zero.

[15] Group Representation Constituency (GRC) MPs show larger effects than Single Member Constituency MPs ($p < 0.001$ vs $p = 0.261$), consistent with multi-member teams diffusing individual accountability and making the vagueness shield more attractive.

[16] $\beta_3$ is the Elected × linear-time-trend interaction in a pre-treatment placebo regression.

carries the primary causal load. The pre-trend is concentrated in the 9th Parliament (2006–2011) and absent in subsequent terms ($p > 0.58$); with topic fixed effects it disappears entirely ($p = 0.83$), suggesting it reflects topic composition shifts over time rather than within-topic behavioural divergence. Under MP $\times$ Year fixed effects, which absorb all MP-specific annual trends, the pre-election coefficient strengthens to $+0.079$ ($p = 0.019$). A placebo test at false election dates produces a null ($-0.008$, $p = 0.63$).[17]

## 5.2 Redistricting Natural Experiment: Primary Causal Identification

A redistricting natural experiment supplies cleaner identification and directly tests the sign-flip property: the EBRC's quasi-exogenous boundary changes (Section 2) shift each MP's effective prior in a known direction, producing the opposing-sign prediction the model derives in Section 3.[18][19] Affected MPs split into two groups: *reduced-prior* (constituency lost reliably favourable voters) and *mixed-prior* (constituency absorbed new voters with different priors).

A uniform increase in vagueness as elections approach is consistent with the shield story but also with broader explanations: heightened scrutiny, strategic withholding, or generic risk-aversion. The redistricting design discriminates between these because the model predicts opposing-signed responses depending on whether prior voters are favourable or unfavourable; alternatives predict uniform responses.

I focus on the 2020 EBRC cycle: in 2011 and 2015 the report was released 72 and 49 days before the election, inside the $\pm 180$-day window, so the pre-treatment baseline overlaps with the pre-election period and the two shocks cannot be separated. The 2020 report was released 119 days before polling, giving a clean pre-EBRC baseline.

[17] Rambachan and Roth (2023) propose bounds on the treatment effect under relaxations of parallel trends; the redistricting design is positioned as primary identification precisely because such bounds on the DiD would plausibly cover zero.

[18] Identification depends only on the direction of each MP's prior shift, signed by 2015 source-constituency vote shares; these are pre-determined and exogenous to post-shock outcomes regardless of how the EBRC selected which constituencies to redraw.

[19] Parliament was dissolved on 23 June 2020, just over three months after the EBRC report, and elections were held on 10 July. The compressed timeline means the post-treatment window is truncated to roughly 100 rather than 180 days, biasing against finding an effect.

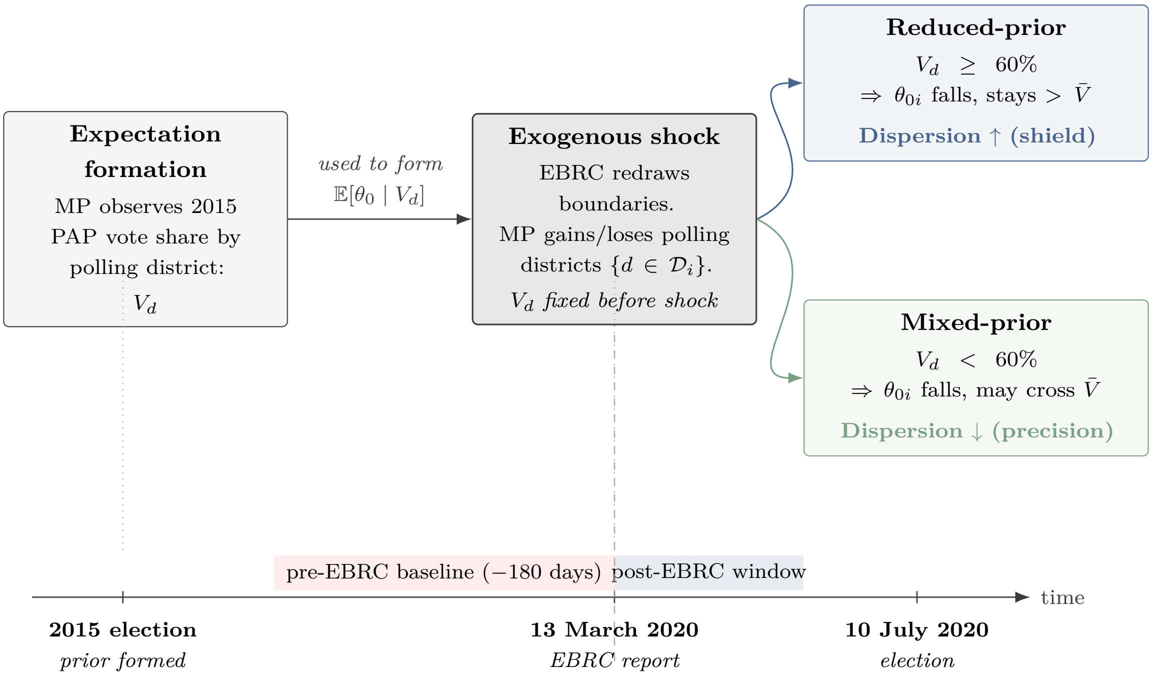


**Figure 6.** Redistricting design and identification logic. The 2015 election forms the prior $\theta_{0i}$ that the EBRC's 13 March 2020 boundary changes then perturb; pre-shock vote shares $V_d$ sign the resulting prior shift exogenously. Opposing predicted responses (shield for reduced-prior, lottery ticket for mixed-prior) cannot arise from any uniform confounder.

**Treatment construction.** The 2020 EBRC produced 14 cross-constituency polling-district transfers across 10 pairs of constituencies (a 15th polling-district change was a within-constituency boundary adjustment). For each transfer I record the PAP vote-share margin in the source constituency at the 2015 general election. The transfers divide cleanly into three bands: safe (margin $\geq$ 40 pp, 4 transfers), lean (20–40 pp, 8 transfers), and competitive ($<$ 20 pp, 2 transfers). A transfer is classified *high-support* if its source margin was at least 20 pp (PAP vote share at or above 60%) and *low-support* otherwise. The 20-pp cut-off corresponds to the natural discontinuity in the 2020 transfer data: the highest competitive margin is 15.0 pp and the lowest lean margin is 24.3 pp, leaving a 9-point gap with no transfers in it, so any cut-off within this range yields the same classification. The absolute threshold approximates the model's MP-specific threshold ($\mathbb{E}[\theta_0 \mid V_d]$ vs. $\theta_{0i}$) because PAP incumbents' baseline priors cluster narrowly. An MP is classified *reduced-prior* if their constituency ceded a high-support polling district, *mixed-prior* if their constituency absorbed a low-support polling district. No MP falls into both groups in 2020. The estimation sample contains 28 reduced-prior MPs and 13 mixed-prior MPs with claims in the ±180-day window, against 84 unaffected elected MPs as the control group. Appendix A lists all 14 cross-constituency transfers and their classifications.

**Specification.** The theory in Section 3 predicts that the two treatments produce opposite-signed responses: vagueness rises for reduced-prior MPs (shield) and falls for mixed-prior MPs (corner solution, Proposition 3), where unfavourable voters pull $\theta_0$ toward $\bar{V}$ and potentially across it, activating the lottery-ticket regime. I test this directly with a joint specification that estimates both interactions simultaneously:

$$\text{Dispersion}_{it} = \alpha_i + \delta_t + \beta_1(\text{Reduced}_i \times \text{Post}_t) + \beta_2(\text{Mixed}_i \times \text{Post}_t) + \mathbf{X}'_{it}\gamma + \varepsilon_{it}, \quad (6)$$

where $\text{Reduced}_i = 1$ for the 28 reduced-prior MPs, $\text{Mixed}_i = 1$ for the 13 mixed-prior MPs, $\text{Post}_t = 1$ for claim dates on or after 13 March 2020, $\alpha_i$ are MP fixed effects, and $\delta_t$ are year fixed effects. Standard errors are HC3 robust rather than MP-clustered.[20] Unaffected elected MPs are the baseline. The sign-flip hypothesis implies $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_1 \neq \beta_2$; the third is testable as a formal sign-flip statistic. If $\beta_1 > 0$ but $\beta_2 = 0$, the shield mechanism survives but the lottery-ticket arm does not: consistent with the model in the safe regime but not in the corner-solution regime. Joint rejection of $\beta_1 = \beta_2$ is the strongest evidence the design can deliver.

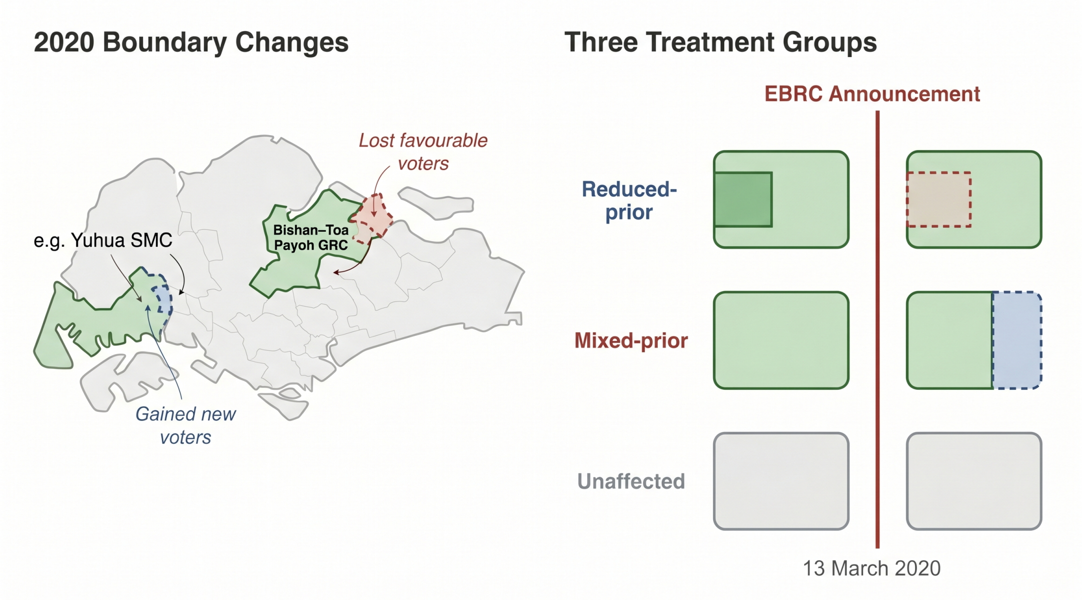


**Figure 7.** Redistricting as a natural experiment. The 2020 EBRC created two treatment groups: *reduced-prior* MPs lost high-support polling districts (e.g. Bishan–Toa Payoh); *mixed-prior* MPs absorbed new voters (e.g. Yuhua). Unaffected MPs are the control. Joint-specification results in Table 5.

**Results.** Table 5 reports the joint estimates. $\hat{\beta}_1 = +0.026$ ($p = 0.009$, HC3) and $\hat{\beta}_2 = -0.023$ ($p = 0.17$). The formal test of $H_0 : \beta_1 = \beta_2$ rejects at $F = 7.89$, $p = 0.005$: direct statistical support for the sign-flip property derived in Section 3. The F-test asks whether the two coefficients differ in the predicted direction, not whether each is individually significant. The reduced-prior coefficient is individually significant; the mixed-prior coefficient is signed as predicted but imprecisely estimated on only 13 treated MPs. Estimated on separate samples with only the relevant treatment group, both coefficients are larger and individually significant (Appendix A): $+0.029$ ($p = 0.003$) and $-0.032$ ($p = 0.05$).

[20] Cluster-robust inference relies on asymptotic approximations that degrade with small cluster counts; with 41 treated MPs, HC3 provides finite-sample correction without that threshold (Cameron and Miller, 2015).

**Table 5.** Redistricting: Joint Sign-Flip Specification

| | (1) Joint specification |
|---|---|
| Reduced-prior × Post ($\beta_1$) | +0.026*** |
| | (0.010) |
| Mixed-prior × Post ($\beta_2$) | −0.023 |
| | (0.017) |
| *Sign-flip test* | |
| $H_0 : \beta_1 = \beta_2$ | $F = 7.89$ |
| | $p = 0.005$ |
| Model prediction | $\beta_1 > 0$, $\beta_2 < 0$ |
| MP FE / Year FE | Yes |
| Observations | 3,354 |
| Treated MPs (reduced / mixed) | 28 / 13 |
| Control MPs (unaffected) | 84 |

*Notes:* 2020 EBRC cycle, ±180-day window around 13 March 2020. Dependent variable is claim-level semantic dispersion. Robust (HC3) standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Any uniform confounder of redistricted MPs (stress, workload, parliamentary procedure, generic election-mode behaviour, COVID-19 disruption to the post-EBRC window) predicts the same directional shift and cannot produce $\beta_1 > 0$ and $\beta_2 < 0$ simultaneously; any alternative must operate in opposite directions across the prior threshold. The most plausible directional alternative is differential media scrutiny, which two features argue against: Singapore's EBRC changes are reported as a national event rather than constituency-level stories, and a joint specification including both redistricting indicators and the pre-election DiD confirms the two sources are separable (reduced-prior coefficient +0.027, $p = 0.007$ conditional on the pre-election indicator).

The reduced-prior coefficient is stable across estimation windows: +0.026 at ±90 days ($p = 0.023$), +0.024 at ±150 days ($p = 0.029$), and +0.029 at ±180 days. Leave-one-MP-out analysis confirms no single MP drives the result, with the coefficient ranging from +0.013 to +0.034 across 28 jackknife iterations. The mixed-prior coefficient is consistently negative across bandwidths but reaches conventional significance only at ±180 days, reflecting the smaller treatment group of 13 MPs.

### 5.3 Testing the Pre-Commitment Assumption

The model requires pre-committed vagueness: the MP chooses a communication *style* (a level of precision $k$) before observing the policy outcome $\theta$. If MPs were only vague when hiding bad outcomes, voters would learn to treat vagueness itself as a signal of poor performance, and the strategy would unravel via standard disclosure logic (see Section 3; Milgrom, 1981).

Pre-commitment therefore generates a testable prediction: vagueness should be persistent within MPs across speeches, shifting only when the strategic environment changes (e.g. an approaching election), not varying claim-by-claim with the content of individual policy outcomes.

I test this with an AR(1) specification that regresses each MP's mean monthly dispersion on its one-month lag, with MP fixed effects (Appendix A, Equation 7). If vagueness were a speech-by-speech reaction to bad news, successive months would show little autocorrelation; if it reflects a stable communication style, the lag coefficient should be substantially positive.

The AR(1) coefficient is $\hat{\rho} = 0.42$ ($p < 0.001$, $N = 4{,}264$ MP-month pairs), confirming that vagueness is a persistent individual trait. Adding a second lag (AR(2)) attenuates $\hat{\rho}_1$ only modestly to 0.36, with a significant second lag of 0.27, indicating multi-month memory rather than one-period noise (Table 7). Pre-election periods produce a level shift ($\hat{\Delta} = +0.032$, $p < 0.001$) on top of this persistence, consistent with MPs adjusting their communication style across the electoral cycle while retaining it within a cycle.

Decomposing persistence by electoral proximity confirms the dynamic logic: more than 24 months from an election, within-MP autocorrelation is high ($\hat{\rho} = 0.56$, $N = 2{,}420$); at 18–24 months it collapses and remains negative through election day, consistent with MPs beginning the shift well before the campaign becomes publicly salient.[21]

[21] Full AR(1), AR(2), and quarterly specifications are reported in Appendix A.

# 6 Robustness

I address three concerns: whether the effect is election-specific, whether the dispersion measure captures the construct, and whether the pre-election shift moves the precision of hedges or their frequency. Table 6 summarises the tests.

**Table 6.** Robustness Summary

| Test | Coeff. | $p$ | Interpretation |
|---|---|---|---|
| Placebo dates | −0.008 | 0.63 | Null at fake elections |
| Window: 6 months | +0.019 | 0.09 | Stable |
| Window: 18 months | +0.038 | 0.02 | Stable |
| BH-adjusted (primary) | — | 0.12 | Above 5% threshold |
| BH-adjusted (redistrict.) | — | 0.007 | Survives 5% FDR |
| Numerical dispersion† | +0.008 | 0.01 | Cross-validates |
| Leave-one-out (cycle) | 0.019–0.038 | — | No single cycle drives |
| Hedge count | +0.003 | 0.71 | Not hedging *more* |
| Hedge density | −0.001 | 0.89 | Hedging more *ambiguously* |

*Notes:* All tests use baseline specification (Table 4, col. 2). BH = Benjamini–Hochberg correction. †Alternative measure constructed from extracted numerical values (see §6).

## 6.1 Election specificity

If the pre-election effect reflects strategic responses to electoral incentives, it should appear at real elections and nowhere else. A placebo test at false election dates produces a null (−0.008, $p = 0.63$). The coefficient is positive and significant across estimation windows (0.019, 0.026, 0.038 at 6, 12, and 18 months) and not driven by any single electoral cycle (0.019 to 0.038 across leave-one-out iterations). Benjamini–Hochberg correction across the nine tests raises the primary DiD's $p$-value to 0.12, above the 5% threshold; the redistricting result survives at a 5% false-discovery rate ($p = 0.007$).

## 6.2 Measurement validity

Semantic dispersion measures vagueness through embedding-space distances; an independent measure should produce concordant results if the construct is valid. Numerical dispersion (an alternative measure constructed from extracted numerical values) produces concordant results (+0.008, $p = 0.01$). Two measurement pipelines sharing no infrastructure beyond the original interpretations agree. The result is not an embedding-space artefact.

## 6.3 Intensive vs. extensive margin

The model predicts adjustment on the intensive margin (precision of each hedge) rather than the extensive margin (frequency of hedging); the two are empirically separable. Hedge count is null (+0.003, $p = 0.71$): MPs do not hedge more often before elections. Hedge density (hedges per word) is also null (−0.001, $p = 0.89$): they do not hedge more densely. The pre-election shift operates on *how* ambiguously MPs hedge, not on *how often*. This confirms the model's specific prediction about which margin the precision parameter $k$ operates on.

# 7 Conclusion

This paper asks whether electoral incentives cause politicians to speak more ambiguously. The answer is yes, when the politician's prior reputation is favourable; when the prior is unfavourable, electoral pressure produces the opposite effect. A signal model predicts this sign-flip: the voter rationally discounts vague messages and falls back on her prior, so vagueness helps the politician when the prior is favourable and hurts her when it is not. A redistricting natural experiment delivers the primary causal evidence: opposing treatment effects from a single shock ($+0.026$ for reduced-prior, $-0.023$ for mixed-prior; sign-flip test $F = 7.89$, $p = 0.005$). Vagueness is a sustained style, not a claim-by-claim reaction to bad news: within-MP persistence is high, with level shifts at electoral transitions. Strategic vagueness is a hidden cost of electoral accountability.

The mechanism does not require Singapore's dominant-party context. Any system in which politicians carry reputations voters can anchor on (Westminster safe seats, US congressional incumbency, German Land-level dominant parties) should exhibit the sign-flip. The mechanism does require sophisticated voters: where receivers fail to discount vagueness rationally, Thaler-style dynamics dominate and vagueness uniformly helps the sender. The effect concentrates among senior MPs, who have spent longer accumulating the reputation that the shield protects. This identifies a channel the term-limits debate has not considered: enforced turnover would prevent the accumulation of high $\theta_0$ that makes shielding effective, eroding the equilibrium in the long run. Existing arguments for term limits centre on responsiveness and incumbent experience; the present results identify a distinct information-quality channel.

A multi-period extension with endogenous prior updating would turn the investing-and-harvesting intuition into quantitative predictions about optimal switching. Cross-country replication would test the mechanism's external generality. Institutional responses (structured debate formats requiring falsifiable predictions; transparency tools flagging hedged language) raise the cost of vagueness directly: the shield works because $C(k)$ is private; making it public shifts the optimum.

As Maskin and Tirole (2004) observe, the distinction between accountable politicians and independent experts is fundamental to institutional design. Canes-Wrone et al. (2001) show that re-election pressure distorts how politicians use private information in policymaking; the present findings suggest that this distortion extends to the language through which representatives communicate. Like the policies they enact, the words politicians choose are chosen to manage what voters can verify.

# A Additional Regression Specifications

### A.1 Persistence.

$$\overline{D}_{i,m} = \alpha_i + \sum_{j=1}^{p} \rho_j \, \overline{D}_{i,m-j} + \delta \, \text{PreElection}_{i,m} + \varepsilon_{i,m}, \tag{7}$$

where $\overline{D}_{i,m}$ is the mean semantic dispersion for MP $i$ in month $m$ and $\alpha_i$ are MP fixed effects. Table 7 reports estimates across lag structures.

**Table 7.** Persistence Across Lag Structures

| | (1) AR(1) | (2) AR(2) | (3) Quarterly |
|---|---|---|---|
| $\hat{\rho}_1$ (1-month lag) | 0.41*** | 0.36*** | — |
| | (0.02) | (0.02) | |
| $\hat{\rho}_2$ (2-month lag) | — | 0.27*** | — |
| | | (0.02) | |
| $\hat{\rho}_1$ (1-quarter lag) | — | — | 0.32*** |
| | | | (0.01) |
| PreElection | +0.025*** | +0.018*** | +0.019*** |
| | (0.003) | (0.004) | (0.003) |
| MP FE | Yes | Yes | Yes |
| Observations | 4,264 | 2,224 | 3,687 |
| Clusters | 211 | 182 | 219 |

*Notes:* Unit is MP-month (cols. 1–2) or MP-quarter (col. 3). All specifications use within-MP demeaning and MP-clustered standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

### A.2 Content controls (PCA interaction).

$$\begin{aligned} D_{it} = \beta_1(\text{Elected}_i \times \text{PreElection}_t) + \sum_{k=1}^{K} \gamma_k \, \text{PC}_{k,it} \\ + \sum_{k=1}^{K} \phi_k \, (\text{PC}_{k,it} \times \text{Elected}_i \times \text{PreElection}_t) + \alpha_i + \delta_t + \varepsilon_{it}. \end{aligned} \tag{8}$$

The $F$-test on the interaction terms ($H_0 : \phi_1 = \cdots = \phi_K = 0$) tests whether the content–dispersion mapping shifts pre-election for elected MPs. Result at $K = 20$: $F_{20} = 2.08$ ($p = 0.005$).

**A.3 Heterogeneity subgroup regressions.** Each column estimates the baseline DiD specification (Equation 5) on the indicated subgroup or with the indicated interaction.

**Table 8.** Heterogeneity in Pre-Election Vagueness Effect

| | (1) Senior (4+ terms) | (2) Budget debates |
|---|---|---|
| Elected × PreElection | +0.041** | +0.096** |
| | (0.018) | (0.038) |
| MP FE / Year FE / Topic | Yes | Yes |
| Observations | 18,402 | 8,741 |
| Clusters | 62 | 198 |

*Notes:* Col. (1) restricts to MPs serving 4+ terms. Col. (2) restricts to claims in budget and fiscal debates. MP-clustered SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

**A.4 Redistricting: separate-sample estimates.** Estimated on separate samples with only the relevant treatment group alongside unaffected MPs, individual coefficients are larger than those from the joint specification (which absorbs the other treatment's effect in its own interaction term): reduced-prior +0.029 ($p = 0.003$, HC3) and mixed-prior −0.032 ($p = 0.05$, HC3). Both signs match the model's predictions; the joint specification in Table 5 is preferred because its control group excludes the other treatment group's post-EBRC response.

**A.5 2020 EBRC transfers.** Table 9 lists all 14 cross-constituency polling-district transfers in the 2020 EBRC cycle, with the source constituency's 2015 PAP margin and the resulting classification.

**Table 9.** 2020 EBRC Polling District Transfers

| Source constituency | Destination constituency | 2015 PAP margin (pp) | Classification |
|---|---|---|---|
| Punggol East | Yuhua | 3.5 | Low-support |
| Fengshan | East Coast | 15.0 | Low-support |
| Sengkang West | Yuhua | 24.3 | High-support |
| Sengkang West | Ang Mo Kio | 24.3 | High-support |
| Marine Parade | Potong Pasir | 28.1 | High-support |
| Jalan Besar | Yuhua | 34.6 | High-support |
| Chua Chu Kang | Yuhua | 35.5 | High-support |
| Chua Chu Kang | West Coast | 35.5 | High-support |
| Marsiling-Yew Tee | Yuhua | 37.4 | High-support |
| Hong Kah North | West Coast | 38.1 | High-support |
| Sembawang | Nee Soon | 44.6 | High-support |
| Pasir Ris-Punggol | Yuhua | 45.8 | High-support |
| Bishan-Toa Payoh | Yuhua | 47.2 | High-support |
| Bishan-Toa Payoh | Yuhua | 47.2 | High-support |

*Notes:* All 14 cross-constituency polling-district transfers in the 2020 EBRC (the 15th is a boundary adjustment with no cross-constituency transfer). The two Bishan–Toa Payoh rows reflect distinct polling districts within the constituency. Margin = PAP vote share minus next-placed party's share at the 2015 election. Classification: margin $\geq$ 20 pp is high-support.

# References


Alesina, A. and Cukierman, A. (1990). The politics of ambiguity. *Quarterly Journal of Economics*, 105(4):829–850.

Aragonès, E. and Neeman, Z. (2000). Strategic ambiguity in electoral competition. *Journal of Theoretical Politics*, 12(2):183–204.

Asirvatham, H., Mokski, E., and Shleifer, A. (2026). GPT as a measurement tool. NBER Working Paper No. 34834.

Ash, E., Morelli, M., and Van Weelden, R. (2017). Elections and divisiveness: Theory and evidence. *Journal of Politics*, 79(4):1268–1285.

Ashworth, S. (2012). Electoral accountability: Recent theoretical and empirical work. *Annual Review of Political Science*, 15:183–201.

Besley, T. (2006). *Principled Agents? The Political Economy of Good Government*. Oxford University Press.

Callander, S. and Wilson, C. H. (2008). Context-dependent voting and political ambiguity. *Journal of Public Economics*, 92(3–4):565–581.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

Canes-Wrone, B., Herron, M. C., and Shotts, K. W. (2001). Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science*, 45(3):532–550.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Di Tella, R., Kotti, R., Le Pennec, C., and Pons, V. (2023). Keep your enemies closer: Strategic platform adjustments during U.S. and French elections. NBER Working Paper No. 31503.

Djourelova, M. (2023). Persuasion through slanted language: Evidence from the media coverage of immigration. *American Economic Review*, 113(3):800–835.

Downs, A. (1957). *An Economic Theory of Democracy*. Harper & Row.

Farkas, R., Vincze, V., Móra, G., Csirik, J., and Szarvas, G. (2010). The CoNLL-2010 shared task: Learning to detect hedges and their scope in natural language text. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning—Shared Task*, pages 1–12.

Gentzkow, M., Kelly, B., and Taddy, M. (2019a). Text as data. *Journal of Economic Literature*, 57(3):535–574.

Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019b). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340.

Glazer, A. (1990). The strategy of candidate ambiguity. *American Political Science Review*, 84(1):237–241.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies*, 66(1):169–182.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kartik, N. (2009). Strategic communication with lying costs. *Review of Economic Studies*, 76(4):1359–1395.

Maskin, E. and Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 94(4):1034–1054.

Mauzy, D. K. and Milne, R. S. (2002). *Singapore Politics under the People's Action Party*. Routledge.

Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12(2):380–391.

Morris, S. (2001). Political correctness. *Journal of Political Economy*, 109(2):231–265.

Ornstein, J. T., Blasingame, E. N., and Truscott, J. S. (2025). How to train your stochastic parrot: Large language models for political texts. *Political Science Research and Methods*, forthcoming:1–18.

Page, B. I. (1976). The theory of political ambiguity. *American Political Science Review*, 70(3):742–752.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591.

Shepsle, K. A. (1972). The strategy of ambiguity: Uncertainty and electoral competition. *American Political Science Review*, 66(2):555–568.

Thaler, M., Toma, M., and Wang, V. Y. (2024). Numbers tell, words sell. CESifo Working Paper No. 11600.

Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34(4):273–286.

Tomz, M. and Van Houweling, R. P. (2009). The electoral implications of candidate ambiguity. *American Political Science Review*, 103(1):83–98.

Vincze, V., Szarvas, G., Farkas, R., Móra, G., and Csirik, J. (2008). The BioScope corpus: Biomedical texts annotated for uncertainty, negation and their linguistic scope. *BMC Bioinformatics*, 9(Suppl 11):S9.

# Data Sources

All data sources are public; the constructed semantic-dispersion scores can be regenerated via the pipeline in Section 4.

- **Singapore Hansard.** Parliamentary debates, 2006–2025.
  `https://www.parliament.gov.sg/parliamentary-business/parliamentary-debates`
- **General-election results.** Constituency- and candidate-level results, 1955–2025. Data.gov.sg, dataset `d_581a30bee57fa7d8383d6bc94739ad00`.
  `https://data.gov.sg/datasets/d_581a30bee57fa7d8383d6bc94739ad00`
- **Electoral Boundaries Review Committee reports.** 2011, 2015, and 2020 cycles, Elections Department of Singapore.
  `https://www.eld.gov.sg/`
- **CoNLL-2010 hedge corpus.** Hedge-detection training data (Farkas et al., 2010).
  `https://rgai.inf.u-szeged.hu/conll2010st/`
- **BioScope corpus.** Hedge-detection training data (Vincze et al., 2008).
  `https://rgai.inf.u-szeged.hu/bioscope.html`
- **GPT-4o-mini.** OpenAI API; all calls use temperature 0, fixed seed 42, and strict JSON-schema output.
  `https://platform.openai.com/docs/models`
- **Sentence embeddings.** `all-mpnet-base-v2`, available on Hugging Face.
  `https://huggingface.co/sentence-transformers/all-mpnet-base-v2`