Television and the Majority-Party Premium in Postwar Congress

Kisoo Kim*

August 25, 2026

Abstract

Broadcast television told postwar American voters almost nothing about their own representative — yet it reshaped House elections. I document that television created accountability without selection. Exploiting television's staggered 1946–1964 rollout across the United States, I find that each year of cumulative exposure raised majority-party House incumbents' probability of returning to the next Congress by 1.6 percentage points relative to minority incumbents, who show no comparable gain. The premium did not reward individual records. It was unrelated to legislative effectiveness, not concentrated among senior members, attached to chamber control rather than the president's party, and followed legislators into and out of the majority. Consistent with collective rather than individual attribution, it was largest when Congress produced least. The mechanism is informational substitution: television replaced candidate-level with party-level knowledge, shifting the House vote onto the majority party's brand. Local journalism's contemporary displacement by national digital media implies the same substitution.

Keywords: television; mass media; U.S. Congress; incumbency advantage; party brands

*Email: kisoo@virginia.edu

1 Introduction

Between 1948 and 1964, broadcast television reached nine in ten American households, the fastest diffusion of any mass medium in U.S. history. Its news was produced nationally: three networks covered Washington and the presidency, not 435 district races, and the medium's arrival displaced the local newspapers and radio that did cover House members (Gentzkow, 2006). The media-accountability literature is accordingly pessimistic about television as an instrument of congressional accountability — media sharpen voters' evaluation of individual representatives only where coverage gives voters something to evaluate them with (Strömberg, 2004; Snyder and Strömberg, 2010; Prat and Strömberg, 2013) — and its prediction for the rollout is an erosion of candidate-level knowledge, not a change in how House elections are decided. A second tradition expected more: television, a medium of faces and personalities, would personalize electoral politics and reward individually visible incumbents (Ranney, 1983; Wattenberg, 1991; Prior, 2006). The era seemed to bear both out: voters' knowledge of their representatives fell, and House results pulled steadily apart from presidential results, with split district-level verdicts climbing from a fifth of districts in 1948 to a third by 1964 (Ornstein et al., 2013). Neither tradition predicts what this article documents: television attached a large and durable electoral premium to majority-party status as such, a reward unconnected to who a representative was or what she had done. How does a medium that tells voters next to nothing about their representative end up organizing the House vote around the party in charge?

The answer this article develops is that television created accountability without selection. Each year of cumulative television exposure raised majority-party House incumbents' probability of returning to the next Congress by 1.6 percentage points; minority incumbents show no comparable gain. Yet a convergent set of tests shows that this premium bears none of the marks of individual selection. It is unrelated to legislative effectiveness once majority status is controlled. It does not concentrate among committee chairs: the chair-eligible senior tail carries the smallest premium of the four seniority quartiles. It follows the same legislator into and out of the majority at nearly identical magnitude: it attaches to the position, not the person. And it tracks control of the chamber, not control of the White House: in the four congresses whose House majority was not the president's party, the premium persists at full strength, and a presidential-co-partisanship interaction carries nothing once majority status enters the model. Each of these facts closes off a different rival account, and together they point to a single conclusion: voters were rewarding what the majority *was*, not what any member did.

A final test makes the collective character of that reward especially vivid: the premium is *largest* when Congress produces the least legislation, the opposite of what credit-claiming for collective output would predict. The mechanism, I show, is an informational substitution: television eroded voters' candidate-specific knowledge while enriching their knowledge of the national party landscape, shifting the basis of the House vote from individual records to the party brand (Snyder and Ting, 2002). A medium that made parties, rather than persons, legible to voters organized House voting around the governing party, resolving the puzzle with which this article began.

Whether a new medium weakens or deepens voters' reliance on party depends on what *kind* of information it conveys. Two cognitively distinct uses of partisan information respond to informational shocks in opposite directions. The *candidate-inference cue* (Type A) uses party as a placeholder for individual-candidate positions; candidate-specific information substitutes for it. The *collective-responsibility cue* (Type B) uses party as an attribution device for chamber outcomes; institutional information about which party governs *enables* it. The distinction follows Snyder and Ting (2002)'s informational theory of party brands, in which the electoral value of a party's label rises with voters' ability to map the party to outcomes it institutionally produces. Broadcast television in this period was a Type B medium: structurally national news content, focused on Washington and the executive branch, oriented to the visible work of the governing party — Speakers, committee chairs, presidential vetoes — rather than to 435 individual House members. A Type B shock should reshape voter knowledge asymmetrically: less about candidates, because the medium also displaced the local press that had carried candidate-level information; more about parties, because its own content taught who governs.

Direct survey evidence from the American National Election Studies bears out this prediction. Identified off the staggered state-level rollout, cumulative television exposure shifted the composition of what voters knew: a summary index of candidate-level knowledge falls by 0.10 standard deviations per year of exposure, while a summary index of party-level knowledge rises by 0.03 standard deviations — the two constructs move in opposite directions under the same shock, and both index estimates are precise. The components tell the same story: candidate-name recall and awareness of whether a candidate is the incumbent decline by three to four percentage points per year, while perceived party difference and engagement with the national outcome rise. Voters in television-exposed regions learned less about candidates and more about parties. A direct voter-level test that interacts each knowledge measure with TV exposure in predicting party-line House voting is consistent with this account, though only partially: all three candidate-knowledge (Type A) interactions are negative, and the party-knowledge (Type B) measure most closely tied to engaged national-outcome voting is positive, while the remaining coefficients are individually indistinguishable from zero (Section 4.1).

The magnitude of the premium is sizable relative to the incumbency advantage of the era. The 1.6-percentage-point gain is in return probability; on the vote-share metric in which incumbency advantages are measured, it is roughly 0.6 percentage points per year of exposure, against Gelman-King-style estimates of two to three percentage points for the era's entire House incumbency advantage (Gelman and King, 1990; Ansolabehere and Snyder, 2002) — about a quarter of the era's advantage per year of exposure. The estimate's scope is majority-versus-minority rather than Democrat-versus-Republican: a joint specification assigns the premium to TV $\times$ majority, leaving TV $\times$ Democrat null, and a Senate replication (Appendix D.4) gives a same-direction headline of essentially identical magnitude.

The identification strategy exploits the staggered FCC license freeze and post-freeze allocation that drove cross-market variation in cumulative television exposure between 1946 and 1964 (Gentzkow, 2006). The premium survives the sharpest test the design permits — absorbing every state-by-congress shock with state $\times$ congress fixed effects, so that no state-level tide or economy can generate it — and beyond that: an event study shows flat (if imprecisely estimated) pre-trends under every binning of event time, a pre-television balance test finds no pre-existing differential (though with limited power), the estimate holds under metropolitan-tide and redistricting controls, across a 72-cell specification grid, under a state-level wild cluster bootstrap, and under leave-one-state-out deletion (Section 3.4).

This paper sits at the intersection of two literatures. It supplies the historical natural-experiment grounding for what Hopkins (2018) and Moskowitz (2021), among others, have documented as contemporary nationalization correlates of local-news decline, and it is theoretically anchored in Snyder and Ting (2002)'s brand model of party labels, with the downstream tightening of party discipline I document consistent with cartel theory (Aldrich, 1995; Cox and McCubbins, 2005).

I make three contributions. Theoretically, the Type A / Type B distinction predicts divergent signs for the same information shock and resolves apparent contradictions in the media-effects literature. Empirically, voter-level evidence of asymmetric substitution and legislator-level evidence of a majority-party premium are estimated under one identification. Substantively, I document the mechanism through which broadcast television reshaped U.S. House electoral outcomes, ruling out meritocratic selection, credit claiming for legislative output, chair-visibility amplification, party-neutral incumbent visibility, and presidential-coattail transmission.

2 Theory

2.1 Two kinds of party cue

The starting point for my argument is a distinction that the literature on party-cue voting has tended to elide. The construct conventionally labeled "the party cue" actually refers to two cognitively distinct uses of partisan information that respond to different kinds of informational environments and predict opposite responses to the same kind of information shock.

The first use, which I call the candidate-inference cue (Type A), is the standard construct of the political behavior literature. When a voter knows little about a specific candidate, she uses the candidate's party label to predict where that candidate is likely to stand. A Democratic candidate is more likely to favor union rights, social spending, and civil rights legislation; a Republican candidate is more likely to favor business interests, lower taxes, and a hawkish foreign policy. The voter who knows nothing else about Smith but knows that Smith is a Democrat can therefore form an educated guess about Smith's positions by appealing to what she knows about Democrats as a class. This kind of inference is the workhorse of the low-information-rationality tradition. It is the mechanism behind Lupia's (1994) demonstration that uninformed California voters could mimic the choices of informed voters on insurance reform initiatives by attending to the endorsements of trusted partisan organizations, and it is the "low-information rationality" Popkin (1991) describes when he argues that voters use partisan shortcuts to economize on cognitive effort; its cognitive substrate — partisans processing political messages through the filter of their identification — is formalized in Zaller's (1992) Receive-Accept-Sample model.

The clearest theoretical implication of the Type A cue is that information about specific candidates is a substitute for it. The voter who learns that Smith voted against her party on a key issue, or that Smith has a particular record on a salient policy area, no longer needs the party label to predict Smith's positions; she can read those positions directly.

This logical implication is the basis for the long-standing intuition that "more information weakens partisan voting," though the empirical record is more complicated than that intuition suggests: Lenz (2012), for instance, finds that political information often strengthens rather than weakens the alignment between voters' partisanship and their expressed positions. The Type A / Type B distinction I develop below explains this empirical complexity by clarifying that "information" shocks have different effects depending on what kind of information they convey.

The second use of partisan information, which I call the collective-responsibility cue (Type B), operates through a different cognitive process. Here the party label is not a substitute for candidate-specific information but an input to a different kind of judgment: the assignment of credit or blame for collective outcomes. Voters know that legislation requires congressional majorities and that the presidency executes laws differently depending on which party holds it. To reward or punish the parties for the state of the country — for the economy, for war, for the productivity of Congress — voters must know which party is institutionally responsible. The party label functions here as an attribution device, allowing the voter to map collective outcomes to a specific organizational target.

This second tradition has its sharpest formal statement in Snyder and Ting's (2002) informational theory of party brands. In their model, the value of a party's label to its individual members is a *state-contingent informational asset*: voters use the label to forecast the party's collective behavior, and the value of doing so depends on the party's institutional position: whether it holds the majority, controls the chamber's agenda, and bears collective responsibility for what the chamber produces. The brand has more informational content, and is therefore more electorally consequential for individual members, when the party occupies a position the voter can map to collective outcomes. Closely allied accounts in the same tradition include Aldrich's (1995) treatment of parties as solutions to collective-action problems within the legislature, Cox and McCubbins's (1993; 2005) demonstration that the majority-party position confers genuine institutional power over the legislative agenda, and Fiorina's (1981) retrospective voting framework. Under all of these accounts, the party label has informational content that depends on the party's institutional position: what the label means depends on whether the party is in the majority, whether it controls the presidency, whether it bears responsibility for current policy outcomes.

The clearest theoretical implication of the Type B cue runs in the opposite direction from the Type A prediction. Information about which party governs is not a substitute for the cue but an enabler of it: the voter who newly learns that the Democrats control Congress can newly assign Democrats credit or blame for what Congress has produced. Information shocks that target institutional knowledge — that teach voters who governs — should therefore amplify, not weaken, the electoral relevance of the party label.

2.2 The cognitive distinction matters

The distinction between Type A and Type B is not a semantic refinement; it implies that the same nominal "party cue" can be eroded or activated by the same nominal "information shock," depending on what kind of information the shock conveys. A piece of information that increases candidate-level knowledge weakens Type A while leaving Type B unchanged. A piece of information that increases institutional knowledge strengthens Type B while leaving Type A unchanged. A medium that delivers both kinds of information would have ambiguous net effects. A medium that delivers asymmetric information — more of one kind than the other — would generate predictable asymmetric responses.

Most empirical work on media and electoral behavior has not made this distinction, and the resulting picture has appeared incoherent. One body of work finds that voters in informationally richer environments rely less on party labels — Snyder and Strömberg (2010) on the historical rollout of local U.S. newspapers being the canonical example. A second, growing body of work finds the opposite: voters in nationalized media environments engage in *stronger* party-line voting. Hopkins (2018) documents that American political behavior has nationalized since 2000, with voters in down-ballot races increasingly relying on national partisan signals rather than local incumbent-specific information. Moskowitz (2021) shows the same relationship from the other side: exposure to local news causally raises voter knowledge of incumbents and weakens party-line voting in down-ballot races, so its erosion removes exactly that counterweight. Darr et al. (2018) document that newspaper closures sharpen partisan polarization in local voting. Martin and McCrain (2019) show that station-ownership consolidation shifts local news content toward nationally homogeneous, national-party-aligned coverage.

The two literatures have appeared contradictory only because they have been read as opposing claims about a single construct. They are not. The first set is consistent with the Type A prediction: local newspapers deliver primarily candidate-level information, so where local newspapers are richer, the Type A cue weakens. The second set is consistent with the Type B prediction: the media that have eroded over the contemporary period (local newspapers) deliver Type A information, while the media that have replaced them (national cable, partisan online content) deliver Type B information — and the documented strengthening of party-line voting tracks the substitution of the latter for the former. My setting takes this contemporary pattern to its source. The 1948-1964 broadcast television rollout displaced local newspapers and rallies in favor of nationally produced political content — the same substitution of national for local information the contemporary literature documents, run for the first time, and with the natural-experiment variation in timing that the contemporary period lacks.

2.3 Why broadcast television was a Type B medium

Three structural features of broadcast television in 1948-1964 made it a much better delivery mechanism for Type B information than for Type A information.

First, the technology was national in its production. Three networks (CBS, NBC, and the early ABC) produced the bulk of news content, which was distributed identically across the markets that received their signals. The local station purchased a national news feed centered on Washington and the executive branch. Whatever was happening in a particular voter's House district — who was running, what they had said, whom they had endorsed — was almost entirely absent from the network nightly news, which by physical and economic constraints could not produce district-specific reporting for hundreds of separate races. What was present was the institutional drama of Congress: which party held the majority, what major bills were being debated, who was being investigated by which committee, whether the President's program was advancing.

Second, broadcast television substituted for media that had been more local in orientation. Gentzkow's (2006) influential entertainment-substitution finding documents that the arrival of television displaced newspaper reading and rally attendance — precisely the media that delivered candidate-level information about House races. Local newspapers in this era covered local races in detail, including endorsements, candidate biographies, and campaign-trail reporting. Rally attendance gave voters direct exposure to candidates and their organizations. The voter whose evening shifted from reading the local political column to watching network news did not simply lose information; she traded one kind of information for another. Type A inputs declined; Type B inputs rose.

Third, the framing of network political coverage emphasized partisan conflict at the institutional level. Network news segments on Congress in this era featured the Speaker, the Majority Leader, the relevant committee chairs — the institutional symbols of party control — far more than they featured rank-and-file members. The Sam Rayburns, Joe Martins, and Lyndon Johnsons of the era were television regulars; the typical House backbencher was not. The cumulative effect was to make "which party is running things" a salient and frequently rehearsed fact, while making "what is one's House member doing" a fact that voters had to seek out elsewhere.

These three features predict the same asymmetric pattern. The Type A cue should weaken, if at all, only because the Type A inputs themselves were eroding — consistent with the Zaller-Lupia tradition. The Type B cue should strengthen because the Type B inputs were being delivered through a new and unprecedented channel that, by the mid-1950s, reached most American households for several hours a day.

2.4 Empirical implications

The argument yields three testable implications. At the voter level, television exposure should increase institutional knowledge about the parties (Type B) without correspondingly increasing candidate-level knowledge (Type A); plausibly the latter declines. At the legislator level, television exposure should raise the electoral fortunes of majority-party incumbents — through higher return probability, larger vote share, and reduced exit — while leaving minority-party incumbents essentially unaffected. And the premium should operate through the party label as such rather than through individual records or institutional roles within the majority: the interaction between television exposure and legislative effectiveness should be null once majority status is controlled, and the effect should not concentrate in the senior, chair-eligible tail of the majority.

The competing accounts make different predictions. Under meritocratic selection, television raises the electoral return to legislative performance, predicting a positive interaction between television and legislative effectiveness on incumbent return. Under chair amplification, television selectively amplifies the institutional visibility of senior majority members who hold committee chairs, predicting that the majority-party effect concentrates in the top quartile of seniority. Under incumbent visibility, the free exposure a new visual medium confers on sitting members advantages incumbents regardless of party (Prior, 2006), predicting a television return premium for minority as well as majority incumbents. Under presidential-coattail transmission, television's presidency-centered content attaches the premium to the president's party rather than to the chamber majority, predicting that the two come apart in the congresses where the White House and the House majority belong to different parties. Under pure entertainment substitution

(Gentzkow, 2006), television depresses all forms of political knowledge and engagement, predicting uniformly negative effects on both Type A and Type B voter outcomes.

These implications are designed to fall together or apart. If the knowledge asymmetry fails to appear at the voter level — if television increases both kinds of knowledge symmetrically or leaves both unchanged — the proposed mechanism for the legislator-level premium is undermined. If the majority-minority asymmetry is absent at the legislator level, there is no premium to explain. If the effect concentrates in chairs or scales with individual legislative effectiveness, the "party label as such" interpretation collapses into a record-based or role-based story.

3 Research Design

3.1 Television introduction as a natural experiment

I exploit the staggered introduction of television stations across U.S. media markets between 1946 and 1964 as a source of plausibly exogenous variation in media access. The staggered rollout was shaped by two supply-side disruptions orthogonal to local political demand. During World War II, the federal government banned construction of new broadcast facilities, halting the nascent television industry. After the war, rapid growth in station applications overwhelmed the FCC's technical allocation system, leading to a freeze on new licenses from September 1948 to July 1952. Markets that had received licenses before the freeze experienced television years earlier than markets that waited for post-freeze allocation. This source of variation, driven by regulatory capacity constraints rather than local demand for political information, follows the identification strategy of Gentzkow (2006) and has been used in Campante and Hojman (2013) and Song (2022) on political-behavior outcomes, and in Fenton and Koenig (2025) on labor-supply outcomes.

I measure treatment at the district level as years of cumulative television exposure: $\max(0, \text{year} - \text{TV introduction year}_d)$, where a district's introduction year averages the first-station dates of its constituent counties' designated market areas (DMAs), using ICPSR study 22720 (Gentzkow and Shapiro, 2008) for first-station dates. Congressional districts often span multiple counties belonging to different DMAs with different introduction dates; the average weights each county part by its allocation share in the congress-matched county-to-district crosswalk (most weights equal one, because most counties lie wholly within a single district), and the crosswalk vintage follows the apportionment map in force (1940 for Congresses 78–82, 1950 for 83–87, 1960 for 88–89). Appendix E.1 states the construction precisely; the headline estimate is invariant to reweighting counties by population, and to fixing the census vintage, at the fifth decimal (Appendix D.5). The treatment variable ranges from zero to approximately nineteen and a half years across my sample, reaching the upper end only in early-station districts that contained pre-freeze metropolitan markets.

The key identifying assumption is that, conditional on district and congress fixed effects, the timing of television introduction is uncorrelated with unobserved determinants of legislator outcomes. Early- and late-adopting districts may differ in baseline levels, but the fixed effects absorb these. I test the design directly with a pre-television balance test: in the pre-television Congresses (the 78th and 79th, 1943–1946), the eventual timing of television's arrival, interacted with majority status, shows no cross-sectional association with incumbent return, though the test's power limits what it can rule out (Appendix A.1). An event study of the majority premium itself shows no differential pre-trend by majority status before television arrives (Figure 2); Section 3.4 reports the identification battery.

3.2 Voter-level data: ANES 1952-1964

For the voter-level analysis I use the American National Election Studies Cumulative Data File 1952-1964, restricted to respondents in the thirty-nine states that overlap with my legislator panel. I construct two families of outcomes.

The Type A (candidate-level knowledge) family draws on three items asked in 1958 and 1964: whether the respondent could recall the names of House candidates running in her district (VCF0972), whether she could state whether either candidate was already in Congress (VCF0977), and whether her identification of the incumbent was correct (VCF0978). The Type B (institutional/party-level knowledge) family draws on three items: whether the respondent perceives an important difference between the parties (VCF0501, available 1952, 1960, 1964), whether she correctly identifies one of the parties as more conservative (VCF0502, available 1960 and 1964), and whether she cares which party wins the presidential election (VCF0311, available throughout; an engagement item — caring about the national party contest — rather than a knowledge test). The American National Election Studies did not ask "which party controls Congress" before 1968, so the Type B measures are indirect proxies for institutional knowledge rather than direct measures.

I estimate

$$Y_{ist} = \alpha + \beta \, \text{years\_of\_tv}_{st} + X_{ist}\gamma + \delta_t + \varepsilon_{ist},$$

where Y_{ist} is one of the knowledge outcomes for respondent i in state s and election year t, X_{ist} contains demographic controls (age, education, income, race), and δ_t is a year fixed effect. Standard errors are clustered at the state level. State fixed effects cannot be added because television penetration was state-level and saturated by 1956: with state fixed effects, years of cumulative exposure becomes nearly collinear with year fixed effects. Identification therefore rests on cross-state variation in first-TV-year, with year fixed effects absorbing national time trends.

3.3 Legislator-level data: 1943-1966 House panel

For the legislator-level analysis I use a panel of all U.S. House members serving in the 78th through 89th Congresses (1943-1966). Ideology measures and the legislator roster are drawn from Voteview (Lewis et al., 2024). The Legislative Effectiveness Score follows the Volden-Wiseman convention (Volden and Wiseman, 2014) as extended to historical Congresses by Chiou and Goplerud (2024). Historical U.S. House election returns combine county-level data for 1840-1972 from ICPSR 8611 (Clubb et al., 2008) with district-level two-party vote tallies for 1946-2012 (Jacobson, 2013) and, for 1944, Clerk of the House returns (Tamas, 2022) (Appendix D.7). The analysis sample, restricted to the thirty-seven non-South states, contains 4,036 legislator-congress observations covering 1,075 unique legislators. The 37-state count excludes the eleven Confederate states from the 48-state legislator panel. The voter-level ANES analysis above retained Southern respondents, because the candidate-knowledge / party-knowledge mechanism I test is not subject to the one-party-South scope condition that motivates excluding the South at the legislator level; its 39 states are the portion of the 48-state panel in which the ANES fielded respondents.

Outcomes include the binary indicator of whether the legislator returns in the next Congress (constructed from Voteview rosters), the legislator's vote share normalized to the share of her own party (so that Democrats and Republicans both have "vote share" measured in the same direction), and the Legislative Effectiveness Score normalized to the congress mean. I define *incumbent return* as the same legislator (`icpsr`) appearing in the next Congress, by whatever route; the variable measures chamber persistence rather than strictly winning reelection, and its complement — failure to return — bundles general-election defeat, primary loss, retirement, resignation, death in office, and candidacy for higher office (Senate or Governor). Appendix F decomposes the headline estimate across these routes. I construct seniority as career House tenure: the count of House Congresses in which the legislator has served through the current one, computed from the full Voteview roster so that service before the panel window counts (a member first elected in 1930 enters the 78th Congress as a seventh-term member, not a freshman; Appendix E.5). Majority-party status is defined per congress: a legislator is "majority" in Congress t if her party held the House majority in Congress t.

I estimate variants of

$$Y_{ict} = \alpha + \beta\,\text{years\_of\_tv}_{ct} \times \text{majority}_{it} + \gamma\,\text{years\_of\_tv}_{ct} + \lambda\,\text{majority}_{it} + Z_{ct}\eta + \mu_c + \delta_t + \varepsilon_{ict},$$

where Y_{ict} is the outcome for legislator i in district c in Congress t, Z_{ct} contains district-level demographic controls (percent urban, percent nonwhite), μ_c is a district fixed effect, and δ_t is a congress fixed effect. Standard errors are clustered at the district-by-decade level, which allows within-district serial correlation across consecutive Congresses while treating decades as independent; results are robust across alternative clustering schemes (state, district-only, two-way state $\times$ congress; Appendix D.3), and a wild cluster bootstrap at the state level — the coarsest level at which introduction timing is assigned, thirty-seven clusters — gives $p < 0.0001$ for the headline interaction (Appendix D.3). The coefficient of interest is the interaction β.

The non-South restriction reflects the structural difference between the two-party North and the one-party South in this period. In Southern districts where general elections were uncontested, the electoral mechanism I describe cannot operate. I report the South-included results as robustness; the South attenuates the effect (Appendix D.1), consistent with the scope condition.

3.4 Identification credibility

The identifying assumption faces two families of threat, and the design answers each with a different instrument. The first family is district-level: early-television districts might have been on different political trajectories for reasons unrelated to television. The second is state-level: because introduction timing varies at the media-market and state level, any state-by-time shock correlated with that timing — state economic conditions, statewide races and their coattails, state party organizations — could in principle load onto the treatment.

The state-level family has a sharp test. Because the estimand is an interaction, I can absorb every state-by-congress shock outright and still identify the premium from the comparison of majority to minority members serving in the same state and congress: with district and state $\times$ congress fixed effects together, the interaction is $+0.015$ ($p < 0.001$, $N = 3{,}983$). No shock common to a state-congress cell — a state economy, a statewide tide — can generate this estimate: each cell's common component is removed by construction, even though those cells absorb 93 percent of the variance in cumulative exposure (Appendix D.5). The absorption does not reach shocks that hit the two parties differently within a state and congress: a statewide race's coattails, a state party organization's mobilization. Those bias the premium only if their partisan incidence tracks television timing, and two results bound them: the joint specification of Section 4.2 finds no TV $\times$ Democrat interaction, and the premium holds at full strength in the four congresses where majority label and president's party diverge (Section 4.3).

The district-level family is addressed in time by the event study of the premium itself (Figure 2): the pre-period interaction coefficients are jointly indistinguishable from zero (Wald $\chi^2(4) = 3.1$, $p = 0.54$) while the premium emerges and grows after arrival, and the verdict holds under all three event-time binnings examined ($p = 0.93$, 0.93, 0.54), of which the published one is the most conservative (Appendix A.2). Two disclosures qualify the test: the pre-period bins are small (72 to 260 observations, against 1,381 in the terminal bin), so individual pre-period intervals do not exclude headline-sized differentials, and the terminal point of Figure 2 pools all congresses five or more after arrival. The event study therefore carries the pre-trend argument jointly with the balance test below, not alone.

The balance test asks the cross-sectional question directly: before any station in the sample operated, were eventually-early-television districts already different? In the 78th and 79th Congresses, eventual introduction timing interacted with majority status shows no association with incumbent return: -0.014 (SE 0.014, $p = 0.32$, $N = 680$). Its limits are stated in Appendix A.1: the regressor is the introduction *year*, so a premium-manufacturing confound would surface as a negative coefficient — the side of zero on which the insignificant point estimate in fact lies — and the minimum detectable effect is 0.039 (two and a half times the headline), so the test is an absence of positive evidence for pre-existing differences, not a bound on them; the specifications with power against confounding are the state $\times$ congress absorption and the event study.

Two further specification checks address composition. The premium is not an artifact of metropolitan, early-television districts being more responsive to national partisan tides: allowing the effect to differ by district urbanization and letting urban districts follow their own congress-specific trajectory leaves the interaction essentially unchanged ($+0.015$, $p < 0.001$). And the premium survives the redistricting that followed the 1950 and 1960 censuses: with district-by-apportionment-cycle fixed effects, so that a district is never compared to its post-redistricting self, the interaction is $+0.019$ ($p < 0.001$). Estimated within single apportionment cycles, the premium attenuates as television matures, consistent with the medium's novelty; the attenuation is an exposure-level phenomenon, not vanishing identification, since the dispersion of introduction timing is stable across every subwindow (Appendix D.5).

Nor is the estimate carried by influential units, specification choices, or the treatment's construction. Dropping any single state moves the headline by at most 8 percent (minimum $|t|$ 5.9 across thirty-seven leave-one-state-out fits); dropping any single congress leaves significance intact. The estimate is positive under every combination of fixed effects, controls, sample window, and missing-covariate handling examined, weakening only in the earliest, lowest-exposure sample window (Appendix D.1). Rebuilding the treatment with population weights, or with census-vintage populations, moves the headline in the fifth decimal (Appendix D.5).

Heterogeneity-robust estimators built for binary absorbing treatments (Goodman-

Bacon, 2021; Sun and Abraham, 2021) do not fit this continuous-exposure design: coarsening exposure to a binary indicator discards the identifying intensity variation, produces degenerate treatment-timing groups under the compressed rollout, and returns uninformative nulls (Appendix A.3).

Randomization inference completes the battery with a division of labor. A design-based test that reassigns states' mean introduction years across states has power against the *level* effect of exposure; for the legislative-effectiveness outcome, the natural stress case, it gives $p = 0.046$ against a conventional clustered 0.020 (Appendix A.2). For the majority premium the same reassignment is uninformative by construction: it preserves the within-state contrasts that identify the interaction, so no verdict can be read from it in either direction (Appendix A.2); the premium's independence from cross-state assignment rests on the state $\times$ congress specification instead.

4 Results

4.1 Voter-level: candidate knowledge declines, party knowledge rises

The American National Election Studies data provide direct micro-level evidence for the asymmetric information substitution that underlies my argument. Because the argument concerns two *constructs* — candidate-level knowledge (Type A) and party-level knowledge (Type B) — rather than any single survey item, Table 1 leads with a summary index for each construct (the standardized average of the available standardized components, following Anderson, 2008), and reports the six components beneath them. All estimates are coefficients on cumulative years of state-level television exposure in the preferred specification with year fixed effects, demographic controls, and state-clustered standard errors. Figure 1 displays the component pattern visually. A first-stage check confirms that the treatment moves actual viewing: each year of cumulative exposure raises the probability that a respondent reports having watched political programs on television by

3.2 percentage points (Table 1), consistent with the established first-stage relationship in Gentzkow (2006).

Table 1: TV exposure and voter knowledge: Type A vs Type B

Outcome	Coef.	SE	p	N	Years
Summary indices (standardized)					
Type A index: candidate-level knowledge	-0.099^{***}	0.020	< 0.001	2,721	1958, 1964
Type B index: party-level knowledge	$+0.029^{***}$	0.010	0.003	5,996	1952–64
Type A components — candidate-level knowledge					
Recall House candidate name	-0.0350^{***}	0.0092	< 0.001	2,679	1958, 1964
Knows if cand was incumbent	-0.0391^{***}	0.0096	< 0.001	2,675	1958, 1964
Correctly IDs the incumbent	-0.0096^{***}	0.0025	< 0.001	1,779	1958, 1964
Type B components — institutional/party knowledge					
Perceives party difference	$+0.0100^{**}$	0.0047	0.032	4,147	1952, 60, 64
Knows which party more conservative	$+0.0092$	0.0073	0.208	2,461	1960, 1964
Cares about presidential outcome	$+0.0102^{**}$	0.0050	0.043	5,915	1952–64
First stage					
Watched political TV programs	$+0.0324^{***}$	0.0058	< 0.001	5,643	1952–64

Coefficients on cumulative years of state TV exposure. All specifications include year fixed effects and respondent-level demographic controls (age, education, income, race). Standard errors clustered at the state level. Each index standardizes its available components and averages them per respondent, then re-standardizes (mean 0, SD 1 in the estimation sample); index coefficients are in standard-deviation units per year of exposure. Item-level inference within the six-component family: the three Type A components survive both Holm and Westfall–Young corrections; the two individually significant Type B components do not (Holm $p = 0.097$ each), which is why the party-level claim rests on the index rather than on any single item. Sample: ANES Cumulative Data File 1952–1964, restricted to 39 states overlapping the legislator panel. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The pattern is stark at the construct level: the Type A (candidate-knowledge) index falls by 0.10 standard deviations per year of exposure while the Type B (party-knowledge) index rises by 0.03 standard deviations, opposite movements under the same shock, each precisely estimated. The components fill in the picture. Each additional year of cumulative television exposure is associated with a 3.5 percentage point reduction in the probability that a respondent can recall her House candidate's name, a 3.9 percentage point reduction in the probability that she reports knowing whether the candidate is the incumbent, and a 1.0 percentage point reduction in the probability that she identifies the incumbent correctly when asked. Two of the three Type B components are individually

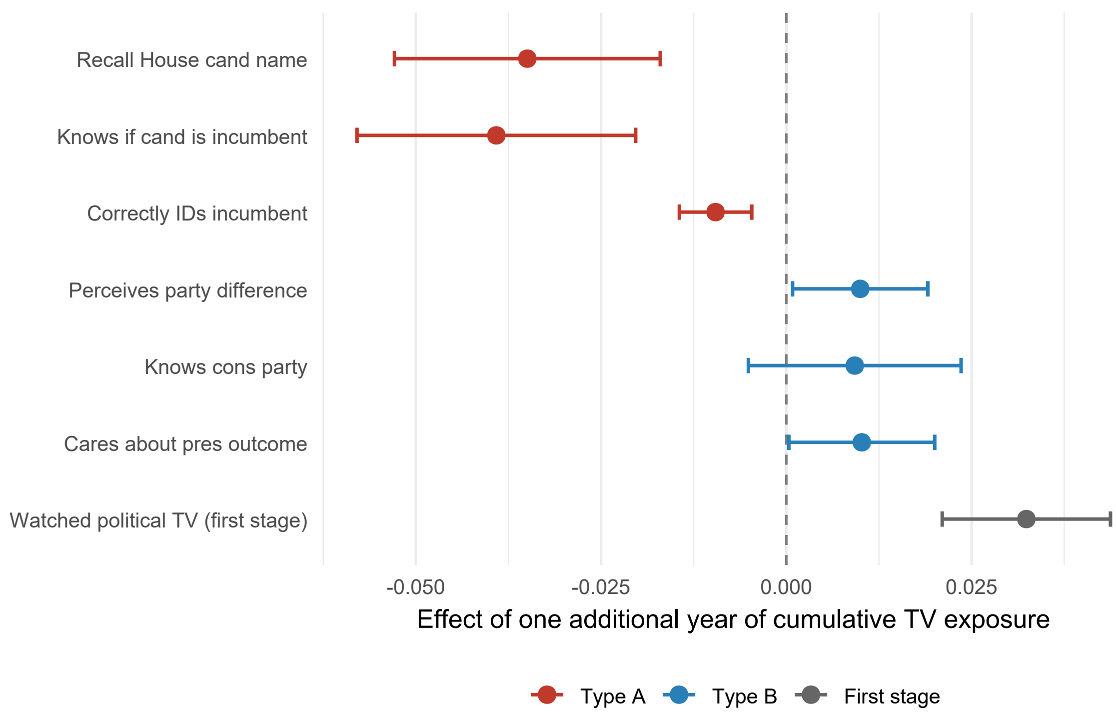

Figure 1: Effect of cumulative TV exposure on voter knowledge, by component. Each row shows the coefficient on years of state-level TV exposure with 95% confidence intervals from the preferred specification (year fixed effects, demographic controls, state-clustered SEs). Type A components (candidate-level knowledge) all decline; Type B components (party-level knowledge) all rise, two of three individually significant — the summary indices in Table 1 carry the construct-level inference. The first-stage relationship between exposure and television viewing is positive as expected.

significant — perceived party difference and caring about the presidential outcome each rise by 1.0 percentage point per year — and the third, knowing which party is more conservative, is positive but imprecise. Held to family-wise corrections across the six-item family, the Type A components all survive while the Type B components individually do not; the party-level half of the asymmetry accordingly rests on the index, not on any single item.

The substantive magnitudes are large, and both patterns are consistent with what survey research has long established about the postwar electorate. Voters have persistently held strikingly little candidate-specific knowledge of their House members: Stokes and Miller found that in the late 1950s most voters knew nothing about either congressional candidate on their ballots (Stokes and Miller, 1962), and Mann, analyzing district surveys two decades later, showed that congressional voting rested on candidate recognition and broad evaluation rather than on detailed recall of records (Mann, 1978). My voter-level results identify a force acting on that low baseline: cumulative television exposure further depressed candidate recall and incumbent identification. The concurrent rise in perceived party difference, in turn, aligns with the growing ideological distinctiveness of the postwar parties documented by Carmines and Stimson (1989) and others. Voters were not becoming uniformly more or less informed; they were trading candidate-specific knowledge for party-level knowledge.

The battery's construction also reveals *where* the Type A decline occurs. The correctness item conditions, by survey design, on respondents who ventured an identification; the other two span the full pool. Re-coding the correctness item onto the full pool more than quadruples its coefficient (to -0.041); conditioning the knows-the-incumbent item on those who ventured an answer shrinks it to an insignificant -0.004. Television's erosion of candidate knowledge thus operates on the extensive margin: it shrank the share of voters who could engage candidate questions at all, while those who still engaged them answered about as accurately as before — the fingerprint of a disappearing information source rather than of noisier information. One qualification attaches: venturing an answer is itself an outcome, so the flat conditional accuracy assumes exposure did not sharply reshape who ventured; the full-pool codings, immune to that selection, carry the decline itself.

The Type A and Type B effects are estimated on different respondent samples (the relevant ANES questions were asked in different election years), so the contrast is at the level of the regression rather than of the individual respondent. Restricting to 1964, the only year in which both question families were asked, the asymmetric pattern survives in cross-section, with smaller magnitudes as the narrower within-year exposure range predicts (Appendix B.2).

To test directly whether the knowledge shift translated into a shift in how voters used that knowledge, I interact each knowledge measure with cumulative TV exposure in predicting party-line House voting — a vote in which the respondent's House choice matches her own partisan identification. The evidence is partial but points in the predicted direction (Appendix B.3). All three candidate-knowledge (Type A) interactions are negative, and the most precise, knowing whether the candidate is the incumbent, is marginally significant (-0.023): the party-line gap between informed and uninformed voters narrows under television exposure. Among the party-knowledge (Type B) measures, caring about the presidential outcome interacts positively with television ($+0.014$), so that voters engaged with national outcomes vote more strongly with their party under television. The remaining four interactions are individually indistinguishable from zero, and two of the Type B measures point the wrong way. The test corroborates rather than decides: the sign pattern is uniform across the three Type A measures and appears again on the Type B measure most closely tied to engaged use of party information, but the individual coefficients are mostly imprecise, and it is the legislator-level evidence that carries the argument's weight.

Two caveats bound the voter-level evidence. First, state fixed effects cannot be added to the preferred specification. By 1956, every state in my sample had at least one operational television station, so the extensive-margin variation in TV access disappears, and the intensive-margin variation in years of cumulative exposure is essentially collinear with the combination of state fixed effects and year fixed effects. Identification therefore rests on cross-state variation in first-TV-year, which may correlate with state-level differences in urbanization, education, or media-market structure that demographic controls can only partially address. Second, the within-1964 first-stage on watched-TV is null, suggesting that within-year cross-state variation in cumulative exposure may capture more than just exposure intensity. The voter-level evidence is therefore best interpreted as direct micro-evidence for the mechanism that produces the legislator-level patterns I report next, rather than as standalone causal estimates of voter-level effects.

4.2 Legislator-level: the majority-party premium

The legislator-level analysis shows the asymmetric pattern that the voter-level evidence predicts. Table 2 reports the headline interaction estimates for incumbent return, vote share, and incumbent exit, estimated on the panel of legislator-congress observations covering the 78th through 89th Congresses with district and congress fixed effects, district-by-decade clustering, and demographic controls. Figure 2 presents the event study of the majority-party premium in incumbent return: the premium is flat before television arrives and emerges only afterward.

The interaction between cumulative television exposure and majority-party status on incumbent return is positive and precisely estimated: $+0.016$ per year. The estimate implies that an additional year of cumulative TV exposure raises a majority-party incumbent's probability of returning to the next Congress by 1.6 percentage points relative to a minority-party incumbent's probability of returning. The minority-party main effect of TV is small and indistinguishable from zero (the separate-sample minority regression in Table 2 reports $+0.001$) — a null that also cuts against party-neutral visibility accounts of

Table 2: TV exposure and majority-party premium

Outcome (per yr of TV exposure)	Coef.	SE	p	N
Headline interactions: TV × majority_party				
Incumbent return	$+0.016^{***}$	0.002	< 0.001	4,031
Own-party vote share	$+0.006^{***}$	0.001	< 0.001	3,923
Incumbent exit (1 − return)	-0.016^{***}	0.002	< 0.001	4,029
Subsample estimates				
TV × majority — R-majority Congresses (80, 83)	-0.001	0.016	0.97	674
TV × majority — within-legislator FE	$+0.019^{***}$	0.003	< 0.001	2,799
Separate samples (no interaction)				
TV — majority sample only, return	$+0.019^{**}$	0.009	0.026	2,012
TV — minority sample only, return	$+0.001$	0.008	0.859	1,984
Within-party marginal effects on vote share (pooled model)				
TV — majority members	$+0.0051^{***}$	0.0017	0.004	3,923
TV — minority members	-0.0005	0.0018	0.785	3,923

All specifications include district and congress fixed effects with demographic controls (% urban, % nonwhite), except the within-legislator row, which substitutes legislator fixed effects. Standard errors clustered at the district-by-decade level; reported N is the estimation sample after fixed-effect singleton removal. Sample: non-South, House congresses 78–89, complete LES + Voteview matching. The two vote-share within-party rows are marginal effects of the pooled model in the headline row and share its estimation sample ($N = 3{,}923$ each); the corresponding separate single-party regressions are uninformative on the same congresses (majority $+0.002$, $p = 0.47$; minority $+0.003$, $p = 0.19$) and are not reported. The within-legislator FE estimate is from the switcher subsample; switcher-subsample estimates with alternative controls range $+0.017$ to $+0.021$, and the full-sample legislator-FE triple specification gives $+0.022$. The R-majority subsample ($N = 674$ across two Republican-majority Congresses) cannot power a separate cross-party test, and the formal triple TV × majority × Republican-Congress is null ($p = 0.93$). Within-legislator FE specification restricts to the 463 House members observed in both majority and minority status. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

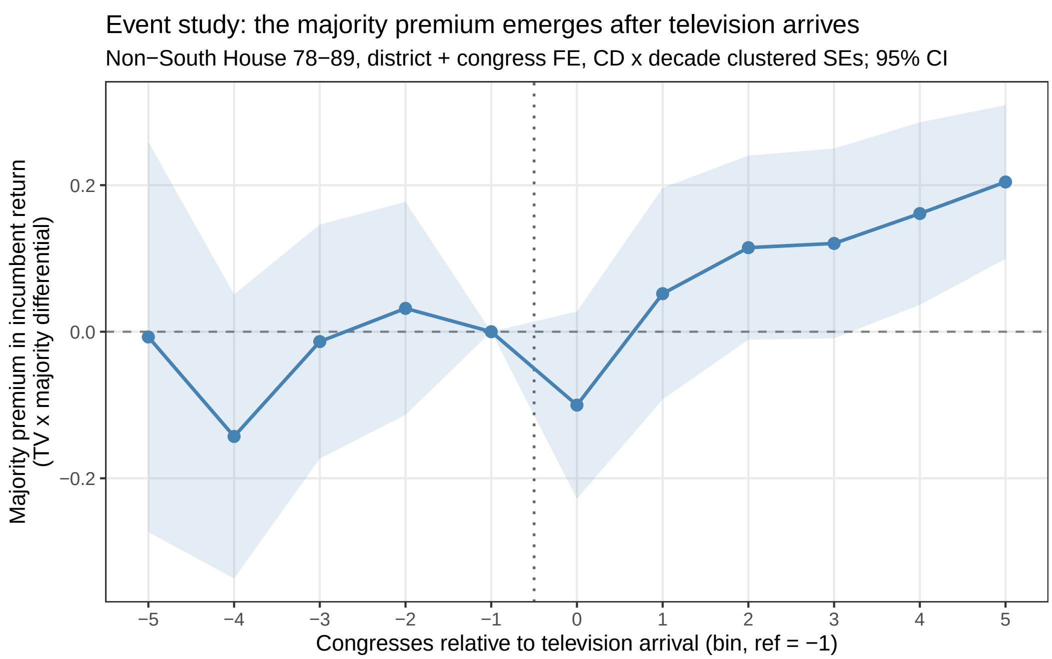

Figure 2: Event study of the majority-party premium in incumbent return. Each point is the majority $\times$ relative-time interaction coefficient — the majority premium at each congress relative to television's arrival — with 95% confidence intervals, from a district- and congress-fixed-effects specification with district-by-decade clustered standard errors; the congress before arrival is the reference. Before television arrives the premium is flat and jointly indistinguishable from zero (Wald test on the pre-period interactions, $\chi^2(4) = 3.1, p = 0.54$); after arrival it emerges and grows. Relative time is binned and clipped at ± 5; the endpoint bins pool all earlier and later congresses (bin -5: $N = 72$; bin $+5$: $N = 1{,}381$). The corresponding event study for the downstream legislative-effectiveness outcome, which is sensitive to event-time binning, appears in Appendix A.2.

television incumbency (Prior, 2006). Benchmarking requires care with units: the premium is in return-probability points, while the incumbency-advantage literature measures vote share. On that common vote-share metric, Gelman-King-style estimates place the U.S. House incumbency advantage at roughly two to three percentage points through the 1950s (Gelman and King, 1990; Ansolabehere and Snyder, 2002), rising to about five to six points around 1960, with the substantial, sustained advantage of later decades emerging only after my window (Cox and Katz, 1996); the premium's own vote-share expression (Table 2) is 0.6 points per year of exposure, so each year of exposure moved a majority incumbent's margin by roughly a quarter of the era's entire advantage. The per-year point estimate is precisely identified; I do not compound it into multi-year magnitudes because the effect attenuates over time[1] and the linear extrapolation extends beyond the within-state variation directly observed.

Whether the asymmetry operates uniformly under both parties' majorities cannot be cleanly tested in this window: the 78–89 sample contains only two Republican-majority Congresses (the 80th and 83rd, combined $N = 674$), and in that small subset the interaction is indistinguishable from zero, as is the formal triple with a Republican-Congress indicator — failure to reject, not evidence of uniformity. Two pieces of evidence nevertheless point to majority status rather than partisan identity. A joint specification including both TV $\times$ majority and TV $\times$ Democrat gives the interaction to majority status ($+0.016$, $p = 0.008$), with TV $\times$ Democrat at $+0.001$ ($p = 0.90$); because the labels separate only in those two Republican-majority Congresses, the Democrat null bounds rather than excludes a party-identity component; the point estimates load the premium entirely on majority status. And an independent Senate replication over the same window (Appendix D.4) yields a majority interaction of $+0.017$, essentially identical to the House's $+0.016$, with a positive (if underpowered) R-majority subset estimate. The discrimination between chamber majority and president's party in Section 4.3 sharpens the same

[1]The premium attenuates with accumulated exposure: the majority $\times$ exposure^2 term is -0.002 ($p < 0.001$).

conclusion on four separated congresses rather than two.

The vote-share results corroborate the return outcome on the same congresses. I measure vote share as the legislator's own party's share of the major-party vote, not the Democratic share, which would conflate party effects with majority effects. The TV-by-majority interaction is $+0.006$ per year of exposure, and the model's within-party marginal effects show the same asymmetry: $+0.0051$ per year for majority members ($p = 0.004$) against -0.0005 for minority members ($p = 0.785$). The premium therefore operates on the margin of victory, not merely on crossing the 50 percent threshold.

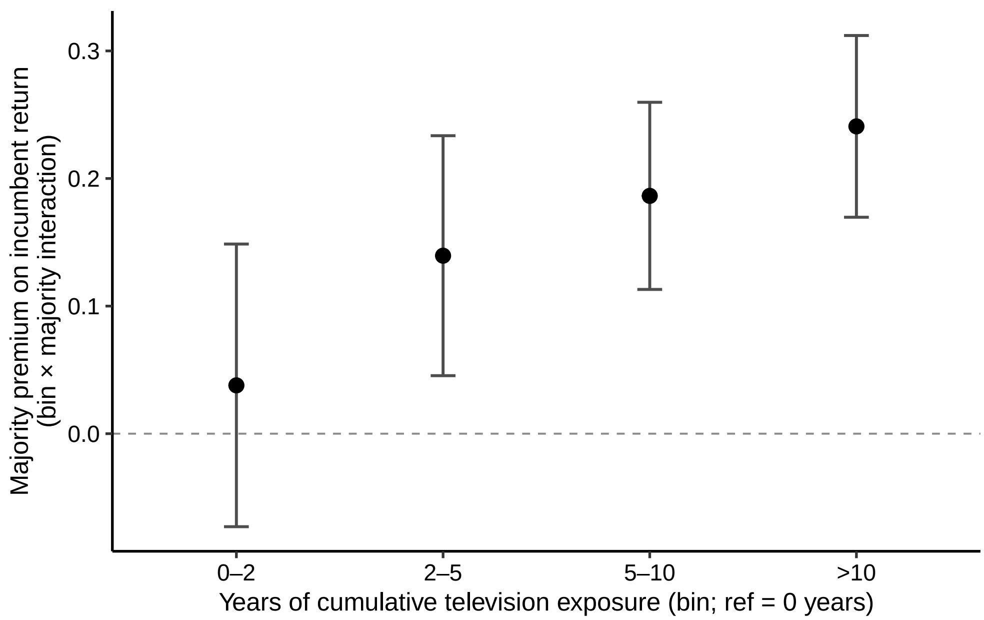

Figure 3: Binned dose-response of the majority-party premium. Each point is the interaction of an exposure bin (years of cumulative television exposure; zero years is the reference category) with majority status, from the district- and congress-fixed-effects specification with district-by-decade clustered standard errors; whiskers are 95% confidence intervals. The premium rises monotonically across the bins; the minority-party main effects across the same bins are flat and never significant.

The premium also behaves like a dose-response rather than a threshold. Replacing the linear exposure term with exposure bins in the incumbent-return specification, the majority premium rises monotonically with accumulated exposure, from an imprecise $+0.04$ in the first two years to $+0.24$ beyond ten (Figure 3), while the minority main effect is flat and insignificant in every bin. The profile is concave in accumulated exposure, consistent with the per-year novelty attenuation noted above; the monotone rise across all four bins is what a genuine cumulative-exposure effect, rather than a coding artifact or a single-period shock, should produce.

The exit margin restates the return effect in terms of who leaves, and locates it. Among majority-party incumbents, each year of TV exposure reduces the probability of exit — defined as not appearing in the next Congress, whether due to electoral defeat, retirement, death, or candidacy for higher office — by 1.9 percentage points. Among minority-party incumbents, the effect is small and insignificant. The interaction between TV and majority on exit is -0.016, the near-mirror of the return effect (the two estimation samples differ by two observations); the exit lens adds interpretation rather than independent evidence. In the exit-route decomposition (Appendix F), defeat at the polls (general-election and primary) accounts for 63.7 percent of the premium; retirement for 12.2 percent, below its 19.5 percent share of exits, on an interaction not distinguishable from zero. The premium is mainly a difference in defeat at the polls; the retirement difference is small and imprecise.

4.3 Mechanism: party label, not individual record

The headline asymmetry between majority and minority incumbents is consistent with several distinct mechanisms that share the same observable signature. Television could amplify the electoral return to legislative effectiveness, with majority members benefiting more because they pass more legislation that voters can credit. It could amplify the visibility of committee chairs and other senior members whose institutional power is concentrated in the majority — visibility that television, with its tendency to focus on institutional leaders, might disproportionately magnify. Or it could operate at the level of the party label itself, rewarding all majority members regardless of record or role. Each alternative would have been the obvious story given some plausible reading of the existing literature, and each makes a sharp, falsifiable prediction.

If television enabled voters to identify and reward more effective legislators — the natural extension of Snyder and Strömberg's (2010) accountability story to the broadcast medium — then the interaction of television exposure with legislative effectiveness should be positive. Voters in television-exposed districts would discriminate among candidates partly on the basis of legislative productivity, and incumbents with strong records would benefit electorally from voters' new ability to evaluate them.

After controlling for majority-party status, the TV-by-LES interaction on incumbent return is null. The triple interaction of TV, majority status, and top-quartile LES is also null. Television does not amplify the electoral premium for legislative productivity. Notably, the absence of an interaction does not reflect television eroding a previously strong reflection of records in electoral outcomes — the baseline LES-return association in this period is itself small and statistically indistinguishable from zero ($+0.003$, district FE without majority control). Television added a record-independent majority-party premium rather than displacing a record-based reward structure. The picture sharpens when I look at the productivity gradient within the majority. Among majority-party members, those in the lowest LES quartile have a point-estimate gain of approximately $+0.040$ from each additional year of TV exposure (not significant at conventional levels in this underpowered cell), while those in the highest LES quartile have a point estimate close to zero (-0.007). Among minority-party members, even those in the highest LES quartile show no comparable gain — though they are by construction the legislators with the strongest individual record to project to voters who might be paying attention to legislative output. Television does not reward effective legislators; if anything, it rewards ineffective majority members at least as much as effective ones. This pattern is inconsistent with any credit-claiming story in which voters use newly acquired information to identify productive members.

A complementary test exploits variation across Congresses in aggregate legislative output. If voters are rewarding majority members for what Congress accomplishes, the majority-party premium should be larger in Congresses that pass more major legislation, where there is more for voters to credit. The opposite is true. The triple interaction of television, majority status, and a high-productivity-Congress indicator is -0.021: the implied premium is $+0.024$ per exposure-year in low-output Congresses against $+0.003$ in high-output ones. Because aggregate output varies only across twelve congresses, formal significance overstates precision here; the split's direction and magnitude carry the claim. When Congress is producing little, majority members benefit more from television exposure, not less. The pattern falsifies credit-claiming for specific accomplishments, whose prediction is a premium that grows with output. Its bearing on collective attribution is narrower: the triple measures how strongly electoral fates track the majority label, not which grade voters assigned; a label whose electoral relevance persists — if anything grows — when output is scarce is what attribution to a standing collective, rather than to itemized accomplishments, implies. A credit-claiming voter has nothing to credit when Congress accomplishes nothing; an attribution voter still knows whom to hold responsible.

A second alternative is that television amplifies the institutional visibility of senior members who occupy committee chairs and other leadership positions. Network news in this era did disproportionately cover institutional leaders. If voters were learning about the work of chairs through television and using that information to reward them, the majority-party premium should concentrate in the senior tail of the majority distribution.

I test this using seniority as a proxy for committee chair status. The proxy is unusually clean for my period: in 1948-1964, House committee chairs were determined by strict seniority within the majority party, with no exceptions for performance, loyalty, or party leadership preference. I measure seniority as career House tenure computed from the full Voteview roster, so that a chair first elected in the 1920s is counted at his true rank rather than as a first-term member of the panel window, and I form quartiles within the majority pool (the top quartile, seven or more terms, is the chair-eligible tail). (Direct committee-assignment data for my period are not available; the Stewart-Woon committee dataset begins in 1993, and the U.S. Party Leadership data have a gap from 1942 to 1994.)

The premium is positive and distinguishable from zero in every seniority quartile, and it is *smallest* in the top quartile, the chair-eligible tail (Table 3). A Wald test does not reject equality across the quartiles at the 5 percent level, and the formal triples with top-quartile and continuous seniority are both null. The chair-amplification story would have predicted concentration at the top; instead the chair-eligible tail carries the smallest premium in the table.

The within-legislator switch evidence reinforces the conclusion: re-estimated among the 463 members who served under both statuses, the seniority triple remains null and negative-signed (-0.006, Table 3) — in the switcher sample, as in the full-sample quartiles, the premium does not concentrate in the senior tail.

A fourth alternative deserves its own test because the theory itself raises it. Network news content centered on Washington and above all on the presidency; if what television taught voters was "which party governs the country," the label carrying the premium could be the *president's* party rather than the chamber majority: presidential coattails delivered through the screen. The two labels usually coincide in this window, but not always: in the 80th Congress the House majority was Republican under a Democratic president, and in the 84th through 86th Congresses the majority was Democratic under President Eisenhower. Four of the twelve congresses separate the labels, twice the variation available for the Democrat-versus-majority comparison.

The chamber label wins outright. In a joint specification including both interactions, majority status keeps the full premium while presidential co-partisanship carries nothing (Table 3). Re-estimated within the four label-opposed congresses (where every majority member is, by construction, not a presidential co-partisan), the premium is $+0.038$ in a subsample refit, and the alignment triple finds no aligned–opposed difference (-0.004, $p = 0.53$): the premium neither requires nor grows with the majority party's control of the White House. A specification with presidential co-partisanship alone finds $+0.009$, which vanishes once majority status enters the model: what looks like a presidential-party effect is the chamber-majority effect wearing the wrong label. The result sharpens the informational reading of the premium: although the medium's content was heavily presidential, the attribution voters performed with that content attached to the institution whose collective output was being reported. Voters used television's institutional information to locate responsibility for Congress in the party that ran Congress.

The cleanest test of the party-label interpretation is the within-legislator switch. The window contains two majority transitions in each direction (the Republican gains of 1946 and 1952, each reversed two years later), so the 463 legislators who served under both majority and minority conditions provide a direct comparison: the same person, observed under both regimes, with all individual characteristics (name recognition, ideology, ability, charisma, district relations) held constant. If the premium is a property of specific legislators rather than of the majority position, the within-legislator effect should be smaller than the cross-legislator effect. If the premium is a property of the position itself, the within-legislator effect should match the cross-legislator estimate.

The within-legislator TV-by-majority interaction on incumbent return is $+0.019$, matching the cross-legislator $+0.016$ within sampling error. When the same legislator's party gains the majority, television exposure increases her probability of returning to the next Congress by approximately the same magnitude I estimate in the cross-section. When her party loses the majority, the gain does not follow her (the minority-side estimates in Table 2 are null). The premium does not travel with the person; it travels with the position. Table 3 collects the key tests.

Each of these tests could have rejected the party-label interpretation; none did. Television's effect on majority-party incumbents is a label phenomenon, not a record phenomenon and not a chair phenomenon. Strategic retirement — a premium produced by who chooses to run rather than by whom voters defeat — is the one rival these tests do not address; the exit decomposition bounds it at 12.2 percent of the premium (Appendix F). The voter-level evidence in Section 4.1 explains the consistency: voters in television-

Table 3: Mechanism tests for the majority-party premium

Test	Estimate	SE	p
Meritocratic selection (TV × LES)			
TV × LES on incumbent return (controlling majority)	-0.0002	0.001	0.84
TV × majority × LES_top (triple)	$+0.002$	0.005	0.74
TV effect, majority Q1 LES (least productive)	$+0.040$	0.024	0.09
TV effect, majority Q4 LES (most productive)	-0.007	0.026	0.79
Cross-Congress productivity gradient			
TV × majority × high-productivity-Congress (triple)	-0.021^{***}	0.005	< 0.001
Chair amplification (seniority)			
Q1 (junior, 1–2 terms)	$+0.019^{***}$	0.006	0.002
Q2 (3–4 terms)	$+0.021^{***}$	0.006	0.001
Q3 (5–6 terms)	$+0.018^{***}$	0.007	0.009
Q4 (senior, 7+ terms; chair-eligible)	$+0.012^{**}$	0.006	0.043
Wald test of equality across quartiles	$\chi^2(3) = 6.83$	—	0.078
TV × majority × senior_top (triple, binary)	-0.001	0.005	0.85
TV × majority × senior_z (triple, continuous)	$+0.003$	0.002	0.23
Presidential-party discrimination			
TV × majority (joint spec. with pres. co-partisan)	$+0.017^{***}$	0.003	< 0.001
TV × presidential co-partisan (same model)	-0.002	0.003	0.51
TV × majority — label-opposed Congresses (80, 84–86)	$+0.038^{***}$	0.008	< 0.001
TV × majority × aligned-Congress (triple)	-0.004	0.006	0.53
Within-legislator switch			
TV × majority within-legislator FE	$+0.019^{***}$	0.003	< 0.001
TV × majority × senior_top (within-legislator)	-0.006	0.006	0.30

Outcome variable: incumbent return to next Congress. All specifications include district and congress fixed effects (the within-legislator rows substitute legislator FE for district FE) with demographic controls. SEs clustered at the district-by-decade level; the presidential-party joint-specification rows are additionally confirmed by a wild cluster bootstrap at the state level (majority $p < 0.0001$; co-partisan $p = 0.46$). Sample: non-South, House congresses 78–89. Seniority is career House tenure from the full Voteview roster; quartiles are defined within the majority-party pool (chair-eligibility threshold ≥ 7 terms). Presidential co-partisanship: the member's party equals the sitting president's party; label-opposed Congresses are those whose House majority was not the president's party. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

exposed regions are not learning more about specific candidates (Type A is declining), so they have less, not more, candidate-level information with which to discriminate within the majority. They are learning more about parties (Type B is rising), so they use the most informationally accessible pieces of party knowledge — including which party holds the majority — to assign collective credit and blame.

These tests also collectively narrow the residual interpretive space. The remaining channels within the party-label family are themselves close cousins: brand valuation in the Snyder-Ting sense, agenda-control awareness in the Cox-McCubbins sense, collective-leadership visibility as a brand attribute (distinct from the individual chair channel the seniority test rejects), and credit/blame attribution as in Fiorina (1981). I cannot empirically separate these; they are bundled in my aggregate "majority-party premium." But all four sit within the Type B / collective-responsibility family that my theory predicts television should amplify, and none requires individual-record information that television does not deliver.

4.4 Downstream consequence: legislative effectiveness rises through majority retention

The empirical regularity that motivated this project — that congressional districts with longer television exposure produced more legislatively effective representatives — survives in point estimate, but it demands a more guarded reading than the premium results above, and I give it one. The within-district effect of cumulative exposure on the current representative's Legislative Effectiveness Score is $+0.042$ with a conventional clustered p of 0.020. Unlike the majority premium, this level effect is identified from cross-state variation in exposure, and its inference is sensitive exactly where that identification is thin: clustering at the state level moves p to 0.060, a wild cluster bootstrap at the state level to 0.100, the design-based permutation test gives 0.046, and dropping New York alone halves the estimate ($+0.021$, $p = 0.20$). Under a population-weighted construction of the treatment the estimate strengthens ($+0.052$, $p = 0.003$), so the sensitivity concerns precision, not sign. The effectiveness results in this subsection are therefore suggestive, and no headline claim rests on them; the decomposition of the point estimate remains informative about mechanism.

A natural decomposition separates two channels. The "replacement" channel captures changes in district-level effectiveness that arise because television districts elect different legislators. The "adaptation" channel captures changes that arise because the same legislators behave differently — work harder, write more bills, manage them more effectively — under television-induced electoral monitoring. The two channels are separable in the panel by comparing district fixed effects (which capture both channels) to legislator fixed effects (which capture only adaptation, holding the legislator's identity constant). Under district fixed effects, I estimate the $+0.042$ total effect. Under legislator fixed effects, the adaptation effect is $+0.010$ with a standard error of 0.028: small and statistically indistinguishable from zero, though imprecisely estimated. The point-estimate difference, roughly three-quarters of the total effect, is the replacement channel.

The replacement channel is real, but its interior is not what the meritocratic-selection account would predict. That account offers two natural routes into the replacement effect. The first is partisan composition: television shifts districts into majority-party hands, and majority members are more effective on average, so district-level effectiveness rises. A mediation decomposition tests this route directly, on the estimation frame with complete majority-share data (total effect $+0.046$, against $+0.042$ in the full panel): the share of the district-level LES effect running through changes in district majority-party share is *negative*, roughly minus nineteen percent of that total. Television does not raise district LES by shifting districts into majority hands; the partisan-composition route, if anything, works against the total. The second route — within-party substitution of less effective members by more effective ones — is the retention-quality question taken up next.

What does drive the replacement point estimate is quality selection in retention. Television districts retain higher-LES members of both parties (within the majority, $+0.073$, with the same inference caveat, wild-bootstrap $p = 0.14$; within the minority, $+0.034$). Combined with the majority-retention finding from Section 4.2, the majority-side selection produces the district-level LES rise through a two-step process: television districts are more likely to retain majority members, and within the majority pool, the retained members are higher-quality. The second step requires some candidate-quality signal to reach voters, and the voter-level evidence permits one: television depressed candidate-level knowledge, it did not extinguish it, and Section 4.1's extensive-margin pattern shows the remaining engaged respondents answering as accurately as before. Majority status operates as the dominant filter; the residual quality differentiation uses the candidate-level information that remained. Why that differentiation registers more strongly in television districts is not something the design resolves — the estimates are imprecise and carry this subsection's provisional reading — and neither step requires generalized meritocratic selection in the broader candidate pool.

One scope note on the decomposition. DW-NOMINATE scores do not vary within a legislator's career by construction, so the legislator-fixed-effects comparison is estimable only for outcomes that vary within career, like LES; the adaptation estimate is identified rather than mechanically zeroed. Its precision is another matter: with a standard error of 0.028, modest adaptation effects would go undetected, so the small estimate is absence of evidence, not a demonstrated behavioral null.

4.5 Downstream consequence: party discipline tightens as the brand becomes more valuable

A second downstream consequence concerns the ideological behavior of majority-party members. The outcome is a member's distance from the majority-party median on roll-call votes, so a negative coefficient means convergence toward the party position. The interaction between television exposure and majority status on this distance is -0.002

per year at the district-fixed-effects level, numerically identical on both ideal-point systems (Table 4); because DW-NOMINATE scores are career-fixed, this district-level estimate compares members serving under different exposure regimes rather than tracking movement within careers. The DW-NOMINATE-based estimate is robust across inference schemes (wild cluster bootstrap at the state level, $p = 0.015$); the Nokken-Poole-based estimate is directionally and numerically identical but less precise under state-level inference (bootstrap $p = 0.13$), so the cross-sectional discipline claim rests on the DW measure. Majority-party members in television-exposed districts sit closer to the party median than their counterparts in non-exposed districts; minority-party members show no comparable convergence.

This pattern is the natural complement of the brand-amplification mechanism. The party label functions as a collective brand whose value to any individual member depends on its consistency. Snyder and Ting's (2002) informational-brand model predicts that as the brand becomes more electorally valuable, parties gain stronger incentives to enforce brand consistency: through nomination, recruitment, and primary discipline that select more party-loyal majority candidates, and potentially through within-member adaptation as well. Separating the channels forces a division of labor: career-fixed DW-NOMINATE carries the cross-sectional claim but cannot enter a within-legislator comparison; congress-varying Nokken-Poole scores can. On the Nokken-Poole measure, the district-fixed-effects coefficient attenuates from -0.0021 to -0.0013 and loses significance under legislator fixed effects in the switcher subsample ($p = 0.18$). The attenuation points to the selection channel — the convergence reflects who serves rather than movement by sitting members — though the switcher subsample is smaller, so adaptation effects below its resolution are not excluded. On this reading, television made the majority-party label more valuable, and the party-discipline tightening that followed operated mainly through *who served* rather than through how the same incumbents voted.

The ideology results in this paper exhibit the same selection-versus-adaptation split, with a twist. At the cross-sectional district-fixed-effects level using DW-NOMINATE (career-fixed ideal points), television exposure is associated with majority Democrats located closer to the chamber center (-0.008 for majority Democrats specifically, on absolute distance from zero) — television districts elect more centrist Democrats. At the within-legislator level using Nokken-Poole congress-varying ideal points, the same Democrats who serve under more cumulative television exposure shift modestly leftward (-0.009 per year). The two channels run in different directions: the replacement channel substitutes more centrist Democrats into majority service, while the surviving Democrats' own positions shift modestly leftward over their careers. The leftward shift is a movement in signed position, not demonstrated convergence to the party median: the direct within-legislator test of distance from the median is null (-0.001, $p = 0.18$; Table 4). The discipline tightening therefore rests on the selection channel, consistent with this subsection's central finding; the within-member drift stands as a descriptive companion rather than a second alignment channel.

The party-discipline tightening complicates simple narratives about television's democratic effects, and its selection character shapes both readings. From one angle, it is a kind of accountability: voters' brand attribution raised the electoral value of a consistent party position, and the majority roster came to consist of members who supplied it — not because sitting members moved toward the party position, but because those who served under television increasingly had voting records that matched the brand. From another, the same process constrains representation: where a district's preferences diverge from the party position, the seat is increasingly held by a member aligned with the party rather than the district. The mechanism that enables collective-responsibility attribution also narrows, seat by seat, the dimensional flexibility of individual representation.

Table 4 collects the downstream-consequence estimates.

Table 4: Downstream consequences: legislative effectiveness and party discipline

Outcome	Coef.	SE	p
Legislative effectiveness (LES, normalized to congress mean)			
TV $\rightarrow$ district LES (district FE, total effect)	$+0.042^{**}$	0.018	0.020
— design-based permutation p (10,000 state-level perms)	—	—	0.046
TV $\rightarrow$ within-legislator LES (legislator FE, adaptation)	$+0.010$	0.028	0.727
implied replacement share	77%	—	—
TV $\rightarrow$ majority-share mediation channel	$\sim -19\%$	—	—
TV $\rightarrow$ within-majority LES (selection within majority)	$+0.073^{**}$	0.036	0.045
Party discipline (distance from majority-party median)			
TV $\times$ majority $\rightarrow$ loyalty (DW-NOMINATE, district FE)	-0.0021^{***}	0.0006	< 0.001
TV $\times$ majority $\rightarrow$ loyalty (Nokken-Poole, district FE)	-0.0021^{**}	0.0007	0.002
TV $\times$ majority $\rightarrow$ loyalty (Nokken-Poole, legislator FE, switchers)	-0.0013	0.0009	0.177
Democrat ideology shift (selection vs adaptation)			
TV $\rightarrow$ DW-NOM Dems (district FE, replacement-side)	$+0.0073^{**}$	0.0031	0.021
TV $\rightarrow$ Nokken-Poole Dems (legislator FE, within-Dem adaptation)	-0.0094^{**}	0.0039	0.015
TV $\rightarrow$ majority Democrat |DW-NOM| (district FE)	-0.0080^{**}	0.0033	0.014

LES regressions: non-South legislator-congress panel, district or legislator FE as indicated, district-by-decade clustered SEs. The district-FE LES total is sensitive to state-level inference (state-clustered $p = 0.060$; wild cluster bootstrap $p = 0.100$; leave-New-York-out $+0.021$, $p = 0.20$) and is presented as provisional; the design-based permutation p uses the state-mean-shift design of Appendix A.2. Loyalty and ideology regressions: same panel, district + congress FE. Negative loyalty coefficient indicates convergence toward party median (less deviation); the DW loyalty estimate's wild-bootstrap p is 0.015, the Nokken-Poole estimate's 0.13. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5 Conclusion

Broadcast television in postwar America did not teach voters about their representatives. It taught voters who was in charge. That asymmetric information substitution — Type A knowledge replaced by Type B knowledge — translated into an electoral premium for majority-party members that did not depend on their individual records or their institutional roles. The premium attached to the majority position itself, indifferent to records, to seniority, and to the president's party; and it ran strongest when Congress produced least, a pattern credit claiming cannot generate and collective attribution can.

The most consequential implication concerns the standard account of how information shocks interact with party-cue voting. A long tradition treats the party label as a cognitive shortcut that voters rely on when candidate-level information is scarce (Popkin, 1991), and a related tradition emphasizes the label as a collective informational asset whose value depends on the party's institutional position (Snyder and Ting, 2002; Grynaviski, 2010). A common reading of the first strand holds that information shocks weaken voters' reliance on the label by supplying the candidate-level information for which the label substitutes. My findings imply that this reading is correctly understood as a statement about a specific kind of information shock. The 1948–1964 television rollout was a powerful information shock that did not increase candidate-level information — it reduced candidate recall and incumbency identification — while increasing institutional information about national two-party politics. The cue that strengthened was not the cue the standard account addresses. There is no contradiction once Type A and Type B cues are distinguished, but the appearance of contradiction has been a recurring source of confusion in empirical work that finds nationalization or brand strengthening in periods of expanding media access (Hopkins, 2018; Moskowitz, 2021). The familiar intuition that more candidate-level information weakens party voting is not wrong; it is incomplete. Information shocks erode party heuristics when they target the candidate-level inputs that party heuristics substitute for. They activate party heuristics when they supply the institutional inputs those heuristics require.

This reading also clarifies where my findings sit relative to Achen and Bartels (2016)'s critique of the heuristics-rescue tradition. The candidate-inference task that the Zaller-Lupia-Popkin tradition emphasized is the task Achen and Bartels (2016) argue voters most clearly fail at, and my data are consistent with that diagnosis: the television return premium does not scale with individual productivity, and the least productive majority members benefit from television at least as much as the most productive. What survives in my setting is a thinner capacity — collective-responsibility attribution, closer in spirit to Fiorina (1981)'s retrospective voting — that operates on the coarse institutional fact of which party holds the gavel rather than on what any individual representative has done. Voters appear to perform collective attribution even where they fail at candidate-level inference.

The result is a paradox of accountability without selection. The conventional intuition holds that media exposure improves democratic accountability through individual selection: voters identify good representatives, retain them, and replace bad ones. Television did increase a kind of accountability. But the reward it created attached to majority status without regard to records: the least productive majority members gained at least as much as the most productive, while minority members — however effective — show no comparable gain, falling behind in relative terms. It produced sharper party brands, tighter party discipline, and stronger collective electoral consequences for governing-party performance, while leaving the already-weak connection between an individual representative's record and her electoral fate unimproved. Whether this is good or bad depends on one's normative theory of representation: advocates of responsible party government will see health, defenders of personal-vote accountability will see loss. The trade-off itself, however, has been largely missed in empirical media-and-democracy work that treats accountability as a unitary good.

The historical setting has acquired renewed contemporary relevance. The decline of local newspapers and the rise of nationally oriented cable and digital media closely resemble the Type A erosion and Type B amplification I document — the same substitution, at greater intensity (Hopkins, 2018). To the extent the mechanism generalizes, the present period should exhibit stronger majority-party premia, tighter party discipline, and weaker candidate-level voter information; nationalized partisan voting (Jacobson, 2015), stronger party-line voting in Congress (Lee, 2016), and declining knowledge of local representatives (Hayes and Lawless, 2021) all point in the predicted direction. Debates about media's democratic role often present "more information" as straightforwardly good. The relevant question is not whether a medium increases information but what kind of information it increases: the medium's structural orientation, not the quantity of political content it produces, determines which side of the accountability trade-off it tilts toward.

Four limits bound these claims. The aggregate majority-party premium bundles closely related channels — brand valuation, agenda-control awareness, leadership visibility, and credit-blame attribution — that my data cannot separate, though all live within the Type B family the theory addresses. The American National Election Studies did not ask which party controls Congress before 1968, so my Type B knowledge measures are indirect proxies for institutional knowledge. The premium itself survives absorbing all state-by-congress shocks, but the voter-level results and the suggestive effectiveness results have no such protection: they rest on cross-state variation in television timing, and unobserved state-level confounds correlated with adoption timing remain a residual concern for those strands. And the 1948–1964 broadcast setting is one information environment: the attenuation of the premium as television aged implies that magnitudes depend on a medium's novelty as well as its structure, the treatment's construction as an average over heterogeneous county markets means measurement shapes them too, and generalization to other media and other legislative systems requires care.

The mechanism I describe likely operates wherever a new medium nationalizes the information environment more than it localizes it, and wherever it transmits institutional positioning more than candidate evaluation. Broadcast television offers a long-run window into how durably such a shift reshapes representation. The contemporary moment, in which local newspapers decline as national platforms expand, suggests the question is not closed.

References

Achen, C. H. and Bartels, L. M. (2016). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press, Princeton.

Aldrich, J. H. (1995). *Why Parties? The Origin and Transformation of Political Parties in America*. University of Chicago Press, Chicago.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training projects. *Journal of the American Statistical Association*, 103(484):1481–1495.

Ansolabehere, S. and Snyder, Jr., J. M. (2002). The incumbency advantage in U.S. elections: An analysis of state and federal offices, 1942-2000. *Election Law Journal*, 1(3):315–338.

Campante, F. R. and Hojman, D. A. (2013). Media and polarization: Evidence from the introduction of broadcast TV in the United States. *Journal of Public Economics*, 100:79–92.

Carmines, E. G. and Stimson, J. A. (1989). *Issue Evolution: Race and the Transformation of American Politics*. Princeton University Press, Princeton, NJ.

Chiou, F.-Y. and Goplerud, M. (2024). Effective lawmaking across congressional eras. *Journal of Politics*, 86(3):997–1012.

Clubb, J. M., Flanigan, W. H., and Zingale, N. H. (2008). Electoral data for counties in the United States: Presidential and congressional races, 1840–1972. ICPSR Study #08611 [distributor]. County-level historical House election returns, 1840–1972.

Cox, G. W. and Katz, J. N. (1996). Why did the incumbency advantage in U.S. House elections grow? *American Journal of Political Science*, 40(2):478–497.

Cox, G. W. and McCubbins, M. D. (1993). *Legislative Leviathan: Party Government in the House*. University of California Press, Berkeley.

Cox, G. W. and McCubbins, M. D. (2005). *Setting the Agenda: Responsible Party Government in the U.S. House of Representatives*. Cambridge University Press, Cambridge.

Darr, J. P., Hitt, M. P., and Dunaway, J. L. (2018). Newspaper closures polarize voting behavior. *Journal of Communication*, 68(6):1007–1028.

Fenton, G. and Koenig, F. (2025). Labor supply and entertainment innovations: Evidence from the U.S. TV rollout. *American Economic Journal: Applied Economics*, 17(4):1–28.

Fiorina, M. P. (1981). *Retrospective Voting in American National Elections*. Yale University Press, New Haven, CT.

Gelman, A. and King, G. (1990). Estimating incumbency advantage without bias. *American Journal of Political Science*, 34(4):1142–1164.

Gentzkow, M. (2006). Television and voter turnout. *Quarterly Journal of Economics*, 121(3):931–972.

Gentzkow, M. and Shapiro, J. M. (2008). Introduction of television to the United States media market, 1946–1960. ICPSR Study #22720 [distributor]. TV introduction years and DMA boundaries used for years_of_tv construction.

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

Grynaviski, J. D. (2010). *Partisan Bonds: Political Reputations and Legislative Accountability*. Cambridge University Press, Cambridge.

Hayes, D. and Lawless, J. L. (2021). *News Hole: The Demise of Local Journalism and Political Engagement*. Cambridge University Press, Cambridge.

Hopkins, D. J. (2018). *The Increasingly United States: How and Why American Political Behavior Nationalized*. University of Chicago Press, Chicago.

Jacobson, G. C. (2013). U.S. House election returns, 1946–2012. Dataset. District-level two-party vote tallies for U.S. House general elections, compiled by Gary C. Jacobson.

Jacobson, G. C. (2015). It's nothing personal: The decline of the incumbency advantage in US House elections. *Journal of Politics*, 77(3):861–873.

Lee, F. E. (2016). *Insecure Majorities: Congress and the Perpetual Campaign*. University of Chicago Press, Chicago.

Lenz, G. S. (2012). *Follow the Leader? How Voters Respond to Politicians' Policies and Performance*. University of Chicago Press, Chicago.

Lewis, J. B., Poole, K., Rosenthal, H., Boche, A., Rudkin, A., and Sonnet, L. (2024). Voteview: Congressional roll-call votes database. `https://voteview.com/`. Data accessed for House congresses 78–89; ideology and roster.

Lupia, A. (1994). Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections. *American Political Science Review*, 88(1):63–76.

Mann, T. E. (1978). *Unsafe at Any Margin: Interpreting Congressional Elections*. American Enterprise Institute, Washington, DC.

Martin, G. J. and McCrain, J. (2019). Local news and national politics. *American Political Science Review*, 113(2):372–384.

Moskowitz, D. J. (2021). Local news, information, and the nationalization of U.S. elections. *American Political Science Review*, 115(1):114–129.

Ornstein, N. J., Mann, T. E., and Malbin, M. J. (2013). *Vital Statistics on Congress*. Brookings Institution Press, Washington, DC. Table 2-16, Ticket Splitting between Presidential and House Candidates, 1900–2012.

Popkin, S. L. (1991). *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. University of Chicago Press, Chicago.

Prat, A. and Strömberg, D. (2013). The political economy of mass media. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics: Tenth World Congress*, volume 2, pages 135–187. Cambridge University Press, Cambridge.

Prior, M. (2006). The incumbent in the living room: The rise of television and the incumbency advantage in U.S. House elections. *Journal of Politics*, 68(3):657–673.

Ranney, A. (1983). *Channels of Power: The Impact of Television on American Politics*. Basic Books, New York.

Snyder, Jr., J. M. and Strömberg, D. (2010). Press coverage and political accountability. *Journal of Political Economy*, 118(2):355–408.

Snyder, Jr., J. M. and Ting, M. M. (2002). An informational rationale for political parties. *American Journal of Political Science*, 46(1):90–110.

Song, B. (2022). Television and economic voting in US presidential elections. *British Journal of Political Science*, 52(4):1622–1634.

Stokes, D. E. and Miller, W. E. (1962). Party government and the saliency of Congress. *Public Opinion Quarterly*, 26(4):531–546.

Strömberg, D. (2004). Mass media competition, political competition, and public policy. *Review of Economic Studies*, 71(1):265–284.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Tamas, B. (2022). U.S. House election results, 1840–1974. District-level returns by candidate, V1; the 1920–1974 span collected from the Clerk of the House of Representatives election statistics.

Volden, C. and Wiseman, A. E. (2014). *Legislative Effectiveness in the United States Congress: The Lawmakers*. Cambridge University Press, Cambridge.

Wattenberg, M. P. (1991). *The Rise of Candidate-Centered Politics: Presidential Elections of the 1980s*. Harvard University Press, Cambridge, MA.

Zaller, J. R. (1992). *The Nature and Origins of Mass Opinion*. Cambridge University Press, Cambridge.

A Appendix A: Identification

A.1 Pre-television placebo

The pre-television balance test asks whether the eventual timing of television's arrival was already associated with electoral outcomes before television existed. In the pre-television Congresses (the 78th and 79th, 1943–1946; $N = 680$), I regress incumbent return on the district's eventual TV-introduction year interacted with majority status, with congress fixed effects and state-clustered standard errors. The interaction is -0.014 (SE 0.014, $p = 0.32$). Interpreting the test requires care with signs: the regressor is the introduction year, so a pre-existing tilt capable of manufacturing the premium (early-television districts already favoring majority members) would appear as a *negative* coefficient. The point estimate sits on that side but is one standard error from zero, and the minimum detectable effect at 80 percent power is 0.039 — about two and a half times the headline in introduction-year units — so the test neither detects nor rules out a differential of the headline's size. I accordingly present it as a null, low-powered check, and rest the identification on the specifications with power against confounding (the state $\times$ congress absorption and the event study of Section 3.4). Implementation: scripts `17_remaining_gaps.R` and `29_audit_placebo_xsec.R`.

A.2 Randomization inference

For the legislative-effectiveness result, I report a design-based permutation test that respects the state-level assignment of introduction timing. Each of 10,000 permutations reassigns the states' mean introduction years across states and shifts every district's introduction year by the difference between its assigned and actual state mean, so the within-state timing structure is preserved and, on the effectiveness estimation sample, the identity assignment reproduces the observed data exactly. The resulting null distribution has standard deviation 0.022 — 1.2 times the conventional clustered standard error of 0.018, the shape a design-based null should have — and the two-sided p is 0.046 (Monte Carlo 95 percent band $[0.042, 0.050]$), against a conventional clustered 0.020. The design-based test therefore places the effectiveness result at the boundary of conventional significance, which is how the main text characterizes it. Applied to the majority-premium interaction, the design is uninformative by construction: the reassignment preserves the within-state contrasts that identify the interaction, so the null distribution collapses to a standard deviation of 0.0007 — one-thirtieth of the clustered standard error — around a mean near the observed coefficient. The nominal two-sided p is 0.15, and neither that number nor the near-invariance carries evidential content about the premium in either direction; the premium's independence from cross-state assignment is shown by the state $\times$ congress fixed-effects specification instead. (On the premium frame, four at-large rows carry a builder fallback fill — Appendix D.5 — so the identity check there holds to within 7×10^{-7} of the coefficient rather than exactly.) Implementation: script `07_randomization_inference.R`.

The headline event study's pre-trend verdict is robust to the event-time binning. Under three binning schemes — floor-based bins clipped to $[-6, 8]$, threshold bins, and the published rounded bins clipped to $[-5, 5]$ — the pre-period joint test returns $p = 0.93$, 0.93, and 0.54 respectively, and the terminal post-arrival coefficient is $+0.23$, $+0.22$, and $+0.20$: every binning passes, and the published one is the most conservative. Implementation: script `31_audit_es_binning.R`.

Figure A.1 shows the corresponding event study for the downstream legislative-effectiveness (LES) outcome. Unlike the majority-premium event study in the main text, its pre-trend verdict is sensitive to how event time is binned: under the reference binning the district-fixed-effects pre-trends are jointly insignificant ($F = 1.48$, $p = 0.19$) and the legislator-fixed-effects pre-trends pass comfortably, while an alternative binning yields a marginal district-FE pre-trend ($F = 2.84$, $p = 0.04$). Because the LES result is a downstream consequence rather than the paper's headline, and because the headline majority-premium event study passes its pre-trend test under every binning, I treat the LES event study as illustrative and lean the identification argument on the battery of Section 3.4. Implementation: script `04_event_study_verification.R`.

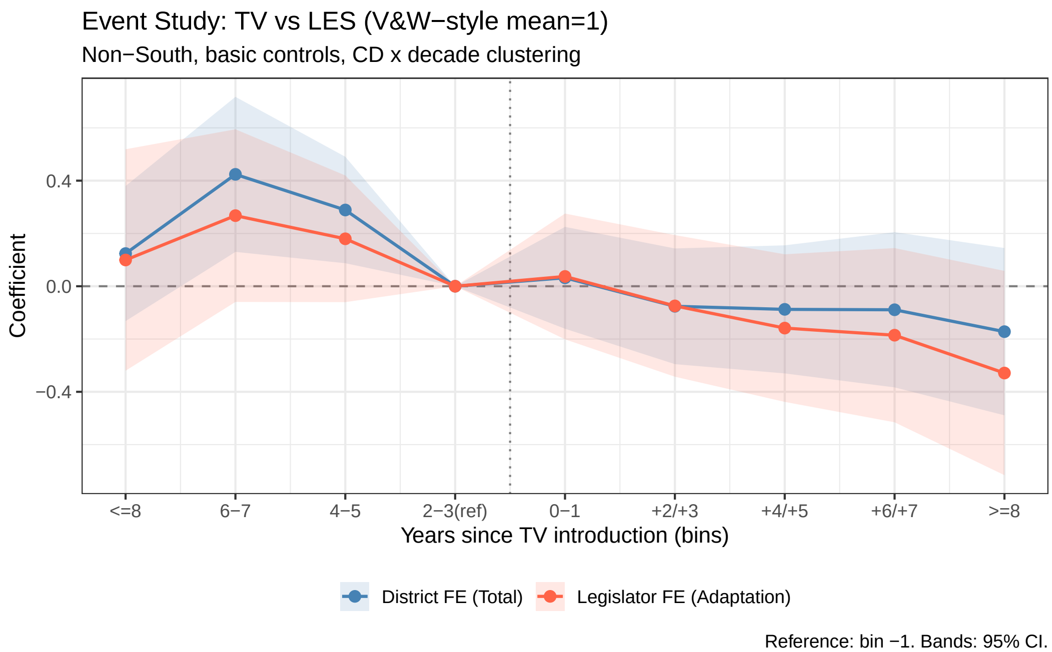

Figure A1: Event study of cumulative television exposure on the Legislative Effectiveness Score (downstream outcome). District- and congress-fixed-effects specification; reference period is the congress before television arrives. Pre-trend tests are reported in the text and are sensitive to event-time binning.

A.3 Goodman-Bacon and Sun-Abraham diagnostics

I attempted Goodman-Bacon (2021) decompositions and Sun and Abraham (2021) interaction-weighted estimators, both designed for binary absorbing treatments rather than the continuous exposure I study. Coarsening exposure to a binary television-presence indicator (any exposure versus none) produces degenerate treatment-timing groups under the compressed postwar rollout: the Goodman-Bacon and Callaway–Sant'Anna routines fail to return an estimate, and the Sun-Abraham interaction-weighted estimator yields an insignificant aggregate ATT (-0.05, $p = 0.72$; the binary two-way fixed-effects analogue is -0.08, $p = 0.45$). This is expected: a binary indicator discards the exposure-intensity variation that identifies my continuous-treatment effect. Details: script `06_staggered_did.R`.

B Appendix B: ANES Type A vs Type B — full specifications

B.1 Robustness across specifications

For each of the six knowledge outcomes, I report estimates under four specifications: (a) bivariate regression on years of TV exposure; (b) plus year fixed effects; (c) plus year fixed effects and demographic controls (preferred); (d) plus state fixed effects (where identification permits).

Specification (d) drops most outcomes because TV penetration was state-level and saturated by 1956; with state fixed effects, years of exposure becomes nearly collinear with year fixed effects. Of the six outcomes, only "cares about presidential outcome" and "perceives party difference" retain identifying variation under state FE; both become statistically indistinguishable from zero, with standard errors an order of magnitude larger than in the preferred specification (cares-about-outcome $+0.007$, SE 0.049; party difference -0.044, SE 0.045), consistent with state fixed effects absorbing the cross-state identifying variation.

B.2 Within-1964 cross-section

The 1964 ANES is the only year in my window with both Type A and Type B knowledge questions. Restricting to 1964 alone, I estimate cross-state effects of cumulative TV exposure on each outcome:

Outcome	Coef.	SE	p	N
Type A: recall House cand name	-0.0225	0.0126	0.074	1,350
Type A: knows if cand is incumbent	-0.0360	0.0128	0.005	1,370
Type A: correctly IDs incumbent	-0.0122	0.0032	< 0.001	947
Type B: perceives party difference	$+0.0094$	0.0085	0.270	1,382
Type B: knows cons party	$+0.0034$	0.0053	0.515	1,371
Type B: cares about pres outcome	$+0.0079$	0.0073	0.278	1,479
First stage: watched political TV	-0.0001	0.0058	0.993	1,381

The within-1964 first-stage on watched-TV is null because by 1964 nearly all respondents in TV-exposed states were watching political TV programs at saturation rates — the cross-state variation in cumulative exposure is no longer driving variation in current viewing intensity. The Type A and Type B effects retain their asymmetric pattern.

B.3 Decision-basis interactions

Section 4.1's decision-basis test asks whether TV exposure changed how voters used each kind of knowledge. For each of the six knowledge measures, I estimate a linear probability model of party-line House voting (the respondent's House choice matching her own partisan identification) on the measure, cumulative TV exposure, their interaction, demographic controls, and year fixed effects, with state-clustered standard errors. The table reports the interaction coefficients:

Knowledge measure × TV exposure	Coef.	SE	p	N
Type A: recall House cand name	-0.0014	0.0088	0.875	611
Type A: knows if cand is incumbent	-0.0231	0.0119	0.052	607
Type A: correctly IDs incumbent	-0.0253	0.0360	0.483	458
Type B: perceives party difference	-0.0032	0.0047	0.497	849
Type B: knows cons party	-0.0201	0.0176	0.253	555
Type B: cares about pres outcome	$+0.0141$	0.0065	0.030	1,159

Each row is a separate regression; samples vary with the availability of the knowledge item and a major-party House vote. Implementation: script `23_anes_decision_basis.R`.

C Appendix C: Mechanism — additional detail

C.1 Marginal effects table (4 quartiles by 2 majority states)

The seniority-by-majority interaction in the chair-amplification test implies a 4×2 matrix of marginal TV effects. Computed from the full triple-interaction specification:

	Q1 (junior)	Q2	Q3	Q4 (senior)
Minority	$+0.011$ (0.006)	$+0.000$ (0.006)	-0.004 (0.007)	-0.003 (0.007)
Majority	$+0.019$ (0.006)	$+0.021$ (0.006)	$+0.018$ (0.007)	$+0.012$ (0.006)

Cell entries: marginal TV effect on incumbent return per year of cumulative exposure (standard error in parentheses). Seniority is career House tenure from the full Voteview roster (Appendix E.5); no minority cell is significant at the 5 percent level. Implementation: script `19_majority_subgroup.R`.

C.2 Open-seat dynamics

A potential concern with the within-legislator switch test is that majority members who lose the majority position may be on a different selection process than incumbents who run continuously. I address this concern with separate analysis of open seats — districts where the previous incumbent did not run for reelection.

Among 1,046 open-seat observations (543 following a majority-party departure, 503 following a minority-party departure), television preserves the departing party's hold on the seat. TV exposure reduces the probability that an open seat switches party by 3.4 percentage points per year ($p = 0.004$; wild cluster bootstrap at the state level $p = 0.022$). Conditioning on which party departed: after a majority-party departure, TV exposure *raises* the probability that the new member is again from the majority party ($+0.036$, $p = 0.028$, $N = 543$); after a minority departure the effect is null (-0.015, $p = 0.35$, $N = 503$), as is the pooled open-seat effect on majority-party occupancy ($p = 0.52$). The open-seat channel therefore works in the same direction as the incumbent-retention channel — the majority label holds seats even when the person departs — which is what a label mechanism, and not a person-specific visibility mechanism, predicts.

This open-seat analysis was also designed to assess whether the small district-level negative association between TV and holding a majority seat (-0.015 per year, $p = 0.008$) reflects systematic flipping of majority-held seats to minority parties at moments of incumbent departure. It does not: at open seats, majority-departure districts are *more* likely to stay majority under TV exposure, not less. The negative district-level coefficient instead reflects the mechanical interaction of incumbent persistence with national majority transitions during 1948–1964 — the same persistent incumbent's majority-status variable flips as the national majority changes across the 80th, 81st, 83rd, and 84th Congresses. I report this analysis for transparency but exclude it from the main text. Implementation: script `20_open_seat_retirement.R`.

C.3 Vote share by competitiveness

Splitting the 78–89 vote-share sample by competitiveness — "competitive" meaning own-party vote share between 0.40 and 0.65, which covers 69.1% of sample rows (2,716 of 3,928) — the TV $\times$ majority interaction is $+0.00351$ per year in competitive districts (SE 0.00046, $p < 0.001$, $N = 2{,}701$) and $+0.00419$ per year in safe districts (SE 0.00225, $p = 0.064$, $N = 1{,}150$); reported N is the estimation sample after fixed-effect singleton removal, the same counting rule as Table 2.

I report this split for completeness and draw no directional conclusion from it, for two reasons. First, the ordering of the two estimates depends on the sample window: on the 79–89 window available without the 1944 returns (Appendix D.7), the competitive estimate is the larger ($+0.00460$, $p < 0.001$, $N = 2{,}455$, against $+0.00403$, $p = 0.092$, $N = 1{,}074$), while on 78–89 the safe estimate is the larger; the difference between the two estimates is not formally tested in either window. Second, the split conditions on the dependent variable: membership in the competitive subsample is defined by own-party vote share itself, so the competitive regression is estimated on an outcome truncated to the 0.40–0.65 band. The truncation is visible in the standard errors — the competitive subsample's standard error is about one fifth of the safe subsample's (0.00046 against 0.00225) in both windows. Implementation: script `14_vote_margin.R`.

C.4 Within-majority quality selection

Among majority-party legislators specifically, the within-majority effect of TV exposure on individual LES is $+0.073$ ($p = 0.045$, $N = 2{,}024$; wild cluster bootstrap at the state level $p = 0.14$, so this estimate carries the same provisional reading as the LES results in the main text); the within-minority counterpart is $+0.034$ ($p = 0.006$, $N = 2{,}012$). Television districts therefore retain higher-LES members of both parties; the majority-side selection is what combines with majority retention to produce the district-level LES rise reported in the main text. I interpret the quality margin as residual local information that survives television's substitution effect — voters who do learn about specific candidates use that information at the margin to retain effective members. Implementation: script `13_majority_mediation.R`.

D Appendix D: Robustness

D.1 Specification curve

I estimate the headline TV × majority interaction on incumbent return across a 72-cell specification grid varying every analytic choice that can move the point estimate: (a) fixed-effects structure (district + congress; district + state × congress; district × apportionment-cycle); (b) controls (none; demographic — % urban + % nonwhite; demographic plus a percent-urban × majority interaction); (c) sample window (congresses 78–89; 78–87; 80–89; 78–82); (d) handling of missing census covariates (recode-to-zero, as in production; complete case). Clustering is deliberately not a grid dimension — it cannot move a point estimate — and is examined separately in Appendix D.3.

The grid contains 55 distinct point estimates (some cells coincide where a dimension is inert, e.g. the missing-covariate choice under a no-controls specification). The interaction is positive in all 72 cells (median $+0.017$, interquartile range $[+0.015, +0.020]$) and significant at the 5% level in 54 of 72 (75%). Every insignificant cell lies in the 1943–1952 window (congresses 78–82), where mean accumulated exposure is 0.84 years — a near-pre-treatment subsample — and the share of significant cells is 100% in each of the three windows that contain the treated era. The fixed-effects and control dimensions leave the share significant unchanged across their levels. Implementation: script `38_audit_multiverse.R`.

I separately report a specification curve for the downstream LES result (script `02_e2_les.R`), which yields a positive median estimate with 25% of specifications significant at the 5% level, consistent with modest per-spec power. A curve for the uncontrolled TV $\times$ LES interaction (script `03_e3_mechanism.R`) is positive in all specifications and significant in half — and that interaction is fully absorbed once TV $\times$ majority enters the model (Table 3), consistent with majority status rather than productivity carrying the premium.

D.2 By-congress heterogeneity

I estimate the TV $\times$ majority interaction separately for each congress in my window. By-congress estimates exhibit substantial volatility consistent with the small per-congress sample sizes (each congress contributes roughly 330 non-South legislator-congress observations); the pooled triple TV $\times$ majority $\times$ Republican-Congress is statistically indistinguishable from zero ($p = 0.93$). The 80th and 83rd Congresses (the only Republican-majority Congresses in this window) are too few in number, and the per-congress samples too small, to support a powered separate test of cross-party generalizability (see Section 4.2 of the main text and the R-majority subsample row in Table 2 of the main text).

D.3 Alternative clustering

I report the headline TV $\times$ majority interaction under four clustering schemes: CD$\times$decade (preferred), state, district, and two-way (state and congress). Under the primary specification (non-South, district + congress fixed effects, demographic controls), the point estimate is $+0.0156$ in all four schemes; the conventional p-value is below 0.001 under CD$\times$decade, state, and district clustering, and is 0.048 under two-way clustering —

significant at the 5% level in all four. Because treatment timing is assigned at the state level and the sample contains only 37 states, I also report a wild cluster bootstrap at the state level (Rademacher weights, 9,999 draws): the headline's bootstrap p is below 0.0001, and the low-productivity triple's is 0.001. Influence diagnostics point the same way: across thirty-seven leave-one-state-out re-estimates the headline ranges from $+0.0147$ to $+0.0168$ with a minimum $|t|$ of 5.9, and across twelve leave-one-congress-out re-estimates it remains significant in every fit. Implementation: scripts `34_audit_wildboot.R` and `33_audit_loso_loco.R`.

D.4 Senate replication

I replicate the headline TV $\times$ majority interaction in the U.S. Senate over the same 78–89 Congress window. The Senate analysis sample contains 942 senator-Congress observations covering 242 senators in 37 non-South states. The regression sample is 788 observations: 154 observations — all of Delaware, Nevada, New Mexico, Vermont, and Wyoming, plus parts of Arizona and North Dakota — drop because their state-aggregated census controls are missing in every congress (these are the single- and two-district states whose district-level census covariates are missing throughout, so the state aggregate is undefined). Applying the House pipeline's convention for missing census covariates instead retains all 37 states and gives $+0.014$ ($p = 0.003$, $N = 942$) — the same conclusion on the full sample. State-level cumulative TV exposure is the population-weighted aggregate of the district-level measure used in the House analysis. Standard errors are clustered at the state level (statewide elections render district-level clustering moot).

The headline result ($+0.017$, $p = 0.003$) is essentially identical in magnitude to the House headline ($+0.016$, $p < 0.001$). The Senate R-majority subset point estimate ($+0.018$, $p = 0.53$, $N = 129$) is positive and points in the same direction as the full-sample interaction, although underpowered. Combined with the House R-majority subset's near-zero estimate (-0.001, $p = 0.97$, $N = 674$), the cross-chamber evidence is consistent with a

Table A1: Senate replication of the majority-party premium

Test	Estimate	SE	p	N
Headline interaction				
TV × majority on incumbent return (state FE)	$+0.017^{***}$	0.005	0.003	788
Cross-party generalizability				
R-majority subset (Cong 80, 83)	$+0.018$	0.029	0.53	129
D-majority subset (Cong 78–89, ex 80, 83)	$+0.015^{**}$	0.006	0.021	659
Triple TV × maj × Republican-Congress	$+0.009$	0.027	0.75	788
Mechanism tests				
Joint spec.: TV × majority (with TV × Democrat)	$+0.024^{*}$	0.014	0.085	788
Joint spec.: TV × Democrat (with TV × majority)	-0.009	0.015	0.58	788
TV × majority on return (controlling LES)	$+0.018^{**}$	0.005	0.002	788
TV × LES on return (controlling majority)	-0.002	0.002	0.39	788
Triple TV × majority × LES_top (binary)	-0.013	0.010	0.18	788
Within-legislator				
TV × majority (legislator FE, full sample)	$+0.008$	0.008	0.28	739
TV × majority (legislator FE, switcher subsample)	$+0.012$	0.008	0.13	547
Party discipline				
TV × maj on |dist from party median| (DW, state FE)	-0.0038	0.0023	0.12	787
TV × maj on |dist from party median| (NP, state FE)	-0.0048^{*}	0.0025	0.066	786
TV × maj on |dist from party median| (NP, legislator FE)	-0.0001	0.0028	0.98	737

Senate sample: 78–89 Congress, 37 non-South states (32 with complete state-aggregated census controls; see text for the five-state drop and the all-37-state robustness estimate). State + congress fixed effects; SE clustered at state level. "Switcher subsample" restricts to the 114 senators observed under both majority and minority status (vs. 463 House switchers). Compare House Tables 2, 3, and 4 of the main text. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Implementation: script `22_senate_replication.R`.

uniform effect across both parties' majorities, though both subsets are failures to reject rather than affirmative demonstrations — the cross-party test is simply underpowered in the available 78–89 window. The triple TV $\times$ majority $\times$ Republican-Congress is null in both chambers (House $p = 0.93$; Senate $p = 0.75$).

The joint specification shows the same pattern as in the House: TV $\times$ majority survives controlling for TV $\times$ Democrat ($+0.024$, $p = 0.085$); TV $\times$ Democrat is null and slightly negative (-0.009, $p = 0.58$). The TV $\times$ LES interaction is null after controlling for majority (-0.002, $p = 0.39$), and the triple TV $\times$ majority $\times$ LES is null ($p = 0.18$) — the same merit-based-selection rejection as in the House.

The party-discipline pattern follows the selection-versus-adaptation split documented for the House. Cross-section party loyalty point estimates are same-signed and larger than the House's under both DW-NOMINATE (-0.0038, $p = 0.12$) and Nokken-Poole (-0.0048, $p = 0.066$) measures of distance from the party median, though only marginally significant given the smaller Senate sample. Under legislator fixed effects on Nokken-Poole, the within-legislator effect is statistically indistinguishable from zero (-0.0001, $p = 0.98$), matching the House finding that the discipline tightening operates through who serves rather than how the same incumbents vote.

The within-legislator FE estimate on incumbent return is statistically insignificant ($+0.008$, $p = 0.28$). Two Senate-specific limits explain the lower power compared to the House. First, six-year Senate terms mean that approximately two-thirds of senator-Congress observations are mid-term, and the binary outcome is mechanically equal to one for those observations (overall sample mean is 0.81). Second, only 114 senators are observed in both majority and minority status during the window (vs. 463 House members), and the switcher-subsample legislator-FE estimate ($+0.012$, $p = 0.13$) is positive but underpowered. The headline state-FE result, however, is essentially identical in magnitude to the House estimate, and the cross-section party-loyalty estimates are same-signed though only marginally significant.

The Senate replication therefore provides cross-chamber consistency on three fronts, estimated on one independent sample: (i) the headline majority-premium magnitude is essentially identical; (ii) the merit-based-selection alternative is rejected the same way; (iii) the party-discipline channel shows the same selection-not-adaptation signature, with the cross-section estimates marginal in the smaller Senate sample. The R-majority point estimate further alleviates the cross-party-uniformity concern that the House subsample alone could not resolve. Implementation: script `22_senate_replication.R`.

D.5 Metropolitan-tide confound and redistricting

Two threats specific to the district-panel design deserve separate treatment.

Metropolitan tide-sensitivity. Early-television districts were disproportionately pre-freeze metropolitan markets, and urban districts may respond more strongly to national partisan tides for reasons unrelated to television. Because the within-window majority variation is driven by the wave elections of 1946–1954, such differential tide-sensitivity could in principle generate the TV $\times$ majority interaction spuriously. I test this by allowing the majority premium to vary with district urbanization (a percent-urban $\times$ majority interaction) and by letting urban districts follow their own congress-specific trajectory (a per-congress urban slope). The interaction is unmoved:

Specification	TV $\times$ majority	SE	p	N
Baseline (Table 2)	$+0.0156$	0.0024	< 0.001	4,031
$+$ percent-urban $\times$ majority	$+0.0156$	0.0025	< 0.001	4,031
$+$ per-congress urban slope	$+0.0149$	0.0024	< 0.001	4,031
$+$ both, and percent-nonwhite $\times$ majority	$+0.0149$	0.0025	< 0.001	4,031

Redistricting. The panel spans the reapportionments that followed the 1950 and 1960 censuses; a district's identifier can persist across a reapportionment even when its geography changes, so a naïve district fixed effect could compare non-comparable districts. I neutralize this by interacting the district fixed effect with the apportionment cycle (1940-census districts, congresses 78–82; 1950-census districts, 83–87; 1960-census and court-redrawn districts, 88–89), so that a district is never compared to its post-redistricting self. The interaction is $+0.019$ ($p < 0.001$, $N = 4{,}030$), if anything larger than the pooled estimate. Restricting instead to a single apportionment cycle attenuates the point estimate — $+0.013$ ($p = 0.38$) in the 1940-census cycle (congresses 78–82) and $+0.008$ ($p = 0.21$) in the 1950-census cycle (83–87). These single-cycle windows are individually underpowered for a reason the data make explicit: each spans only four to five congresses, and the earliest window is nearly pre-treatment (mean accumulated exposure 0.84 years). The dispersion of introduction timing itself is stable across every window (SD of the district introduction year 2.66–2.72 throughout), so the attenuation reflects exposure levels and window length, not vanishing identifying variation. The pattern is consistent with the novelty attenuation documented in the main text, not with a redistricting artifact. Implementation: script `24_headline_robustness.R`.

State-by-congress fixed effects. The decisive specification against state-year confounding absorbs district and state $\times$ congress fixed effects together: the interaction is $+0.0145$ (SE 0.0025 clustered CD$\times$decade, 0.0027 clustered by state; $N = 3{,}983$). A treatment-variance decomposition shows why the specification is demanding and why its survival is informative: state $\times$ congress cells absorb 93.1 percent of the variance in cumulative exposure (district + congress fixed effects absorb 96.0 percent; both together, 97.6 percent), so the interaction is identified from the within-state majority-minority contrast and the residual cross-district timing variation. Implementation: script `28_audit_headline_variants.R`.

Missing census covariates. The production pipeline codes districts with missing percent-urban and percent-nonwhite (128 of the 4,036 headline-frame rows; 142 in the full panel) to zero. Dropping those rows instead leaves the headline at $+0.0157$ ($N = 3{,}906$) and the three metropolitan-tide specifications between $+0.0142$ and $+0.0155$, all $p < 0.001$. Implementation: script `28_audit_headline_variants.R`.

Treatment construction. The exposure measure averages county introduction years with county-part allocation shares (Appendix E.1). Reweighting counties by their 1940 populations, or by census-vintage populations matched to each congress's crosswalk, yields treatments correlated 0.999 with production and moves the headline to $+0.01566$ and $+0.01563$ respectively — the fifth decimal. The legislative-effectiveness level effect strengthens under population weighting ($+0.052$, $p = 0.003$). Four of 4,036 rows carry a stored exposure value that deviates from the construction formula (a state-level fallback fill for at-large districts; maximum deviation 0.26 years); recomputing exposure from the formula moves the headline by 7×10^{-7}. Implementation: scripts `40_audit_alt_treatment.R` and `37_audit_treatment_integrity.R`.

Functional form. The outcome is binary, so I re-estimate the headline interaction by logit and probit with the same district and congress fixed effects and controls. Read as the difference between majority and minority members in the average marginal effect of a year of exposure (the quantity the linear interaction estimates), the logit estimate is $+0.01734$ and the probit estimate $+0.01689$; the average discrete cross-partial, the interaction effect proper in a nonlinear model, is $+0.01658$ under the logit; and a conditional logit with district strata gives a coefficient of $+0.09551$ ($p = 5.5 \times 10^{-10}$). The linear probability model's $+0.01561$ is therefore mildly conservative. The fixed-effect logit and probit drop 269 observations in 41 districts in which every member-congress returned, because the district fixed effect perfectly predicts the outcome there; the conditional logit keeps every row but draws no information from those districts either, and a logit with state fixed effects ($+0.0861$, $p < 0.0001$, $N = 4{,}036$), which drops none, gives the same sign and precision. A pairs cluster bootstrap on the logit marginal-effect difference gives a 95 percent interval of $[+0.0135, +0.0241]$. Implementation: script `50_logit_functional_form.R`.

A district-level test without legislator data. The party-label claim could in principle be tested with no legislator-level variable at all: the Democratic share of the district two-party vote regressed on exposure interacted with an indicator for Democratic-majority congresses, with district and congress fixed effects. The result is uninformative rather than null. The interaction is -0.0023 ($p = 0.15$, $N = 3{,}526$), but rebuilding the same interaction from each of the other 54 of the 55 congress pairs in the sample returns coefficients between -0.0057 and $+0.0069$, 14 of them significant at the 5 percent level. The reason is exposure imbalance: mean exposure in the two Republican-majority congresses (80 and 83) is 2.32 years against 6.90 overall, so the indicator measures congress-specific vote movements in the low-exposure period rather than chamber control. The within-congress contrast between majority and minority members that the member-level design uses does not exist at the district level. Implementation: script `54_district_party_voteshare.R`.

D.6 Heterogeneity probes: market congruence and newspapers

Two district-level moderators probe the mechanism's texture; neither detects significant heterogeneity, and each null is informative. First, *market congruence*: districts whose population is concentrated in a single media market would, under a local-visibility account, show a larger premium (congruent markets can cover their own representative). The triple interaction of TV $\times$ majority with the district's largest-DMA population share is $+0.026$ ($p = 0.053$) — if anything the premium is *larger* in high-congruence districts, the direction opposite the visibility account and consistent with attenuation from treatment measurement error in multi-DMA districts, whose exposure is an average over heterogeneous market timings. Second, *newspaper environment*: interacting the premium with 1940 newspaper circulation per capita gives $+0.001$ ($p = 0.75$; wild bootstrap $p = 0.61$). The ex-ante prediction of the substitution mechanism for this moderator is ambiguous — strong local papers could preserve candidate information (predicting a smaller premium where papers are strong) or could represent exactly the information stock television displaced (predicting a larger one) — and the data detect neither margin. The voter-level evidence of Section 4.1, not district-level heterogeneity, carries the substitution mechanism. Implementation: scripts `44_audit_dma_congruence.R` and `48_audit_newspaper_triple.R`.

D.7 The extended vote-share sample: construction, state-by-congress fixed effects, and pre-trends

The vote-share outcome in Table 2 is estimated on congresses 78–89, the same window as the return outcome. Without the 1944 election the outcome is available only from congress 79, and that restriction is a data-availability artifact rather than a design choice: the district-level returns file (Jacobson, 2013) begins with the November 1946 election, and congress 78's terminal election is November 1944, so congress 78 supplies no vote-share observations. I fill the 1944 election from the Clerk of the House of Representatives' official election statistics, transcribed by Tamas (2022), which supplies own-party vote share for 318 of the 343 non-South congress-78 panel rows (92.7%). The remaining 25 rows are at-large and multimember districts, for which no single incumbent-versus-challenger share exists: building the 1944 table drops 16 multimember candidate rows, and 3 district keys are ambiguous under the at-large coding convention, which rewrites at-large districts to district 1 and so collides with a real district 1 in states holding both; ambiguous keys are dropped rather than resolved by guess. Three construction rules apply: Minnesota's Democratic–Farmer–Labor label counts as Democratic; a candidate co-nominated by several parties is coded major-party if any party slot names that party; candidates carrying both major-party labels are excluded from the two-party total.

The extended variable mixes two sources by construction — the Clerk returns for 1944, the district file for 1946 onward — so agreement between them is verified rather than assumed. Across the 4,615 matched district-elections both sources cover (1946–1966), the two agree exactly (absolute difference below 10^{-6}) on 95.60% of cases, the correlation is 0.99566, the median and 95th percentile of the absolute difference are both 0.000000, and 1.21% of cases differ by more than one percentage point. As a reproduction gate, dropping the 1944 rows from the rebuilt pipeline returns the previous window's estimate exactly: $+0.00624$ (SE 0.00079) on $N = 3{,}605$. On the extended sample the TV $\times$ majority interaction is $+0.00558$ (SE 0.00073, $p < 0.001$, $N = 3{,}923$); both windows print as $+0.006$ (SE 0.001) at Table 2's precision. The standard error falls because congress 78 holds members of both parties (112 majority rows, 206 minority) at uniformly zero TV exposure, against a 79–89 exposure mean of 6.890 years, so the added block identifies the majority–minority difference in levels at zero exposure.

This subsection is also the first place the vote-share outcome receives the two identification checks the paper reports for the return outcome: the state-by-congress fixed-effects specification and the event study with pre-trend tests. Replacing congress fixed effects with state $\times$ congress fixed effects — so that identification uses only variation within a state in a given congress, absorbing statewide tides and any other confound at the state-congress level — gives $+0.00480$ (SE 0.00084, $p = 1.2 \times 10^{-8}$, $N = 3{,}875$); clustering by state instead of district-by-decade raises the standard error to 0.00115. For comparison, the return outcome moves from $+0.01561$ to $+0.01452$ under the same fixed-effect swap.

The vote-share event study on 78–89 passes the joint Wald test on the pre-treatment coefficients under all three binnings of event time: $\chi^2(4) = 3.878$, $p = 0.423$ under the binning of Figure 2, and $\chi^2(5) = 10.572$, $p = 0.061$ and $\chi^2(3) = 5.175$, $p = 0.159$ under the two alternative binnings. The first alternative binning is marginal, and its two pre-treatment bins nearest treatment are individually positive ($+0.0495$, $p = 0.022$; $+0.0412$, $p = 0.015$): the vote-share pre-period is not as clean as the return outcome's. On the 79–89 rows two of the three binnings had rejected; the rejections were a sample-composition artifact, not evidence of a pre-trend. The return outcome, whose published pre-trend tests pass at $p = 0.544$, 0.931, and 0.931, rejected more strongly still when estimated on those same 79–89 rows, and on the extended rows it passes all three binnings ($p = 0.533$, 0.950, 0.879). The far-pre-period cells that produced the rejections held 1 to 25 observations on 79–89; with congress 78 restored they hold 67, 120, 183, and 250. Implementation: scripts `57_voteshare_extend_1944.R` and `14_vote_margin.R`.

E Appendix E: Variable construction details

E.1 Television exposure measure

Years of TV exposure for district d in congress t is constructed in two steps. First, the district's introduction year averages the county-level first-station years:

$$\text{tvyear}_d = \sum_{c \in d} w_c \cdot \text{first\_tv}_c, \qquad \text{years\_of\_tv}_{dt} = \max\bigl(0,\ \text{year}_t - \text{tvyear}_d\bigr),$$

where c indexes the county parts composing district d, first_tv$_c$ is the first year of commercial television availability in county c's designated market area (ICPSR 22720, Gentzkow and Shapiro, 2008), and w_c is the county part's allocation share in the county-to-district crosswalk for the apportionment map in force — the 1940-census crosswalk for congresses 78–82, the 1950 crosswalk for 83–87, and the 1960 crosswalk for 88–89. Because most counties lie wholly within one district, 94 percent of the shares equal one, and the average is effectively an equal-weighted mean over whole counties with fractional counties down-weighted by the fraction inside the district. Truncation at zero is applied once, at the district level, after averaging. Two properties of this construction are documented as robustness rather than assumed: reweighting counties by population (1940 or congress-matched vintages) moves the headline estimate in the fifth decimal (Appendix D.5), and for a small number of at-large districts a builder fallback stores a state-level average whose maximum deviation from the formula is 0.26 years, with no effect on any estimate (Appendix D.5).

E.2 Legislative effectiveness scores

I use Chiou and Goplerud's (2024) extension of the Volden-Wiseman LES framework (Volden and Wiseman, 2014) to historical congresses. The score weights legislative accomplishments by their substantive significance, following the original Volden-Wiseman convention. The source column is the model-based historical score (`impute_LES`), which is on a log scale; I exponentiate it and normalize the exponentiated score to a congress-specific mean of 1, so a value of 1.5 reads as fifty percent above the congress average.

E.3 Majority-party indicator

The majority-party indicator equals 1 if the legislator's party held a majority of House seats in the relevant Congress and 0 otherwise. Majority composition: 78–79 (D), 80 (R), 81–82 (D), 83 (R), 84–89 (D).

E.4 ANES Type A and Type B variables

Variable	ANES item	Asked in	Coding
recall House cand name	VCF0972	1958, 1964	1 if Yes recalled, 0 if No/DK
knows if cand was incumbent	VCF0977	1958, 1964	1 if Yes, 0 if No/DK
correctly IDs incumbent	VCF0978	1958, 1964	1 if correct, 0 if incorrect
perceives party difference	VCF0501	1952, 60, 64	1 if Yes, 0 if No/DK
knows cons party	VCF0502	1960, 64	1 if picks one party, 0 if DK/no diff
cares about pres outcome	VCF0311	1952–64	1 if cares a good deal, 0 otherwise
watched political TV	VCF0724	1952–64	1 if watched, 0 if not

Coding conventions and their consequences: the recall and knows-if items code "don't know" as 0 (uninformed), while the correctness item is defined only for respondents who ventured an identification — the asymmetry examined in the extensive-margin analysis of Section 4.1 (recoding the correctness item onto the full pool gives -0.041; conditioning the knows-if item on those who ventured an answer gives -0.004, insignificant). The cumulative file codes "No" and "don't know" as a single value for VCF0972, so the attempted-only variant cannot be constructed for that item. The two unmatched state codes in the panel merge are the 1954 study's unidentified state code (1,139 respondents) and the District of Columbia (41). Summary indices standardize each available component, average them per respondent, and re-standardize (Anderson, 2008). Implementation: scripts `18_anes_typea_typeb.R` and `36_audit_anes_truth.R`.

E.5 Seniority quartiles

Seniority is career House tenure: the number of House Congresses in which the legislator has served through the current one, computed from the full Voteview roster with no lower congress bound, so that service before the panel window counts and non-consecutive service accumulates. Quartile cuts are computed within the majority-party pool on the analysis sample: Q1 ≤ 2 terms, Q2 ≤ 4 terms, Q3 ≤ 6 terms, Q4 > 6 terms. The 7-plus-term top quartile corresponds approximately to chair-eligibility in the seniority system that prevailed in this period. Implementation: script `19_majority_subgroup.R`.

F Appendix F: Exit-reason decomposition

The outcome in Tables 2 and 3 is chamber persistence: a member who does not appear in the next Congress is coded as not returning whether she lost the general election, lost her primary, retired, died, resigned, or left to run for another office. This appendix classifies every such exit by its route and estimates the majority-party premium separately on each route, so that the share of the headline attributable to voters' verdicts can be read directly rather than assumed.

F.1 Classifying exits from the Biographical Directory

The election files used in the main analysis (ICPSR 8611, the 1946–2012 district returns, and the Clerk of the House returns for 1944) record vote totals but not whether the sitting member was a candidate, so they cannot distinguish a defeated incumbent from one who did not run. The Biographical Directory of the United States Congress states the route of a departure in a fixed vocabulary ("unsuccessful candidate for reelection in 1948 to the

Eighty-first Congress," "unsuccessful candidate for renomination," "was not a candidate for renomination in 1952," "unsuccessful candidate for election to the United States Senate," "resigned"), and it publishes all 13,055 member profiles as a single bulk download. I classify from those profiles the 867 legislator-congress observations in the analysis sample with `incumbent return` equal to zero (806 distinct members; one observation has no Directory identifier).

For each observation the classifier selects the semicolon-delimited clauses of the profile that name the congress's first year, its election year, the following year, or the ordinal name of the current or the next Congress; 92 percent of cases yield between one and three such clauses. Each clause is tested against the Directory's phrase vocabulary for seven routes: election or appointment to higher office (the Senate, a governorship, or the vice presidency), an unsuccessful bid for higher office, defeat for renomination, defeat for reelection, exclusion from membership by House resolution, declining to be a candidate for renomination or reelection, and resignation. Because the clause window spans the whole congress and therefore admits clauses about other terms in a member's career, resolution is two-tier. A clause that names the deciding election year decides first; the wider window decides only when no election-year clause names a route, and there a mid-term resignation (the member never faced the voters) takes precedence over candidacy language that may belong to a different term. Deaths are taken from Voteview's recorded year of death, which separates deaths in office from the fifteen deaths that fall in the transition year between the election and the next Congress's first session; the Directory's own death date is the independent check and agrees on all 109.

The classification yields 389 general-election defeats, 63 primary defeats, 170 retirements, 58 unsuccessful bids for higher office, 35 successful moves to higher office, 94 deaths in office, 15 transition-year deaths, 39 mid-term resignations, 1 exclusion by House resolution, and 3 unclassified cases (0.3 percent). Three checks bear on its accuracy. First, year integrity: of the 715 cases in the five candidacy classes (the two defeats, retirement, and the two higher-office routes), 703 (98.3 percent) have a deciding clause that names the election year itself, and 711 (99.4 percent) are pinned by that year or by the next Congress's ordinal name, leaving 4 cases whose clause could in principle describe another term. Second, agreement with Voteview on the events Voteview itself records: 37 of the 43 Senate moves Voteview identifies are classified as higher-office exits, and 64 of the 80 members whom Voteview shows returning to the House after a gap are classified as defeats, with 10 of the remaining 16 classified as unsuccessful bids for higher office, which also precede a return to the House. Third, a seat-flip check on an observable the classifier never saw: a seat whose incumbent was defeated in the general election should pass to the other party in the next Congress far more often than a seat whose incumbent retired. In Voteview's next-Congress roster, seats classified as general-election defeats change party 94.6 percent of the time and seats classified as retirements 27.0 percent of the time, a 67.6-point gap ($p = 2.3 \times 10^{-61}$); the rates for primary defeat (17.7 percent), higher office won (14.3 percent), and death (23.7 percent) are all closer to the retirement rate than to the defeat rate. Reassigning all three unclassified cases to any single route leaves the combined defeat coefficient reported below between -0.00995 and -0.01019. Implementation: scripts `52_bioguide_exit_classify.R` and `55_exit_classification_verify.R`.

F.2 Decomposition of the majority-party premium

Because the routes partition the exits, estimating the headline specification with each route indicator as the outcome decomposes the headline exactly: the route-specific TV $\times$ majority interactions sum to the negative of the interaction on incumbent return, and the identity holds to 3×10^{-17} on the common estimation sample of 4,031 observations (862 exits; the five observations dropped from the headline as fixed-effect singletons, three general-election defeats and two retirements, account for the difference from the 867 classified cases). Table A2 reports the ten route regressions. A negative coefficient means television made that route less likely for majority-party incumbents relative to minority-party incumbents, and the premium-share column expresses each coefficient as a fraction of the headline $+0.01561$.

Table A2: Exit-reason decomposition of the majority-party premium

Exit route	Coef.	SE	Events	% of premium	% of exits
General-election defeat	-0.00881^{***}	0.00161	386	56.4	44.8
Primary defeat	-0.00114	0.00074	63	7.3	7.3
Retirement (did not seek renomination)	-0.00190	0.00125	168	12.2	19.5
Sought higher office, lost	-0.00247^{***}	0.00069	58	15.8	6.7
Sought higher office, won	-0.00017	0.00054	35	1.1	4.1
Died in office	-0.00087	0.00083	94	5.6	10.9
Died in the transition year	-0.00036	0.00032	15	2.3	1.7
Resigned mid-term	$+0.00025$	0.00060	39	-1.6	4.5
Excluded by House resolution	$+0.00008$	0.00006	1	-0.5	0.1
Unclassified	-0.00025	0.00019	3	1.6	0.3
Sum of exit routes	-0.01561		862	100.0	100.0
Incumbent return (headline)	$+0.01561^{***}$	0.00236			

Each row is a separate linear probability model of an indicator for that exit route on TV exposure interacted with majority-party status, with district and congress fixed effects, demographic controls, and standard errors clustered by district-by-decade, estimated on the identical 4,031 observations. A negative coefficient means television made that route less likely for majority-party incumbents, which contributes positively to the return premium. Because the routes partition the sample, the coefficients sum to the negative of the headline return coefficient by construction; the residual of that identity is 3e-17. Exit reasons are classified from the Biographical Directory (Appendix F.1). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Defeat at the polls accounts for most of the premium. The two defeat routes together give -0.00995 (SE 0.00177, $p < 0.0001$), 63.7 percent of the premium: general-election defeat alone gives -0.00881 (SE 0.00161), 56.4 percent of the premium from a route that makes up 44.8 percent of exits, and primary defeat adds -0.00114 (SE 0.00074, $p = 0.125$), 7.3 percent of the premium against 7.3 percent of exits. Retirement gives -0.00190 (SE 0.00125, $p = 0.128$), 12.2 percent of the premium against 19.5 percent of exits, so the strategic-retirement rival named in Section 4.3 — majority members choosing to stay in television districts and minority members choosing to leave — accounts for at most a small and imprecisely estimated share. The exposure slope on retirement is negative for both parties (-0.00434 for minority members, $p = 0.099$, and -0.00624 for majority members, $p = 0.019$); what is small is the difference between the two slopes. Unsuccessful bids for the Senate or a governorship are a third route: -0.00247 (SE 0.00069, $p = 0.0004$), 15.8

percent of the premium from 6.7 percent of exits. Here only the difference is precise: the majority slope is $+0.00002$ ($p = 0.988$) and the minority slope $+0.00249$ ($p = 0.125$), so the data do not say whether majority members attempted fewer such bids under television or minority members more. Death in office gives -0.00087 ($p = 0.293$), 5.6 percent against 10.9 percent of exits; television should not move that route, and the interaction is not distinguishable from zero, which is what a placebo route should show. The remaining routes (successful moves to higher office, transition-year deaths, resignations, the single exclusion, and the three unclassified cases) each contribute between -1.6 and $+2.3$ percent.

A complementary specification conditions on facing the voters. Restricting the sample to incumbents who returned or lost a general election removes every exit not decided by voters from the outcome and gives $+0.01091$ (SE 0.00185, $p = 4.9 \times 10^{-9}$, $N = 3{,}552$); adding primary losers gives $+0.01177$ (SE 0.00197, $N = 3{,}615$). The restriction conditions on whether the member ran at all, which Table A2 shows television moved differently for majority and minority members, so I report it as a complement to the decomposition rather than a replacement for it. Implementation: script `53_exit_decomposition.R`.