The Decision Window: How Budget Timing Shapes Fiscal Policy

Kerry Siani and Jinyuan Zhang*

March 26, 2026

Abstract

Large organizations receive economic information continuously, but make decisions at discrete points. We study this phenomenon in U.S. state governments using hand-collected budget documents over 1995–2024. The original budget accounts for 73% of within-state spending variation, making the planning stage the key commitment margin for fiscal spending. We develop an empirical beliefs–decision test that distinguishes continuous updating from decision-window updating. Using a shift-share measure of state-specific revenue news and cross-state variation in proposal deadlines, we implement the test and find that only news arriving in the six months before the initial budget proposal—the decision window—updates revenue forecasts, spending plans, and realized spending. Equally informative earlier news is omitted from the proposal and instead predicts forecast errors. These timing differences propagate to the private sector: planning-window news raises employment and wages in government-exposed industries, while off-window news does not.

JEL codes: E62; H30; H61; H72; D83

Keywords: decision window; state budgets; fiscal policy; organizational behavior

*Kerry Y. Siani (ksiani@mit.edu) is at MIT Sloan, Cambridge, MA. Jinyuan Zhang (jinyuan.zhang@anderson.ucla.edu) is at UCLA, Los Angeles, California. All errors are our own.

How do large economic agents incorporate information into spending decisions? In principle, decision makers could continuously revise expenditure plans as new information arrives. To do so can be costly, so organizations often commit to ex ante plans. In practice, institutional constraints—planning cycles, approval processes, and legal commitments—confine meaningful adjustment to narrow decision windows: discrete periods during which information is gathered, expectations are formed, and commitments are set. If decisions are time-gated in this way, the real effects of economic shocks depend not only on their magnitude but also on their timing relative to these institutional decision points.

We study this mechanism in U.S. state government budgets. State and local governments account for roughly 15% of GDP, employ more than 20 million workers, and finance core services in education, healthcare, and public safety. State budgets are also a useful laboratory for expectations and decision-making because expected revenues matter directly for planned spending, budget calendars are transparent and largely predetermined, and the sequence of proposal, enactment, and execution is publicly observed. Using a newly assembled panel of hand-collected budget documents that tracks original budgets, finalized budgets, and realized outcomes for 47 states over 1995–2024, we show that institutional timing fundamentally shapes which economic information translates into public spending—and, through spending, into private economic activity.

To identify the timing with which information enters the budget, we construct a shift-share measure of state-specific revenue news by interacting predetermined state revenue-base shares with national innovations to those bases. Because proposal deadlines differ across states, the same class of national revenue innovations falls inside the active planning window for some states and outside it for others. This variation allows us to compare the same type of information arriving at different points in the budget calendar and to ask whether policymakers incorporate available news continuously or mainly when the organizational calendar forces decisions.

Our first result is that state fiscal policy is largely determined during the initial planning stage. The original budget accounts for 48% of within-state variation in realized revenue growth and 73% of variation in realized expenditure growth. Subsequent revisions and within-year realization contribute comparatively little. This implies that fiscal policy is not characterized by a continuous flow of spending decisions, but rather by a narrow "planning window" in which the fiscal-year budget is largely determined.

Second, revenue information shapes this commitment only when it arrives during the planning window. To identify the information that enters the original budget, we isolate plausibly exogenous revenue news arriving at different points in the budget calendar. Our core instrument is a shift-share measure that interacts predetermined state-level revenue-base shares with national innovations to those bases. Because proposal deadlines vary across states, the same national innovation enters different budget windows across states. This timing variation lets us compare the effects of the same set of news when it arrives before, during, or after active budget planning. We find strong evidence of a planning-window effect: revenue news affects expectations and commitments only when it arrives during the active planning stage—the 0–6 months leading up to the budget proposal deadline. Planning-window shocks strongly predict updates to revenue forecasts and planned spending. However, shocks that occur pre-planning (6-12 months before the initial budget is proposed) have little effect on expectations or plans.

Third, because the original plan largely pins down the fiscal-year budget, planning-window news also predicts realized spending growth. News arriving outside the planning window—before planning, during legislative revision, or during the fiscal year itself—has little effect on realized spending, even though it remains informative about eventual revenues. Only planning-window news systematically enters revenue expectations, shifts expenditure plans, and propagates to realized spending.

We formalize these findings as an availability–belief–action test of organizational updating. The test jointly examines three objects: what information enters the revenue forecast (beliefs), what available information predicts forecast errors (omitted information), and what information moves the expenditure plan (actions). Under continuous updating, all pre-decision news of comparable informational content should enter forecasts equally, forecast errors should be unpredictable, and expenditure plans should respond to each window's news according to the same pass-through. Under decision-window updating, only planning-window news enters forecasts and moves expenditures, while earlier news—despite remaining informative about revenues—is omitted from the forecast and fails to shift the spending plan. The data strongly favor the second pattern: pre-planning news predicts revenue forecast errors, but does not move the expenditure plan, while planning-window news enters both forecasts and spending commitments.

These timing asymmetries have real consequences for the private sector. We link the

same revenue-news measures to private labor-market outcomes using the Quarterly Census of Employment and Wages (QCEW), exploiting cross-industry heterogeneity in exposure to state and local government demand. Planning-window news generates large differential increases in employment and wages in government-exposed industries, while off-window news does not. Local projections confirm that the employment response is absent before the fiscal year, builds during budget execution, and persists beyond fiscal-year end—ruling out anticipation effects and pointing to a procurement-demand channel.

Using a structural model of the budget process estimated to match our reduced-form results, we decompose the attenuation of fiscal pass-through into two channels: institutional frictions and the decision-window bias. The model reveals that the bias is the dominant friction for shocks arriving before the planning window: a pre-planning revenue shock retains only a fraction of its frictionless employment effect, with the bias accounting for the vast majority of the loss. The attenuation is asymmetric: negative pre-planning shocks partially pass through via execution-stage balanced-budget enforcement, while positive shocks are almost entirely lost. By contrast, planning-window shocks pass through nearly fully despite institutional frictions.

For the design of intergovernmental transfers, federal stabilization programs, and monetary policy aimed to stabilize state output, our results suggest that timing relative to state budget calendars may matter as much as magnitude. Moreover, our findings point to a broader lesson about organizational decision-making: large organizations receive information continuously but convert that information into commitments at discrete decision dates, creating a tradeoff between adapting to new information and adhering to plans.

This paper contributes first to the economics of organizations and attention. Large organizations receive information continuously but convert that information into commitments at discrete decision dates, creating a tradeoff between adapting to new information and adhering to plans (**??**). When adapting to information in real time, organizations face communication and coordination costs (**?**), so commitment to ex ante plans can be beneficial, but also creates a distance between decisions and actions. Expectations data is thus useful to study decision-making in real time (**?**), and recent empirical work uses expectations and plans data to study real decisions by large economic agents such as corporations (**?**). We build on this by studying information incorporation and decision-making in the context of state governments, where we can directly observe the expectations concurrently with spend-

ing decisions. We show that information is incorporated in a discrete manner at the decision window, providing micro-evidence supporting **?** and **?**. Our paper therefore identifies a concrete organizational friction, calendar-gated information incorporation, through which institutional decision rules shape fiscal policy.

Second, the paper also contributes to the literature on imperfect information and expectation formation. We build on the literature on sticky-information models with infrequent but complete updating (**??**), rational-inattention models with costly signal processing (**???**), and survey-based tests of noisy or sluggish updating (**????**). Behavioral models find agents may overreact to recent news (**?**), particularly when they have longer forecast horizons (**?**).[1] Our setting directly tests these notions in the context of state government spending. Because the original and final budget proposals are observed updating dates with both expectations and plans, we can ask whether officials fully incorporate information that is already available by that date. We find that they do not: revenue-relevant news arriving during the planning window enters proposal-stage forecasts and spending commitments, while equally informative pre-planning news is not significantly accounted for and instead predicts subsequent forecast errors. This pattern is difficult to reconcile with infrequent but complete updating conditional on an observed review date. More broadly, it suggests that the effective cost of attention is institutionally time-varying and tied to the organizational calendar.

Third, the paper contributes to the literature on fiscal foresight and the timing of policy effects. Prior work shows that the macroeconomic effects of fiscal policy depend on when agents learn about policy relative to when policy is implemented (**????**). In our setting, timing does not just shape when private agents respond to policy; it shapes which information enters policy in the first place. The same class of revenue-base innovations can arrive before, during, or after active budget formation. This institutional timing variation allows us to show that the budget calendar governs whether news is translated into same-year spending commitments at all. Relative to work on state fiscal institutions, balanced-budget rules, and forecast accuracy (**??????**), we identify where in the multi-stage budget process fiscal commitments are set and which signals enter at that stage. In that sense, the paper shifts the focus from whether states forecast well on average to when, in the budget cycle, they process revenue-relevant information.

Finally, the paper contributes to work on procurement demand, production-network

[1] **?** provides a recent survey of the literature on overreaction in macroeconomics.

transmission, and local fiscal effects. Existing work shows that government spending can affect local employment and propagate through procurement and supplier networks (??????). We add an upstream margin that this literature has largely left unmeasured: whether revenue news enters the spending plan before contracts are issued. Planning-window news raises employment and wages in government-exposed industries, whereas earlier news has much smaller and less robust effects, and the response appears mainly during execution rather than before the fiscal year begins. The paper therefore shows that budget timing governs not only the size of fiscal commitments, but also whether revenue news reaches private labor markets at all.

The remainder of the paper proceeds as follows. Section ?? describes the state budget process, defines the decision windows, and introduces the data. Section ?? documents the concentration of fiscal decisions in the planning stage. Section ?? develops the availability–belief–action framework for testing continuous versus decision-window updating. Section ?? introduces the shift-share news measure and implements the test, showing that only planning-window news moves the original budget and propagates to realized expenditures. Section ?? studies real effects in government-exposed industries. Section ?? develops the structural model and presents the counterfactual decomposition. Section ?? concludes.

1 The State Budget Process and Data

1.1 The State Budget Process

State operating budgets typically follow a common process: an initial planning phase in which the governor proposes a budget, a legislative revision phase culminating in enactment, and an execution phase in which spending is implemented during the fiscal year. While procedural details vary across states, the key timing feature—that major spending commitments are set months before the fiscal year begins—is broadly shared.

As an example, Figure ?? illustrates the timeline for the Massachusetts operating budget for fiscal year (FY) 2026 (July 1, 2025 - June 30, 2026). The operating budget for FY 2026 begins development in July 2024, a full year before execution begins. During the fall of 2024, representatives from the state legislature and the Governor's office meet with economists to

establish a consensus revenue projection for the upcoming fiscal year. Using these revenue projections and incorporating the administration's political priorities, the Governor submits an initial budget proposal by the third week of January 2025. Following the Governor's proposal, the Ways and Means Committees in both the Massachusetts Senate and House hold budget hearings, gathering input from various stakeholders. By the end of April, the House votes on its revised version of the operating budget proposal, while the Senate votes on its own version by the end of May.[2] In June, members of the House and Senate Ways and Means Committees negotiate to reconcile their differences, developing a unified budget proposal. This reconciled budget is then submitted to the full legislature for ratification. Upon approval by both the House and Senate, the Governor signs the final operating budget into law around June 30, 2025.

As illustrated in Figure **??**, this process can be divided into three phases: Planning (July 2024–January 2025), Revision (January–June 2025), and Realization (July 2025–June 2026). Initial spending commitments are thus determined six to twelve months before the fiscal year begins, creating substantial temporal distance between decisions and realized outcomes. We manually collect budget calendar information for all states, summarized in Table **??**.

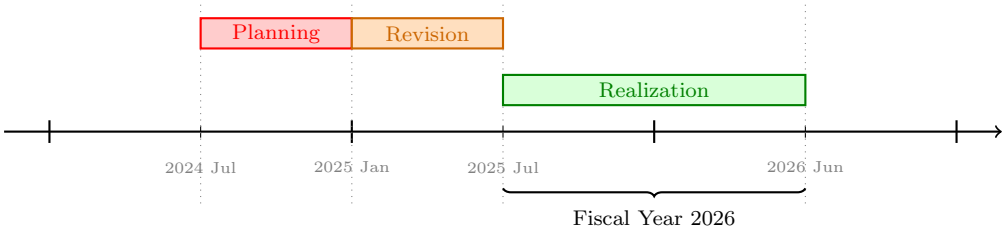


Figure 1: Timing of the Budget Process for a Representative State (FY 2026). This figure illustrates the timing of planning (pre-proposal), revision (proposal to enactment), and realization (fiscal-year implementation) for a representative state.

This transparent institutional timing is typically set by state constitutions and thus generally exogenous to local economic conditions, making state budgets a useful setting for studying how organizations incorporate information into spending decisions.

[2]Although the budgets proposed by the House and Senate must be balanced, the Governor's initial budget proposal is not required to be balanced.

1.2 Data

One of the contributions of this paper is the construction of a comprehensive state-level panel that integrates multiple data sources on planned and realized fiscal outcomes and institutional features spanning several decades. First, we hand-collect the Annual Comprehensive Financial Reports (ACFR) produced by each state from 1995–2024 by downloading PDFs from state websites, contacting state officials by email or phone, and filing formal records requests.[3] We digitize the data using Python and hire undergraduates to manually check the digitization quality.

Our ACFR dataset provides three budget objects for each fiscal year: the initial budget (typically proposed by the governor), the final budget (the enacted budget after legislative revision), and actual outcomes realized during execution, see for example Massachusetts in Figure **??**. We focus on forty-seven states with annual budget cycles; four biennial-budget states (Wyoming, Washington, North Dakota, Oregon) are excluded. The initial and final budgeted data offer a forward-looking window into each state's expected fiscal decisions and allow us to measure how expectations of revenues and borrowing costs translate into planned spending.

We supplement the ACFR data with realized fiscal variables—including debt outstanding and interest expense—from the Census Bureau's Annual Surveys of State and Local Government Finance. We also merge standard state-level controls (e.g., population, income, unemployment) and national macro-financial series used in our baseline specifications, including Treasury yields and measures of aggregate activity, labor-market conditions, and CPI inflation from FRED, as well as S&P 500 returns.

2 Stylized Facts: Decision Windows in State Budgets

How do state governments incorporate information into spending decisions? The institutional framework in Section **??** imposes tight timing and financing constraints, but it is an empirical question whether these constraints bind in practice—and, if so, at which stage of the budget

[3]The availability of ACFR information varies across states. We obtained 84 reports by directly emailing state officials, 11 additional reports from phone calls, and 35 by formal record request; the remainder of the reports were available on state websites.

cycle. Using our panel of state budget plans and realizations constructed from ACFRs, this section establishes two facts that motivate the remainder of the paper: (i) expenditures track revenues closely throughout the budget cycle, and (ii) the initial budget—set months before the fiscal year begins—accounts for most within-state variation in realized spending. Together, these facts imply that the key margin of fiscal adjustment is the initial planning stage, which we refer to as the decision window.

For fiscal year t, the initial budget is typically proposed before fiscal year $t-1$ actuals are realized. We therefore measure planning-stage growth relative to Final_{t-1}, the most recent budget benchmark known at the time of the initial budget. This choice both matches the salient reference point in the budgeting process and cleanly separates new policy intent from prior-year execution surprises. To be concrete, we decompose annual changes in realized outcomes into planning, revision, and realization components using the accounting identity:

(1)
$$\Delta \log(\text{Act}_{s,t}) = \underbrace{\log\left(\frac{\text{Init}_{s,t}}{\text{Final}_{s,t-1}}\right)}_{\text{Planning}} - \underbrace{\log\left(\frac{\text{Act}_{s,t-1}}{\text{Final}_{s,t-1}}\right)}_{\text{Lag realization}} + \underbrace{\log\left(\frac{\text{Final}_{s,t}}{\text{Init}_{s,t}}\right)}_{\text{Revision}} + \underbrace{\log\left(\frac{\text{Act}_{s,t}}{\text{Final}_{s,t}}\right)}_{\text{Realization}}.$$

Yearly changes in actual spending can be thus decomposed into four components: planning, lag realization, revision, and realization. For our empirical analysis, we will refer to the "planning window" as the 0-6 month window before the governor's proposal is published. The 6 months between the governor's proposal and the final budget is passed in the legislature is the "revision window". The "realization" is the 12 months of the actual fiscal year. Because states have different fiscal year start dates, we use the months in which the fiscal year starts and ends to define the planning and realization windows.

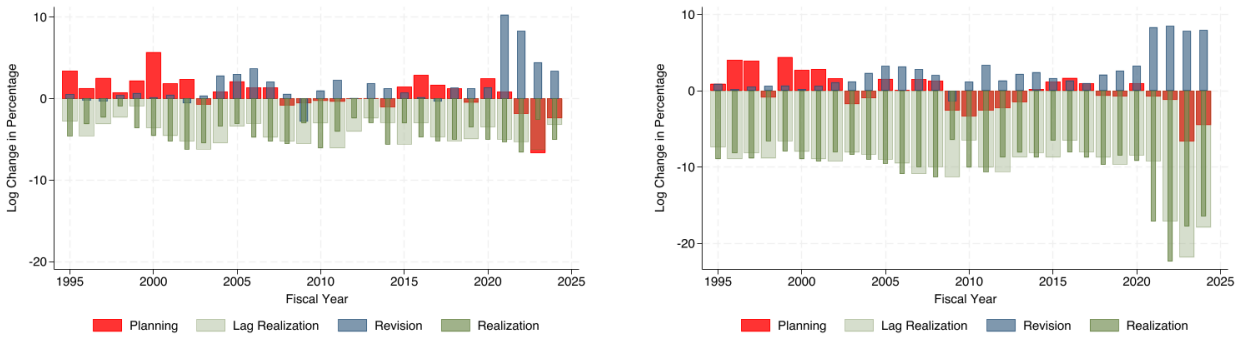
Summary statistics for each of these components are reported in Table **??**. In the median planning window, expenditures increase by 0.4%, while revenues increase by 1.4%. Between the initial budget and the final budget, the median state's expenditures rise by another 0.5%, while median state's revenues decrease by 0.1%. Finally, in the median realization window, spending is 7% lower than the final budget, while revenues are 2% lower. Thus, the median budgeting process overestimates both expenditures and revenues, but overestimates spending more than revenues, likely reflecting a conservativism in the budgeting process. We report the histograms of adjustments in each window in Appendix Figure **??**: budgeted revenues exceed realizations in 75% of state-year observations, and budgeted expenditures

exceed actual outlays in 97%.

Fact 1: Expenditures track revenues throughout the budget cycle. Figure **??** plots average log changes in total real revenues and expenditures across the phases in Equation (**??**). Expenditure adjustments co-move closely with revenue changes in each phase. Both revenues and expenditures typically increase in the planning window, with small further positive adjustments during the revision period. The exceptions are in the years around the 2008 financial crisis and after the COVID pandemic, where both revenues and expenditures decrease in the planning window. However, across the board, realized outcomes are typically negative: actual revenues and, especially, actual expenditures tend to fall below enacted budgets consistently in the time series, with a larger drop in expenditures after the COVID pandemic.

Figure 2: Average Log Changes Over Planning, Revision, and Realization Phases

This figure displays average log changes in total real revenues (Panel a) and total real expenditures (Panel b) across the phases defined in Equation (**??**).



(a) Total Revenue

(b) Total Expenditure

Table **??** quantifies this relationship: revenue changes explain a substantial share of the variation in expenditure adjustments, with particularly strong co-movement during planning and revision. A one-percentage-point reduction in expected revenues reduces expected spending by 0.6 percentage point in the planning window, and 0.7 percentage point in the revision window; during the realization window, a one-percentage-point reduction in expected revenues reduces expected spending by 0.3 percentage point. Appendix Table **??** further shows the asymmetry in realized adjustments: spending adjusts sharply to revenue

shortfalls but responds much less to revenue windfalls.

What happens, concretely, when revenues are higher or lower than expected? For unanticipated revenue windfalls, Appendix Table **??** shows that positive budget surprises flow into rainy day funds: states significantly increase reserve balances when actual surpluses exceed budgeted surpluses, but do not draw down reserves proportionally during deficits. For unexpected revenue shortfalls, the response depends on the size of the shortfall. In 46 states, the executive branch has the authority to unilaterally reduce spending when revenues fall short of expectations. Special legislative sessions can also be called to deal with substantial shortfalls; in 37 states, this requires a supermajority vote of the legislature. The federal government can also help fill the gap in the case of significant shortfalls, such as during Hurricane Katrina, the Great Recession, and COVID. However, given the constraint that nearly all states face to pass a balanced budget, states cannot rely on these mechanisms when setting budgets, and thus rely heavily on expected revenues when setting budgets.

Fact 2: Initial budgets account for most of the variation in realized spending. We next show that the initial budget accounts for the majority of within-state variation in realized expenditures, whereas subsequent revisions and final realization contribute comparatively little. To quantify the contribution of each budgetary phase, we decompose the variance of annual changes in actual outcomes using Equation (**??**). Specifically, we take the covariance of each component with $\Delta \log(\text{Act}_{s,t})$ and scale by $\text{var}(\Delta \log(\text{Act}_{s,t}))$:

(2)

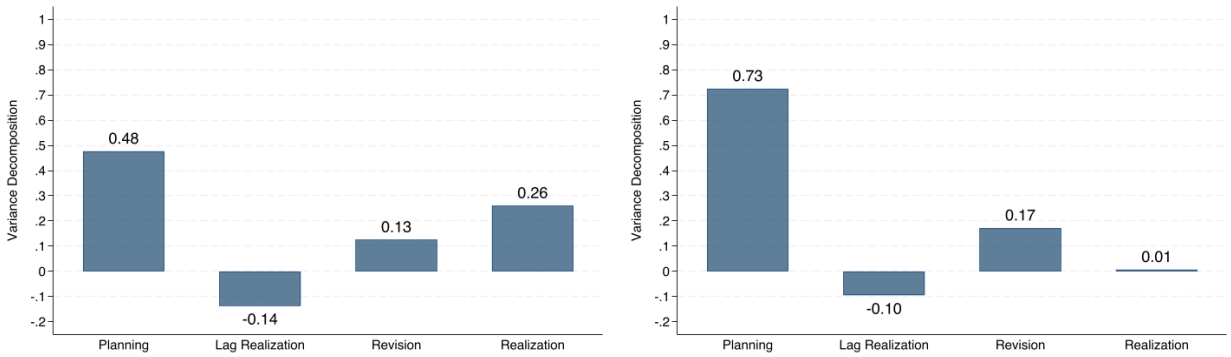
$$1 = \frac{\text{cov}\Big(\log(\text{Init}_{s,t}/\text{Final}_{s,t-1}), \Delta\log(\text{Act}_{s,t})\Big)}{\text{var}\big(\Delta\log(\text{Act}_{s,t})\big)} - \frac{\text{cov}\Big(\log(\text{Act}_{s,t-1}/\text{Final}_{s,t-1}), \Delta\log(\text{Act}_{s,t})\Big)}{\text{var}\big(\Delta\log(\text{Act}_{s,t})\big)} + \frac{\text{cov}\Big(\log(\text{Final}_{s,t}/\text{Init}_{s,t}), \Delta\log(\text{Act}_{s,t})\Big)}{\text{var}\big(\Delta\log(\text{Act}_{s,t})\big)} + \frac{\text{cov}\Big(\log(\text{Act}_{s,t}/\text{Final}_{s,t}), \Delta\log(\text{Act}_{s,t})\Big)}{\text{var}\big(\Delta\log(\text{Act}_{s,t})\big)}.$$

Figure **??** reports the variance decomposition implied by Equation (**??**). The planning component alone accounts for 48% of the variation in revenue growth and 73% of the variation in expenditure growth, indicating that most fiscal adjustment is determined when the initial budget is drafted rather than during revision or execution. Put differently, the decision window is narrow: once initial commitments are set, subsequent phases contribute relatively little to the cross-state, within-year variation in realized spending. Appendix Figure **??** shows that this conclusion is robust to an alternative benchmark that uses two-year-lagged

actual outcomes in the planning stage.

Figure 3: Decomposing the Sources of Variation in Actual Growth

This figure plots the variance-decomposition coefficients implied by Equation (**??**) for total real revenues (Panel a) and total real expenditures (Panel b).



(a) Total Revenues

(b) Total Expenditures

3 A Test of Decision-Window Updating

Section **??** establishes that the initial budget is an important commitment margin for state fiscal outcomes. This raises a natural question: how important is the decision window for actual state spending? More precisely, how different is the impact of news arriving in the decision window versus news arriving outside the decision window? This section develops a conceptual framework that delivers sharp, testable restrictions distinguishing whether organizations update information and decisions continuously or in discrete windows. Section **??** then implements the test empirically.

The framework incorporates two objects that we can directly map to the data: a forecast of an outcome the organization does not directly control over the relevant horizon (revenue) and a revision of an object the organization does control (expenditure). In the state-budget setting, the relationship between these two objects is especially transparent: expected revenue is a direct input into spending commitments, so budget plans provide a clean window into how the organization converts information into action. This is the time-domain analogue of the coordination-versus-adaptation logic in **?** and **?**: organizations may choose not

only which tasks to coordinate ex post, but also when to coordinate information ex post. Our contribution is to make that idea empirically testable at an observed organizational review date. More broadly, the same logic applies whenever the researcher can specify a model-implied pass-through from beliefs about a non-controlled outcome into a controlled decision.

Our framework is distinct in two important ways from the forecast-revision tests in the spirit of **?** and **?**: first, we study forecasts by decision-makers themselves, where the forecasts are directly tied to real outcomes. Second, we observe the relevant review date and can compare information that arrived before versus during that review window. A close analogue is the expectations-and-plans literature (**?**); relative to that literature, we observe both the belief object and the controlled decision object for the same line items at different points in time. The state budget therefore provides a particularly clean laboratory for testing whether organizations update continuously or only when the calendar forces coordination.

3.1 Conceptual Framework

3.1.1 General formulation

The test applies whenever a researcher observes three objects around an organizational decision date: (i) a realized outcome the organization does not directly control over the relevant horizon, (ii) the organization's forecast of that outcome, and (iii) a decision the organization does control. Let Y_{it} denote the non-controlled outcome (e.g., revenues, demand, sales), F_{it} the organization's forecast of Y_{it} at the decision date, and D_{it} the controlled decision (e.g., expenditures, hiring, capex, procurement). Let x_{itw} be news arriving in window w relative to the decision date, where windows are indexed by their timing: pre-window, planning window, and so on.

Because windows may differ in informativeness, we first estimate the predictive content of each window's news for the non-controlled outcome:

$$Y_{it} = \sum_{w} \pi_w \, x_{itw} + u_{it}. \tag{3}$$

The slopes π_w measure how much objective information each window's news contains. To

put all windows on a common footing, we define information-equivalent news $\tilde{x}_{itw} \equiv \hat{\pi}_w \, x_{itw}$, so that a one-unit change in $\tilde{x}_{itw}$ corresponds to the same predicted change in Y, regardless of timing.

Now suppose the organization incorporates each window's information-equivalent news with weight a_w and converts forecasts into decisions with structural pass-through $\lambda > 0$. Let ρ_w denote the loading of forecast errors on window-w news. The three-equation system is:

$$\text{Beliefs:} \quad \Delta F_{it} = \sum_w a_w \, \tilde{x}_{itw} + \varepsilon_{it}, \tag{4}$$

$$\text{Omission:} \quad \underbrace{Y_{it} - F_{it}}_{\equiv \mathrm{FE}^Y_{it}} = \sum_w \rho_w \, \tilde{x}_{itw} + \nu_{it}, \tag{5}$$

$$\text{Action:} \quad \Delta D_{it} = \sum_w \gamma_w \, \tilde{x}_{itw} + \eta_{it}, \tag{6}$$

where all regressions include unit and time fixed effects and controls, suppressed for clarity. The belief equation (**??**) identifies a_w: what information enters forecasts. The omission equation (**??**) identifies ρ_w: what available information was left out and therefore predicts forecast errors. The action equation (**??**) identifies γ_w: what information moves controlled decisions. With the most stringent restrictions, $\rho_w = 1 - a_w$ and $\gamma_w = \lambda a_w$; in practice we estimate a_w, ρ_w, and γ_w as free parameters and test these restrictions directly. The adding-up restriction $\rho_w + a_w = 1$ holds only if the shift-share measure captures all relevant news; in practice, other shocks arrive between the proposal and fiscal-year end, so we estimate ρ_w and a_w freely and treat adding-up as an over-identifying check.

Continuous updating. If the organization fully incorporates all information available by the decision date, three restrictions hold for any window w whose news is known before the decision:

(i) Equal incorporation: a_w is the same across all pre-decision windows (in common Y-units).

(ii) No predictable omission: $\rho_w = 0$ for all pre-decision windows.

(iii) Stable pass-through: $\gamma_w / a_w = \lambda$ for all pre-decision windows.

Decision-window updating. If the organization coordinates information only during an active decision window w^*, then for windows w arriving before w^*:

(i) Selective incorporation: $a_{w^*} > a_w$ for $w \neq w^*$.

(ii) Predictable omission: $\rho_w > 0$ for off-window news; $\rho_{w^*} \approx 0$.

(iii) Attenuated action: $\gamma_w < \gamma_{w^*}$ for off-window news.

Earlier news continues to predict the non-controlled outcome but is omitted from the forecast, so ρ_w is large for off-window periods. Controlled decisions inherit the same omission because they are mediated by the forecast. This is the time-domain analogue of the coordination-versus-adaptation logic in **?** and **?**: organizations choose not only which tasks to coordinate, but when to coordinate information.

The framework above is deliberately agnostic about why the organization gates information at the decision window. The three-equation system identifies the pattern of organizational updating—which windows' news enters beliefs, which is omitted, and which moves actions—but not the underlying mechanism that produces the pattern. The gating could reflect rational attention allocation (**??**), institutional agenda-setting (**??**), or behavioral recency weighting (**??**). These candidate mechanisms generate the same qualitative prediction—selective incorporation of recent news—but differ in auxiliary predictions about signal value, overreaction, and the gradient of attenuation. Section **??** evaluates these candidates against the data and finds that the evidence is most consistent with an institutionally gated attention process that is neither fully rational nor straightforwardly behavioral.

The parameter λ captures the structural mapping from beliefs to actions: conditional on incorporating news into the forecast, how much does the organization adjust its controlled decision? News affects decisions only indirectly, through updated beliefs: $\tilde{x}_{itw} \to \Delta F_{it} \to \Delta D_{it}$. The action-equation coefficient γ_w therefore reflects two margins jointly—how much window-w news enters the forecast (a_w) and how much the decision responds to the forecast (λ). This is why the test has power: under continuous updating, both margins are stable across windows, so $\gamma_w/a_w = \lambda$ for all w; under decision-window updating, earlier news never reaches the decision because it is omitted at the belief stage. The pass-through λ need not be one-for-one. In a corporate setting, λ might be 0.5 if a 1% sales forecast revision moves

investment by 0.5%, as estimated in **?**. The framework requires only that a stable mapping from beliefs to decisions exists.

3.1.2 Application to state budgets

In our setting, $Y_{s,t}$ is realized state revenue, $F_{s,t}$ is the initial revenue forecast embedded in the governor's proposal, and $D_{s,t}$ is the initial expenditure plan. The news measures $x_{s,t,w}$ are the shift-share Bartik shocks from Section **??**, constructed symmetrically across windows. Two windows are available by the proposal date: pre-planning ($w = \text{pre}$, 6–12 months before the proposal) and planning ($w = \text{plan}$, 0–6 months before).

The non-controlled outcome is revenue: the state does not directly control realized tax collections over the fiscal year. Revenue forecast errors therefore isolate omitted information cleanly. The controlled decision is the expenditure plan: the governor's proposal directly sets planned spending. Expenditure forecast errors are less informative in this setting because realized spending is itself constrained by the enacted budget, so the clean controlled-object test is the revision of the expenditure plan, not the forecast error of realized expenditures.

3.2 Empirical Implementation

We operationalize the framework with three state-year regressions estimated on a common sample, using the same notation as equations (**??**)–(**??**). Let $w \in \{\text{pre}, \text{plan}\}$ index the 6–12 month pre-planning window and the 0–6 month planning window. The empirical belief equation is

$$\log\left(\frac{\text{Init}^R_{s,t}}{\text{Init}^R_{s,t-1}}\right) = \alpha_s + \delta_t + \sum_w a_w \,\text{News}_{s,t,w} + \Gamma' X_{s,t} + \varepsilon^F_{s,t}, \tag{7}$$

which asks which news enters the proposal-stage revenue forecast. The empirical omission equation is

$$\text{FE}^R_{s,t} = \alpha_s + \delta_t + \sum_w \rho_w \,\text{News}_{s,t,w} + \Gamma' X_{s,t} + \varepsilon^{FE}_{s,t}, \tag{8}$$

which asks which available news predicts subsequent revenue forecast errors. Finally, the empirical action equation is

$$\log\left(\frac{\text{Init}^{E}_{s,t}}{\text{Init}^{E}_{s,t-1}}\right) = \alpha_s + \delta_t + \sum_{w} \gamma_w \,\text{News}_{s,t,w} + \Gamma' X_{s,t} + \varepsilon^{D}_{s,t}, \tag{9}$$

which asks which revenue news is translated into a controlled fiscal commitment.

These three sets of coefficients map directly to the framework: a_w measures incorporation into beliefs, ρ_w measures omitted information, and γ_w measures translation into action. Under continuous updating, the restrictions are

$$a_{\text{plan}} = a_{\text{pre}}, \qquad \rho_{\text{plan}} = \rho_{\text{pre}} = 0, \qquad \gamma_{\text{plan}} = \gamma_{\text{pre}},$$

and, when a model-implied pass-through λ is available, the sharper restriction is $\gamma_w = \lambda\, a_w$ for each window w. Under decision-window updating, we instead expect

$$a_{\text{plan}} > a_{\text{pre}}, \qquad \rho_{\text{pre}} > \rho_{\text{plan}} \approx 0, \qquad \gamma_{\text{plan}} > \gamma_{\text{pre}}.$$

Our baseline implementation uses the raw shift-share news measures because they are constructed symmetrically across windows. Appendix Table **??** sharpens the test by rescaling each window's news by its predictive slope for realized revenue, placing the coefficients in common ex post revenue units. The qualitative conclusions are unchanged. Section **??** implements this test.

4 Empirical Design and Results

Implementing the availability–belief–action test requires a measure of exogenous revenue news whose timing relative to the proposal date is well defined. The news must be plausibly informative about state revenues—so that it belongs in the information set the organization should, in principle, incorporate—but not driven by the state's own budget choices. We also need the same class of news to arrive inside the planning window for some states and outside it for others, so that we can compare the effects of identical information under different institutional timing.

We construct such a measure by interacting predetermined state revenue-composition shares with national tax-base innovations, following the shift-share logic of **?** and **?**. Because proposal dates differ across states, the same national innovation in a given calendar month enters different budget windows for different states. This cross-state timing variation is the source of identification: it allows us to separate the causal role of institutional timing from the content of the information itself.

4.1 Empirical Design

4.1.1 The Shift-Share Measure

To measure the flow of fiscally relevant information to each state, we exploit cross-state differences in revenue composition. A state that relies heavily on income taxes is more exposed to national income growth, while a state with large severance tax revenues is more exposed to commodity prices. We choose revenue sources k that capture a substantial share of total general revenue, compute shocks to each source over state-specific budget windows $w(s)$, and interact those shocks with predetermined revenue shares to construct state-specific revenue news:

$$\text{News}_{s,t,w} = \sum_{k=1}^{K} \omega_{s,k,t} \times g_{k,t,w(s)}, \tag{10}$$

where $g_{k,t,w(s)}$ is the innovation in national aggregate k over window $w(s)$ and $\omega_{s,k,t}$ is state s's predetermined reliance on revenue source k. We measure

$$\omega_{s,k,t} = \frac{1}{5} \sum_{\tau=t-7}^{t-2} \frac{\text{Revenue}_{s,k,\tau}}{\text{General Revenue}_{s,\tau}}, \tag{11}$$

so the exposure weights are averaged over the five fiscal years ending two years before fiscal year t.

Our baseline uses five own-source revenue categories matched to national tax-base proxies: individual income taxes (real personal income), general sales taxes (real personal consumption expenditures), charges and miscellaneous revenue (real PCE services), corporate income taxes (corporate profits after tax), and severance taxes (WTI crude oil prices). These

five sources account for approximately 62% of general revenue. The remaining revenue consists primarily of intergovernmental transfers, which are driven more by federal policy than by state-specific macroeconomic exposure. Figure **??** shows the substantial cross-state variation in these exposure weights.

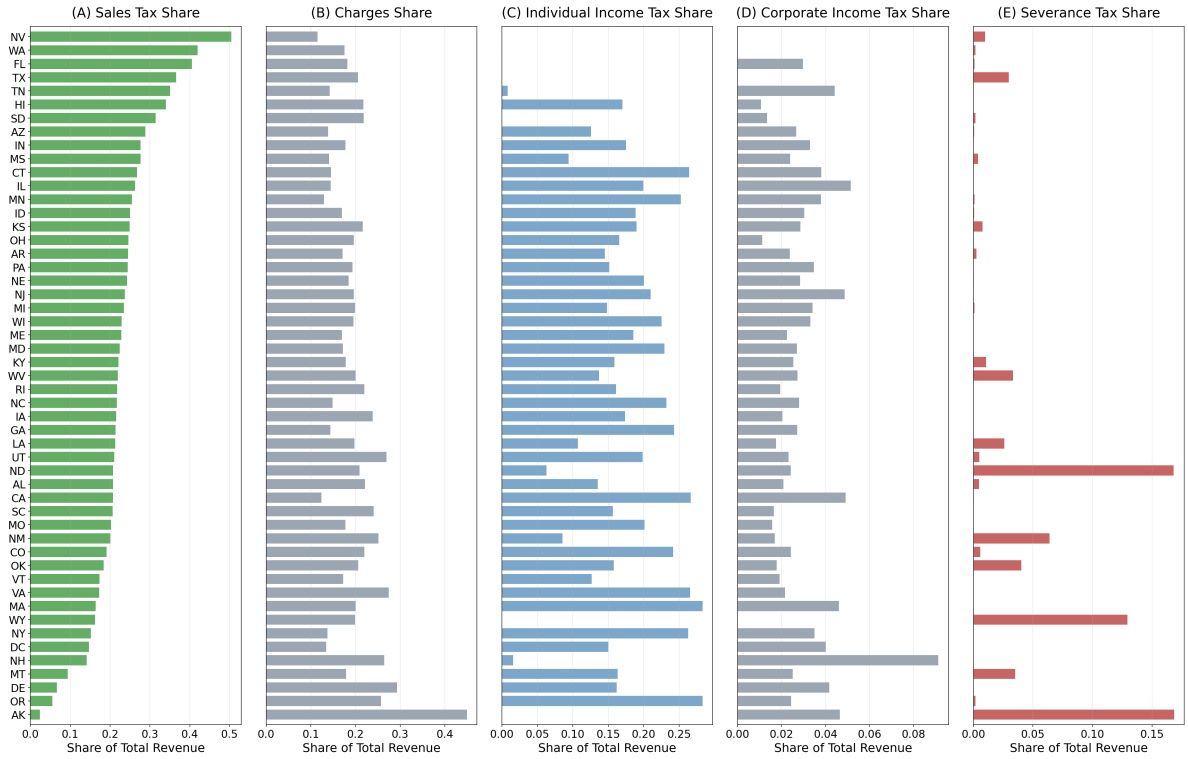


Figure 4: Cross-State Variation in Revenue Composition. Each panel displays the average share of a revenue source in total general revenue for each state over fiscal years 1995–2024. Panel (A): individual income tax; Panel (B): general sales tax; Panel (C): charges and miscellaneous revenue; Panel (D): corporate income tax; Panel (E): severance tax. States are sorted by income tax share in all panels. The five sources constitute the baseline shift-share specification, accounting for approximately 62% of total general revenue. The remaining revenue consists primarily of intergovernmental transfers. Data source: Census Bureau Annual Surveys of State and Local Government Finance.

Because national tax bases are persistent, we purge predictable variation by projecting each cumulative shock on an AR(4) with lagged GDP growth, using only information dated before window w opens. The residual is the component of macroeconomic conditions that was not forecastable from pre-window data: the news that arrives during window w. Full details are in Appendix **??**.

A key feature of the design is that states operate on heterogeneous budget calendars, so

the same national innovation in a given calendar month enters different budget windows for different states. Let $[\tau^0_{sw}, \tau^1_{sw}]$ denote the start and end dates of window w for state s. The raw source-k shock is

$$g^{\text{raw}}_{k,t,w(s)} = \log X_k(\tau^1_{sw}) - \log X_k(\tau^0_{sw}).$$

We define three primary windows: planning (0–6 months before the proposal), revision (proposal to final approval), and realization (during the fiscal year). For the timing tests below, we also construct pre-planning windows—successive six-month intervals prior to the planning window (6–12m, 12–18m, and 18–24m before the proposal). These earlier windows provide the key comparison group: identically constructed news that was already available by the proposal date but arrived before active budget preparation began.

4.1.2 State-Year Timing Specification

The core timing question is whether the proposal responds to all available revenue news or only to news that arrives during the active planning window. We therefore estimate horse-race specifications of the form

$$y_{s,t} = \alpha_s + \delta_t + \beta_{\text{plan}} \text{News}_{s,t,\text{plan}} + \beta_{\text{pre}} \text{News}_{s,t,\text{pre}} + \gamma' X_{s,\text{lags}} + \varepsilon_{s,t}, \tag{12}$$

where $y_{s,t}$ is a fiscal outcome, α_s and δ_t are state and fiscal-year fixed effects, and $X_{s,\text{lags}}$ contains four quarterly lags of real GDP growth, nonfarm payroll employment growth, S&P 500 returns, 5-year Treasury yields, and CPI inflation. These controls are measured to end before the planning window begins, relative to each state's budget calendar. Standard errors are clustered at the state level.

This specification directly implements the availability–belief–action test from Section **??**. We apply it to three outcomes corresponding to the three equations of the conceptual framework: the proposal-stage revenue forecast (equation **??**), the revenue forecast error (equation **??**), and the proposal-stage expenditure plan (equation **??**). Equation (**??**) yields two immediate empirical predictions. If proposal-stage updating is continuous, then planning-window and pre-planning news of comparable informational content should load similarly. If instead the proposal is a discrete review date, the planning-window coefficient should dominate, pre-planning news should predict forecast errors, and expenditure plans should respond disproportionately to planning-window news. Appendix Table **??** sharpens this comparison

by rescaling each window's news into common ex post revenue units; the planning-window premium survives that rescaling, indicating that the contrast reflects differential use of information rather than differential informativeness.

4.1.3 Validation: The Measure Captures State-Relevant Revenue Conditions

Before turning to timing, we verify that the shift-share measure captures economically meaningful variation in realized revenues. Table **??** regresses realized revenue growth on the contemporaneous fiscal-year news measure using equation (**??**). States more exposed to favorable national innovations experience significantly higher realized revenue growth. Appendix Table **??** shows that this relevance is robust to alternative source combinations. The shift-share measure therefore captures state-specific revenue conditions, not merely statistical noise.

4.1.4 Identification

Our shift-share measure has the structure analyzed by **?**, **?**, and **?**, but the object of interest is narrower than in a standard Bartik design. We are not asking whether a single coefficient β is zero. We ask whether the contrast $\beta_{\text{plan}} - \beta_{\text{pre}}$ differs from zero. Because the same national series and the same predetermined exposure weights enter every window, any confounder that is stable across budget windows within a state-year—such as fiscal institutions, forecasting capacity, or persistent political factors—loads symmetrically on both coefficients and therefore differences out of the comparison.

Three features support this interpretation. First, the support of the constructed shocks is similar across planning and pre-planning windows (Table **??**, Panel B; Appendix Figure **??**), so the timing contrast is not driven by mechanically different shock distributions. Second, national tax-base innovations are economy-wide and plausibly orthogonal to any individual state's contemporaneous budget choices. Third, the exposure weights are predetermined, ruling out contemporaneous policy endogeneity in the shares.

4.2 Beliefs: Planning-Window News Dominates Revenue Forecasts

We begin with the proposal-stage revenue forecast—the belief equation in the availability–belief–action system. The dependent variable is the log change in the governor's revenue forecast, $\log\big(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1}\big)$. Using the previous year's initial forecast as the baseline is important: it is fixed before the pre-planning window opens, so earlier news is not mechanically differenced out.

Panel A of Table **??** reveals a sharp timing asymmetry. Planning-window news enters the proposal strongly, with coefficients between 1.09 and 1.32 across specifications, indicating close to unit pass-through of planning-window revenue news into the initial revenue forecast. The coefficient on pre-planning news is essentially zero (-0.01 without controls and -0.10 with controls). The proposal therefore behaves less like a continuously updated forecast and more like a discrete review date: news arriving during active budget preparation moves the forecast, while equally measured news that arrived six to twelve months earlier does not. In the language of Section **??**, $\hat{a}_{\text{plan}} \approx 1$ and $\hat{a}_{\text{pre}} \approx 0$.

Within-Shock Timing Design. A natural concern is that the planning coefficient in the state-year specification reflects the calendar months spanned by the planning window rather than institutional timing per se. Panel B of Table **??** addresses this by moving to a stacked monthly design that compares the same underlying month-level national innovation when it falls inside the planning window for some states and inside the pre-planning window for others.

Let the monthly Bartik shock for state s in calendar month m be $z_{s,m} = \sum_k \omega_{s,k,t} \Delta \log X_{k,m}$, where $\Delta \log X_{k,m}$ is the monthly log change in national aggregate k and the weights are the same predetermined exposure shares used above. For each fiscal year t, we construct a 24-month panel running from January of year $t-2$ through December of year $t-1$, and define a cohort as the set of states that share the same budget calendar. We then estimate

(13)

$$y_{s,m,t} = \gamma_1 z_{s,m} \times \mathbf{1}[\text{Plan}_{s,m,t}] + \gamma_2 z_{s,m} \times \mathbf{1}[\text{PrePlan}_{s,m,t}] + \gamma_0 z_{s,m} + \mathbf{X}'_{s,t}\boldsymbol{\beta} + \alpha_s + \delta_{t\times m\times y} + \varepsilon_{s,m,t},$$

where $\delta_{t\times m\times y}$ denotes cohort $\times$ calendar-month $\times$ year fixed effects. These fixed effects

absorb the national innovation itself, so identification comes from comparing states with different exposure weights to the same month-level shock within the same cohort-month cell.

The within-shock design confirms the state-year result. In the most demanding specification in Panel B of Table **??**, the coefficient on News $\times$ $\mathbb{1}$(Planning) is 1.44, while the coefficient on News $\times$ $\mathbb{1}$(Pre-Plan.) is 0.02. Because the cohort $\times$ calendar-month $\times$ year fixed effects absorb each national innovation entirely, this result cannot be driven by special shocks arriving in special months. It reflects institutional timing.

Figure **??** replaces the two window indicators with a full month-by-month profile around the proposal date. The response is flat from month -12 to month -6, rises sharply once the planning window opens, and peaks near the proposal itself. This profile validates the six-month planning-window definition and makes the organizational interpretation especially transparent: responsiveness turns on when active proposal preparation begins. This specification provides a nonparametric complement testing the belief equation: rather than estimating a_w for two discrete windows, it traces the profile of incorporation of information across all months relative to the proposal, providing a much richer test of the timing predictions.

4.3 Omitted Information: Pre-Planning News Predicts Forecast Errors

Table **??** estimates the omitted-information equation (**??**). The coefficient on planning-window news is small and statistically insignificant, indicating that information arriving during active budget preparation is largely incorporated by the proposal date. The coefficient on pre-planning news is large, positive, and precisely estimated. Favorable revenue news six to twelve months before the proposal leads to realized revenues that exceed the proposal-stage forecast; unfavorable pre-planning news leads to realized revenues that fall short. In the language of the framework, $\hat{\rho}_{\text{pre}}$ is large—roughly 1.08—while $\hat{\rho}_{\text{plan}}$ is close to zero. The omission is strongest for the immediately preceding pre-planning window and largely disappears for still-earlier windows, consistent with signals becoming too stale to predict revenues at the fiscal-year horizon.

Appendix Table **??** shows that this omission is symmetric across positive and negative

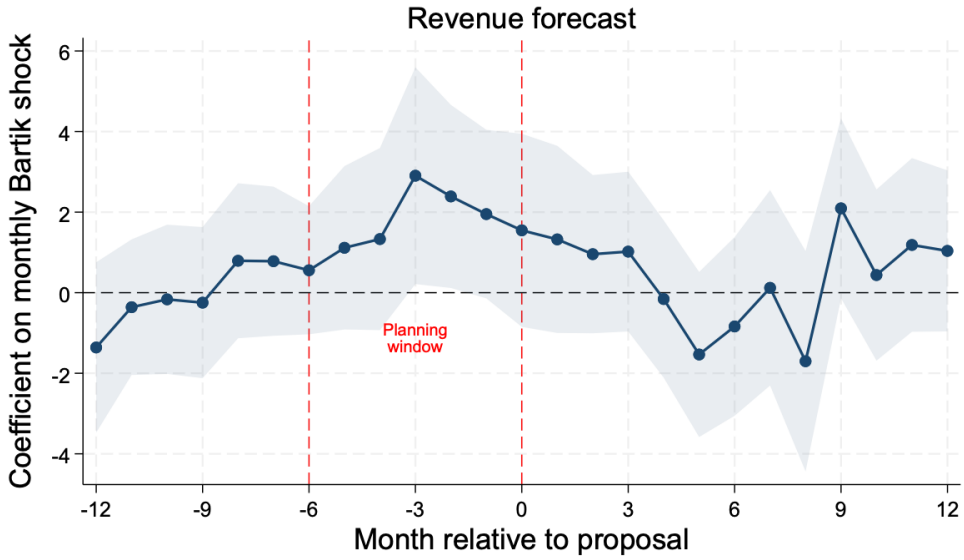


Figure 5: Monthly Profile of Revenue Forecast Updating Around the Proposal Date. The figure plots the coefficient on monthly Bartik news by month relative to the proposal date from the stacked monthly-panel regression with cohort × calendar-month × year fixed effects and lagged macro controls. Coefficients are normalized by subtracting the average coefficient over months -12 to -6. The red dashed lines mark the six-month planning window. Shaded areas denote 95% confidence intervals based on standard errors clustered by state.

news. Appendix Table **??** shows where the omitted information surfaces: the legislature does not systematically correct it, and the omitted signal appears primarily as level miscalibration rather than as a response to current-year volatility.

4.4 Actions: Planning-Window News Moves Controlled Fiscal Decisions

Table **??** estimates the action equation (**??**). Planning-window news not only updates the revenue forecast, but also shifts the initial expenditure plan and, through that plan, realized spending.

The first three columns trace the proposal-stage mechanism. Column (1) reiterates that planning-window news strongly predicts the initial revenue forecast update (coefficient 1.33; Kleibergen–Paap $F = 109$). Column (2) shows that the same news raises the initial expen-

diture plan. Column (3) instruments the expenditure plan with the news-driven component of the revenue forecast and estimates a forecast-mediated pass-through of $\hat{\lambda} = 0.47$: a 1% news-driven increase in the initial revenue forecast is associated with a 0.47% increase in planned spending.

The last two columns show that the proposal-stage decision propagates to realized outcomes. Planning-window news significantly predicts realized expenditure growth, and column (5) implies that a 1% news-driven increase in the revenue forecast is associated with a 0.64% increase in realized spending growth. Pre-planning news does not generate a comparable proposal-stage expenditure response. In the language of the framework, $\hat{\gamma}_{\text{plan}}$ is large while $\hat{\gamma}_{\text{pre}} \approx 0$: the controlled fiscal decision inherits the same timing asymmetry as beliefs.

Panel B of Table **??** confirms the same timing pattern in the stacked monthly panel. Monthly Bartik news interacted with the planning-window indicator predicts both the expenditure plan and realized spending, while the corresponding pre-planning interaction is small and statistically weak.

4.5 Placebo: Later Budget Stages

Table **??** applies the same logic to later stages of the budget cycle. Pre-planning news does not move proposal-stage budgets. Revision-window news predicts later adjustments and realized expenditure growth but has little effect on headline revenue re-forecasting, consistent with legislative bargaining operating more through line-item composition than through aggregate revenue reassessment. Realization-window news predicts revenue over- or under-performance without materially changing realized spending, consistent with limited within-year flexibility. Planning-window news is thus the dominant driver of proposal-stage commitments, while revision and realization windows play narrower roles after the main commitment has been made.

4.6 Robustness

We subject our results to several robustness checks, each targeting a distinct alternative explanation. Full results are in Appendix **??**.

Framework assumptions. The availability–belief–action test in Section **??** rests on two maintained assumptions, each of which we test directly (Appendix **??**). First, the framework treats revenue as a non-controlled outcome. Using Tax Foundation data on state individual income tax top marginal rates (2000–2014), we regress an indicator for any tax rate change on planning-window and pre-planning news. Neither coefficient is statistically significant, confirming that Bartik shocks do not trigger tax legislation (Appendix Table **??**). Second, pre-planning news must contain incremental information that could in principle be used. The correlation between planning-window and pre-planning news is $\hat{\rho} = 0.02$, confirming that the two measures are nearly orthogonal. Pre-planning news strongly predicts revenue forecast errors ($\hat{\rho}_{\text{pre}} = 1.08$, $p < 0.01$) even conditional on planning-window news (Table **??**), ruling out the alternative that earlier information is simply "stale" by proposal time.

Measure construction and outcome definition. The timing result is not an artifact of the particular revenue sources included or the benchmark used to measure updating. Three- and four-source variants produce nearly identical planning-window coefficients and the same sharp drop-off for pre-planning news (Appendix Table **??**). Alternative baselines for proposal updating—including $\text{Init}_t^R/\text{Final}_{t-1}^R$, $\text{Init}_t^R/\text{Final}_{t-2}^R$, and $\text{Init}_t^R/\text{Act}_{t-3}^R$—yield planning-window coefficients between 1.18 and 1.88, with pre-planning news remaining insignificant throughout (Appendix Table **??**).

Calendar composition. Because proposal dates differ across states, the same calendar month falls in the planning window for some states and in the pre-planning window for others. Estimating the horse race separately for early proposers and late proposers yields the same qualitative result in both subsamples: a large planning-window coefficient and an insignificant pre-planning coefficient (Appendix Table **??**). The timing result is therefore not driven by a particular set of calendar months.

Differential informativeness. A different concern is that planning-window shocks are simply more informative about future revenues than earlier shocks. Appendix Table **??** addresses this by rescaling each window's news measure by its predictive slope for realized revenue, placing all windows in common ex post revenue units. The planning-window premium is essentially unchanged. The coefficient gap therefore reflects differential use of information, not just differential informativeness.

Permuted budget calendars. Finally, we randomize the match between each state's institutional calendar and its realized news. Specifically, we reassign shift-share news realizations across states within fiscal year and re-estimate the planning-versus-pre-planning horse race 1,000 times. The actual planning coefficient ($\hat{\beta}_{\text{plan}} = 1.31$) exceeds all but one permuted estimate ($p = 0.001$), and the planning-minus-pre-planning contrast is likewise in the far right tail of the permutation distribution ($p = 0.011$; Appendix Figure **??**). The result depends on the actual match between a state's budget calendar and the timing of the news it faces.

4.7 Relation to Standard Imperfect-Information Models

The availability–belief–action test above is intentionally agnostic about the exact source of limited updating. This subsection asks how the evidence maps into canonical imperfect-information models.

Sticky information. In sticky-information models (**??**), agents update infrequently but, conditional on updating, incorporate their available information set completely. In our setting, the proposal date is the natural observed updating date. Conditional on proposal-stage updating, sticky information predicts that pre-planning and planning-window signals available by that date should both be incorporated. The belief and omission equations reject that prediction: planning-window news enters the forecast, while pre-planning news does not and instead predicts forecast errors.

A standard **?** forecast-revision test is reported in Table **??** as a benchmark, but it is not the most informative statistic in this environment. The revision object spans fiscal years rather than holding the forecast horizon fixed, the relevant updating date is observed rather than latent, and revenue revisions can affect outcomes because budget choices have real consequences. The availability–belief–action system is more tightly aligned with the institutional setting.

Rational inattention. Rational-inattention models (**??**) can rationalize selective updating if pre-planning signals are less informative or less worth processing. The key comparative-static prediction is that earlier news should be used more when it is more informative or when

the cost of forecast errors is higher. We test that prediction by interacting planning-window and pre-planning news with proxies for informativeness (the Herfindahl index of exposure weights, a rolling signal-to-noise ratio) and fiscal stakes (balanced-budget-rule stringency, rainy day fund balances).

Table **??** provides little support for that mechanism. More informative or more consequential pre-planning news is not systematically incorporated more heavily into the proposal. The evidence is difficult to reconcile with the simplest versions of optimal selective attention based solely on signal precision or stakes.

Diagnostic expectations. Diagnostic-expectations models (**?**) imply overweighting of the most recent, salient signal, which would show up as overreaction to planning-window news and a negative coefficient on planning-window news in the forecast-error equation. We do not find that pattern. The planning-window coefficient in the revenue forecast-error regression is small and statistically insignificant. The pattern is underreaction to earlier information, not overshooting on recent information.

Political incentives. A final alternative is that officials strategically ignore earlier information for political reasons. Appendix Table **??** interacts the news measures with governor party and election timing. Neither proxy materially increases the incorporation of pre-planning news into the proposal, and pre-planning news continues to predict forecast errors across political environments. The narrow decision window is difficult to attribute to partisan composition or electoral timing alone.

5 Real Effects of the Decision Window

Section **??** establishes that fiscal-year spending adjusts primarily at the proposal stage: planning-window news drives revenue forecasts and spending commitments, while earlier news is systematically neglected. If the decision window shapes government spending, it should also generate real effects in the private sector. This section tests that prediction by examining employment and wages in government-exposed industries.

5.1 Employment and Wages in Government-Exposed Industries

If planning-window news shifts the fiscal-year spending budget, it should raise labor demand in industries that sell to state and local governments. To identify such industries, we measure exposure to government procurement from BEA input–output accounts. We then test this prediction using the Quarterly Census of Employment and Wages (QCEW), which provides private employment and wage payments at the state–industry level.

Data and outcomes. The unit of observation is a state–industry–fiscal-year, restricted to private establishments and two-digit NAICS industries.[4] Employment is a stock measured at a point in time. For each state s, industry j, and fiscal year t, let $\text{Emp}^{\text{end}}_{s,j,t}$ denote average employment in the last calendar quarter contained in fiscal year t. Our employment outcome is the year-over-year log change,

$$\Delta \log \text{Emp}^{FY}_{s,j,t} \equiv \log\big(\text{Emp}^{\text{end}}_{s,j,t}\big) - \log\big(\text{Emp}^{\text{end}}_{s,j,t-1}\big) .$$

Wages are a flow. We construct the fiscal-year wage bill by summing quarterly wage payments within the fiscal year,

$$\text{WageBill}^{FY}_{s,j,t} \equiv \sum_{q \in FY_t} \text{Wages}_{s,j,q}, \qquad \Delta \log \text{WageBill}^{FY}_{s,j,t} \equiv \log\big(\text{WageBill}^{FY}_{s,j,t}\big) - \log\big(\text{WageBill}^{FY}_{s,j,t-1}\big) .$$

These definitions match the fiscal-year horizon over which budget plans are implemented and allow a direct comparison between planning-window news and within-year labor-market outcomes.

Government exposure. Our identification strategy exploits cross-industry heterogeneity in reliance on state and local government demand. We measure exposure using BEA input–output accounts. For each industry j and calendar year y, we compute the government-purchase share as the ratio of state and local government purchases to total commodity output. To reduce year-to-year noise, we take a five-year trailing average and lag it two

[4] We exclude residual and non-classifiable sectors and drop industries with inconsistent reporting.

years relative to the fiscal year:

$$\text{GovShare}_{j,t} \equiv \overline{\text{GovShare}}^{\text{5yr}}_{j,t-2}.$$

We then define

$$\text{GovExp}_{j,t} = \mathbf{1}[\text{GovShare}_{j,t} \geq 0.05]\,,$$

an indicator for industries whose (smoothed, lagged) government-purchase share exceeds five percent. By construction, $\text{GovExp}_{j,t}$ is predetermined with respect to the budget cycle and the news realizations we study. Results are robust to alternative thresholds, reported in Section **??**.

Empirical strategy. The decision-window mechanism implies a sharp cross-sectional prediction: for a given state-year news realization, industries with $\text{GovExp}_{j,t} = 1$ should respond more strongly than low-exposure industries because only the former are directly tied to government procurement. We implement this logic in a difference-in-differences design that compares exposed and non-exposed industries within the same state and fiscal year. Specifically, we estimate

$$\Delta \log Y^{FY}_{s,j,t} = \alpha_{s,j} + \alpha_{j,t} + \alpha_{s,t} + \sum_{w \in \{\text{pre,plan,rev,real}\}} \gamma_w \big(u_{s,t,w} \times \text{GovExp}_{j,t}\big) + \varepsilon_{s,j,t}, \tag{14}$$

where $Y \in \{\text{Emp}, \text{WageBill}\}$ is measured during the fiscal year and $u_{s,t,w}$ is the shift-share news measure constructed over budget window w as in Section **??**. All regressions are weighted by lagged industry employment to ensure that estimates reflect aggregate employment effects rather than being driven by small state-industry cells. Standard errors are clustered by state.

Identification. Unlike the forecast regressions in Section **??**, where the dependent variable is a government planning choice with no direct feedback to the macroeconomy, the outcomes here are private-sector equilibrium quantities. The shift-share measure must therefore satisfy a stronger exclusion restriction: national tax-base innovations can affect industry-level employment differentially across exposed and non-exposed industries only through the government spending channel. The triple fixed-effect structure in (**??**) largely addresses this concern. State×fiscal-year fixed effects absorb all state-year variation, includ-

ing the level effect of news, contemporaneous state-level shocks, and any state-year policy changes. Industry×fiscal-year fixed effects absorb national industry cycles. Identification therefore comes from within-state-year differential responses across industries: γ_w compares how government-exposed versus non-exposed industries respond to the same state-level news realization. Residual concerns—that the result could be driven by a small number of influential states, a single dominant revenue base, or the particular exposure threshold—are addressed in the robustness checks of Section **??**.

5.2 Results

Table **??** reports the estimates. Planning-window news produces large and precisely estimated differential increases in government-exposed industries: the interaction coefficient γ_{plan} is positive and highly significant for both employment and wages across all specifications. Because planning-window news arrives months before the fiscal year begins, yet the employment and wage outcomes are measured during the fiscal year, the response reflects a transmission lag consistent with the procurement channel: news enters the revenue forecast, shapes the spending plan, and reaches private firms only once contracts are executed.

Pre-planning news, by contrast, has no detectable differential effect, mirroring the budget-stage finding that earlier news does not enter the spending plan. This placebo provides direct evidence for the exclusion restriction: if the differential response reflected general cyclical sensitivity of government-exposed industries rather than procurement demand, earlier news should generate comparable effects. Revision-window news shows a positive interaction with government exposure for wages, consistent with legislatures reallocating line items toward procurement-intensive uses during the deliberation phase.

The concentration of the response in government-exposed industries helps distinguish procurement demand from broader mechanisms. A private-demand channel—for example, income-driven consumption—would raise activity across many industries within a state-year, not differentially in government-exposed sectors. The within-state-year fixed-effect structure absorbs such state-level confounds. The decision window is thus not merely an accounting feature of the budget process—it generates real effects in private labor markets, concentrated in industries that sell to state and local governments. Section **??** reports additional robustness checks targeting influential states, exposure definitions, and shift-share

construction.

Dynamic propagation. The annual results establish that planning-window news affects government-exposed industries, but not when the response materializes. If the mechanism is procurement demand—news enters the forecast, shapes the spending plan, and reaches private firms only once contracts are executed—the employment response should be absent before the fiscal year and build gradually during execution. We test this timing prediction using **?** local projections at the state–industry–quarter level. Let $\text{Emp}^q_{s,j}$ denote average quarterly employment. For each horizon $h = -8, \ldots, +12$ quarters relative to the fiscal-year start, we estimate

$$\log \text{Emp}^q_{s,j,h} - \log \text{Emp}^q_{s,j,-1} = \alpha_{s,j} + \alpha_{j,t} + \alpha_{s,t} + \gamma_h \big(u_{s,t,\text{plan}} \times \text{GovExp}_{j,t}\big) + \varepsilon_{s,j,t,h}, \tag{15}$$

where the baseline is one quarter before the fiscal-year start ($h = -1$), by which point both planning- and pre-planning-window news have been realized. Fixed effects and clustering are identical to (**??**). Plotting $\{\gamma_h\}$ traces out the differential cumulative employment response in government-exposed industries. We estimate (**??**) separately using planning-window news, pre-planning news (placebo), and revision-window news.

Figure **??** traces the dynamic response. Panel A shows that planning-window news has no differential effect on government-exposed employment before the fiscal year begins—the pre-fiscal-year coefficients are indistinguishable from zero. The response emerges only after the fiscal year starts ($h = 0$ to $h \approx 3$) and persists beyond fiscal-year end, consistent with the transmission lag from budget enactment through procurement to private hiring. The delayed onset rules out anticipation effects: industries respond to budget implementation, not to news itself.

Panels B and C provide sharp placebo tests. Pre-planning news—constructed identically but arriving before the decision window—generates no detectable differential response at any horizon, confirming that the real effects require news to enter the spending plan. Revision-window news also generates a positive response, though less precisely estimated, consistent with legislative revisions shifting spending composition toward procurement-intensive uses. Together with the insignificant pre-planning coefficient in Table **??**, the flat pre-trend in Panel A provides direct evidence for the exclusion restriction: government-exposed industries do not differentially respond to macroeconomic news per se, only to news that enters the

spending plan through the decision window.

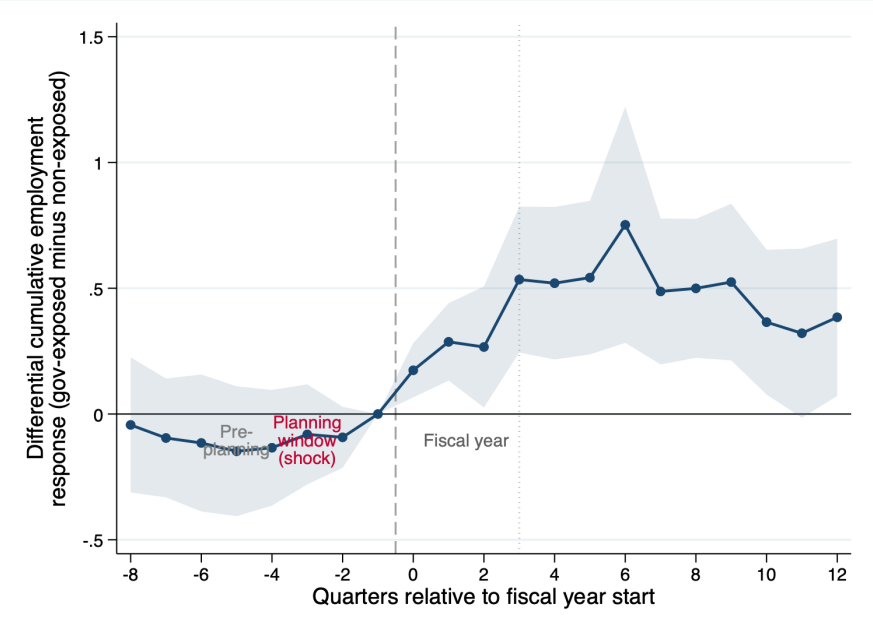
5.3 Robustness

We subject the real-effects estimates to several checks targeting the residual identification concerns flagged above. Full results are in Appendix **??**.

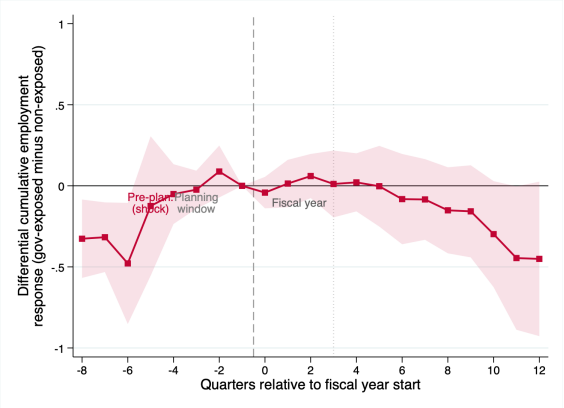
Influential states and revenue bases. A standard concern with shift-share designs is that a small number of units with extreme exposure weights may drive the result (**?**). We address this in two ways. First, we re-estimate the specification dropping one state at a time; the planning-window coefficient clusters tightly around the full-sample estimate across all 47 jackknife samples, with 95% confidence intervals excluding zero in every iteration (Appendix Figure **??**). Second, we drop states whose average severance tax share exceeds 5% of general revenue, since oil-producing states have unusual revenue composition that could dominate the shift-share variation. The planning-window coefficient is essentially unchanged (Appendix Table **??**, column 2).

Exposure definition. The 5% government-purchase threshold could in principle select industries that differ from non-exposed industries in ways correlated with their response to macroeconomic news. Appendix Table **??** shows that the result is robust to using the median government-purchase share as the threshold (column 4), raising the threshold to 6% (column 5), replacing the binary indicator with the continuous government-purchase share (column 6), and removing employment weights (column 3). The planning-window interaction remains significant across all specifications.

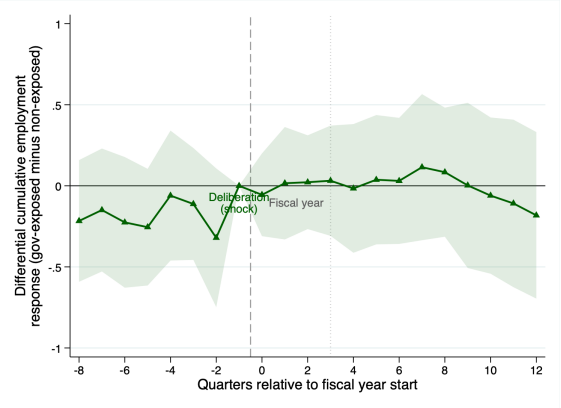
Alternative shift-share constructions. To verify that the result does not depend on a single revenue base entering the shift-share measure, we replace the baseline 5-base construction with 3-base (individual income, severance, and corporate income taxes) and 4-base (adding general sales tax) alternatives. The planning-window coefficient is stable across all three specifications (Appendix Table **??**), consistent with the timing pattern reflecting broad revenue news rather than idiosyncratic movements in any single tax base.



(a) Planning-Window News



(b) Pre-Planning News (Placebo)



(c) Revision-Window News

Figure 6: Impulse Responses: Differential Employment in Government-Exposed Industries. Each panel plots $\{\gamma_h\}$ from local projections (**??**) estimated separately for planning-window news (Panel A), pre-planning news (Panel B), and revision-window news (Panel C). The dependent variable is cumulative log employment change relative to one quarter before fiscal-year start ($h = -1$). Government exposure is defined as industries with government-purchase share exceeding 5%. All regressions are weighted by lagged employment. Shaded areas are 95% confidence intervals based on standard errors clustered by state. All specifications include state×industry, industry×fiscal-year, and state×fiscal-year fixed effects.

Inference. With 47 state clusters, standard cluster-robust standard errors may under-reject (**?**). The leave-one-state-out exercise shows that the coefficient is stable with no outlier-driven inflation, and the triple fixed-effect structure absorbs most potential sources of cross-sectional dependence, making state-level clustering a conservative choice.

6 A Structural Model of the Budget Process

The reduced-form analysis and the private-sector evidence establish that the decision-window bias has real consequences for employment, but they cannot answer quantitative questions. First, how much of the attenuation in fiscal pass-through is due to the decision-window bias versus institutional frictions that dampen even incorporated information? Second, how large is the employment cost of ignoring pre-planning news, and does it differ for positive and negative shocks?

To address these questions, we develop a model that represents the state budget as a sequence of stages—forecasting, proposal, enactment, and execution—each with parameters estimated by matching model-implied moments to their empirical counterparts.

6.1 Model

Fiscal condition. Each fiscal year t, state s faces a fiscal condition $\theta_{s,t}$ (log-deviation from trend) that evolves as

$$\theta_{s,t} = \rho\,\theta_{s,t-1} + z^{\text{pre}}_{s,t} + z^{\text{plan}}_{s,t} + z^{\text{rev}}_{s,t} + z^{\text{real}}_{s,t} + \nu_{s,t}, \tag{16}$$

where z^w are economic shocks arriving during each budget window—pre-planning, planning, revision, and realization—and $\nu_{s,t}$ is idiosyncratic noise. The shocks correspond to the shift-share news measures constructed in the empirical analysis: observed macroeconomic developments (stock returns, employment growth, consumption) interacted with predetermined state revenue-composition weights. Actual revenue equals the fiscal condition: $R^{\text{act}}_{s,t} = \theta_{s,t}$.

Revenue forecasting. Officials form revenue forecasts by selectively incorporating information from each window. The **initial forecast**, set at the proposal stage, is

$$\hat{R}^I_{s,t} = \rho\,\theta_{s,t-1} + K_{\text{pre}}\,z^{\text{pre}}_{s,t} + K_{\text{plan}}\,z^{\text{plan}}_{s,t}. \tag{17}$$

The parameters K_{pre} and K_{plan} govern how much of each window's news enters the initial forecast. These are free parameters that the data will discipline. The **final forecast**, formed at enactment, updates with revision-window news:

$$\hat{R}^F_{s,t} = \hat{R}^I_{s,t} + K_{\text{rev}}\,z^{\text{rev}}_{s,t}. \tag{18}$$

The revenue forecast error relative to the initial forecast is

$$\text{FE}_{s,t} \equiv \theta_{s,t} - \hat{R}^I_{s,t} = (1 - K_{\text{pre}})\,z^{\text{pre}}_{s,t} + (1 - K_{\text{plan}})\,z^{\text{plan}}_{s,t} + z^{\text{rev}}_{s,t} + z^{\text{real}}_{s,t} + \nu_{s,t}. \tag{19}$$

News from any window that is not fully incorporated into the initial forecast passes into the forecast error.

Budget proposal. The initial expenditure plan balances two forces: maintaining prior service commitments and tracking expected revenue under balanced-budget pressure. The governor minimizes

$$\min_{E^I} \;\; \left(E^I - E^F_{s,t-1}\right)^2 \;+\; \phi\left(E^I - \hat{R}^I_{s,t}\right)^2. \tag{20}$$

The first term captures the cost of disrupting existing programs and contractual commitments: changing spending levels requires renegotiating contracts, adjusting staffing, and reallocating across agencies. The second term captures balanced-budget pressure: spending that deviates from expected revenue risks mid-year cuts (if too high) or foregone services (if too low). The parameter $\phi > 0$ indexes the relative strength of these two forces. The first-order condition yields

$$E^I_{s,t} = \frac{E^F_{s,t-1} + \phi\,\hat{R}^I_{s,t}}{1 + \phi}. \tag{21}$$

The proposal pass-through—the marginal effect of the revenue forecast on planned spending—is $\phi/(1+\phi)$. When ϕ is large, balanced-budget pressure dominates and spending closely tracks

the revenue forecast. When ϕ is small, prior commitments anchor spending and the budget responds sluggishly to new revenue information.

Budget enactment. Between the proposal and the start of the fiscal year, the legislature observes revision-window news and updates the budget:

$$E^F_{s,t} = \frac{\alpha\, \hat{R}^I_{s,t} + \alpha_F \Big(\hat{R}^F_{s,t} - \hat{R}^I_{s,t} \Big) + \kappa\, E^I_{s,t}}{1 + \kappa}. \tag{22}$$

The enacted budget is a weighted average of the legislature's desired spending and the governor's proposal E^I, with $\kappa > 0$ indexing the anchoring power of the proposal. The legislature's target has two components: α weights the base revenue forecast $\hat{R}^I$, and α_F weights genuinely new revision-stage information ($\hat{R}^F - \hat{R}^I$). The revision pass-through is $\alpha_F/(1+\kappa)$.

Budget execution. During the fiscal year, the revenue surprise $S_{s,t} \equiv \theta_{s,t} - \hat{R}^F_{s,t}$ is realized. Actual spending deviates from the enacted budget according to

$$E^{\text{act}}_{s,t} = E^F_{s,t} - \bar{u}_{\text{exec}} + \xi_{\text{exec}}\, \max(S_{s,t},\, 0) - c_{\text{neg}}\, \max(-S_{s,t},\, 0). \tag{23}$$

The constant $\bar{u}_{\text{exec}} > 0$ captures systematic under-execution. The parameters ξ_{exec} and c_{neg} govern within-year spending adjustment to revenue windfalls and shortfalls, respectively; the asymmetry reflects that balanced-budget rules make downward adjustments easier than upward ones.

Private-sector response. Government spending transmits to the private sector through procurement demand. Government-exposed employment and wages respond to actual spending:

$$\text{Emp}^{\text{gov}}_{s,t} = \kappa_E\, E^{\text{act}}_{s,t} + \varepsilon^E_{s,t}, \tag{24}$$

$$\text{Wage}^{\text{gov}}_{s,t} = \kappa_W\, E^{\text{act}}_{s,t} + \varepsilon^W_{s,t}. \tag{25}$$

6.2 Estimation and Results

We estimate the model by minimum distance, simulating a panel of 47 states over 30 years and matching model-implied moments to their empirical counterparts. We target 15 moments spanning the main empirical results—the variance decomposition, forecasting regressions, revenue-to-spending pass-throughs, the execution gap, revision-window dynamics, and the private-sector loadings—and evaluate the model on 30+ additional untargeted moments. Full details on the estimation procedure, moment selection, identification, and fit are reported in Appendix **??**.

Parameter estimates. Table **??** reports the estimates. Three features stand out. First, pre-planning information receives essentially zero weight ($K_{\text{pre}} \approx 0$), while planning-window news is incorporated with $K_{\text{plan}} = 1.18$. This is the structural counterpart of the decision-window bias documented in the reduced-form analysis. The estimate $K_{\text{plan}} > 1$ likely reflects that the Bartik news measure captures only part of the revenue-relevant information arriving during the planning window; officials also observe state-specific tax receipt data, budget staff analysis, and legislative testimony that load on the same underlying factors, so the measured news understates the total signal.[5] Second, the proposal pass-through is $\phi/(1+\phi) = 0.67$: for every dollar of news-driven revenue forecast revision, planned spending increases by 67 cents. The revision pass-through is $\alpha_F/(1+\kappa) = 0.72$. Third, the execution stage exhibits asymmetric flexibility: shortfalls pass through at rate $c_{\text{neg}} = 0.22$ while windfalls pass through at $\xi_{\text{exec}} = 0.29$, with systematic under-execution of 7%.

The model matches all 15 targeted moments within 1.4 standard errors and tracks most untargeted moments well (Appendix Tables **??**–**??** and Figure **??**). The planning share of expenditure variance is somewhat underpredicted (0.65 vs. 0.73), reflecting the model's difficulty in fully concentrating variation at the proposal stage; the revision and realization shares are well-matched.

[5] The model's worst-fitting targeted moment, $\beta_{\text{plan,FER}}$, is related: $K_{\text{plan}} > 1$ implies planning news is slightly over-incorporated, predicting a small negative forecast-error coefficient (-0.13), whereas the data show a small positive but insignificant coefficient (0.12, $p > 0.5$). The data cannot reject the model's prediction.

6.3 Counterfactual Decomposition

The estimated model allows us to decompose the attenuation of fiscal pass-through into two channels: information bias ($K_{\text{pre}} \approx 0$, so pre-planning shocks never enter the forecast) and institutional frictions (proposal anchoring, enactment rigidity, and limited execution flexibility, so even incorporated information is only partially passed through). We simulate the employment response to a one-standard-deviation revenue shock under three scenarios:

(A) *No friction*: full information ($K_{\text{pre}} = K_{\text{plan}} = K_{\text{rev}} = 1$) and frictionless budget ($\phi \to \infty$, $\kappa = 0$, $\xi_{\text{exec}} = c_{\text{neg}} = 1$, $\bar{u}_{\text{exec}} = 0$).

(B) *Institutional frictions only*: full information but estimated institutional frictions (proposal anchoring, enactment rigidity, execution constraints).

(C) *Both frictions*: estimated information parameters and estimated institutional frictions.

The comparison A vs. B isolates the cost of institutional frictions when information is fully incorporated. The comparison B vs. C isolates the additional cost of information bias: how much is lost when officials also ignore pre-planning news.

We examine three dimensions of the results.

(a) Pre-planning shocks are dramatically attenuated. Table **??**, Panel A compares the cumulative employment response under all three scenarios. A planning-window shock generates 143% of the frictionless response (column 1): information arriving during active budget preparation passes through fully, with slight amplification from $K_{\text{plan}} > 1$. Pre-planning shocks tell a starkly different story. A positive pre-planning shock retains only 7% of the frictionless response (column 2); a negative pre-planning shock retains 43% (column 3). Equally informative news has dramatically different real effects depending solely on when it arrives: pre-planning shocks are largely lost, while planning-window shocks largely pass through.

(b) Information bias is the dominant friction for pre-planning shocks. The right panel of Figure **??** reveals the source of the attenuation. "Institutional frictions only" (B, red

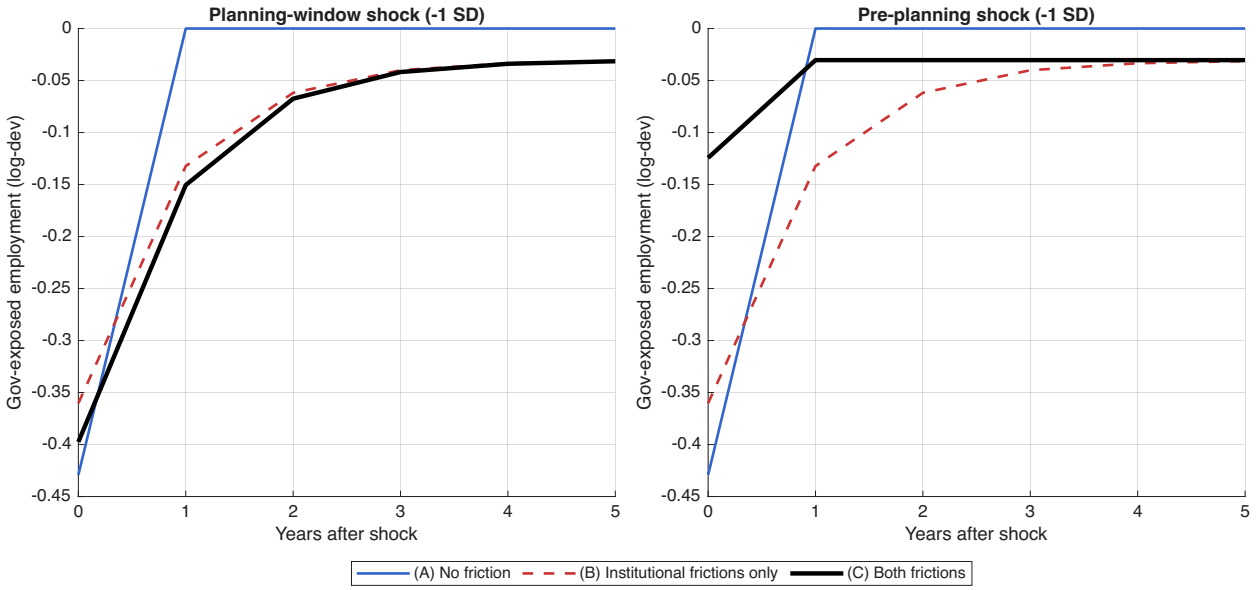


Figure 7: Counterfactual Employment Impulse Responses. Each panel traces government-exposed employment (log-deviation) over six years following a one-standard-deviation negative revenue shock. Left panel: shock arrives during the planning window. Right panel: shock arrives before the planning window. The three scenarios correspond to Table **??**: No friction (A), Institutional frictions only (B), and Both frictions (C). In the right panel, the gap between (B) and (C) shows the cost of information bias: when officials incorporate pre-planning news (B), employment responds strongly despite institutional frictions; when they do not (C), the response is near zero.

dashed) shows a large employment response: when officials incorporate pre-planning news, employment responds strongly even with estimated proposal anchoring, enactment rigidity, and execution constraints. "Both frictions" (C, black) sits near zero: adding information bias on top of institutional frictions collapses the response. The gap between (B) and (C) is the cost of the decision-window bias of ignoring pre-planning news. Panel B of Table **??** quantifies this: for a positive pre-planning shock, information bias ($B-C$) accounts for 86% of the total attenuation and institutional frictions ($A-B$) for only 14%. The pattern reverses for planning-window shocks (left panel), where the information is in the budget pipeline and institutional frictions are the binding constraint.

(c) The bias cushions downturns but blocks expansions. The decision-window bias has strikingly asymmetric consequences for positive and negative pre-planning shocks (Table **??**, columns 2 vs. 3). A negative shock retains 43% pass-through: the revenue shortfall

eventually arrives as an execution-stage surprise, forcing spending cuts through balanced-budget enforcement ($c_{\text{neg}} = 0.22$). A positive shock retains only 7%: the windfall also arrives at execution, but systematic under-execution ($\bar{u}_{\text{exec}} = 0.07$) and limited windfall pass-through ($\xi_{\text{exec}} = 0.29$) mean that most of the surplus is saved rather than spent.

This asymmetry means the decision-window bias partially functions as an unintentional automatic stabilizer for negative shocks—the government does not cut spending when it "should," cushioning employment—while systematically failing to expand spending when revenue conditions improve. The bias is most consequential on the upside: good news that arrives before the planning window generates almost no fiscal or employment response.

Connection to reduced-form evidence. The structural decomposition rationalizes the central reduced-form findings. The near-zero effect of pre-planning news on spending plans and on private-sector employment (Tables 3, 4, and 10) is attributed to $K_{\text{pre}} \approx 0$: the information never enters the budget pipeline. The B-vs-C comparison quantifies what this costs: when officials incorporate pre-planning news (B), employment responds strongly; when they do not (C), the response collapses. Conversely, the strong effects of planning-window news reflect $K_{\text{plan}} \approx 1$ combined with substantial but incomplete institutional pass-through ($\phi/(1+\phi) = 0.67$).

Policy implications. The counterfactual results suggest two conclusions. First, institutional reforms targeting the timing of budget preparation—such as earlier proposal deadlines that would shift pre-planning shocks into the active planning window—could substantially improve fiscal pass-through, especially for positive shocks, where the current retained effect is only 7% compared to 43% for negative shocks. Second, reforms targeting institutional flexibility (e.g., relaxing mid-year spending adjustment rules) would have limited effect for pre-planning shocks because the information was never incorporated into the budget; such reforms would primarily benefit the transmission of planning-window shocks, which already pass through nearly fully.

7 Conclusion

Governments receive a continuous flow of economic information, but fiscal decisions are made at discrete points in the budget calendar. Using a newly assembled panel of hand-collected state budget documents for 47 states over 1995–2024, we show that this timing friction is consequential. The initial budget accounts for 73% of within-state expenditure variation, and only revenue news arriving during the narrow planning window—the months leading up to the governor's proposal—enters the budget. Equally informative news arriving before the planning window is systematically neglected, generating predictable forecast errors that are not easily reconciled with the simplest versions of rational inattention or sticky information. The decision window has real consequences: planning-window news raises employment and wages in government-exposed private industries, while off-window news of similar magnitude does not. A structural model of the budget process indicates that the decision-window bias, not institutional delay, is the primary source of attenuated fiscal pass-through for pre-planning shocks, with asymmetric consequences—negative shocks partially pass through via balanced-budget enforcement, while positive shocks are largely lost.

These findings suggest that the real effects of economic shocks on public spending—and, through spending, on private economic activity—depend not only on their size but also on their alignment with institutional decision cycles. For the design of intergovernmental transfers, revenue-sharing formulas, and federal stabilization programs, the implication is that timing relative to the state budget calendar may matter as much as magnitude. More broadly, our analysis highlights the value of expectations and plans data for studying organizational decision-making (**??**): by observing what officials planned, what they enacted, and what actually happened, we can identify where in the decision process information is lost. Whether similar decision-window dynamics operate in other settings where commitment processes are periodic but information arrives continuously—such as corporate capital budgeting or central bank forecasting—is an open question for future research.

Table 1: Summary Statistics

Panel A reports summary statistics for total actual revenues and expenditures and for their log changes during planning, revision, and realization as defined in Equation (**??**). Panel B reports the distribution of the shift-share news measure across the four budget windows used in the analysis.

Panel A: Fiscal Changes

	Mean	SD	P5	P25	Median	P75	P95
Spending							
Actual total spending (Bn)	19.681	24.066	1.624	4.940	10.952	25.281	66.545
Planning-log($\text{Init}^E_{s,t}/\text{Final}^E_{s,t-1}$)	-0.004	0.100	-0.158	-0.039	0.004	0.041	0.137
Revision-log($\text{Final}^E_{s,t}/\text{Init}^E_{s,t}$)	0.027	0.083	-0.056	-0.011	0.005	0.047	0.192
Realization-log($\text{Act}^E_{s,t}/\text{Final}^E_{s,t}$)	-0.103	0.103	-0.336	-0.124	-0.071	-0.042	-0.019
Revenues							
Actual total revenue (Bn)	20.688	25.328	2.040	5.116	10.999	27.664	68.421
Planning-log($\text{Init}^R_{s,t}/\text{Final}^R_{s,t-1}$)	0.004	0.088	-0.130	-0.026	0.014	0.041	0.109
Revision-log($\text{Final}^R_{s,t}/\text{Init}^R_{s,t}$)	0.019	0.066	-0.051	-0.012	-0.001	0.032	0.138
Realization-log($\text{Act}^R_{s,t}/\text{Final}^R_{s,t}$)	-0.044	0.110	-0.177	-0.058	-0.024	-0.004	0.049

Panel B: Shift-Share News by Budget Window (×100)

	Mean	SD	P5	P25	Median	P75	P95
News Planning (0–6m)	-0.12	1.09	-1.14	-0.38	-0.07	0.25	0.83
News Pre-Plan. (6–12m)	-0.11	0.93	-1.22	-0.41	-0.06	0.25	0.84
News Revision (post-proposal)	-0.10	1.08	-1.09	-0.28	-0.06	0.24	0.91
News Realization (during FY)	-0.25	1.45	-2.18	-0.44	-0.09	0.29	0.99

Table 2: Validation: Shift-Share Measure Predicts Realized State Revenues

The dependent variable is realized revenue growth, $\log(\text{ActR}_{s,t}/\text{ActR}_{s,t-1})$. The shift-share measure is computed over the full fiscal year using the five-base specification (individual income, sales, charges, severance, and corporate income taxes). Columns (1)–(2) use raw fiscal-year log changes in national tax bases; columns (3)–(4) use AR(4)+GDP innovation residuals. Odd columns include no controls or fixed effects; even columns include state and year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Actual Revenue Growth: $\log(\text{ActR}_{s,t}/\text{ActR}_{s,t-1})$			
	Raw Shocks		AR(4)+GDP Innovations	
	(1)	(2)	(3)	(4)
Shift-Share News (Fiscal Year)	1.38**	1.41**	1.65*	1.70**
	(0.67)	(0.70)	(0.83)	(0.84)
State + FYear FE	✓	✓	✓	✓
Macro Controls		✓		✓
Observations	1,061	1,061	1,061	1,061
Adj. R^2	0.08	0.07	0.09	0.08

Table 3: Timing of News and Initial Budget Formation

The dependent variable in both panels is the log change in the initial revenue forecast relative to the prior year's initial forecast, $\log\left(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1}\right)$. Panel A reports the state-fiscal-year specification. "Planning" denotes shift-share news over the 0–6 months before the budget proposal; "Pre-planning" denotes shift-share news over the preceding 6–12 months. Panel B reports the stacked monthly-panel specification, where monthly Bartik shocks are interacted with planning and pre-planning indicators. Across columns, Panel B adds progressively richer fixed effects and macro controls. The shift-share measure uses the five-base specification with AR(4)+GDP innovation shocks. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: State-fiscal-year specification

	Revenue forecast: $\log\left(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1}\right)$			
	(1)	(2)	(3)	(4)
News Planning (0–6m ahead)	1.09***		1.09***	1.32***
	(0.24)		(0.23)	(0.12)
News Pre-Plan. (6–12m ahead)		−0.06	−0.01	−0.10
		(0.32)	(0.31)	(0.34)
State + FYear FE	✓	✓	✓	✓
Macro Controls				✓
Observations	990	990	990	990
Adj. R^2	0.05	0.03	0.05	0.04

Panel B: Monthly panel specification

	Revenue forecast: $\log\left(\text{Init}^R_{t}/\text{Init}^R_{t-1}\right)$			
	(1)	(2)	(3)	(4)
News $\times \mathbb{1}(Planning)$	1.04**	1.13**	1.44***	1.42***
	(0.45)	(0.44)	(0.43)	(0.31)
News $\times \mathbb{1}$(Pre-Plan.)	0.00	-0.00	0.02	0.02
	(0.11)	(0.12)	(0.27)	(0.26)
State FE	✓	✓	✓	✓
Fiscal-Year FE	✓	✓		
Cal.-Month × Year FE		✓		
Cohort × Cal.-Month × Year FE			✓	✓
Macro Controls				✓
Observations	23,760	23,760	23,760	23,760
Adj. R^2	0.098	0.084	0.071	0.084

Table 4: Planning-Window Revenue News Passes Through to Spending Plans and Realized Expenditure Growth

Panel A reports the state-fiscal-year pass-through estimates. Columns (1) and (3) regress (i) the initial expenditure plan, $\log\left(\text{Init}_t^E/\text{Init}_{t-1}^E\right)$, and (ii) realized expenditure growth, $\log\left(\text{Act}_t^E/\text{Act}_{t-1}^E\right)$, on planning-window shift-share news. Columns (2) and (4) rescale the spending response by the component of the initial revenue forecast update predicted by planning-window news. Specifically, $\widehat{\text{Rev Planning}}$ is the fitted value from regressing the initial revenue forecast update on planning-window news with the same fixed effects and controls. Panel B reports the stacked monthly-panel specification for the same two expenditure outcomes. It interacts monthly Bartik news with planning and pre-planning indicators and compares a simple state + fiscal-year fixed-effects specification to the preferred design with cohort × calendar-month × year fixed effects and lagged macro controls. The shift-share measure uses the five-base specification with AR(4)+GDP innovation shocks. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: State-fiscal-year specification

	Revenue forecast: $\log\left(\text{Init}_{s,t}^R/\text{Init}_{s,t-1}^R\right)$	Spending plan response: $\log\left(\text{Init}_{s,t}^E/\text{Init}_{s,t-1}^E\right)$		Spending growth: $\log\left(\text{Act}_t^E/\text{Act}_{t-1}^E\right)$	
	(1)	(2)	(3)	(4)	(5)
News (Planning)	1.33***	0.63***		0.84***	
	(0.13)	(0.16)		(0.15)	
$\widehat{\text{Rev Forecast}}$			0.47***		0.64***
			(0.10)		(0.12)
K-P F-stat			108.89		108.89
State + FYear FE	✓	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓	✓
Observations	990	990	990	990	990
Adj.R^2	0.04	0.05	0.38	0.05	0.08

Panel B: Monthly panel specification

	Spending plan response: $\log\left(\text{Init}_t^E/\text{Init}_{t-1}^E\right)$		Spending growth: $\log\left(\text{Act}_t^E/\text{Act}_{t-1}^E\right)$	
	FE only (1)	Preferred (2)	FE only (3)	Preferred (4)
News ×𝟙(*Planning*)	0.48*	0.61**	0.81*	1.08**
	(0.26)	(0.30)	(0.43)	(0.44)
News ×𝟙(Pre-Plan.)	0.02	-0.20	0.19	0.36
	(0.11)	(0.24)	(0.17)	(0.26)
State FE	✓	✓	✓	✓
Fiscal-Year FE	✓		✓	
Cohort × Cal.-Month × Year FE		✓		✓
Macro Controls		✓		✓
Observations	26,976	26,976	27,504	27,504
Adj. R^2	0.109	0.105	0.107	0.095

Table 5: News from Other Windows Does Not Propagate to Expenditures

Each panel constructs the shift-share measure over a different window relative to the proposal date: pre-planning (6–12 months), revision (proposal to final approval), and realization (during the fiscal year). Column (1) regresses the corresponding revenue object for that window on the shift-share measure; columns (2) and (4) regress planned and realized expenditure outcomes on the same news measure. Columns (3) and (5) relate the expenditure outcomes to the corresponding news-predicted revenue object for that window. All specifications include state and year fixed effects and lagged macroeconomic controls; standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Pre-Initial Window

	Revenue forecast: $\log(\text{Init}^R_t/\text{Init}^R_{t-1})$	Spending plan response: $\log(\text{Init}^E_t/\text{Init}^E_{t-1})$		Spending growth: $\log(\text{Act}^E_t/\text{Act}^E_{t-1})$	
	(1)	(2)	(3)	(4)	(5)
News (Pre-Planning)	-0.18	-0.05		0.45	
	(0.37)	(0.38)		(0.36)	
$\widehat{\text{Rev forecast}}$			0.30		-2.57
			(1.87)		(6.66)
K-P F-stat			0.22		0.22
FEs + Controls	✓	✓	✓	✓	✓
Observations	990	990	990	990	990

Panel B: Revision Window

	Revenue revision: $\log(\text{Final}^R_t/\text{Init}^R_t)$	Spending revision: $\log(\text{Final}^E_t/\text{Init}^E_t)$		Spending growth: $\log(\text{Act}^E_t/\text{Act}^E_{t-1})$	
	(1)	(2)	(3)	(4)	(5)
News (Revision)	0.13	0.41**		0.75***	
	(0.22)	(0.16)		(0.23)	
$\widehat{\text{Rev revision}}$			3.51		6.32
			(6.30)		(11.49)
K-P F-stat			0.29		0.30
FEs + Controls	✓	✓	✓	✓	✓
Observations	1,024	991	991	995	995

Panel C: Realization Window

	Revenue realization: $\log(\text{Act}^R_t/\text{Final}^R_t)$	Expenditure realization: $\log(\text{Act}^E_t/\text{Final}^E_t)$		Spending growth: $\log(\text{Act}^E_t/\text{Act}^E_{t-1})$	
	(1)	(2)	(3)	(4)	(5)
News (Realization)	1.70**	-0.29**		-0.11	
	(0.84)	(0.15)		(0.17)	
$\widehat{\text{Rev realization}}$			-0.17		-0.07
			(0.14)		(0.12)
K-P F-stat		47	4.02		4.11
FEs + Controls	✓	✓	✓	✓	✓
Observations	1,061	1,057	1,057	1,061	1,061

Table 6: Coibion-Gorodnichenko Test: Do Forecast Revisions Predict Forecast Errors?

This table reports the **?** information rigidity test. Under sticky information, noisy information, or rational inattention, forecast revisions should positively predict forecast errors ($\beta > 0$). Columns (1)–(2) regress the revenue forecast error $\log\left(\text{Act}^R/\text{Init}^R\right)$ on the revenue forecast revision $\log\left(\text{Init}_t^R/\text{Init}_{t-1}^R\right)$. All specifications include state and year fixed effects. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue Forecast Error $\log\left(\text{Act}_t^R/\text{Init}_t^R\right)$	
	(1)	(2)
Revenue Revision $\log\left(\text{Init}_t^R/\text{Init}_{t-1}^R\right)$	-0.08	-0.08
	(0.05)	(0.05)
State + FYear FE	✓	✓
Macro Controls		✓
Observations	989	989
Adj.R^2	0.54	0.54

Table 7: Revenue Forecast Efficiency and Decision-Window Bias

The dependent variable in all columns is the revenue forecast error, $\log\left(\text{Act}^R_{s,t}/\text{Init}^R_{s,t}\right)$. Regressors are shift-share news measures constructed over windows relative to each state's proposal date: Planning denotes 0–6 months and Pre-planning denotes 6–12 months. Column (3) includes both planning and pre-planning news; column (4) additionally includes 12–18 and 18–24 month pre-planning windows. All specifications include state and year fixed effects and lagged macroeconomic controls measured prior to the planning window. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue Forecast Error $\log\left(\text{Act}^R_{s,t}/\text{Init}^R_{s,t}\right)$			
	(1)	(2)	(3)	(4)
News Planning (0–6m ahead)	0.06		0.10	0.12
	(0.16)		(0.17)	(0.19)
News Pre-Plan. (6–12m ahead)		1.08***	1.08***	1.25***
		(0.34)	(0.34)	(0.31)
News Pre-Plan. (12–18m ahead)				0.21
				(0.21)
News Pre-Plan. (18–24m ahead)				1.08
				(0.82)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	1,024	1,024	1,024	1,024
Adj.R^2	0.54	0.54	0.54	0.54

Table 8: Rational Inattention: Do More Informative Signals or Higher Fiscal Stakes Attenuate the Planning Premium?

Panel A reports the informativeness tests and Panel B the fiscal-stakes tests. In Panel A, the rows "News Plan × High Info Signal" and "News Pre-Plan × High Info Signal" interact the two news windows with the column-specific informativeness proxy. In Panel B, the rows "News Plan × High Fiscal Stakes" and "News Pre-Plan × High Fiscal Stakes" interact the two news windows with the column-specific stakes proxy listed in the proxy row. In Panel A, the indicators are revenue concentration HHI and a 20-year rolling signal-to-noise ratio. For each state-year and each news window, the rolling SNR is computed from the prior 20 fiscal years ending at $t-2$ as $\widehat{\pi}^2 \mathrm{var}(\mathit{News})/\widehat{\mathrm{var}}(\hat{u})$, where $\widehat{\pi}$ and $\hat{u}$ come from a within-state regression of three-year Census revenue growth, $\log(\mathrm{Rev}_\tau/\mathrm{Rev}_{\tau-3})$, on the corresponding Bartik news measure and a linear time trend. In Panel B, the indicators are a strict balanced-budget-rule indicator and fiscal stress, defined as the negative of the rainy day fund balance as a share of expenditure, lagged two fiscal years. All continuous indicators are standardized to mean zero and unit variance. All specifications include state and fiscal-year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Signal Informativeness

	Revenue Forecast: $\log(\mathrm{Init}^R/\mathrm{Init}^R_{t-1})$		Revenue Forecast Error: $\log(\mathrm{Act}^R/\mathrm{Init}^R)$	
	(1)	(2)	(3)	(4)
News Planning	1.24***	1.06***	0.56	0.23
	(0.35)	(0.23)	(0.36)	(0.23)
News Pre-Plan.	-0.06	0.23	1.23*	1.11***
	(0.50)	(0.34)	(0.65)	(0.28)
News Planning × High Info Signal	0.04	-0.43	-0.30**	0.16
	(0.15)	(0.32)	(0.14)	(0.31)
News Pre-Plan. × High Info Signal	-0.06	-0.01	0.07	-0.63
	(0.19)	(0.41)	(0.31)	(0.44)
Informativeness proxy	HHI	20y rolling SNR (3y)	HHI	20y rolling SNR (3y)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	990	955	1,024	989
Adj. R^2	0.05	0.05	0.55	0.54

Panel B: Fiscal Stakes

	Revenue Forecast: $\log(\mathrm{Init}^R/\mathrm{Init}^R_{t-1})$		Revenue Forecast Error: $\log(\mathrm{Act}^R/\mathrm{Init}^R)$	
	(1)	(2)	(3)	(4)
News Planning	1.48*	0.90***	0.54	-0.42
	(0.79)	(0.33)	(0.70)	(0.89)
News Pre-Plan.	-0.50	-0.06	1.26	1.32**
	(1.28)	(0.65)	(1.03)	(0.51)
News Planning × High Fiscal Stakes	-0.76	-0.05	0.30	-0.10
	(0.80)	(0.05)	(0.69)	(0.13)
News Pre-Plan. × High Fiscal Stakes	0.45	-0.02	-0.01	0.03
	(1.09)	(0.08)	(1.15)	(0.07)
Stakes proxy	Strict rule	Fiscal stress	Strict rule	Fiscal stress
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	878	931	910	964
Adj. R^2	0.04	0.06	0.29	0.54

Table 9: Real Effects of Decision-Window News in Government-Exposed Industries

The unit of observation is state–industry–fiscal-year (QCEW, private sector). Dependent variables are year-over-year log changes in employment (columns 1–3) and the wage bill (columns 4–6). Shift-share news measures are constructed over budget windows as in Section **??**. Government exposure is a binary indicator equal to one for industries whose government-purchase share (from BEA input–output accounts) exceeds 5%, using a five-year trailing average lagged two years. All regressions are weighted by lagged employment. All specifications include state×industry, state×fiscal-year and industry×fiscal-year fixed effects. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	$\Delta \log(\text{Private Employment})^{FY}$			$\Delta \log(\text{Private Wages})^{FY}$		
	(1)	(2)	(3)	(4)	(5)	(6)
News Pre-Plan. × Govt Exp.		0.01	-0.01		0.03	0.05
		(0.11)	(0.13)		(0.11)	(0.11)
News Planning × Govt Exp.	0.54***		0.55***	0.42***		0.45**
	(0.15)		(0.15)	(0.16)		(0.17)
News Revision × Govt Exp.			0.07			0.26**
			(0.18)			(0.10)
State×Industry FE	✓	✓	✓	✓	✓	✓
Industry×FYear FE	✓	✓	✓	✓	✓	✓
State×FYear FE	✓	✓	✓	✓	✓	✓
Observations	22,326	22,326	22,326	22,326	22,326	22,326
Adj.R^2	0.79	0.79	0.79	0.75	0.75	0.75

Table 10: Structural Parameter Estimates

Minimum-distance estimates using 15 targeted moments and diagonal inverse-SE weighting. Panel: 47 states × 30 years. 15 multistart runs; all converge to objectives in $[4.60, 4.97]$.

Parameter	Symbol	Estimate	Interpretation
Forecasting			
Pre-planning incorp.	K_{pre}	≈ 0	Pre-planning info ignored
Planning incorp.	K_{plan}	1.18	Planning news incorporated
Revision incorp.	K_{rev}	0.53	Partial revision-stage updating
Proposal			
Balanced-budget pressure	ϕ	2.00	Pass-through: $\phi/(1+\phi) = 0.67$
Enactment			
Base revenue response	α	2.07	
Revision-news response	α_F	10.0	
Legislative anchoring	κ	12.8	Pass-through: $\alpha_F/(1+\kappa) = 0.72$
Execution			
Windfall pass-through	ξ_{exec}	0.29	Asymmetric within-year adjustment
Shortfall coefficient	c_{neg}	0.22	
Under-execution	$\bar{u}_{\text{exec}}$	0.071	7% systematic underspend
Dynamics and private sector			
Persistence	ρ	≈ 0	
Employment loading	κ_E	0.43	
Wage loading	κ_W	0.46	
Minimized objective Q		4.60	

Table 11: Counterfactual Decomposition: Cumulative Employment Response

Cumulative government-exposed employment response over 3 years to a one-standard-deviation revenue shock. Panel A reports the response under all three scenarios. Panel B decomposes the total attenuation ($A - C$) into institutional frictions ($A - B$, the cost of frictions when information is fully incorporated) and information bias ($B - C$, the additional cost of ignoring pre-planning news).

Panel A: Employment Response by Scenario

	(1) Planning (−1 SD)	(2) Pre-plan. (+1 SD)	(3) Pre-plan. (−1 SD)
(A) No friction	−0.429	+0.429	−0.429
(B) Institutional frictions only	−0.554	+0.372	−0.554
(C) Both frictions	**−0.615**	**+0.031**	**−0.185**
Retained effect (C as % of A)	143%	7%	43%

Panel B: Decomposing Attenuation of Pre-Planning Shocks

	Pre-plan. (+1 SD)		Pre-plan. (−1 SD)	
Total attenuation ($A - C$)	0.398		0.244	
Institutional frictions ($A - B$)	0.057	(14%)	−0.125	(−51%)
Information bias ($B - C$)	0.341	(86%)	0.369	(151%)

Appendix

A Data and Institutional Background

This appendix provides additional detail on the data sources, institutional features, and descriptive patterns that underlie the analysis in the main text.

A.1 State Budget Calendar

Table **??** reports the budget calendar for each state in our sample. Columns include the fiscal year start date, the typical month in which the governor submits the initial budget proposal, and the statutory deadline for legislative approval. These dates define the planning, revision, and realization windows used throughout the paper. The key source of cross-state identification is the variation in proposal dates: early-proposing states (e.g., those with October–January proposal months) and late-proposing states (e.g., February–March) face the same national economic conditions but assign them to different budget windows.

Table A.1: State Budget Timeline and Data Sources

"FY Start" denotes the beginning of the fiscal year. "Initial Proposal" refers to the typical period in which governors submit budget proposals to legislatures, while "Final Approval" denotes the deadline by which the budget must be enacted.

State	FY Start	Initial Proposal	Final Approval	State	FY Start	Initial Proposal	Final Approval
AL	Oct 1	Feb	May	MS	Jul 1	Nov	Mar–Apr
AK	Jul 1	Dec	Jun	MO	Jul 1	Jan	May
AZ	Jul 1	Jan–Feb	Jun	MT	Jul 1	Nov	May
AR	Jul 1	Nov–Jan	Apr–May	NE	Jul 1	Jan	Jun
CA	Jul 1	Jan	Jun	NV	Jul 1	Jan	Jun
CO	Jul 1	Nov	Apr–Jun	NH	Jul 1	Feb	Jun
CT	Jul 1	Feb	Jun	NJ	Jul 1	Feb	Jun
DC	Oct 1	Mar	Jun	NM	Jul 1	Jan	Mar–Apr
DE	Jul 1	Jan	Jun	NY	Apr 1	Jan	Mar
FL	Jul 1	Jan	Mar–May	NC	Jul 1	Mar	Jun
GA	Jul 1	Jan	May	OH	Jul 1	Feb	Jun
HI	Jul 1	Dec	May	OK	Jul 1	Feb	May–Jun
ID	Jul 1	Jan	Mar–Apr	PA	Jul 1	Feb	Jun
IL	Jul 1	Feb	Jun	RI	Jul 1	Jan	Jun
IN	Jul 1	Jan	Apr–May	SC	Jul 1	Jan	May
IA	Jul 1	Feb	May	SD	Jul 1	Dec	Mar–Apr
KS	Jul 1	Jan	May	TN	Jul 1	Feb	May–Jun
KY	Jul 1	Jan	May	TX	Sep 1	Jan	May
LA	Jul 1	Feb	Jun	UT	Jul 1	Dec	Mar–Apr
ME	Jul 1	Jan	Jun	VT	Jul 1	Jan	May–Jun
MD	Jul 1	Jan	Apr–May	VA	Jul 1	Dec	Apr–May
MA	Jul 1	Jan	Jun	WV	Jul 1	Jan	Apr
MI	Oct 1	Feb	Jul	WI	Jul 1	Jan	Jun
MN	Jul 1	Jan	May				

A.2 Primary Data Source: Annual Comprehensive Financial Reports

Our primary data source is the Annual Comprehensive Financial Report (ACFR), which each state publishes after the close of the fiscal year. ACFRs contain audited financials and, crucially, comparisons of budgeted and actual figures for both revenues and expenditures. Figure **??** reproduces a representative page from the Massachusetts ACFR for fiscal year 2002. The top panel shows the budgetary comparison statement, which reports the original budget, the final revised budget, and actual amounts for each major revenue and expenditure category. The bottom panel illustrates the reconciliation between budgetary-basis and GAAP-basis reporting. We hand-collected and digitized analogous tables from 47 states over fiscal years 1995–2024, extracting the initial budget, final budget, and actual amounts used

to construct the decomposition in Equation (??).

Budgeted Funds

Combined Schedule Of Revenues, Expenditures And Changes In Fund Balances - Statutory Basis - Budget and Actual

Fiscal Year Ended June 30, 2002

(Amounts in thousands)

	Original Budget	Final Budget	Actual	Variance over/(under)
REVENUES AND OTHER FINANCING SOURCES				
Revenues:				
Taxes	$ 14,930,000	$ 14,930,100	$ 13,622,710	$ (1,307,390)
Assessments	-	2,700	129,354	126,654
Federal grants and reimbursements	4,126,300	4,126,300	4,334,934	208,634
Departmental	1,672,300	1,422,389	1,205,629	(216,760)
Miscellaneous	-	92,483	265,492	173,009
Total revenues	**20,728,600**	**20,573,972**	**19,558,119**	**(1,015,853)**
Other financing sources:				
Fringe benefit cost recovery	-	-	161,524	161,524
Lottery reimbursements	-	87,341	87,341	-
Lottery distributions	-	857,700	871,136	13,436
Caseload increase mitigation transfer	-	-	4,878	4,878
Tobacco settlement transfer	-	154,728	154,728	-
Operating transfers in	1,062,000	221,911	515,531	293,620
Transfer for transitional escrow	-	-	579,215	579,215
Stabilization transfer	-	-	1,030,000	1,030,000
Capital projects funds year end transfer	-	-	51,369	51,369
Federal reimbursement transfer in	-	-	1,902	1,902
Transfer for tax reduction	-	-	33,605	33,605
Total other financing sources	**1,062,000**	**1,321,680**	**3,491,229**	**2,169,549**
Total revenues and other financing sources	**21,790,600**	**21,895,652**	**23,049,348**	**1,153,696**
EXPENDITURES AND OTHER FINANCING USES				
Expenditures:				
Legislature	88,227	72,850	58,909	13,941
Judiciary	584,123	587,686	580,057	7,629
Inspector General	2,019	1,975	1,973	2
Governor and Lieutenant Governor	5,611	5,655	5,542	113
Secretary of the Commonwealth	35,088	35,028	34,693	335
Treasurer and Receiver-General	187,159	121,024	119,114	1,910
Auditor of the Commonwealth	15,380	15,351	15,347	4
Attorney General	35,687	36,335	35,674	661
Ethics Commission	1,415	1,412	1,411	1
District Attorney	79,566	81,253	81,056	197
Office of Campaign & Political Finance	1,089	1,088	1,440	(352)
Sheriff's Departments	193,220	215,292	211,973	3,319
Disabled Persons Protection Commission	1,616	1,714	1,708	6
Board of Library Commissioners	7,251	7,270	6,984	286
Comptroller	8,352	8,415	8,410	5
Administration and finance	1,429,507	1,386,683	1,344,343	42,340
Environmental affairs	226,182	234,966	221,717	13,249
Communities and development	122,269	123,018	116,790	6,228
Health and human services	4,241,579	4,403,293	4,312,159	91,134
Transportation and construction	129,104	151,990	145,917	6,073
Education	464,899	462,232	412,846	49,386
Higher education	1,014,147	1,039,047	1,029,516	9,531
Public safety	850,519	945,382	917,982	27,400
Economic development	46,362	46,241	34,008	12,233
Elder affairs	276,755	291,729	273,759	17,970
Consumer affairs	46,909	48,552	46,872	1,680
Labor	83,542	89,750	59,319	30,431
Direct local aid	5,191,844	5,193,357	5,189,116	4,241
Medicaid	4,985,458	5,282,868	5,259,275	23,593
Pension	931,029	797,064	795,782	1,282
Debt service:				
Principal retirement	659,802	710,674	692,001	18,673
Interest and fiscal charges	691,889	641,017	612,734	28,283
Total expenditures	**22,637,599**	**23,040,211**	**22,628,427**	**411,784**

Figure A.1: Example: ACFR for Massachusetts. This figure illustrates the structure of data from Massachusetts's ACFR for fiscal year 2002. The top panel shows the budgetary comparison statement reporting original budget, final budget, and actual amounts. The bottom panel shows the reconciliation to GAAP-basis financials.

A.3 Anchoring of the Initial Budget

A key assumption underlying our decomposition is that the prior year's final budget, $\text{Final}_{s,t-1}$, is the relevant reference point when the initial budget is drafted. This is natural because the initial budget for fiscal year t is typically proposed before fiscal year $t-1$ actuals are known. Table **??** provides direct evidence. We regress the log initial budget on $\log(\text{Final}_{t-1})$ and $\log(\text{Act}_{t-2})$, the two most recent available benchmarks. The initial budget is far more strongly anchored to Final_{t-1} than to Act_{t-2}, and the latter adds little explanatory power conditional on the former. This supports the use of Final_{t-1} as the baseline in the planning component and confirms that the initial budget is set relative to the most recent enacted plan rather than past realizations.

Table A.2: Anchoring of the Initial Budget

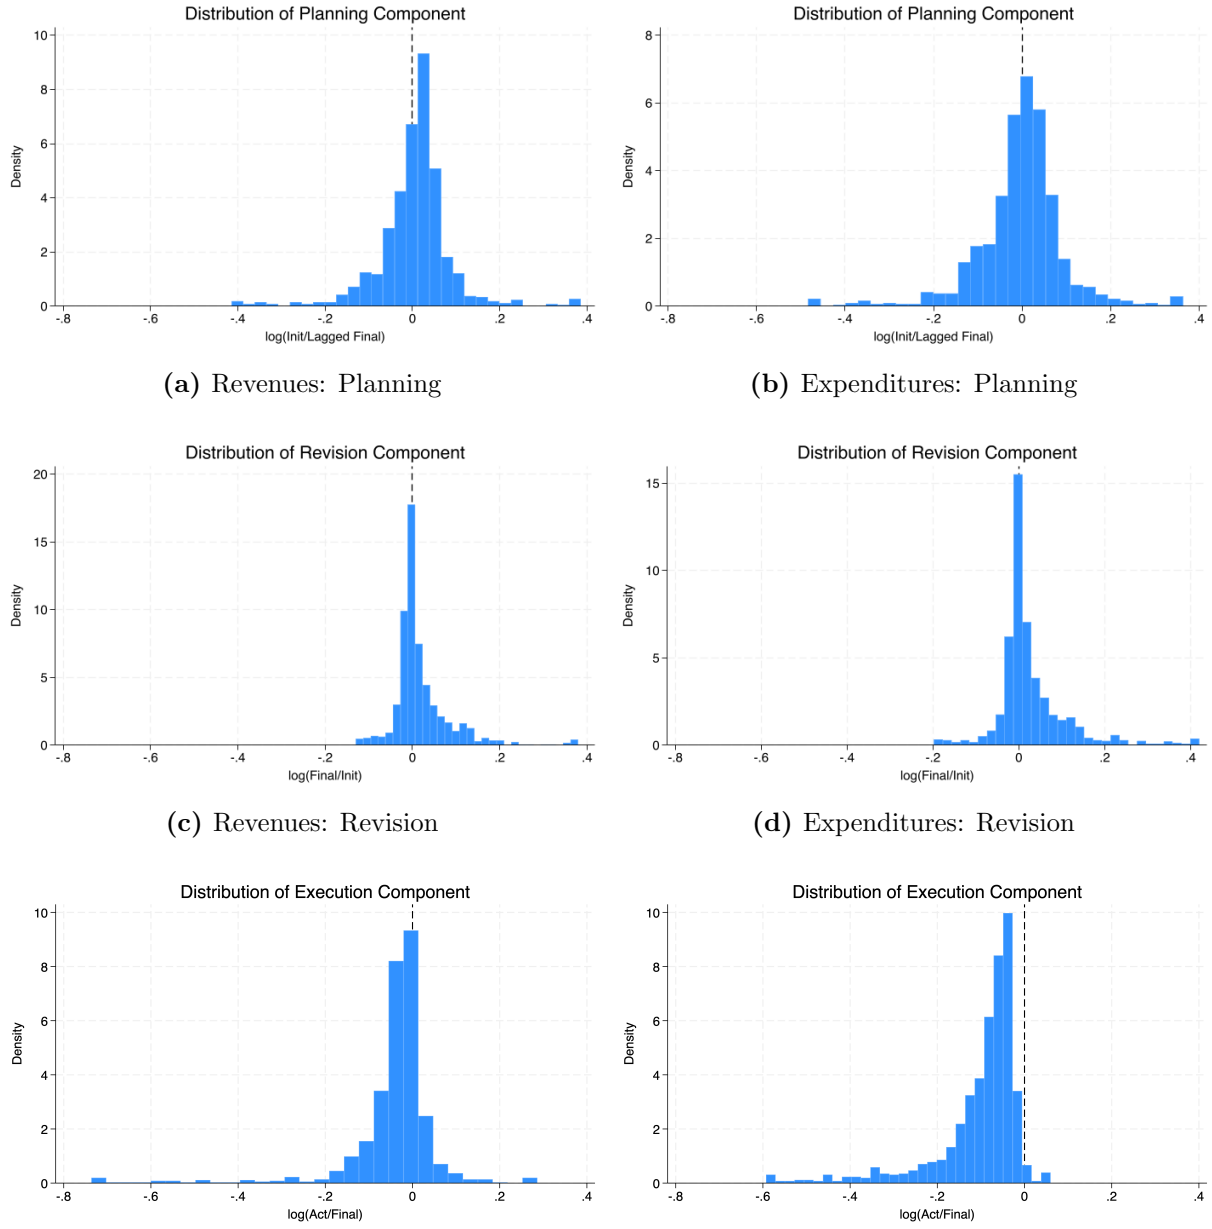
The dependent variable is $\log(\text{Init}_t^R)$ in columns (1)–(3) and $\log(\text{Init}_t^E)$ in columns (4)–(6). Regressors include $\log(\text{Final}_{t-1})$ and $\log(\text{Act}_{t-2})$. All specifications include state and year fixed effects. Standard errors are reported in parentheses and are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue: $\log(\text{Init}_t^R)$			Spending: $\log(\text{Init}_t^E)$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(\text{Final}_{t-1})$	0.80***		0.78***	0.74***		0.59***
	(0.03)		(0.06)	(0.04)		(0.06)
$\log(\text{Act}_{t-2})$		0.60***	0.02		0.67***	0.20***
		(0.09)	(0.07)		(0.05)	(0.07)
State + Year FE	✓	✓	✓	✓	✓	✓
Observations	989	958	954	1,122	1,088	1,084
Adj.R^2	0.992	0.988	0.992	0.986	0.984	0.986

A.4 Distribution of Budget-Cycle Adjustments

Figure **??** displays the empirical distributions of revenue and expenditure changes across the three phases of the budget cycle defined in Equation (**??**). Several patterns are noteworthy. First, the planning-phase distributions are centered near zero with moderate dispersion, indicating that initial budgets typically project modest growth relative to the prior year's final budget. Second, revision-phase adjustments are tightly concentrated near zero, consistent with the main-text finding that the legislative process makes limited changes to the governor's proposal. Third, the realization-phase distributions are left-skewed: budgeted revenues exceed actual revenues in approximately 75% of state-year observations, and budgeted expenditures exceed actual outlays in 97%. This systematic overestimation reflects the

conservatism inherent in balanced-budget requirements and executive spending controls.



(a) Revenues: Planning

(b) Expenditures: Planning

(c) Revenues: Revision

(d) Expenditures: Revision

(e) Revenues: Execution

(f) Expenditures: Execution

Figure A.2: Distribution of State Revenue and Expenditure Changes over the Budget Cycle. The figure shows the distributions of the planning, revision, and execution (realization) components in Equation (**??**) for both revenues and expenditures.

A.5 Revenue Constraints and Expenditure Responses

Table **??** quantifies the relationship between revenue and expenditure changes across the three budget phases. Revenue changes explain a substantial share of the variation in expenditure adjustments, with particularly strong co-movement during planning and revision. A one-percentage-point reduction in expected revenues reduces expected spending by 0.6 percentage points in the planning window and 0.7 percentage points in the revision window; during the realization window, the pass-through is smaller at 0.3 percentage points. These estimates are consistent with the balanced-budget constraint binding most tightly during budget formation, when policymakers must align planned expenditures with projected revenues.

Table A.3: Revenue Constraints and Expenditure Responses Across Budget Phases

This table relates phase-by-phase expenditure adjustments to phase-by-phase revenue changes using the decomposition in Equation (**??**). All specifications include the fixed effects and controls reported in the table; standard errors are reported in parentheses. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Total Spending					
	Planning $\log(\text{Init}^E_{s,t}/\text{Final}^E_{s,t-1})$		Revision $\log(\text{Final}^E_{s,t}/\text{Init}^E_{s,t})$		Realization $\log(\text{Act}^E_{s,t}/\text{Final}^E_{s,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)
Rev (Planning)	0.68***	0.60***				
	(0.08)	(0.07)				
Rev (Revision)			0.83***	0.73***		
			(0.07)	(0.07)		
Rev (Realization)					0.57***	0.27***
					(0.08)	(0.08)
State FE		✓		✓		✓
Obs.	989	989	1,024	1,024	1,040	1,040
Adj.R^2	0.36	0.51	0.40	0.59	0.42	0.63

A.6 Asymmetric Spending Adjustment

Table **??** documents asymmetry in how state expenditures respond to revenue surprises during the realization phase. States adjust spending sharply downward when actual revenues fall short of the final budget—consistent with balanced-budget enforcement and executive impoundment authority—but do not increase spending proportionally when revenues exceed budgeted amounts. This asymmetry reinforces the rigidity of spending commitments docu-

mented in the main text: once the spending plan is set, upside revenue surprises flow into reserves (see Table **??** below) rather than triggering supplemental appropriations.

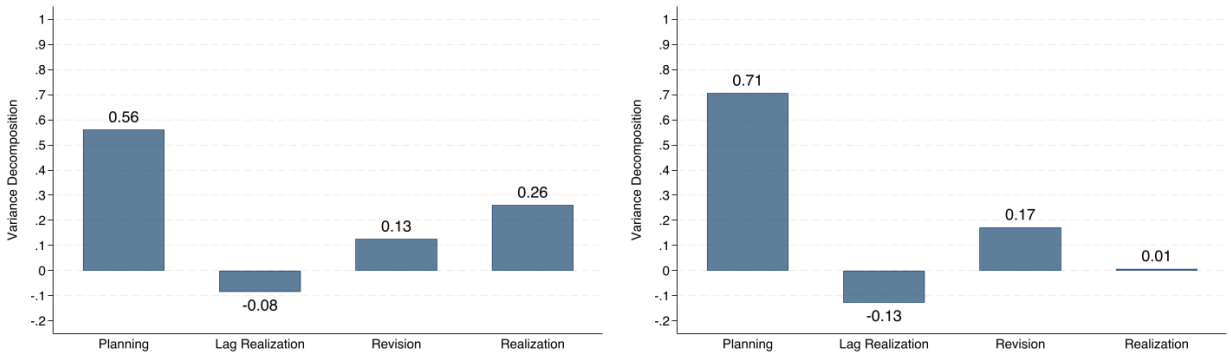
Table A.4: Asymmetric Spending Adjustment to Revenue Shortfalls and Windfalls

This table documents asymmetry in the realization phase. The dependent variable is the expenditure realization term, $\log\left(\text{Act}^E_{s,t}/\text{Final}^E_{s,t}\right)$. The key regressor interacts the realization-phase revenue change with an indicator for negative realization (a shortfall relative to the final budget). All specifications include the fixed effects and controls reported in the table. Standard errors are reported in parentheses and are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Realization $\log(\text{Act}^E_{s,t}/\text{Final}^E_{s,t})$	
	(1)	(2)
Rev (Realization) × Negative Realization	0.97***	0.72***
	(0.10)	(0.17)
Rev (Realization)	-0.28***	-0.26**
	(0.08)	(0.11)
Negative Realization	-0.00	0.00
	(0.01)	(0.01)
State FE		✓
Obs.	1,040	1,040
Adj.R^2	0.51	0.66

A.7 Variance Decomposition: Alternative Benchmark

The main-text variance decomposition in Figure **??** benchmarks the planning component to the prior year's final budget, $\log(\text{Init}_{s,t}/\text{Final}_{s,t-1})$. An alternative is to benchmark to two-year-lagged actual outcomes, $\log(\text{Init}_{s,t}/\text{Act}_{s,t-2})$, which avoids any mechanical dependence on the prior year's final budget. Figure **??** replicates the variance decomposition under this alternative benchmark. The planning component continues to account for the majority of the variation in realized spending, confirming that the concentration of fiscal adjustment in the initial budget is not an artifact of the benchmark choice.



(a) Total Revenues **(b)** Total Expenditures

Figure A.3: Decomposing the Sources of Variation in Actual Growth (Alternative Benchmark). This figure replicates Figure **??** using two-year-lagged actual outcomes, $\log(\text{Init}_{s,t}/\text{Act}_{s,t-2})$, as the planning-stage benchmark.

A.8 Rainy Day Fund Response to Budget Surprises

What happens to unanticipated revenue windfalls? Table **??** examines how rainy day fund (RDF) balances respond to budget surprises. States significantly increase reserves when actual surpluses exceed budgeted surpluses, but do not draw down reserves proportionally during deficits. This asymmetric pattern is consistent with the rigid spending commitment documented in the main text and with the asymmetric realization response in Table **??**: revenue windfalls accumulate in fiscal buffers rather than triggering additional spending within the fiscal year.

Table A.5: Rainy Day Fund Response to Budget Surplus and Forecast Errors

The dependent variable is $\Delta \log(\text{RDF}_{s,t})$. Columns 1–3 use the actual surplus normalized by lagged expenditure as the regressor; columns 4–6 use the surplus forecast error. Columns 2 and 5 allow separate slopes for positive and negative values; columns 3 and 6 add state GDP growth as a control. All regressors are winsorized at the 1st and 99th percentiles. Sample: 47 states, 1995–2024. Standard errors clustered by state in parentheses. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Dependent Variable: $\Delta \log(\text{RDF})$					
	Actual Surplus			Surplus Forecast Error		
	(1)	(2)	(3)	(4)	(5)	(6)
Surplus / Lag Exp	0.424					
	(0.284)					
Surplus × 1(In Surplus)		0.617**	0.571**			
		(0.294)	(0.283)			
Surplus × 1(In Deficit)		0.079	-0.052			
		(0.589)	(0.488)			
Surplus FE / Lag Exp				0.229**		
				(0.101)		
Surplus FE × 1(Good News)					0.255**	0.209**
					(0.111)	(0.097)
Surplus FE × 1(Bad News)					-0.965	-1.474
					(2.072)	(1.992)
State GDP Growth			3.222**			3.283**
			(1.452)			(1.481)
State FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
GDP Control			✓			✓
Observations	843	843	843	843	843	843
Adj. R^2	0.073	0.073	0.080	0.072	0.071	0.079

B Shift-Share Measure: Construction Details

This appendix provides complete details on the construction of the shift-share news measure introduced in Section **??**. We describe the revenue source pairings used to match state revenue bases to national indicators, the definition of budget windows, the innovation extraction procedure, and the construction of exposure weights. We also report the empirical distribution of the resulting news measures and show that planning-window and pre-planning shocks have comparable support.

B.1 Details of Measurement Construction

Revenue source pairings. We consider six revenue sources, each paired with a national macroeconomic series tracking the underlying tax base (Table **??**). The logic of each pairing is straightforward. Individual income tax revenue depends on the trajectory of personal income; sales tax revenue tracks consumer spending; severance taxes are driven by commodity prices (oil for the majority of severance-tax-dependent states); corporate income taxes follow corporate profitability; charges and fees correlate with services consumption; and federal transfers reflect the trajectory of federal grants-in-aid to state and local governments.

Our baseline specification in the main text uses five sources: charges & misc. revenue, individual income taxes, general sales taxes, severance taxes, and corporate income taxes. These constitute the major own-source revenue categories with clear, well-measured national indicator pairings. Appendix Table **??** evaluates alternative combinations.

Table B.1: Revenue Source–National Indicator Pairings

Sources 1–5 constitute the baseline specification. Source 6 is included in the six-base robustness specification only. Personal income and consumption are deflated by the PCE price index. Oil prices and corporate profits are in nominal terms; the AR(4)+GDP residualization absorbs predictable inflation components.

Revenue Source	National Indicator	FRED Series
1. Individual income tax	Real Personal Income	RPI
2. General sales tax	Real Personal Consumption Exp.	PCE/PCEPI
3. Charges & misc. revenue	Real PCE Services	PCESV/PCEPI
4. Corporate income tax	Corporate Profits After Tax	CP
5. Severance tax	WTI Crude Oil Price	DCOILWTICO
6. Federal intergov. transfers	Federal Grants-in-Aid	FGSL

Budget windows. States operate on heterogeneous budget calendars. We define four windows based on each state's budget proposal date and fiscal year start: (i) planning: 0–6 months before; (ii) revision: proposal to final legislative approval; (iii) realization: during the fiscal year. For each window, the raw shock to national indicator k is the cumulative log-change:

$$g_{k,t,w}^{(s),\text{raw}} = \log X_k(\tau_{sw}^1) - \log X_k(\tau_{sw}^0), \tag{B.1}$$

where $[\tau_{sw}^0, \tau_{sw}^1]$ are the start and end dates of window w for state s in fiscal year t. Because states have different proposal dates, the same calendar-month shock enters different windows for different states. Table **??** reports the budget proposal dates and fiscal year start dates used for each state.

Innovation extraction. The raw cumulative shocks $g_{k,t,w}^{(s),\text{raw}}$ contain both a predictable component (based on recent macroeconomic trends) and an unpredictable innovation. We extract the innovation by residualizing against an AR(4) model augmented with lagged real GDP growth, estimated separately for each budget window.

For each national indicator k and each state-year-window observation (s, t, w), we estimate:

$$g_{k,t,w}^{(s),\text{raw}} = \alpha + \sum_{\ell=1}^{4} \rho_\ell \, g_{k,w}^{\text{lag}_\ell} + \sum_{\ell=1}^{4} \phi_\ell \, \Delta \log \text{GDP}_{w}^{\text{lag}_\ell} + u_{k,t,w}^{(s)}, \tag{B.2}$$

where $g_{k,w}^{\text{lag}_\ell}$ is the ℓ-th quarterly lag of the cumulative shock to indicator k and $\Delta \log \text{GDP}_{w}^{\text{lag}_\ell}$ is the ℓ-th quarterly lag of real GDP growth, both measured backward from the start of window w. The lag structure is anchored to each window's start date, so the conditioning set differs across windows even within the same state-year. The innovation is the residual: $g_{k,t,w}^{(s)} = u_{k,t,w}^{(s)}$. We prefer cumulative residualization over summing monthly innovations because budget-makers observe the net change in conditions over the decision window; the relevant surprise is the gap between this cumulative change and what was forecastable from pre-window data. Figure **??** plots the raw and residualized series for each indicator, showing that the procedure removes predictable trends while preserving sharp innovations around recessions.

Exposure weights. The exposure weight $w_{s,k,t}$ measures state s's reliance on revenue source k:

$$w_{s,k,t} = \frac{1}{5} \sum_{\tau=t-7}^{t-2} \frac{\text{Revenue}_{s,k,\tau}}{\text{General Revenue}_{s,\tau}}. \tag{B.3}$$

Revenue by source is from the Census Bureau's Annual Survey of State Government Finances. The denominator is total general revenue— including sources outside the baseline— so the weights do not sum to one. The sum $S_s = \sum_k w_{s,k,t}$ varies across states, capturing total exposure to included revenue bases; state fixed effects absorb time-invariant differences in S_s. Figure **??** displays the full revenue composition by state, showing the relative magnitude of each source and the share of general revenue not captured by the baseline categories.

Shock magnitudes across windows. Table **??** (Panel B) summarizes the empirical distribution of the shift-share news measures across budget windows. Figure **??** plots the distributions of planning- and pre-planning-window shocks directly. The two distributions overlap substantially, indicating that the main timing results are not driven by planning-window

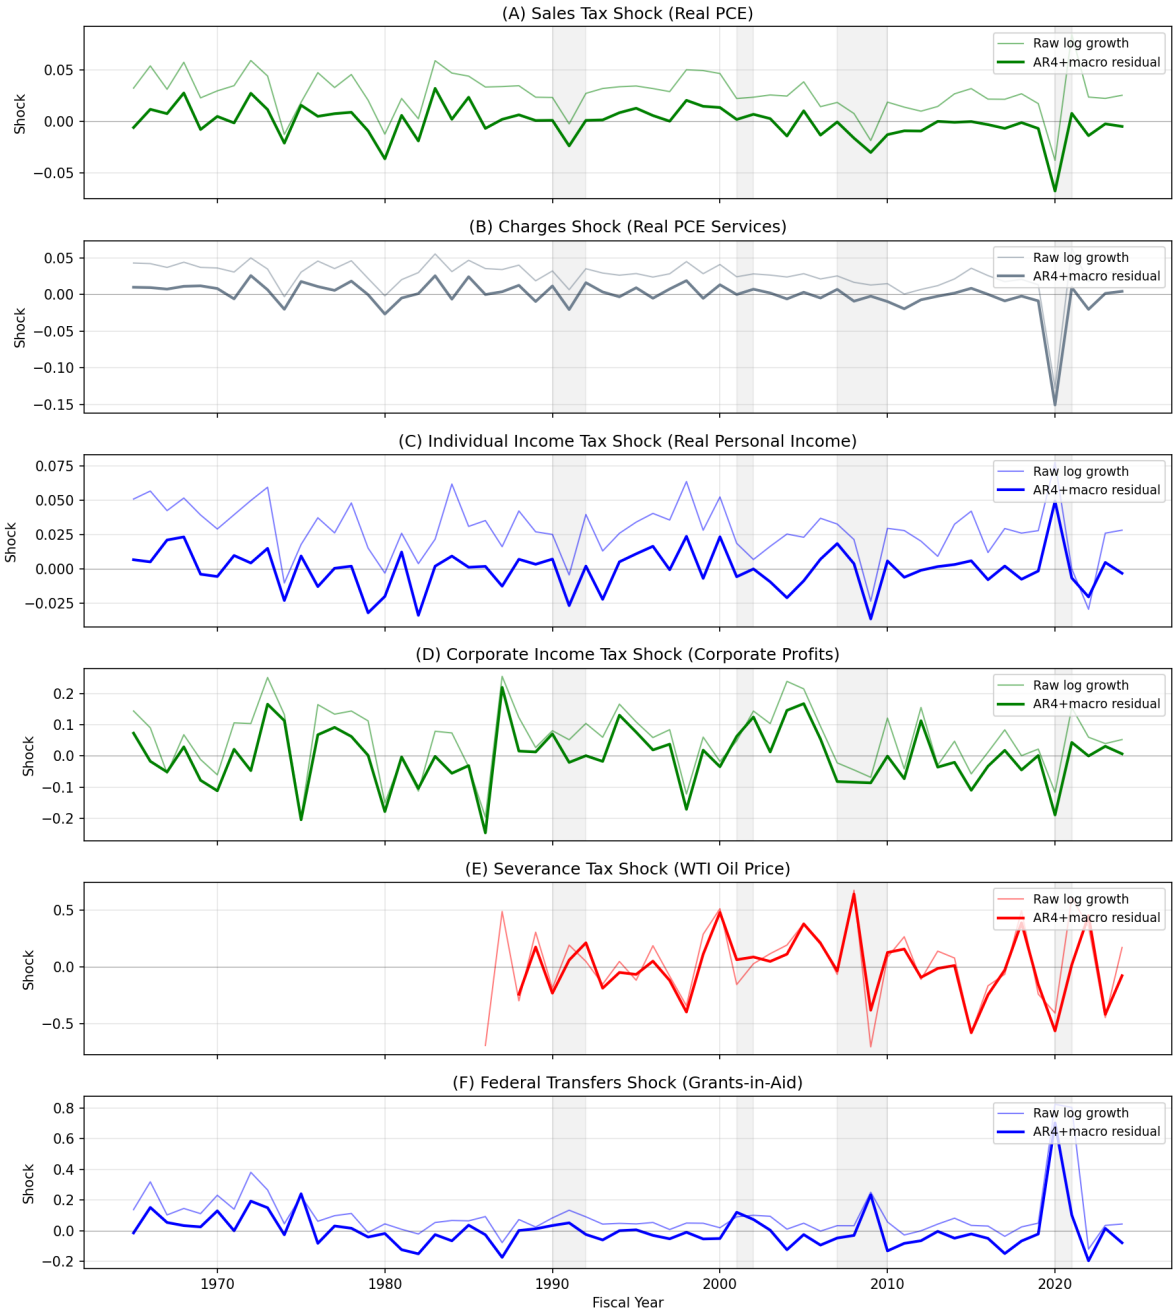


Figure B.1: Raw Shocks and AR(4)+GDP Innovations. Each panel plots the raw cumulative log-change (thin line) and the AR(4)+GDP innovation residual (thick line) for one national indicator over the fiscal year. Panel (A): real personal consumption expenditures (sales tax); Panel (B): real PCE services (charges); Panel (C): real personal income (income tax); Panel (D): corporate profits (corporate income tax); Panel (E): WTI crude oil (severance tax); Panel (F): federal grants-in-aid. Panels (A)–(E) correspond to the five baseline sources. The innovation is the residual from projecting each cumulative shock on four own lags and four lags of real GDP growth. Shaded areas indicate NBER recession dates.

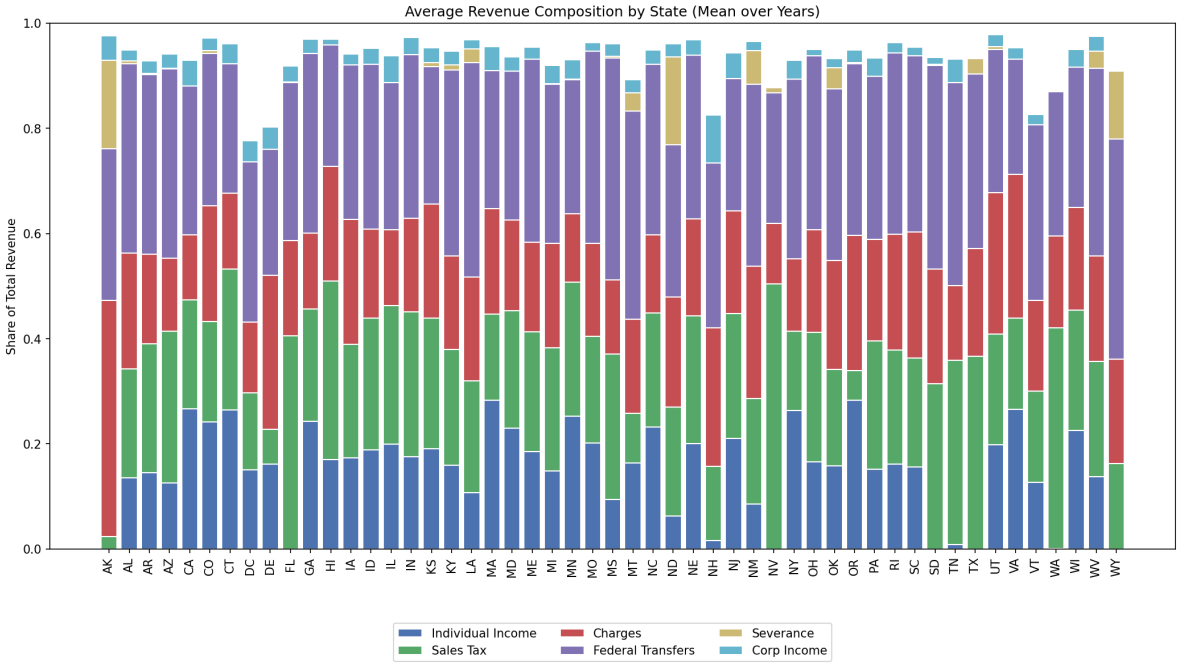


Figure B.2: Full Revenue Composition by State. Stacked shares of six revenue sources in total general revenue, averaged over fiscal years 1995–2024. The five baseline sources (individual income, sales, charges, corporate income, and severance taxes) account for approximately 62% of general revenue. Federal transfers account for most of the remainder. The gap between bar tops and 1.0 reflects other revenue sources (license fees, transfer from local governments etc.) included in the denominator of the exposure weights but not in the shift-share measure. Data source: Census Bureau Annual Surveys of State and Local Government Finance.

shocks simply having a wider support or mechanically larger dispersion than pre-planning shocks.

B.2 Within-Planning Gradient

The main text documents a sharp contrast between the planning and pre-planning windows at the six-month level. Figure **??** further shows that the response rises continuously as the proposal date approaches. Table **??** formalizes this gradient by splitting the planning window into the near quarter (0–3 months before the proposal) and the far quarter (3–6 months before the proposal). The near-quarter coefficient is larger than the far-quarter coefficient for both revenue forecasts and expenditure plans. For expenditure plans in particular, the far-quarter coefficient is small and statistically insignificant, while the near-quarter coefficient is large and precisely estimated. This pattern suggests that information incorporation intensifies as the

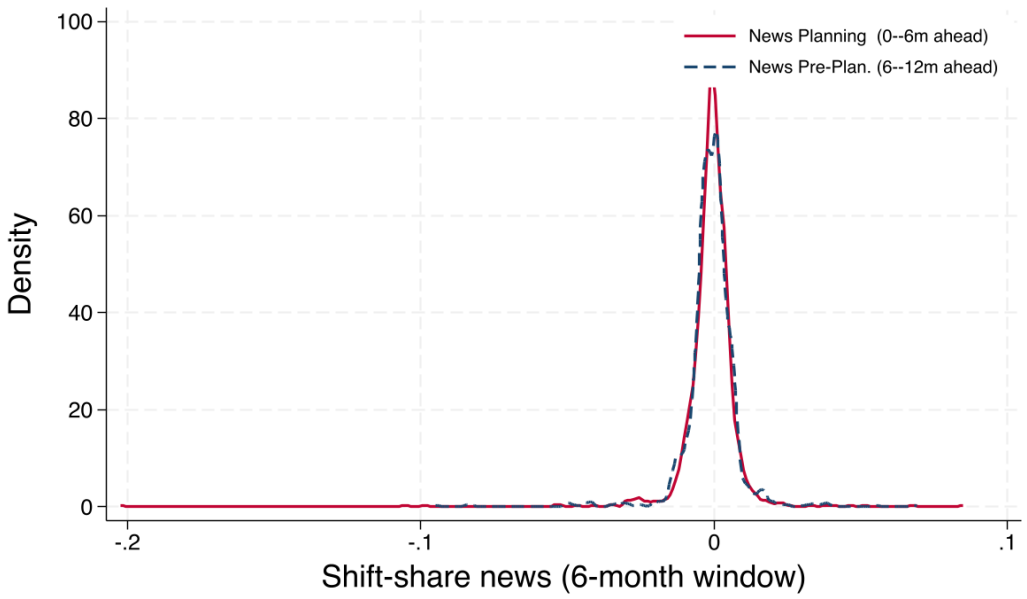


Figure B.3: Distribution of Planning and Pre-Planning Shift-Share News. The figure plots the empirical distributions of the preferred five-base AR(4)+GDP shift-share measure in the planning and pre-planning windows. The substantial overlap shows that the timing results are not driven by planning-window shocks occupying a distinct region of support.

proposal date approaches, consistent with attention concentrating around the institutional decision point.

C Robustness of Main Results

This appendix verifies that the decision-window findings from Sections **??** and **??** are robust along four dimensions. We first confirm that the results do not depend on how the shift-share news measure or the planning-stage outcome is constructed (Sections **??** and **??**). We then rule out calendar-month confounds using cross-state variation in proposal deadlines (Section **??**). Finally, we show that the planning-window premium cannot be attributed to differences in the informativeness of earlier versus later news (Section **??**). Throughout, the core pattern is stable: planning-window news strongly predicts the initial budget, while pre-planning news does not.

Table B.2: Within-Planning Recency: Preferred Five-Base Specification

This table splits the planning window into the near quarter (0–3 months before the proposal) and the far quarter (3–6 months before the proposal) using the preferred five-base AR(4)+GDP shift-share measure. Columns (1) and (2) report revenue forecasts; columns (3) and (4) expenditure plans. Odd columns include state and fiscal-year fixed effects; even columns add lagged macroeconomic controls measured prior to the planning window. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue forecast $\log(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1})$		Spending plan response $\log(\text{Init}^E_{s,t}/\text{Init}^E_{s,t-1})$	
	(1)	(2)	(3)	(4)
News Planning Near (0–3m ahead)	1.64***	1.81***	0.55***	0.56***
	(0.24)	(0.19)	(0.12)	(0.14)
News Planning Far (3–6m ahead)	1.15**	1.43**	0.69	0.74*
	(0.55)	(0.57)	(0.45)	(0.42)
State + FYear FE	✓	✓	✓	✓
Macro Controls		✓		✓
Observations	990	990	1,124	1,124
Adj. R^2	0.05	0.05	0.05	0.05

C.1 Alternative Revenue Base Combinations

The baseline specification uses five own-source revenue bases to construct the shift-share measure. We verify that the results are not sensitive to this choice by examining three-, four-, and six-base alternatives.

Relevance. Table **??** reports the relevance test from Section **??** across alternative revenue base combinations. The shift-share measure significantly predicts actual revenue growth using three, four, or five own-source bases with either raw or AR(4)+GDP shocks (Columns 1–4). Point estimates are stable across these specifications, ranging from 1.4 to 2.0. The six-base specification, which adds federal transfers, yields small and insignificant coefficients (Columns 5–6), consistent with federal grants responding to appropriations rather than macroeconomic conditions. This motivates the exclusion of federal transfers from the baseline.

Timing and recency. Table **??** confirms that the recency pattern is robust to alternative revenue base combinations. In Panel A, the planning-window coefficient remains large and highly significant across both the three-base and four-base specifications, with

Table C.1: Validation: Alternative Revenue Base Combinations

This table verifies that the relevance of shift-share news documented in Table **??** is robust to alternative combinations of own-source revenue bases. The dependent variable is actual revenue growth ($\log \text{ActR}_{s,t}/\text{ActR}_{s,t-1}$). "3-Base" includes individual income, severance, and corporate income taxes. "4-Base" adds general sales tax. "6-Base" adds charges and federal intergovernmental transfers. Odd columns use raw log-change shocks; even columns use AR(4)+GDP innovation residuals. All specifications include state and year fixed effects and lagged macroeconomic controls. Standard errors are two-way clustered by state and year. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Actual Revenue Growth: $\log(\text{ActR}_{s,t}/\text{ActR}_{s,t-1})$					
	3-Base		4-Base		6-Base	
	Raw	AR(4)	Raw	AR(4)	Raw	AR(4)
	(1)	(2)	(3)	(4)	(5)	(6)
Shift-Share News (Fiscal Year)	1.45**	2.01**	1.43*	1.86**	0.14	0.38
	(0.71)	(0.82)	(0.72)	(0.85)	(0.28)	(0.45)
State + FYear FE	✓	✓	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓	✓	✓
Observations	1,061	1,061	1,061	1,061	1,061	1,061
Adj. R^2	0.07	0.09	0.07	0.09	0.05	0.05

point estimates of 1.35–1.36 for revenues and 0.57–0.60 for expenditures—closely matching the baseline five-base results in Table **??**. Pre-planning windows are uniformly insignificant, with the exception of the most distant window (18–24 months), which enters with a marginally significant negative sign for expenditures, likely reflecting mean reversion in macroeconomic conditions rather than a genuine behavioral response. Panel B shows that the within-planning-window recency gradient is equally stable: the near-planning coefficient (0–3 months) dominates the far-planning coefficient (3–6 months) for both revenue base combinations, and the expenditure response to far-planning news remains small and insignificant regardless of specification.

C.2 Alternative Planning Benchmarks

The baseline timing regressions measure the planning-stage forecast update as $\log\left(\text{Init}_t^R/\text{Init}_{t-1}^R\right)$, benchmarking the current initial budget to the prior year's initial budget. A potential concern is that this choice of denominator could affect the timing result—for example, if the prior initial budget itself reflects planning-window news in a way that mechanically generates the asymmetry. We address this by re-estimating the horse-race specification using three alternative benchmarks that span progressively longer horizons.

Table C.2: Recency: Alternative Revenue Base Combinations

This table verifies that the recency pattern documented in Table **??** is robust to alternative combinations of own-source revenue bases. The dependent variable is the log change in the initial budget relative to the prior year's final budget. Panel A includes all four budget windows simultaneously; Panel B splits the planning window into two quarters. "3-Base" includes individual income, severance, and corporate income taxes. "4-Base" adds general sales tax. The shift-share measure uses AR(4)+GDP innovation shocks throughout. All specifications include state and year fixed effects. Panel B additionally includes lagged macroeconomic controls (four quarterly lags of GDP growth, payroll employment growth, S&P 500 returns, 5-year Treasury yields, and CPI inflation, measured before each state's planning window). Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Planning vs. pre-planning windows

	Revenue forecast $\log(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1})$		Expenditure plan response $\log(\text{Init}^E_{s,t}/\text{Init}^E_{s,t-1})$	
	3-Base	4-Base	3-Base	4-Base
	(1)	(2)	(3)	(4)
Planning (0–6m)	1.23***	1.15***	0.63***	0.58***
	(0.18)	(0.22)	(0.17)	(0.16)
Pre-planning (6–12m)	−0.07	−0.02	0.23	0.24
	(0.34)	(0.37)	(0.32)	(0.33)
State + FYear FE	✓	✓	✓	✓
Observations	990	990	1,124	1,124
Adj. R^2	0.05	0.05	0.05	0.05

Panel B: Within the planning window

	Revenue forecast $\log(\text{Init}^R_{s,t}/\text{Init}^R_{s,t-1})$		Expenditure plan response $\log(\text{Init}^E_{s,t}/\text{Init}^E_{s,t-1})$	
	3-Base	4-Base	3-Base	4-Base
	(1)	(2)	(3)	(4)
News Planning Near (0–3m ahead)	1.86***	1.83***	0.54***	0.55***
	(0.15)	(0.20)	(0.13)	(0.14)
News Planning Far (3–6m ahead)	1.38**	1.46**	0.80*	0.75*
	(0.61)	(0.63)	(0.40)	(0.43)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	990	990	1,124	1,124
Adj. R^2	0.05	0.05	0.05	0.05

Table **??** reports the results. Column (1) reproduces the baseline. Column (2) benchmarks to the prior year's enacted budget, Final^R_{t-1}, which is the reference point used in the Section **??** decomposition. Column (3) uses the two-year-lagged enacted budget, Final^R_{t-2}, which avoids any overlap with the current budget cycle. Column (4) uses three-year-lagged actuals, Act^R_{t-3}, a purely realized quantity with no planned component.

The planning-window coefficient is large and highly significant across all four specifications (1.18–1.88), confirming that the timing result is not an artifact of the benchmark choice. Pre-planning news remains statistically insignificant in every column. The point estimates on pre-planning news rise with longer benchmarks (from −0.10 in column 1 to 1.28 in column 4), but so do the standard errors, reflecting the additional variation absorbed by benchmarks that span longer horizons. In no case can we reject that pre-planning news has zero effect on the initial revenue forecast.

Table C.3: Robustness to Alternative Planning Benchmarks

Each column regresses the initial revenue forecast, measured relative to a different benchmark, on planning-window and pre-planning shift-share news entered jointly. Column (1) uses the baseline benchmark $\log\left(\text{Init}^R_t/\text{Init}^R_{t-1}\right)$. Column (2) uses the prior year's enacted budget, $\log\left(\text{Init}^R_t/\text{Final}^R_{t-1}\right)$. Column (3) uses the two-year-lagged enacted budget, $\log\left(\text{Init}^R_t/\text{Final}^R_{t-2}\right)$. Column (4) uses three-year-lagged actuals, $\log\left(\text{Init}^R_t/\text{Act}^R_{t-3}\right)$. All specifications include state and fiscal-year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	$\frac{\text{Init}^R_t}{\text{Init}^R_{t-1}}$	$\frac{\text{Init}^R_t}{\text{Final}^R_{t-1}}$	$\frac{\text{Init}^R_t}{\text{Final}^R_{t-2}}$	$\frac{\text{Init}^R_t}{\text{Act}^R_{t-3}}$
	(1)	(2)	(3)	(4)
News Planning (0–6m ahead)	1.32***	1.41***	1.18***	1.88***
	(0.12)	(0.15)	(0.13)	(0.38)
News Pre-Plan. (6–12m ahead)	−0.10	0.18	1.03	1.28
	(0.34)	(0.32)	(0.70)	(1.17)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	990	989	951	922
Adj. R^2	0.04	0.18	0.13	0.34

C.3 Cross-Deadline Horse Race: Early vs. Late Proposers

The previous two subsections show that the timing result is stable across measurement choices. A distinct concern is that the planning-window effect reflects something special about the calendar months that happen to fall in the planning window rather than the institutional timing per se. We address this by exploiting the fact that different states have different proposal deadlines, so the same calendar month falls in the planning window for some states and the pre-planning window for others.

We classify states by their modal proposal month into early proposers (October–January) and late proposers (February–March) and estimate the horse-race specification separately for each group. Table **??** confirms that both subsamples exhibit the planning-window pattern: the planning-window coefficient is large and highly significant for both early and late proposers, while pre-planning news is insignificant. Because a national shock in, say, November falls in the planning window for early proposers but the pre-planning window for late proposers, this result cannot be driven by calendar-month composition.

Table C.4: Cross-Deadline Horse Race: Early vs. Late Proposers

States are classified by modal proposal month into early proposers (October–January) and late proposers (February–March). Planning-window and pre-planning news are entered jointly. All specifications include state and fiscal-year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue forecast $\log(\text{Init}_t^R/\text{Init}_{t-1}^R)$		Expenditure plan response $\log(\text{Init}_t^E/\text{Init}_{t-1}^E)$	
	Early (1)	Late (2)	Early (3)	Late (4)
News Planning (0–6m ahead)	1.40***	3.44**	0.52***	2.08**
	(0.11)	(1.58)	(0.18)	(0.93)
News Pre-Plan. (6–12m ahead)	-0.27	-1.50	-0.01	-2.76
	(0.36)	(2.33)	(0.29)	(2.01)
Sample	Early (Oct–Jan)	Late (Feb–Mar)	Early (Oct–Dec)	Late (Jan–Mar)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	672	318	783	341
Adj. R^2	0.090	-0.055	0.075	-0.014

C.4 Information-Equivalent News Test

A remaining defense of the baseline timing pattern is that pre-planning news may simply be less informative than planning-window news about fiscal-year revenues. If so, a stronger proposal-stage response to planning-window news need not reflect decision-window bias; it could instead reflect rationally greater weight on more informative signals.

To assess this possibility, we construct information-equivalent versions of the window-specific shift-share shocks. The goal is not to recover structural signal precision, but to place news arriving in different windows into comparable ex post revenue units. We then re-estimate the proposal-stage timing regression using these rescaled shocks.

Construction. Let $News^w_{s,t}$ denote the shift-share news for state s, fiscal year t, and window w. To measure how informative each window's news is for fiscal-year-t revenues, we estimate:

$$\log\left(\frac{\text{Act}^R_{s,t}}{\text{Final}^R_{s,t-1}}\right) = \alpha_s + \delta_t + \sum_w \pi_w \, News^w_{s,t} + \Gamma' X_{s,t} + u_{s,t}, \tag{C.1}$$

where $X_{s,t}$ contains the same macro controls as in the main timing regressions, and the sum runs over four six-month windows spanning the two years before the proposal.

To avoid in-sample contamination, we estimate (**??**) leaving state s out and define the rescaled shocks as

$$\widetilde{News}^w_{s,t} \equiv \hat{\pi}_{w,-(s)} \, News^w_{s,t}.$$

By construction, a one-unit increase in $\widetilde{News}^w_{s,t}$ corresponds to the same ex post predicted change in the fiscal-year-t revenue outcome $Y_{s,t}$, regardless of when the news arrived.

Test. We re-estimate the proposal-stage regression replacing raw shocks with $\widetilde{News}^w$. This exercise is best interpreted as a refinement of the main timing results in Section **??**. Section **??** shows that planning-window news enters the initial budget while pre-planning news does not. The present exercise asks whether that asymmetry can be explained simply by differences in ex post informativeness. Under complete updating, once shocks are in comparable revenue units the proposal should respond equally to planning and pre-planning news ($\beta_{\text{plan}} \approx \beta_{\text{pre}}$). Under decision-window bias, planning-window news should retain a premium ($\beta_{\text{plan}} > \beta_{\text{pre}}$).

Results. Table **??** reports the results. The left panel reproduces the baseline timing regression using un-scaled shift-share shocks; the right panel uses information-equivalent shocks

constructed via leave-one-state-out estimation of the predictive slopes $\hat{\pi}_w$ in (**??**). Across all three outcomes—revenue forecasts, expenditure plans, and actual expenditures—the planning-window coefficient remains large and highly significant after rescaling (0.63–1.03), while pre-planning coefficients remain small and statistically insignificant. The planning-window premium survives the rescaling, confirming that the timing asymmetry reflects selective incorporation of information rather than differences in the informativeness of earlier versus later news.

Table C.5: Proposal-Stage Response: Un-scaled vs. Information-Equivalent Shocks

The left panel reports coefficients from regressing each outcome on un-scaled shift-share shocks by window. The right panel uses information-equivalent shocks $\widetilde{News}^{w}_{s,t} = \hat{\pi}_{w,-(s)} News^{w}_{s,t}$, where $\hat{\pi}_{w,-(s)}$ is the predictive slope from regressing realized revenue $\log\left(\text{Act}^{R}_{s,t}/\text{Final}^{R}_{s,t-1}\right)$ on all window-specific shocks jointly, estimated leaving state s out. Standard errors clustered by state in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

	Un-scaled Shift-Share Shocks			Information-Equivalent Shocks		
	Rev. Forecast	Exp. Plan	Exp. Actual	Rev. Forecast	Exp. Plan	Exp. Actual
	(1)	(2)	(3)	(4)	(5)	(6)
News Planning (0–6m)	1.40***	0.71***	0.85***	1.03***	0.52***	0.63***
	(0.16)	(0.19)	(0.16)	(0.15)	(0.16)	(0.14)
News Pre-Plan. (6–12m)	0.09	0.38	0.48	0.08	0.35	0.40
	(0.33)	(0.49)	(0.38)	(0.39)	(0.54)	(0.38)
News Pre-Plan. (12–18m)	-0.09	0.12	-0.07	-0.03	-0.17**	0.03
	(0.19)	(0.14)	(0.19)	(0.10)	(0.07)	(0.06)
News Pre-Plan. (18–24m)	-0.54	-0.49*	-0.18	0.06	0.09	-0.21***
	(0.36)	(0.27)	(0.39)	(0.06)	(0.06)	(0.07)
Info-equiv. rescaling	No	No	No	Yes (LOO)	Yes (LOO)	Yes (LOO)
State + FYear FE	✓	✓	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓	✓	✓
Observations	989	1,122	1,146	989	1,122	1,146
Adj. R^2	0.18	0.27	0.05	0.18	0.27	0.05

C.5 Pseudo-Window Placebo

A natural concern is that the planning-window effect reflects something mechanical about the six-month aggregation window or the particular national shocks it captures, rather than the institutional budget calendar per se. We address this with a randomization test that permutes budget calendars across states while keeping outcomes, revenue composition, and national shocks fixed.

Specifically, we randomly reassign shift-share news realizations across states within each fiscal year, so that state B's forecast outcome is paired with the news constructed for state A (reflecting state A's revenue-composition weights and budget-calendar window). Since cross-state variation in the news measure arises primarily from differences in proposal timing—which calendar months fall in the planning versus pre-planning window—this permutation tests whether the result requires the actual institutional calendar. We repeat this procedure 1,000 times and compare the actual planning-window coefficient (and the planning-minus-pre-planning contrast) to the permutation distribution.

Figure **??** reports the results. The actual planning coefficient ($\hat{\beta}_{\text{plan}} = 1.31$) exceeds all but one of the 1,000 permuted estimates ($p = 0.001$). The planning-minus-pre-planning contrast (1.53) is likewise in the far right tail of the permutation distribution ($p = 0.011$). The permutation distribution is centered near zero, confirming that randomly assigned budget calendars do not generate the timing asymmetry. The planning-window effect is tied to the actual institutional calendar.

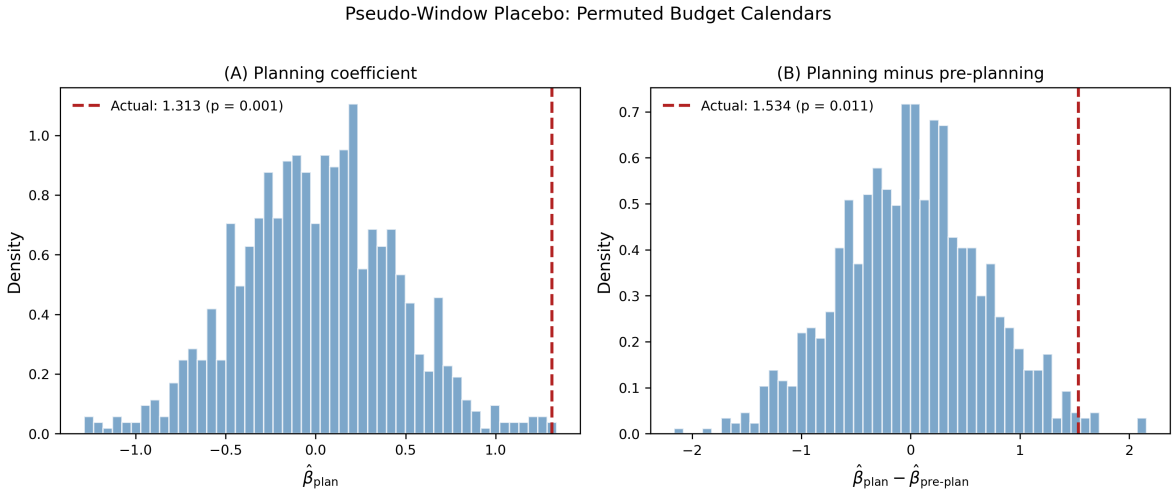


Figure C.1: Pseudo-Window Placebo: Permuted Budget Calendars. Panel (A) plots the distribution of $\hat{\beta}_{\text{plan}}$ from 1,000 permutations that randomly reassign shift-share news across states within each fiscal year, breaking the link between a state's actual budget calendar and its forecast outcome. Panel (B) plots the distribution of $\hat{\beta}_{\text{plan}} - \hat{\beta}_{\text{pre-plan}}$. Dashed red lines mark the actual estimates. The dependent variable is $\log\left(\text{Init}^{R}_{s,t}/\text{Init}^{R}_{s,t-1}\right)$. All regressions include state and fiscal-year fixed effects.

D Robustness of Real-Effects Results

This subsection reports robustness checks for the government-exposed employment and wage results in Section **??**. Table **??** varies the definition of government exposure and the sample

of states; Table **??** uses alternative shift-share constructions. Throughout, the planning-window interaction coefficient remains positive and significant, while pre-planning news has no detectable differential effect.

Leave-one-state-out. Figure **??** reports the distribution of the planning-window interaction coefficient γ_{plan} from re-estimating equation (**??**) (saturated specification with all windows) dropping one state at a time. For both employment and wages, the coefficients cluster tightly around the full-sample estimate, confirming that no single state drives the result.

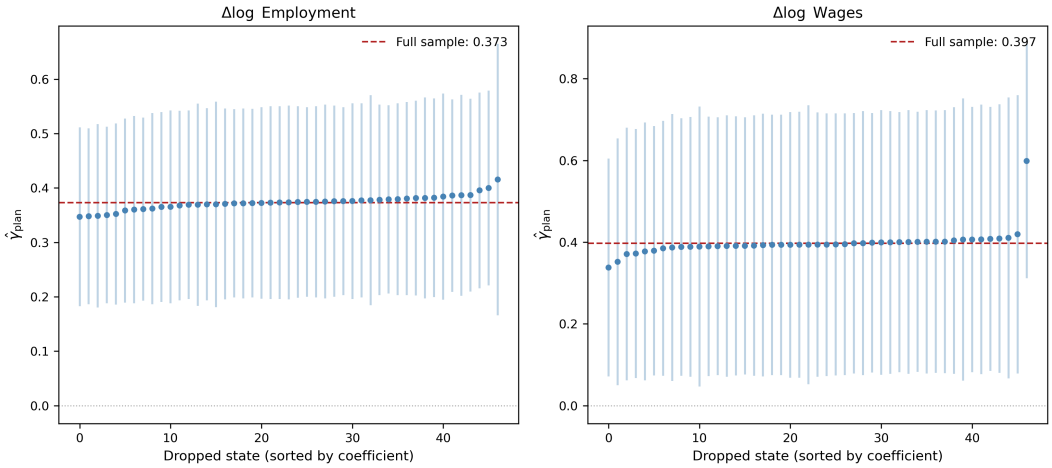


Figure D.1: Leave-One-State-Out: Planning-Window Interaction Coefficient. Each dot represents $\hat{\gamma}_{\text{plan}}$ from the saturated specification (equation **??**) estimated after dropping one state. The dashed line marks the full-sample estimate. Panel (A): Δ log Employment; Panel (B): Δ log Wages. All specifications include state×industry, industry×fiscal-year, and state×fiscal-year fixed effects with standard errors clustered by state.

Excluding resource-dependent states. States with large severance tax revenues (e.g., Alaska, Wyoming, North Dakota) have unusual revenue composition and may dominate the shift-share variation through oil price movements. Column (2) of Table **??** drops states whose average severance tax share exceeds 5% of general revenue. The planning-window coefficient is essentially unchanged in both panels.

Alternative government exposure definitions and weighting. The baseline uses employment-weighted regressions with a binary indicator for industries whose government-purchase share exceeds 5%. Table **??** varies both choices. Column (3) removes employment weights; column (4) uses the median government-purchase share as the exposure threshold; column (5) raises the threshold to 6%; column (6) replaces the binary indicator with the

continuous government-purchase share. The planning-window interaction remains positive and significant across all specifications, while the pre-planning interaction is insignificant.

Table D.1: Robustness of Real Effects: Sample and Exposure Definition

This table reports robustness checks for the government-exposed employment and wage results in Table **??**. The baseline (column 1) uses the 5-base AR(4)+GDP shift-share, a binary government-exposure indicator at the 5% threshold, and employment weights. Column (2) drops states whose average severance tax share exceeds 5% of general revenue. Column (3) removes employment weights. Column (4) uses the median government-purchase share ($\sim$ 3%) as the exposure threshold. Column (5) uses a 6% threshold. Column (6) replaces the binary indicator with the continuous government-purchase share. Each panel row reports the specification's weighting scheme and exposure definition. All specifications include state×industry, industry×fiscal-year, and state×fiscal-year fixed effects. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Panel A: $\Delta \log(\text{Private Employment})^{FY}$					
	(1)	(2)	(3)	(4)	(5)	(6)
News Plan. × GovExp	0.552***	0.709***	0.244***	0.415***	0.474***	3.344**
	(0.146)	(0.154)	(0.066)	(0.119)	(0.136)	(1.325)
News Pre-Plan. × GovExp	-0.008	-0.073	0.170*	0.014	0.154	0.215
	(0.129)	(0.132)	(0.093)	(0.126)	(0.094)	(1.087)
Specification	Baseline	Drop Sev.	Unwtd	Med Cut	6% Cut	Cts Share
FE (S×I, I×Y, S×Y)	✓	✓	✓	✓	✓	✓
Observations	22,326	21,414	22,407	22,326	22,326	22,326
Adj.R^2	0.791	0.792	0.572	0.791	0.791	0.791

	Panel B: $\Delta \log(\text{Private Wages})^{FY}$					
	(1)	(2)	(3)	(4)	(5)	(6)
News Plan. × GovExp	0.448**	0.652***	0.183	0.363**	0.391***	3.275**
	(0.171)	(0.137)	(0.125)	(0.147)	(0.141)	(1.318)
News Pre-Plan. × GovExp	0.053	-0.024	0.162	0.109	0.203*	0.906
	(0.108)	(0.100)	(0.125)	(0.119)	(0.120)	(1.094)
Specification	Baseline	Drop Sev.	Unwtd	Med Cut	6% Cut	Cts Share
FE (S×I, I×Y, S×Y)	✓	✓	✓	✓	✓	✓
Observations	22,326	21,414	22,407	22,326	22,326	22,326
Adj.R^2	0.748	0.749	0.446	0.748	0.748	0.748

Alternative shift-share constructions. Table **??** replaces the baseline 5-base shift-share measure with 3-base (individual income, severance, and corporate income taxes) and 4-base (adding general sales tax) alternatives. The planning-window interaction coefficient is stable across all three specifications for both employment and wages.

Table D.2: Robustness of Real Effects: Alternative Shift-Share Constructions

This table replaces the baseline 5-base AR(4)+GDP shift-share measure with alternative revenue-base combinations. "3-Base" includes individual income, severance, and corporate income taxes. "4-Base" adds general sales tax. All specifications use the saturated interaction design (equation **??**) with state×industry, industry×fiscal-year, and state×fiscal-year fixed effects. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Panel A: $\Delta \log(\text{Private Employment})^{FY}$		
	5-Base	3-Base	4-Base
	(1)	(2)	(3)
News Plan. × GovExp	0.552***	0.736***	0.731***
	(0.146)	(0.178)	(0.165)
News Pre-Plan. × GovExp	-0.008	-0.083	-0.079
	(0.129)	(0.161)	(0.141)
FE (S×I, I×Y, S×Y)	✓	✓	✓
Observations	22,326	22,326	22,326
Adj.R^2	0.791	0.792	0.792

	Panel B: $\Delta \log(\text{Private Wages})^{FY}$		
	5-Base	3-Base	4-Base
	(1)	(2)	(3)
News Plan. × GovExp	0.448**	0.571***	0.602***
	(0.171)	(0.146)	(0.155)
News Pre-Plan. × GovExp	0.053	-0.033	0.023
	(0.108)	(0.140)	(0.127)
FE (S×I, I×Y, S×Y)	✓	✓	✓
Observations	22,326	22,326	22,326
Adj.R^2	0.748	0.748	0.748

Revenue Non-Controllability

The availability–belief–action framework in Section **??** treats realized state tax revenue as an outcome the government does not directly control over the fiscal-year horizon. In practice, states can alter tax rates, adjust enforcement, or modify the tax base. If Bartik shocks trigger legislative tax policy responses, realized revenue would partly reflect the state's own policy choices, contaminating the forecast-error equation.

We test this using Tax Foundation data on state individual income tax top marginal rates for 2000–2014. For each state-year, we construct indicator variables for whether the top marginal rate changed, increased, or decreased relative to the prior year. We then

regress each indicator on planning-window and pre-planning news using the same state and fiscal-year fixed effects, macro controls, and clustering as the main specification.

Table D.3: Do Bartik shocks predict state tax policy changes?

	(1) Any change	(2) Increase	(3) Decrease
Planning news	0.346	-0.154	0.500
	(0.947)	(0.369)	(1.100)
Pre-planning news	-0.158	-0.033	-0.126
	(1.193)	(0.386)	(0.973)
Observations	627	627	627
Adj. R^2	0.163	0.070	0.127

State + FY FE; macro controls; SEs clustered by state.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column regresses a tax-policy-change indicator on planning-window and pre-planning shift-share news. Tax rate data from the Tax Foundation (top marginal state individual income tax rate, 2000–2014). State + FY FE; macro controls; SEs clustered by state.

Neither planning-window nor pre-planning news predicts tax rate changes (Table **??**). All coefficients are small, statistically insignificant, and economically negligible. The result holds separately for increases and decreases. Revenue news therefore does not trigger tax legislation over this sample, supporting the assumption that realized revenue is non-controlled.

E Forecast Error Decomposition

The main text establishes that pre-planning news predicts revenue forecast errors (Table **??**), indicating that information available at the proposal date is not fully incorporated into the initial budget. A natural follow-up question is: where does this ignored information eventually surface? Does the legislature correct the governor's omission during the revision phase, or does the neglected news instead materialize as unplanned within-year revenue surprises? And is the forecast error driven by a miscalibrated budget level or by unpredictable current-year volatility?

This appendix addresses these questions by decomposing the forecast error along two dimensions. We also examine whether the same patterns appear for expenditure forecast errors.

Expenditure forecast errors. For completeness, we also estimate the baseline forecast-efficiency regression using expenditure forecast errors, $\log\left(\text{Act}^E_{s,t}/\text{Init}^E_{s,t}\right)$, as the dependent variable. Unlike revenue forecast errors, expenditure forecast errors show little systematic predictability from either planning-window or pre-planning news, consistent with realized spending being constrained by the enacted budget.

Table E.1: Expenditure Forecast Errors and Budget Timing

The dependent variable in all columns is the expenditure forecast error, $\log\left(\text{Act}^E_{s,t}/\text{Init}^E_{s,t}\right)$. Regressors are shift-share news measures constructed over windows relative to each state's proposal date: Planning denotes 0–6 months and Pre-planning denotes 6–12 months. Column (3) includes both planning and pre-planning news; column (4) additionally includes 12–18 and 18–24 month pre-planning windows. All specifications include state and year fixed effects and lagged macroeconomic controls measured prior to the planning window. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Expenditure Forecast Error $\log\left(\text{Act}^E_{s,t}/\text{Init}^E_{s,t}\right)$			
	(1)	(2)	(3)	(4)
News Planning (0–6m ahead)	0.02		0.02	0.00
	(0.14)		(0.13)	(0.14)
News Pre-Plan. (6–12m ahead)		-0.02	-0.02	0.01
		(0.24)	(0.24)	(0.24)
News Pre-Plan. (12–18m ahead)				-0.08
				(0.15)
News Pre-Plan. (18–24m ahead)				0.15
				(0.42)
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	1,159	1,159	1,159	1,159
Adj.R^2	0.68	0.68	0.68	0.68

Revision vs. execution. We decompose the forecast error using the final approved budget as pivot:

$$\text{FE}^R_{s,t} = \underbrace{\log\left(\frac{\text{FinalR}_{s,t}}{\text{InitR}_{s,t}}\right)}_{\text{Legislative revision}} + \underbrace{\log\left(\frac{\text{ActR}_{s,t}}{\text{FinalR}_{s,t}}\right)}_{\text{Execution surprise}}. \tag{E.1}$$

If the legislature corrects the governor's blind spot, pre-planning shocks should predict the revision component. If the ignored information instead materializes as unplanned fiscal-year surprises, it should predict the execution component.

Level miscalibration vs. current-year volatility. We also decompose using the prior-year actual as pivot:

$$\text{FE}^R_{s,t} = \underbrace{\log\left(\frac{\text{ActR}_{s,t-1}}{\text{InitR}_{s,t}}\right)}_{\text{Prior actual gap}} + \underbrace{\log\left(\frac{\text{ActR}_{s,t}}{\text{ActR}_{s,t-1}}\right)}_{\text{Actual growth}}. \tag{E.2}$$

This asks: is the forecast error driven by a miscalibrated budget level (the initial budget doesn't reflect the prior revenue trajectory) or by unpredictable current-year revenue volatility?

Results. Table **??** traces where the ignored pre-planning information eventually surfaces. Panel A decomposes the revenue forecast error into a legislative revision component, $\log\big(\text{Final}^R/\text{Init}^R\big)$, and an execution surprise, $\log\big(\text{Act}^R/\text{Final}^R\big)$. Pre-planning news (6–12 months) strongly predicts the execution component (coefficient 1.000, $p < 0.01$) but not the revision component (0.254, insignificant). The legislature does not systematically correct the governor's neglect of earlier information; instead, the ignored news materializes as unplanned within-year revenue surprises. For expenditures, neither component is significantly predicted by pre-planning news, consistent with the enacted spending plan disciplining actual outlays.

Panel B uses the prior-year actual as pivot, decomposing the forecast error into a prior actual gap, $\log\big(\text{Act}^R_{t-1}/\text{Init}^R_t\big)$, and actual revenue growth, $\log\big(\text{Act}^R_t/\text{Act}^R_{t-1}\big)$. All three pre-planning windows strongly predict the prior actual gap (coefficients 1.238, 0.778, and 1.343, all significant at the 1% level) but not actual revenue growth. This pattern reveals that the forecast error is driven by level miscalibration: the initial budget fails to reflect the revenue trajectory already established by past economic conditions, rather than by unpredictable current-year volatility. Officials set the baseline revenue level as if earlier conditions had not occurred, producing a systematic gap between the budget starting point and realized revenues.

Asymmetric forecast errors: positive vs. negative pre-planning news. A natural question is whether the neglect of pre-planning information differs for favorable versus adverse news. If officials are selectively optimistic—ignoring bad early news while incorporating good news—or selectively pessimistic, the forecast-efficiency violation would be concentrated in one tail. To test this, we decompose pre-planning news into positive and negative components, $\text{News}^+_{s,t,\text{pre}} \equiv \max(\text{News}_{s,t,\text{pre}}, 0)$ and $\text{News}^-_{s,t,\text{pre}} \equiv \min(\text{News}_{s,t,\text{pre}}, 0)$,

Table E.2: Decomposing the Forecast Error: Where Does Pre-Planning Information Emerge?

Panel A decomposes the revenue forecast error using the final approved budget as pivot: the revision component is $\log\left(\text{Final}^R/\text{Init}^R\right)$ and the execution surprise is $\log\left(\text{Act}^R/\text{Final}^R\right)$. Panel B uses the prior-year actual as pivot: the prior actual gap is $\log\left(\text{Act}^R_{t-1}/\text{Init}^R_t\right)$ and actual growth is $\log\left(\text{Act}^R_t/\text{Act}^R_{t-1}\right)$. Analogous decompositions are shown for expenditures. All specifications include state and year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Decomposition using final budget as pivot

	Revenue		Expenditure	
	Revision $\log\left(\text{Final}^R/\text{Init}^R\right)$ (1)	Execution $\log\left(\text{Act}^R/\text{Final}^R\right)$ (2)	Revision $\log\left(\text{Final}^E/\text{Init}^E\right)$ (3)	Execution $\log\left(\text{Act}^E/\text{Final}^E\right)$ (4)
News Pre-Plan. (6–12m ahead)	0.22	0.87***	0.23	-0.25
	(0.22)	(0.31)	(0.27)	(0.15)
State+FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	1,024	1,040	1,159	1,175
Adj.R^2	0.466	0.600	0.459	0.789

Panel B: Decomposition using prior-year actual as pivot

	Revenue		Expenditure	
	Prior Actual $\log\left(\text{Act}^R_{t-1}/\text{Init}^R_t\right)$ (1)	Actual Growth $\log\left(\text{Act}^R_t/\text{Act}^R_{t-1}\right)$ (2)	Prior Actual $\log\left(\text{Act}^E_{t-1}/\text{Init}^E_t\right)$ (3)	Actual Growth $\log\left(\text{Act}^E_t/\text{Act}^E_{t-1}\right)$ (4)
News Pre-Plan. (6–12m ahead)	1.09**	-0.11	-0.15	0.44
	(0.41)	(0.39)	(0.48)	(0.34)
State+FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	995	1,061	1,128	1,146
Adj.R^2	0.438	0.043	0.526	0.041

and estimate

$$\text{(E.3)}\quad \text{FE}^R_{s,t} = \alpha_s + \delta_t + \beta^{pre+}\text{News}^+_{s,t,\text{pre}} + \beta^{pre-}\text{News}^-_{s,t,\text{pre}} + \beta^{\text{plan}}\text{News}_{s,t,\text{plan}} + \gamma' X_{s,t} + \varepsilon_{s,t}.$$

Under symmetric neglect, $\beta^{pre+} \approx \beta^{pre-}$; under selective optimism, $|\beta^{pre-}| > |\beta^{pre+}|$; under motivated over-optimism, $|\beta^{pre+}| > |\beta^{pre-}|$.

Table **??** reports the results. Point estimates on positive and negative pre-planning news are similar in magnitude (0.80–1.02 vs. 1.11–1.18), and an F-test cannot reject symmetry in any specification ($p > 0.76$). The negative component is more precisely estimated, likely reflecting greater variance in adverse shocks rather than a true asymmetry. Planning-window news remains insignificant throughout, whether entered symmetrically (column 2) or split by sign (column 3). The results indicate that the decision-window bias reflects blanket inattention to pre-planning information regardless of its sign, rather than strategic revenue shading or motivated reasoning.

Legislative-stage information rigidity test. The main text applies the **?** test to the governor's proposal and finds no evidence of underreaction ($\beta \approx 0$). We apply the same logic to the legislative stage: does the legislative revision, $\log\big(\text{Final}^R_{s,t}/\text{Init}^R_{s,t}\big)$, predict the subsequent execution error, $\log\big(\text{Act}^R_{s,t}/\text{Final}^R_{s,t}\big)$? Under information rigidities, $\beta > 0$: the legislature's revision should predict same-sign execution surprises.

Table **??** reports the results. The legislative revision coefficient is negative and significant (-0.20, $p < 0.05$), indicating that larger upward legislative revisions are associated with actual revenue falling further below the enacted budget. Importantly, column (3) shows that pre-planning news continues to predict execution-stage revenue surprises (0.90, $p < 0.01$) even after controlling for the legislative revision, confirming that the legislature does not correct the governor's neglect of earlier information.

F Alternative Theories: Derivations and Additional Results

This appendix provides formal derivations for the heterogeneity tests in Section **??**. We first derive the rational inattention predictions that motivate the informativeness and fiscal-stakes tests, then show that the rolling signal-to-noise ratio used in the paper is a valid proxy for structural signal precision. We conclude with additional results on political incentives.

Table E.3: Asymmetric Forecast Errors: Positive vs. Negative News

The dependent variable is the revenue forecast error, $\log\left(\text{Act}^R_{s,t}/\text{Init}^R_{s,t}\right)$. Pre-planning news is decomposed into positive (> 0) and negative (< 0) components. Column (1) includes asymmetric pre-planning news only; column (2) adds symmetric planning-window news; column (3) splits both windows by sign. The F-test reports the test of $\beta^{pre+} = \beta^{pre-}$. All specifications include state and year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue Forecast Error $\log\left(\text{Act}^R_{s,t}/\text{Init}^R_{s,t}\right)$		
	(1)	(2)	(3)
News Planning (0–6m)		0.09	
		(0.19)	
News Planning (0–6m, > 0)			1.42
			(1.11)
News Planning (0–6m, < 0)			-0.44
			(0.53)
News Pre-Plan. (6–12m, > 0)	0.96	1.02	0.80
	(0.66)	(0.68)	(0.81)
News Pre-Plan. (6–12m, < 0)	1.18**	1.13**	1.11**
	(0.52)	(0.54)	(0.53)
F-test: $\beta^{pre+} = \beta^{pre-}$	0.05	0.01	0.09
p-value	0.820	0.911	0.767
State + FYear FE	✓	✓	✓
Macro Controls	✓	✓	✓
Observations	1,024	1,024	1,024
Adj.R^2	0.54	0.54	0.54

Table E.4: Legislative-Stage Coibion-Gorodnichenko Test

This table applies the **?** information rigidity test to the legislative stage. Columns (1)–(3) regress the revenue execution error, $\log\left(\text{Act}_t^R/\text{Final}_t^R\right)$, on the legislative revision, $\log\left(\text{Final}_t^R/\text{Init}_t^R\right)$. Column (3) additionally includes pre-planning news (6–12 months before the proposal). All specifications include state and year fixed effects. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue Execution Error $\log\left(\text{Act}_t^R/\text{Final}_t^R\right)$		
	(1)	(2)	(3)
Legislative Revision $\log\left(\text{Final}_t^R/\text{Init}_t^R\right)$	-0.21**	-0.20**	-0.20**
	(0.08)	(0.08)	(0.08)
News Pre-Plan. (6–12m)			0.90***
			(0.31)
State + FYear FE	✓	✓	✓
Macro Controls		✓	✓
Observations	1,024	1,024	1,024
Adj.R^2	0.61	0.60	0.61

F.1 Rational Inattention: Formal Predictions

Setup. Let $\theta_{s,t}$ be the latent fiscal condition, with realized revenue $R_{s,t} = \bar{R}_s + \kappa\,\theta_{s,t} + \xi_{s,t}$ where $\xi_{s,t} \sim (0, \sigma_\xi^2)$. Before the proposal date, policymakers observe two signals $x_{s,t}^w = \theta_{s,t} + \varepsilon_{s,t}^w$ for $w \in \{pre, plan\}$, with precision $q_w \equiv 1/\sigma_w^2$. Under full Bayesian updating, the optimal weight on signal w is $\omega_w^* = q_w/(q_0 + q_{pre} + q_{plan})$.

Predictions. Under rational inattention, attention weights are endogenous: $a_w = a_w(q_w, L_{s,t})$, where $L_{s,t}$ indexes the cost of forecast errors. The first-order condition from a quadratic loss objective implies

$$a_w = \frac{\omega_w^*}{1 + c_w/(\kappa^2 L_{s,t})},$$

where c_w is the marginal processing cost. This yields two testable monotonicity conditions:

$$\frac{\partial a_{pre}}{\partial q_{pre}} > 0, \qquad \frac{\partial a_{pre}}{\partial L} > 0. \tag{F.1}$$

Higher signal precision or higher forecast stakes should increase the weight on pre-planning news. In the proposal regression, this implies a positive interaction $\theta_{pre} > 0$; in the forecast-

error regression, a negative interaction $\rho_{pre} < 0$.

F.2 Rolling Signal-to-Noise Ratio as a Precision Proxy

The heterogeneity tests require a proxy for the structural precision q_w. We show that the rolling signal-to-noise ratio constructed in Section **??** is monotonically increasing in q_w.

Regress realized revenue on the signal from window w:

$$R_{s,t} = \alpha + \pi_w \, x^w_{s,t} + u^w_{s,t}. \tag{F.2}$$

The population slope is

$$\pi_w = \frac{\kappa \, \sigma_\theta^2}{\sigma_\theta^2 + \sigma_w^2} = \frac{\kappa \, \sigma_\theta^2 \, q_w}{1 + \sigma_\theta^2 \, q_w}, \tag{F.3}$$

which is strictly increasing in q_w. The residual variance is

$$\mathrm{var}(u^w_{s,t}) = \sigma_\xi^2 + \kappa^2 \frac{\sigma_\theta^2 \, \sigma_w^2}{\sigma_\theta^2 + \sigma_w^2}, \tag{F.4}$$

which is decreasing in q_w. The population signal-to-noise ratio

$$SNR_w \;\equiv\; \frac{\pi_w^2 \mathrm{var}(x^w_{s,t})}{\mathrm{var}(u^w_{s,t})} = \frac{\kappa^2 \sigma_\theta^4}{\sigma_\xi^2(\sigma_\theta^2 + \sigma_w^2) + \kappa^2 \sigma_\theta^2 \sigma_w^2} \tag{F.5}$$

is therefore strictly increasing in q_w.

The empirical proxy used in the paper is the direct sample analogue. For each state s, fiscal year t, and window $w \in \{pre, plan\}$, we estimate

$$\log\left(\frac{\mathrm{Rev}_{s,\tau}}{\mathrm{Rev}_{s,\tau-3}}\right) = a_{s,t,w} + \pi_{s,t,w} \, x^w_{s,\tau} + b_{s,t,w} \, \tau + u^w_{s,\tau} \tag{F.6}$$

using a 20-year rolling window $\tau \in \{t-21, \ldots, t-2\}$ and construct $\widehat{SNR}^{\,w}_{s,t} \equiv \widehat{\pi}^2_{s,t,w} \mathrm{var}(x^w_{s,\tau}) / \widehat{\mathrm{var}}(\hat{u}^w_{s,\tau})$. Because the estimation window ends at $t-2$, the proxy uses only information available before the proposal is formed.

Revenue concentration, measured by the Herfindahl index $\sum_k w^2_{s,k,t}$, provides a complementary proxy: states with more concentrated revenue bases are more exposed to identifiable national shocks, so a given innovation maps more tightly into total revenues.

F.3 Political Incentives

Section ?? notes that the decision-window pattern could reflect strategic forecasting rather than informational frictions. Table ?? interacts news measures with governor party and election timing.[6] Neither proxy significantly increases the incorporation of pre-planning news, and pre-planning news continues to predict forecast errors regardless of political environment.

Table F.1: Political Incentives: Do Governor Party and Election Timing Matter?

Each column interacts planning-window and pre-planning news with a political-environment proxy. Columns (1) and (3) use governor party (Democrat = 1). Columns (2) and (4) use governor election year. Columns (1)–(2) report the proposal forecast regression; columns (3)–(4) the forecast-error regression. All specifications include state and fiscal-year fixed effects and lagged macroeconomic controls. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue Forecast: $\log(\text{Init}^R/\text{Init}^R_{t-1})$		Revenue Forecast Error: $\log(\text{Act}^R/\text{Init}^R)$	
	(1)	(2)	(3)	(4)
News Planning	1.58***	1.30***	0.11	-0.24
	(0.13)	(0.15)	(0.11)	(0.34)
News Pre-Plan.	-0.04	0.21	1.33***	1.82***
	(0.61)	(0.36)	(0.43)	(0.33)
News Planning × Indicator	-0.96**	0.05	0.13	1.70*
	(0.42)	(0.42)	(0.38)	(0.89)
News Pre-Plan. × Indicator	0.45	-0.69	-1.05	-3.25***
	(0.84)	(1.01)	(0.92)	(0.98)
Indicator	Gov. party (D)	Election year	Gov. party (D)	Election year
State + FYear FE	✓	✓	✓	✓
Macro Controls	✓	✓	✓	✓
Observations	941	942	974	976
Adj. R^2	0.05	0.05	0.49	0.56

G Structural Estimation: Details and Model Fit

This appendix provides details on the estimation of the structural model in Section ??. We describe the estimation procedure (Section ??), the targeted moments and the variation that identifies each parameter (Section ??), the fit to targeted and untargeted moments (Section ??), and sensitivity diagnostics (Section ??).

[6] Governor party is coded as a binary indicator (Democrat = 1, Republican = 0) using the State Partisan Balance dataset (?) for years before 2009 and National Conference of State Legislatures reports for 2009–2026. Governor election year is coded from the same sources.

G.1 Estimation Procedure

We estimate the model by minimum distance, choosing parameter values to minimize the weighted distance between empirical moments and their model-implied counterparts. The procedure has three steps.

Step 1: Simulation. For a given parameter vector θ, we simulate the model (equations **??**–**??**) for a panel of $N = 47$ states over $T = 30$ years with a 3-year burn-in period. We draw common random numbers (seed 12345) for the four shock windows ($z^{\text{pre}}, z^{\text{plan}}, z^{\text{rev}}, z^{\text{real}}$) and idiosyncratic noise (ν). The simulation produces panel data on initial and final revenue forecasts ($\hat{R}^I, \hat{R}^F$), initial and enacted expenditure plans (E^I, E^F), actual spending (E^{act}), and government-exposed employment and wages.

Step 2: Model moments. On the simulated panel, we compute model-implied moments $\mathbf{m}(\theta)$ using the same procedures as in the empirical analysis. For regression-based moments (e.g., the coefficient of planning-window news on the initial revenue forecast), we run two-way fixed-effects regressions demeaning by state and year. For the variance decomposition, we compute the covariance of each budget-stage growth component with realized expenditure growth, matching the decomposition in equation (**??**) of the main text. This ensures that model and empirical moments are computed using identical procedures.

Step 3: Optimization. The estimator minimizes

$$Q(\theta) = \left[\mathbf{m}^{\text{data}} - \mathbf{m}(\theta)\right]' W \left[\mathbf{m}^{\text{data}} - \mathbf{m}(\theta)\right], \tag{G.1}$$

where $W = \text{diag}(1/\text{SE}_j^2)$ is the diagonal inverse-standard-error weighting matrix. Moments without standard errors (variance shares, execution-gap medians) use the median SE of the remaining moments. We use 15 multistart runs with randomized initial conditions and logit-bounded parameter transforms to avoid boundary solutions. All 15 runs converge to objectives in $[4.60, 4.97]$, with the best 7 in $[4.60, 4.63]$, suggesting a well-defined global minimum. Table **??** reports the parameter bounds, starting values, and estimates.

G.2 Identification

Each parameter is disciplined by a distinct set of empirical moments. Table **??** reports the 15 targeted moments, their empirical values and standard errors, and the primary parameters each moment identifies. The moments span five blocks of the empirical analysis:

Table G.1: Parameter Bounds, Starting Values, and Estimates

[†]At upper bound. The revision-news response α_F is large because the legislature's desired revision response is strong but heavily dampened by legislative anchoring: the net revision pass-through is $\alpha_F/(1+\kappa) = 10.0/13.8 = 0.72$.

Parameter	Symbol	Lower	Upper	Start	Estimate
Pre-planning incorp.	K_{pre}	0	2.5	0.00	≈ 0
Planning incorp.	K_{plan}	0.5	2.5	1.14	1.18
Revision incorp.	K_{rev}	0.1	2.5	0.80	0.53
Balanced-budget pressure	ϕ	0.1	20	1.50	2.00
Base revenue response	α	0.3	3.0	1.65	2.07
Revision-news response	α_F	0.3	10.0	3.65	10.0[†]
Legislative anchoring	κ	1.0	15.0	3.99	12.8
Windfall pass-through	ξ_{exec}	0	0.5	0.01	0.29
Shortfall coefficient	c_{neg}	0.1	2.0	0.50	0.22
Under-execution	$\bar{u}_{\text{exec}}$	-0.2	0.3	0.071	0.071
Persistence	ρ	0	0.99	0.00	≈ 0
Employment loading	κ_E	0	3.0	0.43	0.43
Wage loading	κ_W	0	3.0	0.48	0.46

- **Variance decomposition** (Section **??**, Figure **??**). The planning share of expenditure variance (0.73) and the revision share (0.17) discipline the proposal and enactment anchoring parameters ϕ and κ: higher anchoring concentrates more variation in the initial plan.

- **Revenue forecasting** (Section **??**, Tables **??** and **??**). The planning-window coefficient for the initial revenue forecast (1.32) identifies K_{plan}. The near-zero pre-planning coefficient (-0.10) and the large pre-planning forecast-error coefficient (1.08) jointly pin $K_{\text{pre}} \approx 0$: pre-planning news is informative about future revenue but excluded from the forecast.

- **Revenue-to-spending pass-through** (Section **??**, Table **??**). The planning-stage pass-through (0.60) identifies ϕ; the revision-stage pass-through (0.73) identifies the ratio $\alpha_F/(1+\kappa)$; the realization-stage pass-through (0.27) identifies ξ_{exec} and c_{neg}.

- **Execution** (Section **??**, Table **??**). The median execution gap (-0.071) directly pins $\bar{u}_{\text{exec}}$.

- **Private sector** (Section **??**, Table **??**). The planning-window interaction coefficients for government-exposed employment (0.37) and wages (0.40) directly identify κ_E and κ_W.

Table G.2: Targeted Moments and Identifying Variation

Moments are drawn from the main empirical tables as indicated in the text. The "Primary parameters" column lists the parameters most directly disciplined by each moment; in practice, all parameters are jointly identified through the model's cross-equation restrictions.

Moment	Data	SE	Primary parameters
Variance decomposition			
Planning share of exp. variance	0.730	—	ϕ, κ
Revision share of exp. variance	0.170	—	κ
Revenue forecasting			
Plan. news $\to$ init. revenue forecast	1.320	0.12	K_{plan}
Pre-plan. news $\to$ init. revenue forecast	-0.100	0.34	K_{pre}
Plan. news $\to$ forecast error	0.120	0.19	K_{pre}, K_{plan}
Pre-plan. news $\to$ forecast error	1.080	0.34	K_{pre}
Spending responses			
Plan. news $\to$ init. expenditure	0.630	0.16	ϕ
Plan. news $\to$ actual expenditure	0.840	0.15	ϕ, κ, ξ_{exec}
Revenue $\to$ spending (planning)	0.600	0.07	ϕ
Revenue $\to$ spending (revision)	0.730	0.07	α_F, κ
Revenue $\to$ spending (realization)	0.270	0.08	ξ_{exec}, c_{neg}
Rev. news $\to$ spending revision	0.410	0.16	K_{rev}, α_F, κ
Execution			
Median log(ActE/FinalE)	-0.071	—	$\bar{u}_{\text{exec}}$
Private sector			
Plan. news $\times$ gov-exp. $\to$ employment	0.370	0.09	κ_E
Plan. news $\times$ gov-exp. $\to$ wage bill	0.400	0.16	κ_W

G.3 Model Fit

Targeted moments. Table **??** reports the fit to targeted moments. All 15 moments are matched within 1.3 standard errors. The model closely reproduces the revenue-to-spending pass-through coefficients (planning: 0.58 vs. 0.60, revision: 0.73 vs. 0.73, realization: 0.25 vs. 0.27), the execution gap, and the private-sector loadings. The largest tension is in the planning-news forecast-error coefficient (std. miss = -1.33), which reflects the slight over-incorporation of planning news ($K_{\text{plan}} = 1.18 > 1$) discussed in Section **??**.

Untargeted moments. Table **??** evaluates the model on moments not used in estimation. The model correctly predicts that revision-window news affects actual expenditure (0.75 in data, 0.51 in model—somewhat underpredicted) and that pre-planning news predicts revenue execution surprises (0.87 in data, 0.98 in model). The main tension is in execution-stage distributional moments: the model underpredicts the frequency with which actual spending

Table G.3: Targeted Moment Fit

Std. miss = (model − data)/SE. Moments without standard errors are matched in levels. All standardized misses are within 1.33 SE.

Moment	Data	SE	Model	Std. miss
Planning share of exp. variance	0.730	—	0.649	—
Revision share of exp. variance	0.170	—	0.180	—
Plan. news → init. rev. forecast	1.320	0.12	1.179	−1.18
Pre-plan. news → init. rev. forecast	−0.100	0.34	0.021	0.36
Plan. news → init. expenditure	0.630	0.16	0.769	0.87
Plan. news → actual expenditure	0.840	0.15	0.864	0.16
Plan. news → forecast error	0.120	0.19	−0.132	−1.33
Pre-plan. news → forecast error	1.080	0.34	0.971	−0.32
Median log(ActE/FinalE)	−0.071	—	−0.071	—
Revenue → spending (planning)	0.600	0.07	0.578	−0.32
Revenue → spending (revision)	0.730	0.07	0.730	0.01
Revenue → spending (realization)	0.270	0.08	0.252	−0.22
Rev. news → spending revision	0.410	0.16	0.389	−0.13
Plan. news × gov-exp. → emp.	0.370	0.09	0.370	0.00
Plan. news × gov-exp. → wages	0.400	0.16	0.400	0.00

falls below the enacted budget (0.97 vs. 0.58) and the mean execution gap (−0.10 vs. −0.03). These reflect the simplified execution rule, which captures the median gap but not the full distributional asymmetry.

Visual summary. Figure **??** plots model-implied moments against their empirical counterparts for all targeted and untargeted moments. The targeted moments (blue circles) cluster tightly around the 45-degree line. Most untargeted moments (red diamonds) are also well-matched; the main outliers are the execution-stage distributional moments noted above.

G.4 Sensitivity

Jacobian. Figure **??** displays the numerical Jacobian of targeted moments with respect to parameters ($\partial m_j/\partial \theta_k$). Each column corresponds to a parameter; dark entries indicate moments that are sensitive to that parameter. The block structure confirms the identification logic in Table **??**: forecasting moments load on K_{pre} and K_{plan}; spending moments load on ϕ, α_F, and κ; execution moments load on ξ_{exec}, c_{neg}, and $\bar{u}_{\text{exec}}$; and private-sector moments load on κ_E and κ_W.

Table G.4: Selected Untargeted Moments (Out-of-Sample Validation)

These moments were not used in estimation. The model is evaluated at the estimated parameter vector. Most regression-based moments are matched within 1 SE. The main tensions are in execution-stage distributional moments, which reflect the simplified execution rule.

Moment	Data	Model	Std. miss
Revision-stage moments			
Rev. news → actual expenditure	0.750	0.507	−1.05
Rev. news × gov-exp. → emp.	0.300	0.219	−0.58
Rev. news × gov-exp. → wages	0.320	0.236	−0.77
Pre-planning moments			
Pre-plan. news → execution surprise	0.870	0.977	0.34
Pre-plan. × gov-exp. → emp.	0.170	0.110	−0.54
Pre-plan. × gov-exp. → wages	0.130	0.120	−0.07
Execution and distribution			
Rev. forecast → actual spending	0.640	0.757	0.97
Freq. $E^{\text{act}} < E^F$	0.970	0.575	—
Freq. $R^{\text{act}} < \hat{R}^F$	0.750	0.501	—
Mean log(ActE/FinalE)	−0.103	−0.034	—

Weakly identified parameters. Three parameters warrant discussion. First, the pre-planning incorporation rate K_{pre} is pinned at near zero, but the objective function is flat for values in $[0, 0.05]$—the data strongly reject $K_{\text{pre}} = 1$ but cannot distinguish $K_{\text{pre}} = 0$ from $K_{\text{pre}} = 0.03$. This has negligible consequences for the counterfactual decomposition, which depends on $K_{\text{pre}} \approx 0$ but not on its exact value. Second, persistence ρ is weakly identified and estimated near zero; this has minimal effects on the core results because the model's cross-sectional variation is driven by contemporaneous shocks. Third, α_F is estimated at the upper bound of its range (10.0). However, the economically relevant quantity is the net revision pass-through $\alpha_F/(1+\kappa) = 0.72$, which is interior (well within $[0, 1]$) and precisely identified by the revision-stage revenue-constraint moment (data: 0.73, model: 0.73, std. miss = 0.01). The individual parameters α_F and κ are not separately well-identified—only their ratio matters for the model's predictions—but this does not affect the counterfactual decomposition, which depends on the pass-through rate rather than on the individual values.

Figure G.1: Data vs. Model Moments

Each point represents one empirical moment plotted against its model-implied counterpart. Blue circles: 15 targeted moments; red diamonds: untargeted moments. The dashed line indicates perfect fit.

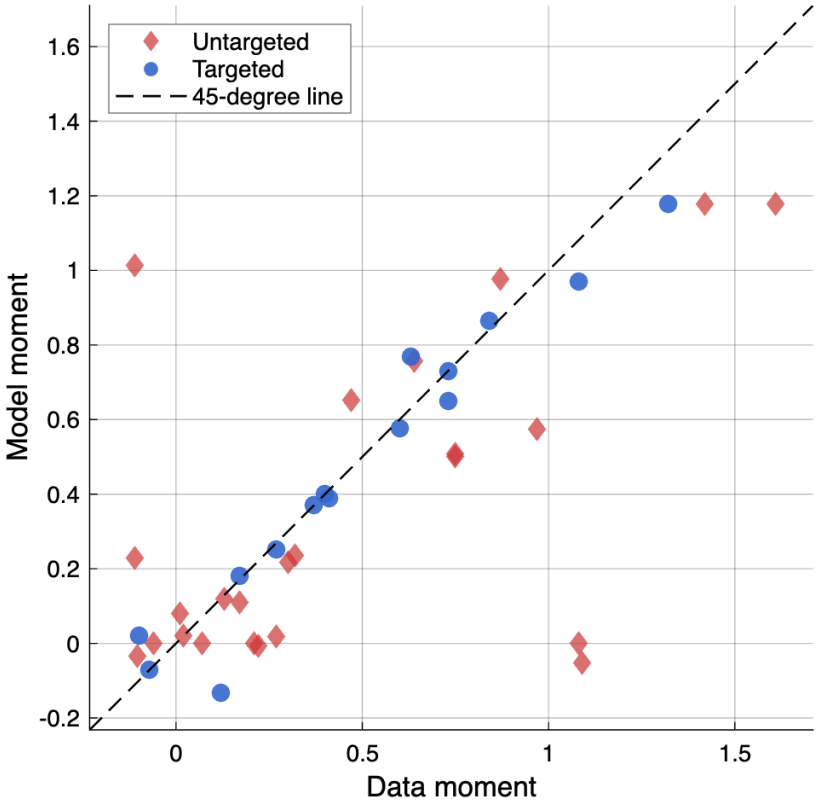


Figure G.2: Jacobian of Targeted Moments with Respect to Parameters

Each cell shows $\partial m_j / \partial \theta_k$ evaluated at the estimated parameter vector. Dark entries indicate high sensitivity. The block-diagonal structure confirms that each parameter block is identified by distinct moments.

