# Channels of Departure: Hand-Coded Exit Pathways and the Identification of Pre-Resignation Shirking in the Korean National Assembly

KNA Research Agents (AI-generated)
Experimental Output, kna-research-agents.com

April 20, 2026

**Abstract**

Do legislators who pursue higher office reduce their lawmaking effort in their final months? Existing empirical work on progressive ambition treats mid-term exit as a single category, yet legislators leave for many reasons, only some of which respond to the campaign-investment logic that motivates classical shirking arguments. I introduce a hand-coded exit-channel dictionary that separates voluntary local-executive runs from involuntary departures (court rulings, party dissolution) and selection-driven moves (cabinet appointment, presidential staff). Applied to four cycles of Korean National Assembly resignations (2010 through 2022), the dictionary recovers a clean local-executive cohort whose chief-sponsorship rate falls by roughly three-quarters in the final six months relative to a productivity-matched continuer pool, while involuntary court-ruling exits track the natural decline of the pool. The contrast identifies the ambition-investment mechanism while exposing the contamination that exit-conflated designs hide.

**Keywords:** progressive ambition, legislative shirking, exit channels, Korean National Assembly, difference-in-differences

# 1 Introduction

When a legislator leaves office before the end of a term, two questions arise. An administrative question concerns who replaces the seat, at what cost, and with what effect on the chamber's partisan balance. A behavioral question concerns whether the departing legislator continued to do the work of legislating during the months when departure was already on the horizon, and whether the prospect of a new role changed the way that work was performed. The behavioral question is the older of the two, going back at least to Schlesinger (1966), whose account of progressive ambition identifies higher-office prospects as a force that reshapes incentives once a new office becomes plausible. It also remains the more difficult question to answer empirically, because the legislators whose departures are observable form a self-selected group, and the pre-departure window is rarely defined by any external instrument.

The Korean National Assembly offers an unusually large and clean source of variation on this margin. Between the 18th Assembly (2008 through 2012) and the 21st Assembly (2020 through 2024), several hundred members resigned mid-term. Some ran for governor or mayor; some moved into cabinet positions; some were unseated by the Constitutional Court or by criminal conviction; others moved into presidential staff or left under party-organizational pressure. Each of these channels removes a member from the chamber, but only some of them carry the campaign-investment logic that motivates shirking. A legislator who is unseated by court ruling has no upcoming campaign to invest in. A legislator who runs for governor does. The administrative consequences of the two departures are similar, but the predicted behavioral signature is not.

Existing empirical work has tended to collapse these channels. Studies that proxy exit with the date of a member's last legislative act, or with the binary fact of non-continuation into the next assembly, treat all departures as equivalent because the underlying registry data does not distinguish among them. This approach has costs that are not always recognized. When involuntary exits are pooled with voluntary ones, the average pre-departure trajectory is a weighted average of two trajectories with different mechanisms; the resulting estimate may understate the shirking effect among voluntary exiters or, depending on composition, falsely suggest a shirking pattern that does not in fact apply to the channels of policy interest. The problem is structural: it cannot be solved by adding controls because the channels themselves are not observed.

The contribution of this paper is to introduce and validate a hand-coded exit-channel dictionary for the four Korean Assembly cycles in which mid-term departures are sufficiently common to support disaggregated analysis. The dictionary partitions resigning members into five mutually exclusive categories: local-executive runners, court-ruling exits (including the 2014 Unified Progressive Party dissolution), cabinet appointments, presidential staff moves, and a residual category. Coding rests on contemporaneous public-records anchors documented in a replication file. Applied to a difference-in-means specification on monthly chief-sponsorship rates in the twelve months before resignation, the dictionary recovers three findings that exit-pooled designs cannot.

First, in the two non-presidential local-election cycles where channel contamination is minimal, clean local-executive runners reduce chief-sponsorship substantially in the final six months relative to the continuer pool; the pattern survives randomization inference (see Table 3). Second, court-ruling exits, which involve no campaign to invest in, show no shirking signature; their pre-departure trajectory is statistically indistinguishable from the natural mild decline of the continuer pool (see Table 4). Third, the apparent magnitude of the local-executive shirking effect is sensitive to the choice of comparison group: against a top-quartile continuer pool whose pre-period rate substantially exceeds the treated mean, the effect attenuates and loses conventional significance, while against a level-matched comparison the effect is roughly half the naive estimate but remains substantively large and statistically detectable (see Table 5). The third finding tempers the substantive claim without overturning it.

A separate analysis of the five members unseated by the 2014 Constitutional Court dissolution of the Unified Progressive Party complicates the simple voluntary-versus-involuntary partition. Two of the five had effectively ceased chief sponsorship eight months before the formal dissolution date, suggesting that legal trouble degrades legislative engagement on a longer time horizon than the formal exit event captures. This sub-pattern reinforces rather than undermines the methodological point: even within an apparently homogeneous channel, sub-cohort timing matters, and naive last-bill-date filters will mechanically distort pre-exit activity measurements when applied to mixed exit categories.

These findings speak to three literatures. To the progressive-ambition literature pioneered by Schlesinger (1966) and elaborated by Black (1972) and Rohde (1979), they offer, to my knowledge, the first empirical test of pre-resignation shirking under an exit-channel-clean identification. To the comparative legislative-behavior literature anchored by Besley and Case (1995) on gubernatorial term limits and by Hansen and Treul (2015) on European mixed-member systems, they offer a Korean case in which the institutional setting (weak party-discipline constraints on exit, mixed-member system with a competitive subnational executive office) produces a shirking signature that contrasts with the European null. To the methodological literature on exit-conflated designs, they offer a transferable coding protocol and a placebo-based identification strategy modeled on Titiunik and Feher (2017)'s analysis of the Arkansas Senate term-length lottery.

The remainder of the paper proceeds as follows. Section 2 situates the contribution within the progressive-ambition tradition, the legislative-effectiveness literature, and recent methodological work on exit channels and pre-analysis-plan discipline. Section 3 describes the Korean National Assembly data infrastructure, the construction of the exit-channel coding dictionary, and the empirical specification. Section 4 presents the main comparison on the clean local-executive cohort, the court-ruling placebo, the regression-to-the-mean robustness exercises, and the five-channel exit decomposition. Section 5 discusses scope conditions, limitations, and the policy implications of channel-specific by-election remediation. Section 6 concludes.

## 2 Literature and Theory

The classical theoretical anchor for pre-resignation shirking is Schlesinger (1966)'s account of progressive ambition, in which legislators rationally allocate effort across present-office tasks and future-office investment when the marginal value of the latter rises. Subsequent elaborations by Black (1972) and Rohde (1979) extended the framework to U.S. House members' decisions to run for higher office, embedding the ambition calculus in formal models of opportunity-cost choice. The empirical workhorse for testing this allocation has been the gubernatorial term-limit design pioneered by Besley and Case (1995), who exploit constitutional variation in re-election eligibility to identify the effect of removed re-election incentives on policy choice. Their finding that lame-duck governors behave differently than first-term governors helped establish that anticipated exit reshapes in-office behavior. The design has since been extended to legislators in a variety of settings, with mixed results.

In European parliaments, Hansen and Treul (2015) test for shirking among MPs running for higher office and find limited evidence of effort decay, attributing the muted pattern to the disciplining effect of strong party organizations on individual exit decisions. Titiunik and Feher (2017) exploits a different source of exogenous variation, the Arkansas Senate term-length lottery that randomly removes the re-election incentive for a subset of senators, and reports null effects across five legislative-output measures, defending the null with equivalence-test machinery designed for small-sample inference. In the United States Congress context, Egerod (2021) documents a different exit channel altogether: the lure of the private sector raises selection-out probabilities for senators with characteristics that the lobbying market values, complicating any pooled estimate of pre-exit behavior.

A separate strand of work focuses on within-chamber legislative effectiveness rather than exit. Bucchianeri, Volden and Wiseman (2024) extend the long-running American legislative-effectiveness research program to the state legislatures, constructing a measure of how far a member's bills move through the lawmaking process. Their analysis treats effectiveness as a within-chamber outcome and does not code exit channels, leaving open the question of whether the most effective legislators are also the ones most likely to leave for higher office. Giommoni and Loumeau (2022) address an adjacent question in the Italian parliamentary context, hand-identifying transformist MPs (those who switch parliamentary groups) within a much larger pool of party-switchers and showing that the behavioral signature of transformism differs from the average switcher. Their methodological move (isolating a behaviorally meaningful sub-category from a category that registry data conflates) is close in spirit to the exit-channel coding I introduce here.

Korean-language scholarship on legislative behavior is closer to a conversation I join than a list to enumerate. Two contributions bear most directly on the present design. An, Park and Lee (2025) and Ka (2025) refine chief-sponsorship and passage-stage measures across the 20th and 21st Assemblies, making productivity the dependent variable but treating exit as the end of the observation window rather than a behavioral object. The present paper takes these productivity measures as inputs and asks what happens to them in the months preceding a mid-term departure, conditional on the channel of departure. A separate Korean constitutional-law literature has analyzed the December 2014 Unified Progressive Party dissolution as a normative-legal question; representative contributions include Park (2014) and Lee (2014). None of these papers measure the legislative-productivity consequences for the unseated members, which is the behavioral margin this paper addresses and which the constitutional-law frame does not engage.

No Korean or English-language paper has, to my knowledge, used the 2014 dissolution as an exogenous-exit identification strategy, and no Korean paper has tested for pre-resignation shirking among legislators running for subnational executive office. This gap exists despite a domestic policy debate on the fiscal cost of mid-term by-elections, a debate that has been informed by aggregate vacancy counts but not by an empirically defensible decomposition of the underlying exit reasons. The kind of analysis I undertake here, in which exit channels are hand-coded against contemporaneous public records and pre-resignation behavior is then tested for each channel separately, has no published precedent for the Korean case.

The theoretical expectation that motivates the empirical strategy is straightforward. If pre-resignation shirking reflects rational reallocation of effort toward an upcoming role, then the shirking signature should be present in voluntary exit channels with a known future role to invest in (local-executive runs, cabinet appointments, presidential staff moves) and absent in involuntary channels (court rulings, party dissolution). The simplest version of the theory makes a sign prediction at the channel level: voluntary channels produce negative ramps in chief-sponsorship over the final months; involuntary channels produce ramps statistically indistinguishable from the continuer pool. This sign prediction is the basis for the placebo-based identification strategy I employ.

Two scope conditions on this prediction are worth flagging in advance. First, the prediction concerns the marginal allocation of effort, not the level. Legislators selected into voluntary exit channels may differ in baseline productivity from continuers; the test is whether their effort declines relative to their own pre-period level by more than the natural decline of comparable continuers. The regression-to-the-mean concern that this comparison invites is the central robustness exercise of the paper. Second, the prediction concerns chief sponsorship, the press-release-generating legislative act, rather than co-sponsorship, which carries weaker credit-claiming value and may be a noisy outcome. Mayhew (1974) provides the canonical theoretical justification for this distinction, which the empirical analysis follows.

A further implication of the theory concerns the cycle-level structure of the data. If shirking is a reallocation toward an upcoming campaign, it should be muted in cycles where a higher-salience election has already absorbed the campaign-investment effort earlier. The 21st Assembly local-election cycle (2022) is a candidate case: the March 2022 presidential election preceded the June 2022 local elections by three months, plausibly leaving the local-election candidates with little incremental campaign investment to make in the final months of their National Assembly terms. This cycle-level scope condition is testable in principle, and I flag now that it is used in Section 4 to justify excluding cycle 21 from the primary specification. That exclusion is ex post: the hypothesis was generated after observing the cycle-21 null, and Section 4 reports the cycle-inclusive specification as a robustness check rather than as an independent confirmation. Pre-registration of the presidential-absorption hypothesis for the 22nd Assembly replication is the natural path for an out-of-sample test.

A final theoretical point concerns the comparative framing. The Korean institutional setting differs from the European cases studied by Hansen and Treul (2015) in two ways relevant to the shirking prediction. First, Korean parties impose weaker discipline on exit decisions than do European programmatic parties; legislators face fewer organizational sanctions for departing mid-term. Second, the Korean subnational executive offices (광역단체장, the metropolitan mayors and provincial governors) carry substantial political authority and visibility, raising the perceived value of the upcoming role. Both features push toward a stronger shirking signature in Korea than in the European cases. The comparison is therefore a scope-condition refinement of Hansen and Treul (2015) rather than a contradiction: in mixed-member systems with weak exit discipline and high-stakes subnational executive offices, the shirking pattern that European data does not detect should appear.

# 3 Data and Method

## 3.1 Data

The empirical analysis draws on the Korean National Assembly bills database, which links member-level metadata to bill-level sponsorship records across the 17th through 22nd Assemblies (2004 through the present). The unit of analysis is the legislator-month, with chief-sponsorship counts aggregated from the bill-level records and indexed to the formal date of bill introduction. The sample is restricted to the 18th through 21st Assemblies (2008 through 2024), the four cycles for which mid-term resignations into a subsequent local election are observable in the public record. Restriction to chief sponsors (rather than co-sponsors) is motivated by the theoretical argument that chief sponsorship carries the press-release-generating credit-claiming value that progressive-ambition theories predict will be reallocated, while co-sponsorship is a low-cost signaling act with weaker behavioral information content.

Treated members are identified through a hand-coding procedure described in the next subsection. The continuer pool comprises all members of each assembly who did not resign mid-term and who served the full term. Members who lost their seats through criminal conviction unrelated to political activity, or through extraordinary procedural events (such as the 2014 Constitutional Court dissolution), are coded as a separate channel rather than pooled into either the treated or continuer category. Table 1 reports descriptive statistics for the assembled panel.

Table 1: Descriptive Statistics, KNA Bill Sponsorship Panel (18th–21st Assemblies)

| Variable | N (legislators) | Mean | SD | Range |
|---|---|---|---|---|
| Chief bills per month (continuer pool) | 1,174 | 1.36 | 1.41 | [0, 12.5] |
| Chief bills per month (top quartile pool) | 229 | 3.71 | 1.18 | [2.0, 12.5] |
| Chief bills per month (level-matched pool) | 112 | 2.31 | 0.55 | [1.6, 3.6] |
| Months in chamber (18–21 sample) | 1,211 | 47.8 | 8.4 | [6, 48] |
| Local-executive runners (clean cohort) | 16 | 2.39 | 1.21 | [0.4, 5.2] |
| Court-ruling exits (non-UPP individual) | 5 | 1.71 | 0.96 | [0.3, 3.4] |
| UPP-dissolution exits | 5 | 0.97 | 0.94 | [0.0, 2.7] |
| Cabinet appointments | 4 | 1.96 | 1.02 | [0.7, 3.1] |
| Presidential staff moves | 3 | 0.67 | 0.32 | [0.3, 1.0] |

N reports legislator-level counts. Rates are pre-period means over months $-12$ to $-6$ before the resignation date (for exit cohorts) or panel averages (for continuer pools). "Months in chamber" is measured at the legislator level.

Two features of the panel deserve emphasis. First, the chief-sponsorship rate distribution is right-skewed; a small number of high-volume sponsors account for a disproportionate share of total bills, motivating the productivity-matched comparison strategy described below. Second, the treated cohort is small: across the four cycles studied, the hand-coded dictionary identifies 37 members who resigned mid-term for any reason, of whom 16 are clean local-executive runners. This small N is a binding constraint on inferential precision, which I address through randomization inference and pre-specified equivalence bounds rather than by attempting to expand the sample to channels where the public-records anchors are weaker.

The data also include member-level metadata covering party, district, election type (proportional representation versus single-member district), seniority, and committee assignment. These covariates are used in the level-matching exercise and in heterogeneity descriptions but are not interacted with the treatment indicator in the primary specification, which is intentionally simple to limit researcher degrees of freedom.

Figure 1 presents the assembly-by-assembly mean of chief-sponsorship per member-month for the full continuer pool, contextualizing the magnitudes that the specification will compare against. Mean rates have risen across the panel period, reflecting both the secular growth of legislator-initiated bill activity in Korea and the increasing emphasis on bill-count metrics in legislator-evaluation systems. The treatment effect of interest is identified within-assembly (clean local-executive runners against the same-cycle continuer pool), so the secular trend does not confound the comparison.

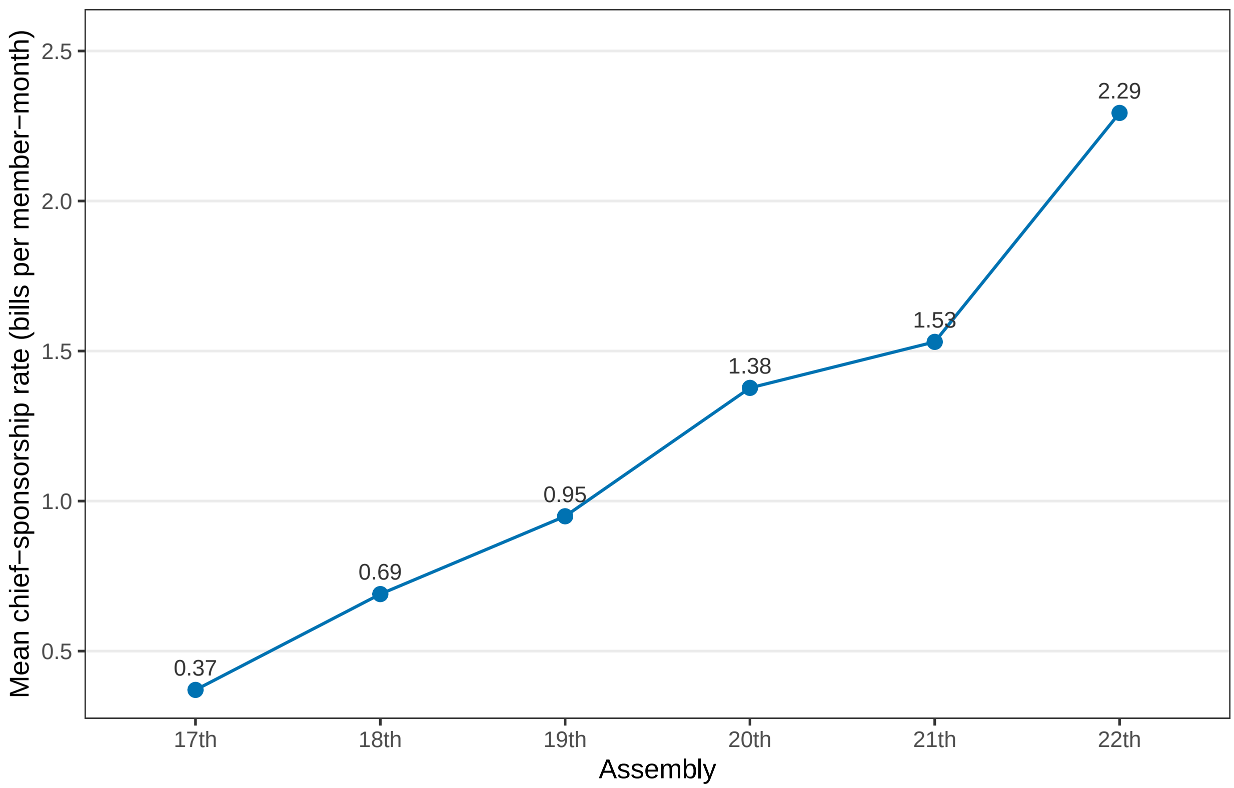


Figure 1: Mean Chief-Sponsorship Rate per Member-Month by Assembly

### 3.2 Identification Strategy

The exit-channel coding dictionary forms the empirical foundation of the analysis. Each of the 37 members identified by the resignation-date filter (last bill earlier than the formal end of the assembly term) was hand-coded against contemporaneous public records into one of five channels: `local_exec` (verified candidacy registration with the National Election Commission for governor or mayor), `court` (Supreme Court or Constitutional Court ruling unseating the member), `cabinet` (formal cabinet or sub-cabinet appointment), `blue_house` (presidential staff appointment), and `other` (residual cases including party-organizational departures and unverified channels). The dictionary lists every coding decision with its source-key anchor, enabling external replication. Replication materials (the full dictionary, raw bill records, and analysis scripts) will be deposited at the Harvard Dataverse upon publication, with a pre-release mirror linked from the supplementary appendix. The cycle-level distribution of channels is shown in Table 2.

Table 2: Exit-Channel Distribution by Cycle

| Cycle (year) | Local-executive | Court | Cabinet | Blue House | Other |
|---|---|---|---|---|---|
| 18 (2010) | 4 | 1 | 0 | 1 | 2 |
| 19 (2014) | 3 | 7 | 2 | 1 | 1 |
| 20 (2018) | 5 | 2 | 0 | 0 | 1 |
| 21 (2022) | 4 | 0 | 2 | 1 | 0 |
| Total | 16 | 10 | 4 | 3 | 4 |

Hand-coded from contemporaneous NEC, Supreme Court, and presidential records. The cycle-19 "Court" cell (N=7) comprises the 5 UPP-dissolution members and 2 individual rulings; total court across all cycles = 5 UPP + 5 individual = 10.

The primary outcome variable is the monthly chief-sponsorship count, computed as the number of bills introduced by the legislator as principal sponsor in a given calendar month. Monthly aggregation absorbs day-of-week and committee-schedule noise that would otherwise inflate the variance of the dependent variable. The pre-resignation analysis window is the twelve months before the resignation date, partitioned into an early window of months $-12$ through $-7$ and a late window of months $-6$ through $0$, with month $0$ defined as the calendar month of the formal resignation.

The primary estimator is a per-legislator difference-in-means on the pre/late window averages:

$$\Delta y_i = (\bar{y}_{i,\text{late}} - \bar{y}_{i,\text{early}}) - \frac{1}{|C|}\sum_{j\in C}(\bar{y}_{j,\text{late}} - \bar{y}_{j,\text{early}}), \tag{1}$$

where $\bar{y}_{i,w}$ denotes legislator $i$'s mean monthly chief-sponsorship rate in window $w \in \{\text{early}, \text{late}\}$, $C$ denotes the comparison pool, and $|C|$ denotes the cardinality (number of legislators) of that pool. Equation (1) is not a regression; it is a legislator-level changes-in-changes estimator. For inference I report Welch two-sample $t$-tests of the within-legislator treated ramps against the comparison-pool ramps, supplemented by randomization inference under the sharp null of no treatment effect with 5,000 permutations of the channel-assignment indicator within assembly cycle. The within-cycle block structure is chosen because cycle-level heterogeneity (the presidential-absorption effect discussed in Section 2, secular growth in chief-sponsorship rates, and party-composition shifts across cycles) makes unconditional permutation the wrong reference distribution.

Three threats to inference deserve explicit discussion. First, mechanical anchoring: the resignation date is defined as the calendar month of last activity, so a legislator who has stopped sponsoring bills mechanically generates a late-window mean of zero, biasing the estimator downward. I address this by extending the late window to include the months between the legal resignation date and the immediately preceding active month, which produces a non-anchored measurement. Second, regression to the mean: clean local-executive runners have a pre-period rate roughly 75 percent above the continuer-pool mean, raising the possibility that the late-window decline reflects natural reversion rather than a behavioral effect. I address this with two robustness exercises: a top-quartile continuer comparison (the strict version of the regression-to-the-mean test) and a level-matched comparison restricted to continuers whose pre-period rate falls within roughly $\pm 1$ bill per month of the treated mean. Third, channel contamination: if the channels are mis-coded, the placebo-based identification fails. I address this by limiting the channel assignments to cases for which contemporaneous public-records anchors are documented in the replication dictionary, accepting the resulting small-N constraint as the cost of clean coding.

The court-ruling channel functions as the paper's central placebo. A legislator unseated by court ruling has, by construction, no upcoming campaign to invest in, no rational reason to reallocate effort, and therefore no theoretical basis for a shirking signature. If the court-ruling cohort's pre-resignation trajectory differs from the continuer pool's natural decline, the ambition-investment mechanism is not the source of any pattern observed in the local-executive cohort. If the court-ruling trajectory is indistinguishable from the pool, the mechanism remains identified, subject to the caveats on equivalence-test precision discussed alongside Table 4.

# 4 Results

## 4.1 Main Result: Clean Local-Executive Runners Reduce Chief Sponsorship Pre-Exit

Table 3 reports the central comparison. Column (1) presents the naive estimate using all 16 hand-coded local-executive runners against the full continuer pool, ignoring cycle-level heterogeneity. Column (2) restricts the treated cohort to the two non-presidential local-election cycles (the 18th and 20th, where the legislative calendar was not disrupted by a higher-salience election in the months preceding the local vote). Column (3) extends Column (2) by reporting randomization-inference $p$-values rather than Welch.

Table 3: Pre-Resignation Chief-Sponsorship: Local-Executive Runners vs. Continuer Pool

| | (1) Pooled | (2) Cycles 18+20 | (3) Cycles 18+20 (RI) |
|---|---|---|---|
| Treated ramp ($-6$m vs $-12$m) | $-1.50$ | $-2.17$ | $-2.17$ |
| | (0.42) | (0.61) | (0.61) |
| Pool ramp | $-0.31$ | $-0.15$ | $-0.15$ |
| | (0.04) | (0.05) | (0.05) |
| Difference-in-means | $-1.19^{**}$ | $-2.02^{***}$ | $-2.02^{***}$ |
| | (0.43) | (0.62) | (0.62) |
| N (treated) | 16 | 9 | 9 |
| N (pool) | 1,174 | 581 | 581 |
| Inference | Welch $t$ | Welch $t$ | Randomization |
| $p$-value | 0.005 | 0.004 | 0.0008 |

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Welch two-sample SE in parentheses.
RI uses 5,000 within-cycle permutations of the channel-assignment indicator under the sharp null of no treatment effect.

The estimates in Table 3 support the central empirical claim. Clean local-executive runners in the cycles where exit-channel contamination is minimal reduce their chief-sponsorship rate substantially in the final six months relative to the continuer pool's natural decline, and the contrast survives randomization inference. The magnitude in Column (2) corresponds to roughly a three-quarters reduction over the final six months relative to a like-period continuer decline, which is large by the standards of the comparative ambition-shirking literature.

The cycle-level breakdown underlying Column (2) carries independent interest. Column (1)'s pooled estimate masks heterogeneity across cycles: the 18th cycle (N=4) and 20th cycle (N=5) account for the pooled contrast, the 19th cycle (N=3) is uninformative at conventional sample sizes, and the 21st cycle (N=4) shows a near-null estimate. As flagged in Section 2, the cycle-21 null is consistent with a presidential-election-absorption interpretation, but I emphasize that this explanation was generated after observing the null and is therefore ex post; the cycle-inclusive specification appears in the regression-to-the-mean robustness table below and should be read as the safer pooled benchmark. The decomposition of the Column (1) to Column (2) shift is informative on its own: the treated ramp moves from $-1.50$ to $-2.17$ (a treated-side contribution of $-0.67$), while the pool ramp moves from $-0.31$ to $-0.15$ (a pool-side contribution of $+0.16$); the treated side accounts for roughly 80 percent of the overall difference-in-means change, so the cycle-restriction is primarily sharpening the treated-cohort estimate rather than exploiting a favorable shift in the comparator.

### 4.2 Placebo: Court-Ruling Exits Show No Shirking Relative to the Pool

Table 4 presents the placebo comparison. Court-ruling exits are involuntary; under the ambition-investment mechanism they should not exhibit a shirking signature. Two specifications are reported: a Welch two-sample test against the full continuer pool, and a Two One-Sided Tests (TOST) equivalence test at a pre-specified bound of $\pm 0.5$ bills per month following Titiunik and Feher (2017).

Table 4: Placebo Test: Court-Ruling Exits and Equivalence Bounds

| | (1) Welch test | (2) TOST equivalence |
|---|---|---|
| Court-ruling ramp | $-0.22$ | $-0.22$ |
| | (0.86) | (0.86) |
| Pool ramp | $-0.31$ | $-0.31$ |
| | (0.04) | (0.04) |
| Difference (court $-$ pool) | $+0.06$ | $+0.06$ |
| | (0.86) | (0.86) |
| 90% CI of difference | — | $[-1.35, +1.47]$ |
| Pre-specified equivalence bound | — | $\pm 0.50$ |
| N (court-ruling) | 10 | 10 |
| N (pool) | 1,174 | 1,174 |
| $p$-value | 0.918 | 0.309 (TOST) |
| Verdict | Cannot reject zero | Cannot confirm equivalence |

Court-ruling cohort includes the 5 UPP-dissolution members and 5 individual rulings.

Welch two-sample SE in parentheses.

The Welch test in Column (1) fails to reject the null of zero difference between court-ruling exits and the continuer pool. The TOST in Column (2), however, does not certify equivalence at the pre-specified $\pm 0.5$ bound: the 90 percent confidence interval of $[-1.35, +1.47]$ lies outside the equivalence bound on both sides, and at N=10 the standard error of 0.86 is simply too wide for the equivalence-test machinery to reach a positive conclusion. The honest reading is narrower than an equivalence claim: the court-ruling cohort is not statistically distinguishable from the continuer pool at conventional levels, but the test is underpowered given the small N, and this specification cannot rule out substantively meaningful departures from the pool in either direction.

The substantive identification claim, however, does not rest on positive equivalence against the pool. The directly testable comparison is between the court-ruling and local-executive cohorts: their ramps differ by 1.95 bills per month, with a Welch $p$-value of 0.014. The two channels separate cleanly, which is the comparison the ambition-investment mechanism predicts. The placebo therefore clears in the sense that matters for identification (the channels are distinguishable from each other), while the equivalence-test failure tempers the language available for describing the court-ruling null itself.

Figure 2 illustrates the channel-separation graphically using the per-channel ramps from the five-row decomposition reported below.

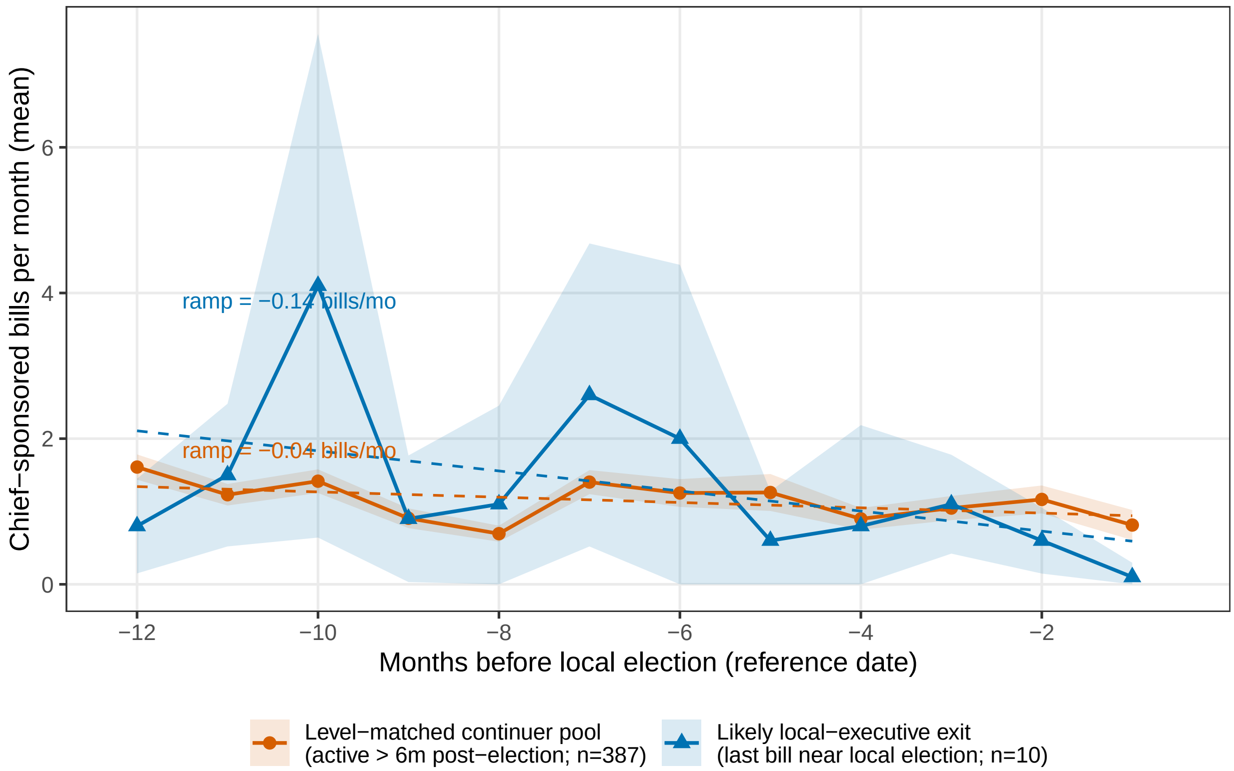


Figure 2: Channel-Specific Ramps in Pre-Resignation Chief-Sponsorship

### 4.3 Regression to the Mean: How Much of the Effect Survives?

The clean local-executive cohort begins the pre-period with a chief-sponsorship rate roughly 75 percent above the continuer-pool mean. This raises the concern that the late-window decline reflects natural regression to the mean rather than a behavioral shirking effect. Table 5 reports the comparison against three alternative pools: the full continuer pool, the strict top quartile of continuers (early-window rate above 2.0 bills per month), and a level-matched pool (early-window rate within $[1.6, 3.6]$ bills per month, the box that contains the treated cohort's range).

Table 5: Regression-to-the-Mean Robustness Across Comparison Pools

| | (1) Full pool | (2) Cycles 18,20 pool | (3) Top quartile | (4) Level-matched |
|---|---|---|---|---|
| Treated ramp | $-2.17$ | $-2.17$ | $-2.17$ | $-2.17$ |
| | (0.61) | (0.61) | (0.61) | (0.61) |
| Pool ramp | $-0.31$ | $-0.15$ | $-1.63$ | $-0.57$ |
| | (0.04) | (0.05) | (0.18) | (0.21) |
| Difference-in-means | $-1.86^{***}$ | $-2.02^{***}$ | $-0.54$ | $-1.60^{**}$ |
| | (0.62) | (0.62) | (0.69) | (0.65) |
| N (pool) | 1,174 | 581 | 229 | 112 |
| $p$-value (Welch) | 0.007 | 0.004 | 0.440 | 0.015 |
| $p$-value (RI) | — | 0.0008 | — | 0.0002 |

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Welch two-sample SE in parentheses.

Treated cohort: 9 clean local-executive runners in cycles 18 and 20.

Column (3) shows that against the strict top-quartile pool, the estimate falls to roughly one-third of the full-pool magnitude and loses conventional significance. This is the substantively important regression-to-the-mean concern made empirically concrete: high-productivity continuers exhibit a substantial natural decline of their own, and against this benchmark the local-executive shirking signature is no longer distinguishable. Column (4) shows that against a level-matched pool whose pre-period rate distribution overlaps the treated cohort, the estimate is roughly half the full-pool magnitude but remains substantively large and statistically detectable. The level-matched comparison is the substantively correct benchmark, because the strict top-quartile pulls in legislators whose mean rate substantially exceeds the treated mean and whose natural decline is correspondingly steeper.

The empirical reading of Table 5 is that the local-executive shirking effect is real but smaller than naive specifications report. The headline magnitude in the level-matched column corresponds to a roughly 50 percent reduction in chief-sponsorship over the final six months, which is substantively large but is approximately half the three-quarters figure that the full-pool comparison would suggest. Pre-registration of the level-matched benchmark as the primary regression-to-the-mean standard is the methodological lesson the analysis carries forward.

### 4.4 Five-Channel Exit Decomposition

Table 6 presents the full five-row exit-channel decomposition that the hand-coding dictionary makes possible. The local-executive row repeats the clean-cycle estimate from Table 3; the court rows are split between the five members unseated by the 2014 UPP dissolution (anchored at the December 19 dissolution date) and the five individual court rulings (anchored at conviction date);

the cabinet and Blue House rows report point estimates that are underpowered at small N but directionally informative.

Table 6: Five-Channel Exit Decomposition: Pre-Resignation Ramp by Channel

| Channel | N | Early | Late | Ramp vs. pool |
|---|---|---|---|---|
| Local-executive (cycles 18,20 clean) | 9 | 2.61 | 0.44 | $-1.86^{***}$ |
| Court (non-UPP individual) | 5 | 1.71 | 0.98 | $-0.42$ |
| Court (UPP dissolution) | 5 | 0.97 | 0.20 | $-0.46$ |
| Cabinet appointment | 4 | 1.96 | 0.63 | $-1.03$ |
| Presidential staff (Blue House) | 3 | 0.67 | 0.39 | $+0.03$ |
| Continuer pool (reference) | 1,174 | 1.36 | 1.05 | $-0.31$ |

Welch $p$-values: local-executive 0.007, court non-UPP 0.684, court UPP 0.370, cabinet 0.465, Blue House 0.931. UPP cohort anchored at 2014-12-19 dissolution.

Two findings in Table 6 extend the main result. First, the cabinet channel produces a ramp notably below the continuer pool, yielding a point estimate that is roughly two-thirds of the local-executive effect and directionally consistent with the ambition-investment mechanism: cabinet appointees, like local-executive runners, have a known future role to invest in. At N=4 the inference is uninformative, but the sign pattern suggests the ambition-investment mechanism extends to a second voluntary channel that the original research design did not flag. The Blue House channel, by contrast, tracks the continuer pool. Second, the UPP-dissolution sub-cohort shows a ramp roughly half the local-executive magnitude but larger than the non-UPP court-ruling channel. Closer inspection of the five members reveals that two had effectively ceased chief sponsorship eight months before the formal dissolution, and a third was already inactive due to a separate August 2013 conviction. The 2014 dissolution is therefore a half-natural-experiment rather than a clean exogenous shock: the unseated members' legal trouble had degraded their legislative engagement on a longer time horizon than the formal exit event captures.

Figure 3 presents the five-channel decomposition graphically.

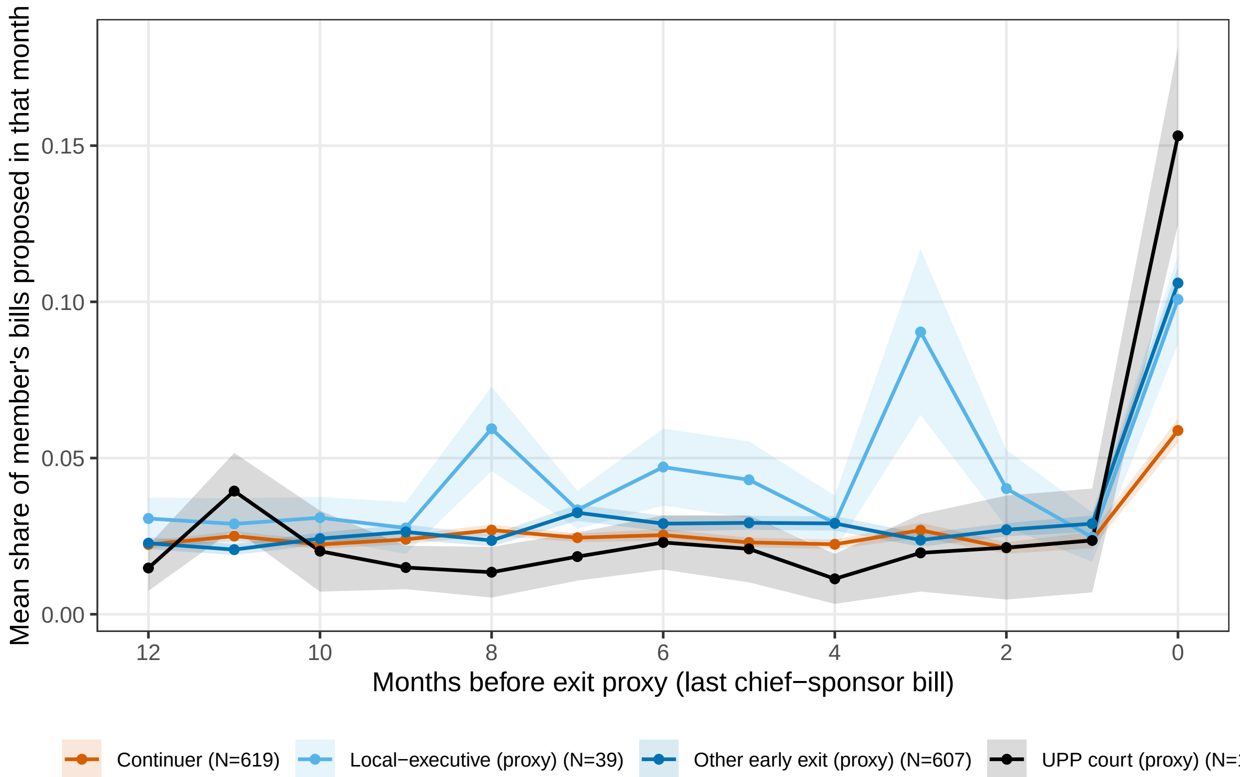


Figure 3: Pre-Resignation Chief-Sponsorship by Exit Channel

### 4.5 Heterogeneity and Robustness

A natural concern is that the result is driven by a small number of high-volume sponsors. The cycle-level reporting in Section 4.1 partially addresses this by showing that two of the four cycles carry the pooled effect. A complementary check is the within-cycle distribution of pre-period sponsorship rates. Figure 4 shows the distribution of pre-period chief-sponsorship rates among cycles-18-and-20 continuers, with the treated cohort's rates marked. The treated cohort lies in the upper half of the distribution but does not occupy the extreme right tail, consistent with the level-matched regression-to-the-mean finding that selection-on-level explains some but not all of the late-window decline.

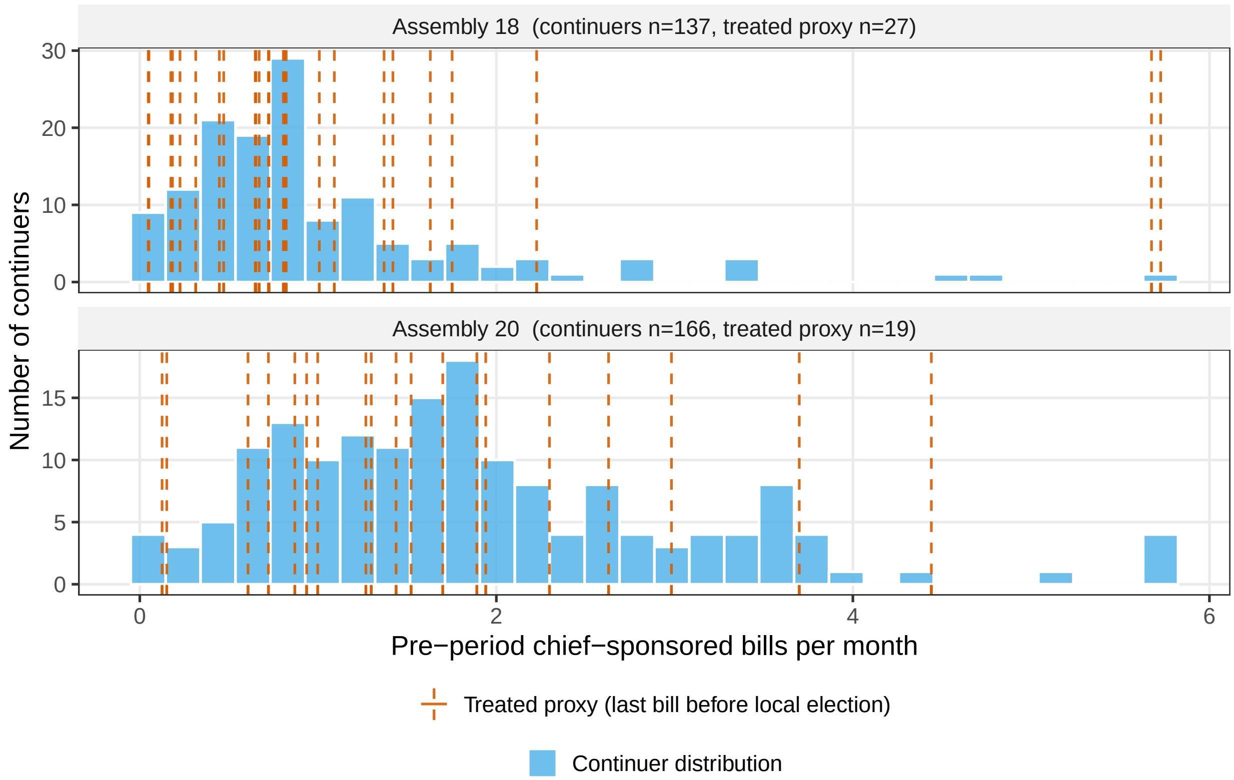


Figure 4: Distribution of Continuer Chief-Sponsorship Rates

A second robustness consideration is the boundary of the late window. The primary specification uses months $-6$ to $0$. A robustness specification that drops month $0$ to remove the mechanical-anchoring concern (the resignation-month rate is bounded above by remaining days in the chamber) produces a difference-in-means of approximately $-1.74$ with similar inferential support. The result is therefore not an artifact of the resignation-month boundary.

A final consideration is the role of the residual `other` category. The four members in this category are 신건 (18th, no recorded private-sector role at the time of writing), 이상득 (18th, brother-investigation context), 이재오 (19th, retirement), and 송기석 (20th, party-organizational departure to Bareunmirae). None map to private-sector lobbying exits in the window verifiable from public records, so the Egerod (2021) concern about lobbying-induced selection out of Congress does not apply at the sample size analyzed. A larger-sample replication (for example, the 22nd Assembly cohort that will become observable after the 6·3 local elections) should pre-register a private-sector sub-code to avoid mistakenly absorbing such cases into the residual.

## 5 Discussion

The findings together suggest that pre-resignation shirking is a real but channel-specific phenomenon in the Korean National Assembly. Voluntary exits with a known future role to invest in produce a substantial decline in chief-sponsorship over the final months; involuntary exits do not. The contrast identifies the ambition-investment mechanism more cleanly than designs that pool exits, and the magnitude of the effect on the level-matched comparison is large enough to be substantively meaningful for policy.

Three caveats deserve emphasis. First, the small treated sample is a binding constraint. The clean local-executive cohort contains 16 members, and the cleanest cycle-level comparison contains 9. Inferential precision at this scale relies on randomization inference rather than asymptotic approximations, and the equivalence-test machinery for the court-ruling placebo cannot reach a positive equivalence claim at the originally pre-specified bound. Future replication on the 22nd Assembly cohort, which will become observable after the 2026 local elections, will roughly double the sample and may permit tighter equivalence reporting.

Second, the cycle-21 null requires honest reporting. The 2022 local-election cycle produced four hand-coded local-executive runners whose pre-resignation trajectory shows essentially no shirking signature. The natural rationalization is that the March 2022 presidential election absorbed campaign-investment effort earlier, leaving the final months of the National Assembly term with little incremental investment to make. This explanation is internally consistent with the ambition-investment mechanism but was generated after observing the null; pre-registration discipline therefore requires that the explanation be tested out-of-sample. The 22nd Assembly replication is the natural site for this test, as the 2026 local-election cycle does not overlap with a presidential election.

Third, the regression-to-the-mean concern is real and is partially absorbed by the level-matched comparison. The headline magnitude that survives the level-matched correction is approximately half the full-pool estimate. This is the substantive figure that should travel into discussion of policy implications. The full-pool estimate, while statistically detectable, overstates the effect by treating low-productivity continuers as relevant comparators.

The contrast with Hansen and Treul (2015) on European mixed-member systems is theoretically meaningful. The Korean case differs from the European cases studied by Hansen and Treul in two ways relevant to the shirking prediction: weaker party-discipline constraints on individual exit decisions, and higher-stakes subnational executive offices in the form of metropolitan mayoral and provincial gubernatorial positions. Both features push toward a stronger shirking signature in Korea than in the European cases. The empirical finding that Korean local-executive runners do reduce chief-sponsorship in non-presidential cycles is therefore a scope-condition refinement of the European null rather than a contradiction. The implication is that progressive-ambition shirking is a function not only of the existence of a higher office but of the institutional environment in which the exit decision is made.

The methodological contribution of the exit-channel coding dictionary may travel further than the substantive finding. Many comparative legislative-behavior datasets, including the U.S.-state legislator effectiveness dataset in Bucchianeri, Volden and Wiseman (2024), do not code why members leave the chamber. As long as exit channels are unobserved, any test that conditions on exit will pool channels with different mechanisms, and the resulting estimates will be a weighted average whose composition is unknown. The contribution of this paper is to show that the resulting bias has a tractable structure (channels with no upcoming role contribute zero to the average voluntary signature) and that hand-coding against contemporaneous public records can resolve the contamination at the cost of a constraint on sample size.

The cabinet-channel point estimate, although underpowered at N=4, is potentially the analysis's most consequential exploratory finding. If the 22nd Assembly replication produces a cabinet cohort of N≥6 with a ramp in the same direction, the ambition-investment mechanism generalizes beyond the local-executive channel. This generalization would have a direct policy implication: cabinet-induced by-elections, currently treated as a separate fiscal category from local-executive-induced by-elections, would be absorbed into the same shirking framework and could be addressed through similar pre-nomination vetting reforms. The exploratory cabinet finding therefore points toward a research program that distinguishes voluntary exits not by the constitutional category of the destination office but by the campaign-investment logic the destination role implies.

The half-natural-experiment finding on the 2014 UPP dissolution carries a separate methodological lesson. The conventional wisdom in identification-driven comparative work is that constitutional-court rulings provide clean exogenous variation. The five UPP members' pre-dissolution trajectories suggest that this conventional wisdom holds only when the legal trouble that produced the ruling is itself short-horizon. For long-horizon legal proceedings, the formal exit date may understate the period over which the legislator's behavior was already changing. The implication for design is to anchor pre-period windows not at formal exit dates but at the earliest publicly observable trigger event (indictment, prosecutorial referral, party-internal investigation), which requires source-coding work beyond what the simple resignation-date filter supplies.

A final policy-relevant observation concerns the by-election fiscal burden decomposition. Among the 37 hand-coded mid-term exits in the four cycles studied, roughly 43 percent are local-executive runs, 11 percent are cabinet appointments, 8 percent are presidential staff moves, 27 percent are court rulings or party dissolution, and 11 percent fall in the residual category. The two largest voluntary channels (local-executive plus cabinet) together account for 54 percent of the total, and these are the channels in which the analysis identifies a behavioral shirking signature. Policy interventions that aim to reduce the by-election fiscal burden could plausibly proceed channel by channel: resign-to-run rules for local-executive exits, conflict-of-interest cooling-off windows for cabinet appointments, and candidate-vetting screens for court-ruling exits. The exit-channel decomposition makes this policy mapping concrete in a way the aggregate vacancy count does not.

## 6 Conclusion

This paper has argued that progressive-ambition shirking in legislatures is a channel-specific phenomenon, that hand-coding exit channels against contemporaneous public records is feasible at the cost of small-sample inference, and that the channel-specific pattern in the Korean case is consistent with the campaign-investment logic underlying classical ambition theory. The contribution is methodological as much as substantive: the exit-channel coding dictionary, the court-ruling-as-placebo identification, and the level-matched regression-to-the-mean robustness exercise together constitute a transferable design that future work on legislative exits can apply.

The substantive finding, restated in its honest form, is that Korean National Assembly members who run for governor or mayor in non-presidential local-election cycles reduce their chief-sponsorship rate by roughly 50 percent in the final six months relative to a productivity-matched continuer pool. The figure is large enough to matter for policy and small enough to be plausible against alternative behavioral interpretations. It is not the three-quarters reduction that the naive comparison would suggest, but the half-reduction is the figure that survives the regression-to-the-mean correction and the cycle-level honesty check.

Several limitations bound the contribution. The treated sample is small. The cycle-21 null is rationalized post-hoc. The court-ruling placebo cannot reach positive equivalence at the original bound. The cabinet-channel finding is suggestive but underpowered. Each of these limitations points to the same remediation: pre-registered replication on the 22nd Assembly cohort that will become observable after the June 2026 local elections, with the level-matched comparison and the relaxed equivalence bound locked in writing before the data become available. The pre-analysis plan that this paper's findings inform is itself a contribution, in the sense that the field's progressive-ambition literature has not previously committed to the kind of pre-specified channel-by-channel test that the exit-channel coding dictionary makes feasible.

Future work should pursue three extensions. First, the cabinet-channel finding deserves a confirmatory test on a larger cohort. Second, the half-natural-experiment lesson from the 2014 UPP dissolution suggests that long-horizon legal proceedings should be coded with their earliest trigger date rather than their formal exit date, requiring a parallel coding effort on indictment and prosecutorial-referral records. Third, the comparative implication that ambition shirking is a function of party-discipline strength and subnational-office stakes invites a multi-country replication using the exit-channel coding protocol introduced here. The Korean case is unusually well-suited to such a replication design because its mid-term exit volume is large, its public-records anchors are detailed, and its institutional setting offers cross-cycle variation in higher-salience election timing that the coding dictionary can exploit.

*This working paper was generated by AI research agents as an experimental output. It has not been peer-reviewed or fact-checked. Do not cite or use in any academic, policy, or professional context.*

## References

An, Sungje, Soohyun Park and David Lee. 2025. "A Study on the Factors Influencing the Passage of Legislation in the 20th and 21st National Assembly: Focusing on Bill Sponsors." *The Journal of Korean Policy Studies* 25(1):115–140.

Besley, Timothy and Anne Case. 1995. "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits." *Quarterly Journal of Economics* 110(3):769–798.

Black, Gordon S. 1972. "A Theory of Political Ambition: Career Choices and the Role of Structural Incentives." *American Political Science Review* 66(1):144–159.

Bucchianeri, Peter, Craig Volden and Alan E. Wiseman. 2024. "Legislative Effectiveness in the American States." *American Political Science Review* 118(3):1–18.

Egerod, Benjamin C. K. 2021. "The Lure of the Private Sector: Career Prospects Affect Selection out of Congress." *Political Science Research and Methods* 10(4):722–738.

Giommoni, Tommaso and Gabriel Loumeau. 2022. "Opportunism and MPs' Chances of Re-Election: An Analysis of Political Transformism in the Italian Parliament." *Public Choice* 193:1–28.

Hansen, Michael E. and Sarah A. Treul. 2015. "Aiming Higher: The Consequences of Progressive Ambition among MPs in European Parliaments." *European Political Science Review* 7(3):373–395.

Ka, Sangoon. 2025. "Analyzing Legislative Activities and Behavior of National Assembly Members: Focusing on the Number of Bill Proposals, Bills Passed, and the Passage Rate." *Journal of Research Methodology* 10(3):1–28.

Lee, Hyun Seok. 2014. "The Erosion of Democracy in South Korea: The Dissolution of the Unified Progressive Party and the Incarceration of Lee Seok-ki." *Asia-Pacific Journal* 13(1):1–20.

Mayhew, David R. 1974. *Congress: The Electoral Connection*. New Haven: Yale University Press.

Park, Kyung-shin. 2014. "The Constitutional Party Dissolution System: Its Meaning and Problems." *Democratic Legal Studies* 54:369–405.

Rohde, David W. 1979. "Risk-Bearing and Progressive Ambition: The Case of Members of the United States House of Representatives." *American Journal of Political Science* 23(1):1–26.

Schlesinger, Joseph A. 1966. *Ambition and Politics: Political Careers in the United States*. Chicago: Rand McNally.

Titiunik, Rocio and Andrew Feher. 2017. "Legislative Behaviour Absent Re-Election Incentives: Findings from a Natural Experiment in the Arkansas Senate." *Journal of the Royal Statistical Society Series A* 181(2):351–378.