The Decision Window: How Budget Timing Shapes Fiscal Policy

Kerry Siani and Jinyuan Zhang*

July 13, 2026

Abstract

Whether economic news affects fiscal policy depends on when it arrives. Using hand-collected budget documents for 47 U.S. states over 1995–2024, we show that enacted state budgets lock in 79 percent of within-state spending variation, and we identify when those budgets incorporate information. Exploiting staggered proposal dates, whereby the same national shock falls inside one state's budget-writing months and outside another state's, we find that news arriving in the five months before a state's proposal date—the "decision window"—is incorporated into revenue forecasts and spending, while similarly informative news arriving outside the window is not incorporated, and instead predicts forecast errors. Budget timing gates the transmission of economic news to the real economy: decision-window news moves employment in government-dependent industries by four times as much as off-window news, with larger effects for bad news. Monetary policy affects state spending more when it arrives during the state's decision window, suggesting an institutional source of "long and variable lags."

*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. This paper subsumes an earlier circulated draft entitled "The Subnational Fiscal Channel of Monetary Policy." All errors are our own.

Economic news arrives continuously, but government policy is decided at discrete dates, based on budget calendars. This mismatch could make the effect of a shock depend on when it arrives. We document this calendar dependence in U.S. state budgeting, where states form spending budgets based on their revenue expectations. Revenue-relevant news enters the enacted budget primarily during a five-month window before the initial budget proposal date—the "decision window"—while news arriving before or after that window is left out of the current budget. Because incorporated news changes government spending and affects private employment, the budget calendar determines not only when fiscal policy responds, but also how economic shocks propagate to the real economy via state governments.

To establish this, we hand-collect the budget documents of 47 U.S. states from 1995 to 2024, observing for the same state-year the government's enacted spending budget, corresponding revenue forecast, and realized revenues and spending. We show that policy is tightly anchored to the budget plan. The enacted spending budget accounts for 79 percent of within-state variation in realized spending growth, and the budget itself is set against the revenue forecast, a link that balanced-budget rules make binding (Poterba, 1994). The forecast is therefore the gate between news and policy: what it misses, fiscal policy misses. Which news becomes fiscal policy thus includes a question of timing: when does information enter the forecast?

Consider New Jersey and Virginia, which both begin their fiscal year on July 1 but propose budgets at different times: Virginia in December, New Jersey in February. In January 2016, oil prices fell sharply and renewed depreciation of the renminbi pushed the S&P 500 down 5 percent in a month.[1] The shock arrived inside New Jersey's decision window but after Virginia's had closed. New Jersey revised its FY2017 revenue forecast down nearly 9 percent and cut planned spending by about 9 percent; Virginia, whose forecast was set before the shock, planned for revenue growth of 6 percent and accordingly increased its budgeted spending by 6 percent.

We generalize this example across all states for our identification. Each state budgets on a fixed annual calendar anchored by three dates: the governor's budget proposal, the legislature's enactment, and the start of the fiscal year. Proposal dates range across states from November to March, and are sticky: they are set by constitutions or long-standing procedure. Thus, a national shock falls at different points in each state's budget cycle. We combine this cross-state identification with a state-specific monthly revenue shock, which we

[1] See Bloomberg News, "Equity Slump Set to Continue in Asia as Crude Lingers Around $34," January 6, 2016, https://www.bloomberg.com/news/articles/2016-01-06/equity-slump-set-to-continue-in-asia-as-crude-lingers-around-34.

construct by weighting national innovations to five revenue bases by each state's lagged revenue composition. Fiscal-year-by-calendar-month fixed effects absorb whatever is common to a month, so identification comes from where news falls relative to each state's proposal date.

This design delivers three sets of results. First, we show that the decision window exists: state budgets incorporate revenue news only when it arrives inside the window, even though off-window news, both before and after it, also predicts revenues. Second, the decision window matters for the real economy: the same news moves employment persistently in government-exposed industries when it arrives on-window and barely at all off-window, a contrast that lets us quantify the real effects of routine budget adjustment. Third, the decision window effect applies to other shocks: monetary policy surprises affect state spending significantly more when they arrive in-window, providing an institutional source of the "long and variable lags" of monetary policy.

We establish the first result with a forecast-incorporation test. Although the enacted revenue forecast is not finalized until enactment, months after the governor's proposal, only news arriving in the five months before the proposal enters it. Off-window news, both before and after the window, instead predicts forecast errors, and predicts realized revenues in future fiscal years. Each side of this exclusion is surprising. Pre-window news is months old and freely available by the time the forecast is written, yet it systematically stays out. This is the opposite of sticky information, under which older news is caught up at the next update (Mankiw and Reis, 2002). Deliberation news, arriving between the proposal and the start of the fiscal year while the budget is still being revised, is the newest information available, yet it too stays out. This is the opposite of overreaction to recent news or fading memory, under which the newest signals carry the greatest weight (Bordalo et al., 2018; Nagel and Xu, 2022; Afrouzi et al., 2023; Azeredo da Silveira et al., 2024). Incorporation is thus tracking the budget calendar rather than the age or salience of information.

The forecast, in turn, translates into planned and, ultimately, realized state spending. When we instrument the revenue forecast with decision-window news, a 1 percent news-induced increase in the forecast raises planned spending by 0.37 percent and realized spending by 0.56 percent; off-window news, instead, does not move the revenue forecast.

Next, we show that the decision window matters for the real economy. Firms in industries that sell more to state and local governments are more exposed to the state budget through procurement. Comparing more- and less-exposed industries within the same state

and year, decision-window news that raises the state's revenue forecast by 1 percent raises private employment in exposed industries by 0.10 percentage points and wages by 0.14 percentage points, while the same news arriving off-window has little effect. The differential response has little anticipatory effect, appearing significantly only once the fiscal year begins, and persisting for several years. We further test for state-level employment and wage bill responses to on- and off-decision-window news, and find that state-level private employment and wages respond roughly four times more strongly to decision-window news than to off-window news.

The employment response is asymmetric. Negative decision-window news produces large and persistent contractions in exposed industries, while positive news generates smaller and more transient expansions. The asymmetry does not appear to come from the budget: we cannot reject that spending plans pass positive and negative news through symmetrically. We interpret it instead as a firm-level adjustment friction. When bad news arrives inside the window, states convert it into same-year spending cuts that exposed firms have little time to absorb (Barrot and Nanda, 2020). In contrast, bad news arriving off-window enters a later budget cycle, so the eventual cuts come with over year of notice, allowing firms with fiscal foresight to adjust gradually (Leeper et al., 2013).

Beyond these industry-level responses, the on/off design also lets us estimate directly something that standard fiscal-multiplier designs cannot: the real effects of the systematic response of budgets to economic conditions. If we write spending as a rule plus a shock, $G = G(X) + \varepsilon$, the credible designs in the literature identify the effects of ε via military procurement, federal stimulus, or innovations relative to expectations, precisely because these are plausibly unrelated to local conditions. We instead identify the real effects of a key margin of the rule's routine response to news. Dividing the aggregate on- versus off-window employment premium by the corresponding premium in realized spending, we estimate that a 1 percent increase in state spending generated by routine budget adjustment raises private employment by 0.34 percent and private wages by 0.46 percent, roughly $36,000 of induced spending per private job-year, of the same order as estimates from the cross-region multiplier literature (Chodorow-Reich, 2019). Similar to much of the literature, these are relative, cross-state estimates; they do not capture aggregate general-equilibrium effects.[2]

[2]The mapping from cross-state estimates to aggregate multipliers is not one-for-one. Cross-state comparisons difference out the monetary and federal-tax responses that dampen aggregate effects in normal times (Nakamura and Steinsson, 2014), while trade across states and positive spillovers could make them a rough lower bound on the aggregate no-monetary-response, deficit-financed multiplier (Chodorow-Reich, 2019; Auerbach et al., 2020).

Finally, we document that the decision window matters for monetary policy transmission. We use the same month-fixed-effects specification to show that a contractionary monetary policy surprise lowers a state's budgeted and realized spending when it arrives during that state's decision window, relative to the same surprise arriving outside it. Consistent with this state-level result, the aggregate time series displays a seasonal pattern: because many states' decision windows overlap the fourth calendar quarter, Q4 tightening surprises are followed by declines in national state-and-local spending that build over two to three years, with a trough response of up to one-fifth of the corresponding total output response. Meanwhile, surprises in the first three quarters have negligible effects, even though Q4 and non-Q4 surprises are statistically indistinguishable in size and distribution. These results provide an institutional narrative to help explain the "long and variable" lags of monetary policy effects (Friedman, 1961): how long policy takes to reach state spending depends on where in the budget cycle it arrives.

One caveat to our results is that budgets can still change during the fiscal year, as states can pass supplemental appropriations and executives can withhold spending in the case of emergencies. Importantly for our exercise, the original enactment accounts for 79% of within-state variation in realized spending (Section 1.3). In the case of very significant crises, like the 2008 financial crisis, the hardest-hit states may attract a larger share of federal stimulus, which offsets these states' own spending cuts and biases our estimated effects toward zero. We conduct a series of robustness checks to show our baseline results hold when excluding the Global Financial Crisis (GFC).

This paper contributes to four literatures. We contribute first to the empirical literature on the real effects of fiscal policy. Because government spending is inherently political and can be endogenous to fundamentals like productivity or growth expectations, credible estimates of fiscal effects typically rely on rare and unusual spending variation precisely because it is plausibly exogenous (Ramey, 2011; Nakamura and Steinsson, 2014; Chodorow-Reich et al., 2012; Wilson, 2012; Suárez Serrato and Wingender, 2016); see Chodorow-Reich (2019) for a survey.[3] Our estimates resemble that literature in mapping employment responses to government spending, but differ by identifying the systematic, revenue-expectations-driven component of budget policy rather than deviations from it.[4] Ramey (2011) shows that the

[3]Fiscal policy can also move asset prices, as shown in Wiegand (2025).

[4]A related literature shows that fiscal multipliers vary over the business cycle (Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018); our results add that they vary over the budget cycle. Because our spending variation is generated by the budget rule rather than by windfall transfers, it embeds the state's own financing constraint; Clemens and Miran (2012) show that such rule-induced, deficit-financed variation yields smaller multipliers than windfall-financed variation, consistent with Ricardian effects. Appendix G

timing of fiscal news is central to identifying the effects of spending shocks; we show that timing also determines whether news enters the budget at all. Within state public finance, Poterba (1994) show that budget institutions affect the speed at which states adjust their budgets, and Clemens and Miran (2012) show that balanced budget requirements promote procyclical policy. We differ by holding the rules fixed and exploiting variation in budget calendars, thereby establishing that the timing of news, not just the stringency of the institutions, determines whether procyclical adjustment occurs. The distinction matters for stabilization policy: aid or forecasts that arrive after budgets are committed may prevent later cuts, but they are unlikely to change the same-year spending plans that govern procurement and employment. The relevant policy horizon is the decision window, not the fiscal year.

Second, we contribute to work on expectations formation and its real consequences (Manski, 2004; Coibion and Gorodnichenko, 2012, 2015; Gennaioli et al., 2016; Bordalo et al., 2020). Our test builds on Coibion and Gorodnichenko (2012) by using forecast errors regressed on economic shocks to understand expectations formation, with a few important differences: our forecasts are not an average across many professionals, but governments forecasting their own revenues over which they have some control. Standard sticky-information, rational-inattention, and overreaction models (Mankiw and Reis, 2002; Sims, 2003; Reis, 2006; Bordalo et al., 2018; Azeredo da Silveira et al., 2024) can each match a piece of the pattern we document, but none matches its defining feature: a decision window based on each state's constitutionally set proposal date, where informative news is ignored outside the window but incorporated fully within it.[5] We build on these papers to show that expectations formation determines the passthrough of shocks via fiscal institutions to the real economy.

Third, our real-effects evidence speaks to work on government demand and firm outcomes. State budgets reach the private economy largely through procurement, and our industry- and firm-level results locate the response among government-exposed vendors, consistent with evidence that government demand shapes firm employment and dynamics (Barrot and Nanda, 2020; Cox et al., 2024; Ferraz et al., 2015). Our industry-exposure design is similar in spirit to Nekarda and Ramey (2011), which map government-purchasing shares

formalizes the object our design identifies.

[5] A related literature studies the bias and rationality of revenue forecasts themselves: state revenue forecasts tend to be biased downward (Feenberg et al., 1989), while national governments' official forecasts tend toward over-optimism (Frankel, 2011; Frankel and Schreger, 2013); Mikesell and Ross (2014) study how consensus forecasting builds political acceptance of the forecast. Our focus is on what information these forecasts incorporate, and the real consequences.

to industry outcomes; however our identification focus is on cross-state variation in budget timing. The asymmetry we document, where negative news contracts exposed employment by more than positive news expands it, is consistent with credit-constrained government vendors (Barrot and Nanda, 2020). A closely related paper is Liebman and Mahoney (2017), who show that expiring federal appropriations distort the timing of procurement outlays, as agencies rush lower-quality spending into the final weeks of the fiscal year. We differ in three respects. First, we study the decision margin rather than the execution margin; second, we observe the full chain from revenue forecast to enacted commitment to realized spending, rather than outlays alone; third, we document the real effects via private sector employment outcomes. More broadly, our setting makes observable a friction emphasized in organizational economics: institutions receive information continuously but revise commitments only at discrete, calendar-fixed decision points (Dessein and Santos, 2006; Dessein et al., 2016; Grenadier et al., 2016; Malenko, 2024). The same friction should operate in any organization that commits to plans on a fixed cycle, from corporate capital budgets to central-bank projections.

Finally, we contribute to the literature on monetary transmission and its lags. A large body of work identifies monetary shocks from narrative records and high-frequency asset-price responses (Romer and Romer, 2004; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018; Bauer and Swanson, 2023) and traces their propagation to real activity (Christiano et al., 1999, 2005; Miranda-Agrippino and Ricco, 2021), finding that transmission strength varies with economic conditions (Tenreyro and Thwaites, 2016). A closely related paper is Bouscasse and Hong (2026), who show how federal fiscal policy responds to monetary policy. We differ from this body of work in two ways: we focus on state spending as an outcome, and add an institutional margin: the same monetary surprise moves a state's spending significantly more when it arrives inside that state's decision window than outside it. Olivei and Tenreyro (2007) document that policy surprises have a stronger effect in Q1-Q2 due to wage rigidities; our evidence isolates a distinct, institutional channel that leads to larger spending effects for Q4 shocks; moreover our effect is also identified in the cross section of states based on budget institutions. Our mechanism parallels Selgrad and Siani (2025), who show that monetary policy transmits through firms' long-horizon investment plans; we document the government analogue, operating through revenue-expectations incorporation and the planned spending. Together, these results point to planning calendars as one institutional source of the "long and variable lags" of monetary policy (Friedman, 1961): how long a policy shock takes to reach spending depends on where in the budget cycle it lands.

The rest of the paper is organized as follows. Section 1 describes the state budget pro-

cess, data, and discusses the importance of the original enacted budget for actual spending. Section 2 develops the forecast-error test used to identify when the budget process incorporates revenue-relevant information. Section 3 constructs the revenue-news measure and identifies the budget decision window. Section 4 estimates the real effects of the decision window and shows the implications for monetary policy transmission. Section 5 concludes.

1 The State Budget Process and Data

1.1 The State Budget Process

State operating budgets follow a common sequence. Governments first enter a planning stage in which the executive prepares a revenue forecast and budget proposal, then a deliberation stage in which the legislature revises and enacts the budget, and finally an execution stage in which spending is implemented during the fiscal year. Figure 1 illustrates that budgeting process for Massachusetts fiscal year (FY) 2026, which runs ends June 30, 2026. Budget preparation begins second half of 2024, then the governor submits a budget "proposal" in January 2025, and then the legislature deliberates before reconciling a final enacted budget around June 2025. The budget incorporates revenue forecasts from meetings with economists and legislators, as well as spending plans that incorporates the political priorities of the governor and legislative bodies.[6] Across states, procedural details differ, but the core timing fact is shared: major same-year spending is determined several months, and in some cases years, before the fiscal year begins.

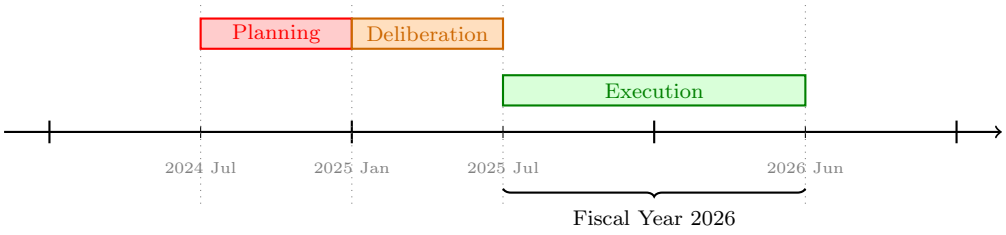


Figure 1: **Timing of the Budget Process for a Massachusetts (FY 2026).** This figure illustrates the timing of planning (pre-proposal), deliberation (proposal to enactment), and execution (fiscal-year implementation) for a Massachusetts.

[6]For revenue forecasts, 28 states develop a consensus forecast, often with participation by economic advisors; in 12 states, both executive and legislative branches draft independent forecasts, while in 10 states, the executive branch develops the forecast on its own (National Association of State Budget Officers, 2021).

This creates an institutional friction: revenue news arrives continuously, but same-year spending commitments are revised only at discrete points in the calendar. The proposal date is the first time that states combine a revenue forecast with a spending plan. Proposal dates differ meaningfully across states, ranging from November to March, even though most fiscal years begin on July 1. The budget timing for each state is typically determined by constitutions, statutes, and long-standing budget procedures, thus is plausibly exogenous to fundamentals. In contrast, budget enactment dates can vary based on economic conditions (Andersen et al., 2012).

1.2 Data and Sample

We analyze a new, hand-collected state-level fiscal panel covering 1995–2024. Our primary fiscal source is the Annual Comprehensive Financial Report (ACFR), which we gather from state websites, direct correspondence with state officials, and formal records requests.[7] We digitize these reports using Python and manually verify the extracted data. We rely on ACFR for budget data rather than the commonly used National Association of State Budget Officers (NASBO) Fiscal Survey of States for two reasons. First, while NASBO is constructed from a semiannual survey collected in simultaneously for all states, the ACFR reports the budget dated according to each state's own budget cycle, which is central to our analysis. Second, the two sources measure different fiscal aggregates, and the ACFR is more standard across states. Appendix A.3 describes the differences between the two data sources in detail.

We measure state budget calendars from archival publications of the NASBO, which report proposal dates, enactment dates, fiscal-year start dates, and budget-cycle rules. We construct a piecewise-constant calendar from the NASBO waves, assigning each fiscal year the most recent prior NASBO snapshot.[8] Proposal dates are highly persistent across waves, so this procedure mainly records stable institutional calendars rather than frequent timing changes. Figure 2 shows the cross-state variation in proposal timing using the 2021 NASBO snapshot. Proposal months vary substantially across states and do not follow a simple geographic pattern.[9]

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

[8] The NASBO budget-process snapshots used in the paper are 1987, 1989, 1992, 1997, 1999, 2002, 2008, 2015, and 2021. For each fiscal year, we use the most recent snapshot observed at or before that year and carry it forward until a newer snapshot is available. This avoids using future calendar information.

[9] Appendix Table A.1 reports each state's full budget-cycle timing in the 2021 snapshot.

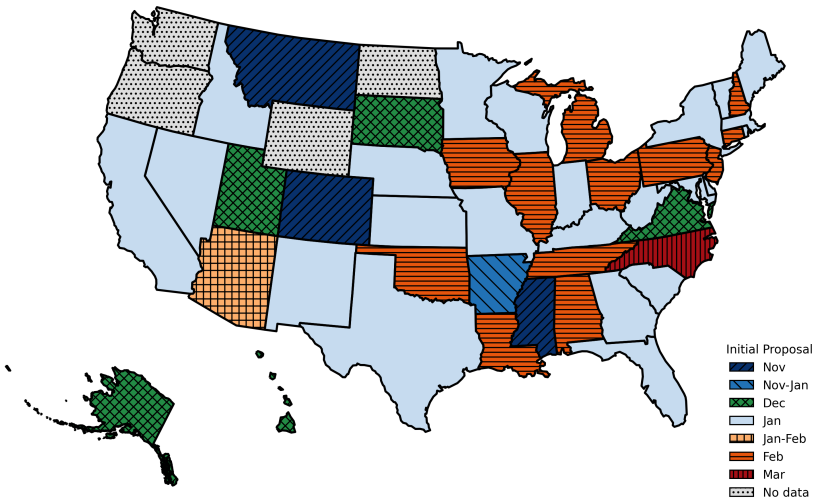


Figure 2: **Cross-State Variation in Original Budget Proposal Month.** The map shades each state by the typical month in which its governor submits the original budget proposal to the legislature, taken from Table A.1 in Appendix A. Colors range from November (dark blue) through January (light blue), February (orange), and March (dark red); December is shown in green, and states with two-month proposal windows (Nov–Jan, Jan–Feb) use intermediate colors. Hatch patterns are included so categories remain distinguishable in black-and-white printing. States shown in light grey (North Dakota, Oregon, Washington, Wyoming) are not in our analysis sample.

Under GASB Statement No. 34, each state's ACFR reports a budgetary comparison statement with three fiscal-year objects: the "original budget," the "final budget," and actual outcomes (see Appendix Figure A.1). The original budget is the first complete appropriated budget adopted for the fiscal year, including legally authorized changes made before the fiscal year begins.[10] Thus, the ACFR original budget is enacted prior to the fiscal year start, and is distinct from the governor's initial proposal.

This distinction is important for interpretation. The proposal date is the first point in the budget year that states pair a statewide revenue forecast with a spending plan, so we use it to organize the budget calendar. However, we only observe the enacted original budget, which is finalized later. It is therefore an empirical question whether the enacted original budget reflects information that arrives during legislative deliberation, or instead remains anchored to the forecast assembled before the proposal.

In our analysis, we do not use the "final budget", because it incorporates within-year

[10]GASB 34, paragraph 131, defines the original budget as "the first complete appropriated budget," adjusted for reserves, transfers, and other legally authorized changes made before the fiscal year begins.

legally authorized changes whose timing is unclear. We focus instead on the original-budget revenue forecast and expenditure plans, which we notate with $\text{Orig}^R_{s,t}$ and $\text{Orig}^E_{s,t}$.

Balanced-budget requirements are widespread at the budget stage. 45 states require the governor to submit a balanced budget, 44 require the legislature to pass a balanced budget, and 41 require the budget signed by the governor to be balanced. Ex-post requirements are less uniform: 35 states require the executed budget at year end to be balanced (National Association of State Budget Officers, 2021). Consistent with this requirement, state expenditures broadly track revenues over the budget cycle (see Appendix B.3 for details).[11]

1.3 Original Budget Explains Most Spending Variation

We first show that the original budget is the key commitment margin for state fiscal policy. To quantify its importance, we decompose annual changes in realized outcomes into three components: growth in the original budget, the current-year execution deviation from that budget, and the lagged execution deviation. For fiscal year t, we benchmark original-budget growth to the prior year's original budget, $\text{Orig}_{s,t-1}$, and write:[12]

$$\Delta \log(\text{Act}_{s,t}) = \underbrace{\log\left(\frac{\text{Orig}_{s,t}}{\text{Orig}_{s,t-1}}\right)}_{\text{Budget}} + \underbrace{\log\left(\frac{\text{Act}_{s,t}}{\text{Orig}_{s,t}}\right)}_{\text{Execution}} - \underbrace{\log\left(\frac{\text{Act}_{s,t-1}}{\text{Orig}_{s,t-1}}\right)}_{\text{Lag execution}}. \tag{1}$$

Summary statistics for these components are reported in Table 1; Appendix Figure B.1 plots their distributions.

Figure 3 reports the variance decomposition implied by Equation (1).[13] The planning component alone accounts for 51% of the within-state variation in revenue growth and 79% of the within-state variation in expenditure growth. The contrast between the two is notable:

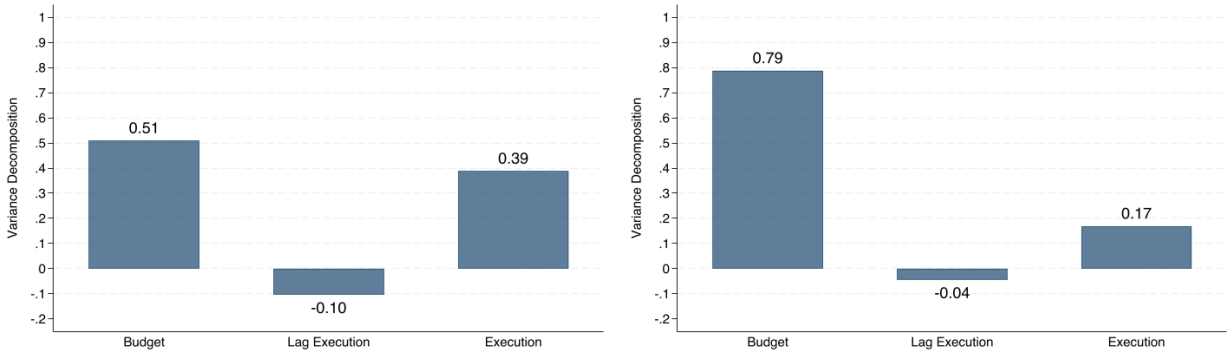
[11] States can manage prospective deficits through spending cuts, drawing down on rainy day funds, deferred payments, inter-fund transfers, short-term borrowing, and in rarer cases, revenue increases or accounting changes. See Poterba (1996) for a detailed discussion.

[12] This benchmark treats the prior year's original budget as the reference point from which the FY-t budget is drafted. Appendix Table B.1 supports this choice: when the log original budget is regressed on the two most recent available benchmarks, the prior original budget dominates, while two-year-lagged actual outcomes add little conditional on it.

[13] We take the covariance of both sides of Equation (1) with respect to log changes in actual outcomes, then divide by the variance of log changes in actual outcomes. Then we compute each term using an OLS regression with state fixed effects. See Appendix B.2 for more details on how the variance decomposition is estimated. Appendix Figure B.2 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 (1), for total real revenues (Panel a) and total real expenditures (Panel b). Components are: Budget $\log(\text{Orig}_{s,t}/\text{Orig}_{s,t-1})$, Execution $\log(\text{Act}_{s,t}/\text{Orig}_{s,t})$, and Lag Execution $\log(\text{Act}_{s,t-1}/\text{Orig}_{s,t-1})$.



(a) Total Revenues

(b) Total Expenditures

while revenues often deviate from the forecast, spending does little. Thus, the original budget is important: once those commitments are set, execution deviations in spending contribute relatively little to the annual variation in realized spending with state.

The dominance of the original budget also shows up in line-item adjustments. Appendix B.3 shows that, at the budget stage, positive and negative revenue-forecast updates pass through to spending plans in broadly symmetric ways across categories. The asymmetry appears later, during execution: when realized revenue falls short of the budget, states cut mainly education spending; when revenue exceeds the budget, windfalls flow primarily to infrastructure and general government. Appendix B.3-B.4 reports additional evidence on revenue–expenditure comovement, execution-stage asymmetry, and rainy-day-fund adjustment.

In summary, we find that the original budget is important for actual fiscal policy, and revenue forecasts affect spending plans across a wide range of line items. We next examine how these revenue forecasts are formed and translated into spending plans.

2 Framework

We start by studying how the original budget is formed: which information enters the budget, and when does that information arrive? To do so, we develop a forecast-incorporation test, to identify the period of the budget cycle whose revenue news is embedded in the original budget. We call this period the "budget decision window". Because state spending plans are built based on revenue forecasts, identifying when forecasts update also identifies when revenue news can become current fiscal policy.

2.1 The Budget Forecast as the Fiscal State Variable

The state chooses original-budget spending $G_{s,t}^{\text{budget}}$ by trading off expected fiscal balance against a baseline $\bar{G}_{s,t}$, which summarizes service needs, statutory obligations, and political commitments. Most states face balanced budget constraints at initial and final budgeting, so expected revenues based on the information set at the time of the budget $R_{s,t}^{\text{budget}} \equiv \mathbb{E}\Big[R_{s,t}^{\text{actual}} \mid \mathcal{A}_{s,t}^{\text{budget}}\Big]$ enters into the state's problem as well.

$$G_{s,t}^{\text{budget}} = \arg\min_{G} \left\{ \tfrac{\psi}{2}\big(G - R_{s,t}^{\text{budget}}\big)^2 + \tfrac{\phi}{2}\big(G - \bar{G}_{s,t}\big)^2 \right\} \qquad \text{s.t.} \qquad G \leq R_{s,t}^{\text{budget}}. \tag{2}$$

The first term reflects the cost of fiscal imbalance, where spending less than expected revenues can be politically costly, and the second term is the cost of moving spending away from baseline commitments. The constraint is the balanced-budget requirement at enactment: the budget cannot commit spending beyond the revenue the state forecasts on its own books (Poterba, 1994).

Solving the problem gives

$$G_{s,t}^{\text{budget}} = \min\Big\{ \lambda\, R_{s,t}^{\text{budget}} + (1-\lambda)\,\bar{G}_{s,t},\ R_{s,t}^{\text{budget}} \Big\}, \qquad \lambda \equiv \frac{\psi}{\psi + \phi} \in (0, 1]. \tag{3}$$

When the balanced-budget constraint is slack, spending is a convex combination of the revenue forecast and baseline commitments; when it binds, spending is capped by the forecast and forecast updates pass through one-for-one. Either way, news that moves expected revenues $R_{s,t}^{\text{budget}}$ can affect the spending budget.

2.2 Testing Forecast Incorporation

We establish a framework to test whether institutions incorporate new information into their expectations continuously or discretely based on their institutional calendars. Suppose we observe news $z_{s,t,m}$ arriving in month m that could positively affect state s's revenue in fiscal year t. We first test if the news affects actual revenues $R_{s,t}^{\text{actual}}$, and then test if the news is incorporated into the revenue forecast $R_{s,t}^{\text{budget}}$. Equivalently, we can test if the news predicts forecast errors, where $R_{s,t}^{\text{fe}} \equiv R_{s,t}^{\text{actual}} - R_{s,t}^{\text{budget}}$. All tests take the form:

$$R_{s,t}^{j} = \sum_{m} \beta_n^j z_{s,t,m} + u_{s,t}^{j}. \tag{4}$$

where $j \in \{\text{actual}, \text{budget}, \text{fe}\}$. Notice that for any month, any news that is informative but not incorporated into the forecast will predict the forecast error: $\beta_m^{\text{fe}} = \beta_m^{\text{actual}} - \beta_m^{\text{budget}}$.

The full-information benchmark. The benchmark is full incorporation of all information available before the original budget is finalized. This benchmark requires the revenue forecast $R_{s,t}^{\text{budget}}$ reported in the ACFR original enacted budget to incorporate any informative news before enactment. Forecast errors should therefore be unpredictable with respect to such information:

$$\beta_m^{\text{budget}} = \beta_m^{\text{actual}}, \qquad \beta_m^{\text{fe}} = 0, \qquad \text{for all } m \text{ before finalization.}$$

Budget Decision Windows. The decision-window hypothesis is that the relevant information set is defined by when the budget forecast is actively built. In our setting, the governor's proposal is the first date at which a statewide revenue forecast is paired with a spending plan. Even though the original enacted budget may be finalized later, the information-processing window may close earlier.

We call the active forecast-building period the budget decision window, $\mathcal{W}_{s,t}^{\text{Decision}}$. The hypothesis is that

$$\mathcal{A}_{s,t}^{\text{budget}} = \{z_{s,t,m} : m \in \mathcal{W}_{s,t}^{\text{Decision}}\}.$$

This yields a clear test. Pre-decision news is already available when the forecast is being built, so we can directly test if older information is relevant and incorporated. After-decision news tests whether the enacted budget continues to update after the proposal. The decision-

window hypothesis predicts that only news arriving in the window enters the current original budget. For any month with informative news ($\beta_m^{\text{actual}} \neq 0$), the hypothesis implies:

$$\begin{aligned} &\beta_m^{\text{budget}} = \beta_m^{\text{actual}}, \qquad \beta_m^{\text{fe}} = 0, \quad m \in \mathcal{W}_{s,t}^{\text{Decision}}, \\ &\beta_m^{\text{budget}} = 0, \qquad \beta_m^{\text{fe}} > 0, \qquad\quad m \text{ outside } \mathcal{W}_{s,t}^{\text{Decision}}. \end{aligned} \tag{5}$$

Differencing across windows collapses both predictions into a pair of sufficient statistics, the decision-window premium (where "on" means $m \in \mathcal{W}_{s,t}^{\text{Decision}}$ and "off" means $m \notin \mathcal{W}_{s,t}^{\text{Decision}}$):

$$\theta^{\text{budget}} \equiv \beta_{\text{on}}^{\text{budget}} - \beta_{\text{off}}^{\text{budget}} > 0 \qquad \text{and} \qquad \theta^{\text{fe}} \equiv \beta_{\text{on}}^{\text{fe}} - \beta_{\text{off}}^{\text{fe}} < 0. \tag{6}$$

Alternative models. The "decision-window" hypothesis is distinct from alternative stories of information incorporation. In a sticky-information model (Mankiw and Reis, 2002), updating is infrequent but complete: when the forecast is revised, it incorporates all information available on that date, so β^{fe} for older news (pre-window) should be closer to 0. Overreaction and diagnostic expectations (Bordalo et al., 2018; Afrouzi et al., 2023) instead overweight the most recent signals, implying β^{fe} approaches zero for news arriving closest to the forecast date. Rational inattention (Sims, 2003; Maćkowiak and Wiederholt, 2009) attenuates responses to signals that are costly to process or low in salience. Appendix Table C.1 summarizes these contrasts.

2.3 From Incorporation to Spending

We care about state revenue forecasts because they transmit directly to fiscal policy via the optimal budget rule (3). Specifically, original-budget spending responds to month-m news only through the incorporated component,

$$\frac{\partial G_{s,t}^{\text{budget}}}{\partial z_{s,t,m}} = \underbrace{\frac{\partial G_{s,t}^{\text{budget}}}{\partial R_{s,t}^{\text{budget}}}}_{\text{passthrough}} \beta_m^{\text{budget}}, \tag{7}$$

where the passthrough equals λ when the balanced-budget constraint is slack, and equals one when it binds. State fiscal spending therefore should inherit the incorporation pattern of revenue forecasts. If comparable news has the same direct effect on the economy whenever it arrives, the on- versus off-window premium in any state-level outcome identifies the

component that is mediated by state spending.

3 Identifying the Budget Decision Window

The framework delivers two testable predictions: forecast incorporation should be concentrated in the budget decision window, and incorporated forecast updates should pass through to spending. To test these predictions, we exploit cross-state variation in proposal dates to compare similar news at different proposal-relative times.

3.1 Research Design

States write their budgets at different times, so a common national shock lands at different points in different states' budget calendars. Consider New Jersey and Virginia: both begin their fiscal year in July, but Virginia proposes its budget in December, while New Jersey proposes in February. In January 2016, oil prices fell sharply and renewed depreciation of the renminbi pushed the S&P 500 down 5 percent in a month.[14] New Jersey was still drafting its budget, and revised its FY2017 revenue forecast down nearly 9 percent and cut planned spending by about 9 percent; Virginia, whose forecast was set before the shock, planned for revenue growth of 6 percent and raised planned spending by 6 percent.

Our core identification builds on this example and exploits variation in the timing of state-specific revenue news relative to each state's budget calendar. Because proposal dates are staggered, news arriving in a given month falls inside the decision window for some states and outside it for others, and we compare how it is incorporated into forecasts. The key assumption is that, conditional on state and time fixed effects and controls, this timing is as good as randomly assigned: whether a given news realization falls inside or outside a state's decision window is uncorrelated with the state's unobserved fiscal fundamentals. The main threat is endogenous proposal timing, if states moved budget dates in response to shocks; we address this concern in Section 3.4.3 and find that proposal dates, typically set by statute, rarely move.

[14] See Bloomberg News, "Equity Slump Set to Continue in Asia as Crude Lingers Around $34," January 6, 2016, https://www.bloomberg.com/news/articles/2016-01-06/equity-slump-set-to-continue-in-asia-as-crude-lingers-around-34.

3.2 Measuring State-Specific Revenue News

Implementing the design requires a monthly, state-specific measure of revenue news. However, there is no readily availably measure of monthly state-specific shocks to economic conditions for our full sample. Thus, we adapt a shift-share strategy built from national series and state-level exposures for constructing state-month level shocks.

To build the shift-share shock, we begin with five key own-source revenue bases k: personal income, general sales, charges and miscellaneous revenue, corporate income, and severance taxes. Together, they account for 62% of general revenue.[15] For each base, we identify a national shock $z_{k,m}$, compute a predetermined state exposure weight $\omega_{s,k,t}$, and estimate a national scaling $\hat{\eta}_k$:

$$\text{News}_{s,t,m} = \sum_{k=1}^{K} \hat{\eta}_k \, \omega_{s,k,t} \, z_{k,m}. \tag{8}$$

The shock $z_{k,m}$ is the standardized innovation in month m of base k's national source series, taken as the residual from an AR(4) in monthly log changes so that it captures news rather than the predictable component. Because the source series are national, all cross-state variation in $\text{News}_{s,t,m}$ comes from the exposure weights. For example, a severance-reliant state loads on oil-price innovations, an income-tax-reliant state on national income. This composition sets the magnitude of a state's exposure to each shock, while the decision-window effect is identified from the timing of the shock in the budget calendar. The weight is state s's revenue share in base k, averaged over five years ending two years before fiscal year t,

$$\omega_{s,k,t} = \frac{1}{5} \sum_{\tau=t-6}^{t-2} \frac{\text{Revenue}_{s,k,\tau}}{\text{General Revenue}_{s,\tau}}, \tag{9}$$

so it is predetermined relative to the FY-t budget. The scaling $\hat{\eta}_k$ is the slope from a national regression of aggregate revenue growth in base k on the corresponding fiscal-year shock,

$$\Delta \log \text{Rev}^{US}_{k,\tau} = \eta_k \, \text{Shock}^{US}_{k,\tau} + u_{k,\tau}; \tag{10}$$

it puts the bases in positive common predicted-revenue units and has no cross-state variation.

[15] The baseline national source series are national personal income for personal income taxes, personal consumption expenditures for sales taxes, PCE services for charges and miscellaneous revenues, industrial production for corporate income taxes, and crude petroleum prices for severance taxes; Appendix Table D.1 lists the exact series mnemonics. Intergovernmental transfers (about 32% of general revenue) are excluded because they are stable and driven largely by federal policy rather than state tax-base exposure.

Construction details are in Appendix D.

Two properties make the measure suitable for our design. First, it is predetermined: neither decision-window nor pre-decision news predicts changes in state top marginal income-tax rates (Appendix Table E.2), so revenue composition does not respond to the shocks. Second, it is relevant: fiscal-year news predicts realized revenue growth (Appendix Table E.1), though this relationship weakens once fiscal-year fixed effects absorb the large common national component.

Although the news measure is a shift-share object, the exogeneity of its exposure shares is not the key identifying condition in our design. We compare news inside versus outside a state's decision window; a state's exposure is common to both sides of that comparison. The identifying assumption is therefore not that the shift-share measure is an exogenous fiscal-policy shock, but that any non-budget relationship between the news measure and the outcome does not change discretely at the decision-window boundary (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022).

3.3 Locating the Decision Window

We first use the original-budget revenue forecast to locate the timing of forecast incorporation. Because the news measure is monthly, we build a state–fiscal-year–calendar-month panel: for each state–fiscal-year budget round we stack the monthly news observations over a 30-month event window, from July of calendar year $t-3$ to December of $t-1$, repeating the annual outcome across those months. We anchor event time at the budget proposal date, letting $\ell_{s,t,m}$ denote the number of months between calendar month m and state s's FY-t proposal, so $\ell < 0$ precedes the proposal and $\ell \geq 0$ follows it. We then regress original-budget revenue forecast growth, $y_{s,t} = \log(\text{Orig}^R_{s,t}/\text{Orig}^R_{s,t-1})$, on news arriving at each event-time distance from the proposal:

$$y_{s,t} = \sum_h \rho^y_h \left(\text{News}_{s,t,m} \times \mathbf{1}\{\ell_{s,t,m} = h\}\right) + \mathbf{X}'_{s,t}\Pi^y + \alpha_s + \delta_{t,m} + \varepsilon^y_{s,t,m}, \tag{11}$$

The key coefficient ρ^y_h is the loading on news arriving h months from the proposal. State fixed effects α_s absorb persistent cross-state differences; fiscal-year-by-calendar-month fixed effects $\delta_{t,m}$ absorb the common national component of each innovation within a budget round, so identification comes from proposal-relative timing across states, scaled by exposure. The controls $\mathbf{X}_{s,t}$ are four lagged quarterly changes of each national macro series (CFNAI, payroll

employment, the S&P 500, the five-year Treasury yield, and CPI), the most recent ending four quarters before the budget; the Borusyak et al. (2022) aggregate-share control; and two-year-lagged state rainy-day-fund balances and GDP growth. Standard errors are two-way clustered by state and calendar month-year.

A positive ρ_h^y means that news arriving h months from the proposal is incorporated into the original-budget revenue forecast. Figure 4 plots the month-specific coefficients from (11). Expected revenue growth is affected by news primarily in the five months before the proposal date; coefficients outside this interval are small and statistically indistinguishable from zero. Visually, the timing profile identifies $\ell \in [-5, -1]$ as the candidate budget decision window: the months in which monthly revenue news is reflected in the original-budget revenue forecast.

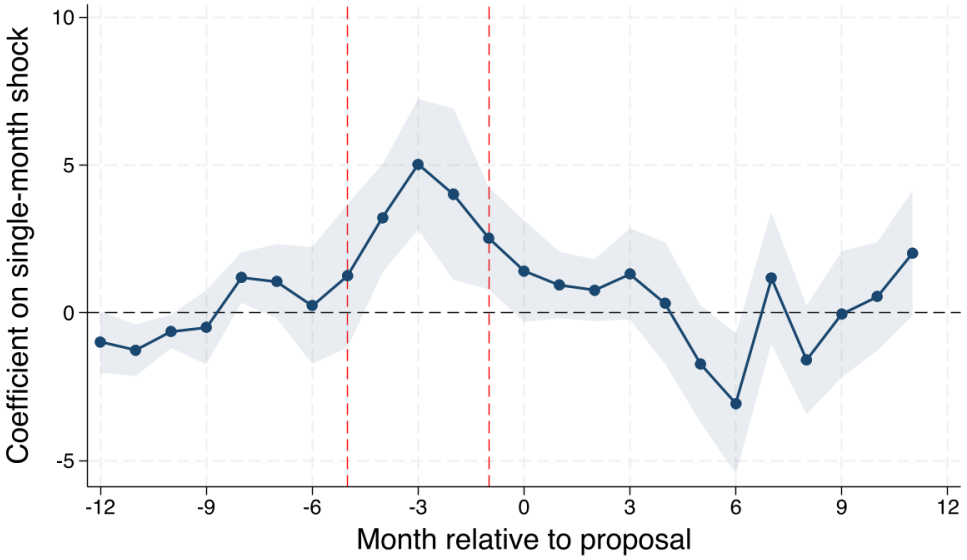


Figure 4: **Monthly Profile of Revenue Forecast Updating Around the Proposal Date.** The figure plots coefficients on the state-specific shift-share news measure interacted with single-month proposal-relative indicators. The dependent variable is one-year original-budget revenue forecast growth. The specification includes state fixed effects, fiscal-year-by-calendar-month fixed effects, predetermined controls, and the aggregate-share control. Shaded areas denote 95 percent confidence intervals. The dashed lines mark the budget decision window used in the window-level specifications.

To check that this pattern is not an artifact of anchoring event time at the proposal date, we repeat the exercise anchoring instead at the enactment date (Appendix Figure E.1). The response shifts correspondingly earlier relative to enactment, consistent with the proposal-date window identified here.

We accordingly partition each budget round's pre-fiscal-year months into three regions: "pre-decision" ($\ell \in [-11, -6]$), the "decision window" ($\ell \in [-5, -1]$), and "deliberation" (the months after the proposal, before the fiscal year begins). Table 1, Panel B, reports summary statistics for the monthly news measure within each window.[16]

3.4 Forecast Incorporation and the Decision Window

While the monthly profile locates where forecast updating occurs, the framework in Section 2 requires a sharper decomposition: revenue-relevant off-window news should not enter the original-budget forecast and should instead reappear in forecast errors. We restrict the sample to the three proposal-relative windows defined in Section 3.3 and allow the news loading to differ across windows:

$$y_{s,t} = \sum_{W \in \{\text{Pre,Decision,Delib}\}} \beta_W^y \left(\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^W\}\right) + \mathbf{X}_{s,t}'\Pi^y + \alpha_s + \delta_{t,m} + \varepsilon_{s,t,m}^y, \quad (12)$$

where $\mathcal{W}_{s,t}^{\text{Pre}}$, $\mathcal{W}_{s,t}^{\text{Decision}}$, and $\mathcal{W}_{s,t}^{\text{Delib}}$ denote the pre-decision, decision, and deliberation windows. We estimate (12) separately for realized revenue, the original-budget revenue forecast, and the forecast error, so the coefficients map directly into $\beta_{\mathcal{W}}^{\text{actual}}$, $\beta_{\mathcal{W}}^{\text{budget}}$, and $\beta_{\mathcal{W}}^{\text{fe}}$ in Equation (5) of the framework.

To test the decision-window premium in Equation (6), we also estimate an in/out specification that pools the off-window months:

$$y_{s,t} = \gamma^y \text{News}_{s,t,m} + \theta^y \left(\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}\right) + \mathbf{X}_{s,t}'\Pi^y + \alpha_s + \delta_{t,m} + \varepsilon_{s,t,m}^y, \quad (13)$$

where $\mathcal{W}_{s,t}^{\text{Decision}}$ is the decision window. Because we restrict the sample to end before the fiscal year begins γ^y is the pooled off-window response and θ^y is the decision-window-minus-off-window premium. When the outcome variable is revenue forecast, the test shows whether news is incorporated more when it arrives during the decision window. When the outcome variable is the revenue forecast error, then the test shows whether revenue-relevant news is differentially incorporated in- versus out-of-window.

[16] Decision-window news has a smaller standard deviation than the other windows in Panel B. The gap reflects COVID: the pandemic-era innovations of 2020–2021, concentrated in the spring of 2020, fall in calendar months covered by the pre-decision and deliberation windows, while decision windows sit mostly in August–January. Appendix Table D.2 reproduces the panel excluding calendar years 2020–2021; the standard deviations are then comparable across windows (0.23–0.26, ×100 scale). Decision windows therefore faced smaller shocks on average over the full sample, which works against finding a premium.

The formal counterpart of the timing assumption in Section 3.1 is the moment condition

$$\mathbb{E}\left[\left(\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}\right) \varepsilon_{s,t,m}^{y} \mid \mathbf{X}_{s,t}, \alpha_s, \delta_{t,m}\right] = 0.$$

3.4.1 Main results

Table 2 reports the results. Panel A estimates the separate-window loadings for the original-budget revenue forecast, the forecast error, and realized revenue. All columns share the same sample, so the realized-revenue loadings in Column (3) equal the sum of the forecast and forecast-error loadings, coefficient by coefficient. Column (1) shows that decision-window news moves the original-budget revenue forecast strongly, with $\beta_{\text{Decision}}^{\text{budget}} = 2.71$, while pre-decision and deliberation news do not: their coefficients are 0.02 and 0.55, respectively, and statistically insignificant.

Column (2) shows the mirror image: pre-decision and deliberation news predict forecast errors, while decision-window news, already incorporated, does not. This is a substantive anchoring result. Deliberation news arrives after the proposal but before enactment, so it could in principle affect the original enacted budget; its forecast-error loading shows that this information is not fully incorporated. Pre-decision news provides the complementary test: it is already available when the decision-window forecast is being built, yet it also remains outside the original-budget forecast and reappears in forecast errors. Together, the two off-window results show that incorporation of news is governed not by availability but by timing relative to the window.

Finally, Column (3) closes the identity with realized revenue. Both decision-window and deliberation news predict realized revenue, but the decision-window coefficient is larger. The larger decision-window coefficient, 2.54, reflects a potential multiplier effect of spending plans: news written into the budget can also move the tax base through the resulting fiscal choices.[17]

Panel B collapses the contrast into the in- versus out-of-window premium. Two key results are consistent with the decision window prediction: the forecast premium is large and significantly positive, and the forecast-error premium is negative and rejected against

[17] Appendix C formalizes this point. If realized revenue contains a same-year fiscal-feedback component, realized-revenue loadings combine fundamental revenue news with any spending response induced by the budget. The forecast-error decomposition remains the clean incorporation test: under consistent scoring of the enacted plan, the fiscal-feedback term cancels from $R^{\text{actual}} - R^{\text{budget}}$.

zero at the 10 percent level one-sided ($p = 0.071$). Taken together, the evidence validates that the five months before the proposal is the budget decision window: the period in which revenue news is written into the original budget.

In summary, incorporation of news is concentrated in the decision window leading up to the proposal date, and does not continue systematically through legislative deliberation, even though deliberation occurs before the original enacted budget is finalized. The decision window is therefore an empirical feature of how the enacted budget is formed, rather than a mechanical consequence of when it is written.

3.4.2 Alternative stories

We can distinguish the decision-window mechanism from theories of information frictions, because incorporation is not monotone in the age of news. Under sticky-information updating, pre-decision news should be caught up when the forecast is updated; instead, it does not enter the forecast and predicts forecast errors. Under recency or overreaction, deliberation news should receive the greatest weight because it is the newest pre-fiscal-year information; instead, it is left out and reappears in forecast errors. Nor is there evidence of gradual updating within the decision window: splitting the window into early months $\ell \in [-5, -3]$ and late months $\ell \in [-2, -1]$ yields forecast loadings of 2.81 and 2.63, which are statistically indistinguishable (Appendix Table E.3). Finally, an endogenous-attention model predicts that unusually salient pre-decision shocks should be more likely to enter the forecast. They are not: interacting pre-decision news with its absolute magnitude yields a small and insignificant coefficient (Appendix Table E.4). Incorporation therefore tracks calendar position rather than news salience.

3.4.3 Robustness

Addressing the budget calendar. One concern is that the decision-window premium reflects state fundamentals correlated with budget calendars rather than the budget calendar itself. Proposal dates help address this concern because they are set by constitutions, statutes, and legislative calendars and change infrequently. [18] Since the empirical strategy applies these predetermined calendars to national innovations realized later, the identifying

[18] Andersen et al. (2012) show that budget enactment can be delayed by economic conditions; our design uses proposal dates, not enactment dates.

assumption is that states cannot move long-standing proposal months in response to future oil-price, income, or sales shocks.

The data support this assumption in three ways. First, proposal months are highly persistent: forty of fifty-one states never change their reported proposal month across the NASBO waves covering our sample period, and state fixed effects absorb 93 percent of panel variation in proposal months (Appendix Table E.7). Second, the result is robust to using a fixed calendar vintage: holding every state's calendar fixed at its most recent 2021 NASBO wave changes only seven states' proposal months and leaves the decision-window loading essentially unchanged (Appendix Table E.8). Third, a pseudo-window placebo that relocates each state's decision window to random positions away from the proposal does not generate the result. Across 500 relocations, the estimated premium is centered near zero ($p < 0.002$; Appendix Figure E.2), supporting the conclusion that the premium is generated by the actual budget calendar.

Biennial budgets and legislative calendars. Some states have biennial budget cycles, as discussed in Section 1. For these states, the second year of the biennium has no official budget proposal of its own, so we test whether the baseline result depends on how those state-years are treated. Appendix Table E.9 re-estimates the baseline three ways: interacting each window term with a second-year indicator, dropping second-year cohorts, and reassigning their months to the post-proposal period. The decision-window coefficient remains between 2.7 and 2.8 across the three treatments. The interaction is also informative about the mechanism: news assigned to a second-year "decision window" is not incorporated into the forecast (implied loading -0.57) and instead predicts the forecast error (3.11, significant at 5 percent). Appendix Table E.10 further shows that dropping strictly biennial state-years, or keeping only states with annual legislative sessions or annual budget cycles, leaves the decision-window coefficient between 2.6 and 2.9, significant at the 1 percent level in all three samples.

Other concerns. Appendices E.3–E.5 report the remaining robustness analysis. The decision-window pattern is stable across controls and fixed effects, nearby definitions of the window, an FY2002+ sample that addresses changes in ACFR budgetary reporting around GASB 34, leave-one-out exercises, and a transparent severance-only specification based on oil-price innovations. It also survives shift-share inference checks based on month-year wild bootstrapping and exposure-vector permutation, and is not explained by exposure-specific

calendar seasonality. These checks support the interpretation that the premium is generated by the actual budget calendar rather than by a particular specification or influential state-year.

3.5 Updated Revenue Forecasts Affect Spending Plans

Next, we test whether forecast updates induced by decision-window news translate into spending decisions. Equation (7) implies a pass-through from the original-budget revenue forecast to spending. We estimate this pass-through by instrumenting the revenue forecast update with news from each proposal-relative window. Table 3 reports the results.

The first stage exists only in the decision window. The decision-window instrument has a Kleibergen–Paap cluster-robust first-stage Wald F-statistic of 22.0, while the pre-decision and deliberation instruments have essentially no first stage ($F = 0.0$ and 1.3). Using the decision-window instrument, a 1 percent news-induced increase in the original-budget revenue forecast raises the original spending plan by 0.37 percent and realized spending by 0.56 percent. Because the first stage is moderate, Table 3 also reports weak-instrument-robust Anderson–Rubin confidence sets. These sets are bounded and exclude zero for the decision-window instrument; for the two off-window placebo instruments they are unbounded, reflecting the absence of forecast first-stage variation.

Through the lens of the framework, the planned-spending coefficient is the pass-through from revenue forecasts to the original spending plan. The estimate below one indicates that forecast updates are only partially translated into planned spending, consistent with reserves, spending rigidities, or other budget constraints absorbing part of the revenue news. The realized spending coefficient is larger, 0.56, showing that forecast updates also predict actual spending over the fiscal year. Incorporation is therefore fiscally consequential: revenue news affects spending only when it first enters the budget forecast.

4 Real Effects: Budget Calendar Gates Transmission

The decision window governs which revenue news enters the budget; this section asks whether it also governs which news reaches the real economy. We find that it does. Using the same on-versus off-window comparison of similar shocks, together with state-industry private employ-

ment data, we show that revenue news arriving inside the decision window moves employment at firms in government-exposed industries, and does so persistently, while revenue-relevant off-window news has at most small, transient effects. A given shock that arrives during the state's decision window impacts private sector employment by roughly four times more than if it arrives outside the decision window. We then show that the budget calendar shapes how national monetary policy transmits to state fiscal policy.

4.1 Employment Effects of the Decision Window

We test the real effects of decision windows in three steps: first, is there a differential response to revenue news in industries exposed to state and local government demand? Second, does this differential response depend on the budget cycle? Finally, does the same pattern appear in aggregate state employment and wages?

4.1.1 Effects Concentrated in Government-Exposed Industries

We first test differential private sector responses to revenue news using QCEW state–industry employment and wage data and measure exposure to government demand using BEA input–output accounts. The intuition is that firms in industries that sell relatively more to state and local government demand are more likely to be affected by changes in the state's operating budget via procurement.

The design is a triple difference: we compare government-exposed and less-exposed industries within the same state-year, and ask whether that exposure difference is larger for revenue news arriving on- versus off-window. To measure industry-level exposure to state and local government demand, we compute the state-and-local government purchase share from national input–output data for each two-digit NAICS industry and fiscal year. To keep the exposure variation predetermined with respect to the budget cycle, the measure is a five-year trailing average lagged two years. For robustness, we test two alternative exposure measures: a high-exposure indicator based on the median average state–industry exposure, and a standardized continuous exposure measure.

Let $Y_{s,j,t}$ denote private employment or wages in state s, industry j, and fiscal year t,

facing revenue news in each month m leading up to the fiscal year t. The full specification is

$$\begin{aligned}\Delta \log Y_{s,j,t} = \kappa \left(\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}_{j,t}\right) \\ + \alpha_{s,t} + \alpha_{j,t} + \alpha_{s,j} + \delta_{t,m} + \mathbf{Z}'_{s,j,t,m}\Lambda + \varepsilon_{s,j,t,m},\end{aligned} \tag{14}$$

where $\mathbf{Z}_{s,j,t,m}$ includes the relevant lower-order interactions. State-by-fiscal-year fixed effects, $\alpha_{s,t}$, absorb all state-year shocks, including the level effect of cumulated revenue news, contemporaneous state policies, and the state-level controls used in specification (13). Industry-by-fiscal-year fixed effects, $\alpha_{j,t}$, absorb national industry cycles, and state-by-industry fixed effects, $\alpha_{s,j}$, absorb persistent local industry composition. Fiscal-year-by-calendar-month fixed effects, $\delta_{t,m}$, absorb common monthly shocks within each fiscal-year cohort as in our previous analysis of state-level revenue news. Identification therefore comes from whether government-exposed industries respond differentially to decision-window news within the same state-year.

Table 4 supports the procurement interpretation across exposure definitions and fixed-effect specifications. The news measure has no natural units of its own, so we express the magnitudes in budget units using the paper's fiscal first stage. The news variable raises the enacted revenue forecast with coefficient 2.71 (Table 2), so dividing the real-effect coefficients by this pass-through estimate converts them into effects per one-percent news-induced revenue forecast revision. In the preferred full fixed-effect specification of Panel A, Columns (2) and (4), decision-window news that raises the state's revenue forecast by 1 percent raises private employment in high-government-exposure industries by 0.10 percentage points (0.26/2.71) and private wages by 0.14 percentage points relative to less-exposed industries in the same state-year. Using the continuous exposure measure in Panel B, the same news raises employment in industries with a one-standard-deviation higher government-purchase share by 0.06 percentage points and wages by 0.07 percentage points.

The response does not appear to operate through public payroll. We detect no significant effect of decision-window news on public-sector employment or wages (Appendix Table F.6). Combined with the concentration of private employment effects in government-exposed industries, this evidence is consistent with transmission through procurement and other non-payroll spending margins.

The stringent fixed-effects specification rules out many potential alternative explanations. For example, state-year fixed effects imply that the pattern cannot be explained by government-exposed industries being located in states with better economic conditions. The

result is also not driven by the 2008–2009 financial crisis: Appendix Table F.2 reports similar magnitudes excluding FY2010, whose planning period for many states overlaps or closely follows the September 2008 fall of Lehman Brothers. A leave-one-state-out exercise shows that no single state drives the estimate (Appendix Figure F.1). Finally, randomization inference that permutes government exposure across two-digit industries places the observed employment coefficient near the upper tail of the placebo distribution ($p = 0.059$; Appendix Table F.5).

4.1.2 Persistent Employment Effects

Next, we test when the differential private employment response materializes, and how long it persists. We estimate quarter-by-quarter versions of specification (14), replacing the dependent variable with cumulative employment growth from the pre-fiscal-year base to quarterly horizon h, $\Delta_h \log Y_{s,j,t} = \log Y_{s,j,t+h} - \log Y_{s,j,t-1}$, using the high-government-exposure indicator and the same fixed effects. Figure 5 plots the differential employment response of high- relative to low-exposure industries, separately for off-window news and for decision-window news; the vertical distance between the two series is the decision-window premium.

The dynamics show how the budget calendar translates revenue news shocks to the real economy. Prior to the fiscal year, there is no differential response in government-exposed industries, showing that we cannot reject parallel pre-trends. Once the fiscal year begins, and the spending from the enacted budget begins, the employment in government-exposued industries responds positively to on-window news, while the response to revenue-relevant off-window news is small throughout. Because decision-window news has already arrived before the fiscal year starts, a purely informational or anticipatory response would appear when the news arrives. Instead, the response tracks budget execution, consistent with the procurement channel rather than a contemporaneous reaction to the news. Appendix Figure F.2 shows similar dynamics using the continuous exposure measure.

Importantly, the decision-window employment response persists for years beyond the fiscal year whose budget the news affects, whereas the off-window response stays small and transient throughout (Figure 5). This gap is not driven by a difference in informativeness of news arriving on- versus off-window. We show in table 5 that pre-decision news strongly and significantly predicts realized revenue in $t + 2$, over and above the year-t forecast, so it carries genuine revenue information. Moreover, the budget does not ignore this information:

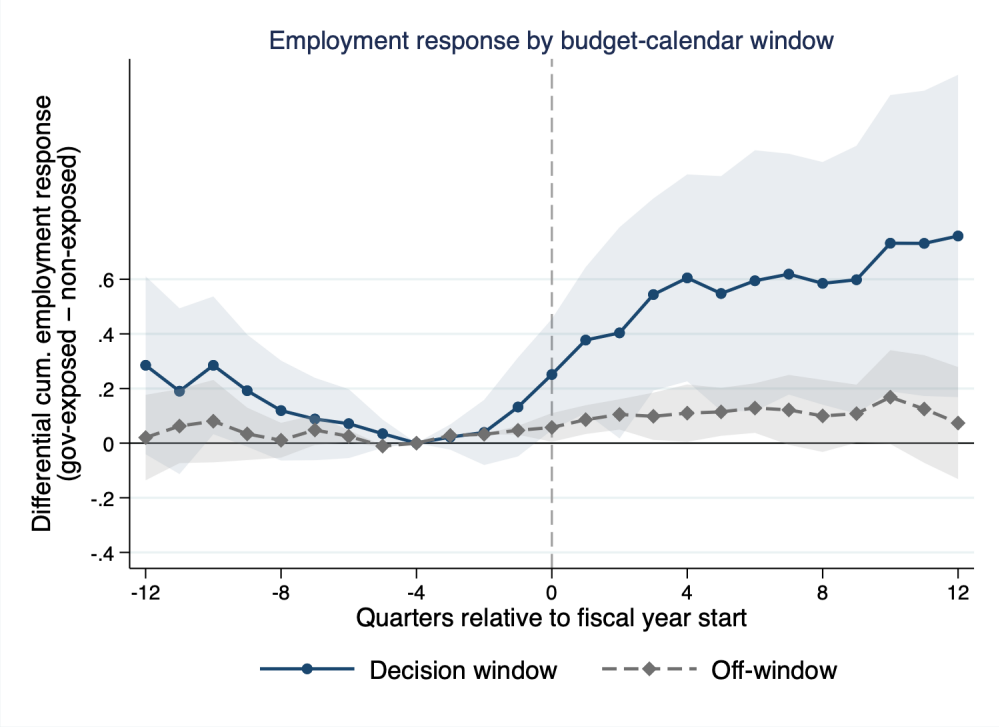


Figure 5: **Dynamic Employment Response by Budget-Calendar Window.** The figure plots horizon-by-horizon local projections where the dependent variable is cumulative private employment growth from the pre-fiscal-year base to horizon h, $\Delta_h \log Y_{s,j,t} = \log Y_{s,j,t+h} - \log Y_{s,j,t-1}$. The off-window series is the coefficient on $\text{News}_{s,t,m} \times \text{GovShare}^{high}$, measuring the differential response of high- versus low-exposure industries to off-window news. The decision-window series adds the triple interaction $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}^{high}$, so the vertical distance between the two series is the decision-window premium. Government exposure is a predetermined high-exposure indicator based on the state-and-local-government purchase share. Regressions include state-by-industry, industry-by-fiscal-year, and state-by-fiscal-year fixed effects. Shaded areas are 90 percent confidence intervals based on clustered standard errors.

pre-decision news enters the revenue forecast itself by $t + 2$.[19] Off-window news is therefore neither noise nor permanently ignored; it is simply incorporated into the budget later.

Moreover, we show that the private employment response to decision-window news is asymmetric and concentrated in negative news. Running the same specification as in Figure 5, we plot the "decision-window premium"—the differential response of high- versus low-exposure industries to on- versus off-decision-window news—separately for positive and negative news (Figure 6). The persistent response comes almost entirely from negative

[19] In the full sample, the decision-window row carries large negative later-budget coefficients; these are driven by the FY2010 crisis round, in which the September 2008 shock entered many states' decision windows, and they are no longer distinguishable from zero once FY2010 is excluded; see Appendix Table F.3.

decision-window news, while positive news generates a smaller and less persistent expansion. We verify that this asymmetry does not necessarily come from the budget. In Appendix Table F.1 we test whether state budgets respond differentially to positive versus negative decision-window news and cannot reject symmetric pass-through.

We interpret these results as evidence of firm-level adjustment frictions. Decision-window news shifts government demand within the fiscal year, giving exposed firms little notice, whereas off-window news reaches spending through later budgets that firms can anticipate and respond to in advance (Leeper et al., 2013). Because budgets pass positive and negative news through symmetrically, the asymmetric employment response supports a firm-level friction, consistent with credit-constrained government vendors who cannot bridge abrupt demand shortfalls, leading to layoffs that outlast the shock itself (Barrot and Nanda, 2020). The budget calendar thus converts similar news into different real outcomes, less by changing what states ultimately spend than by changing how much time firms have to adjust.

4.1.3 Aggregate State Responses

Next, we test if the decision-window effect propagates to aggregate state outcomes. We estimate equation (13) with $Y_{s,t}$ equal to state-wide private employment or private wages growth. The specification uses the same news measure, fixed effects, controls, and restricted pre-fiscal-year sample as the budget regressions.

Table 6 shows that the decision-window premium is positive and precisely estimated. The interaction coefficient on news times the decision window indicator is $\theta^Y = 0.41$ for private employment and $\theta^Y = 0.55$ for private wages. The pooled off-window response is much smaller, 0.11 for employment and 0.19 for wages. These estimates show that when the budget channel is active, state-level private employment and wages also respond by more: roughly 4 times the off-window response. Appendix Table F.4 shows that the aggregate premia are slightly larger when FY2010 is excluded.[20]

In summary, we find that fiscal decision-windows not only determine what shocks enter state spending, but also how shocks transmit to the local real economy. We interpret this differential aggregate response to on- versus off-decision-window news as arising from a com-

[20]The 2008 financial crisis began during the planning period for many states; those states that were more negatively impacted were also more likely to receive federal stimulus, which counters our effects. Thus, excluding FY2010 sharpens our results.

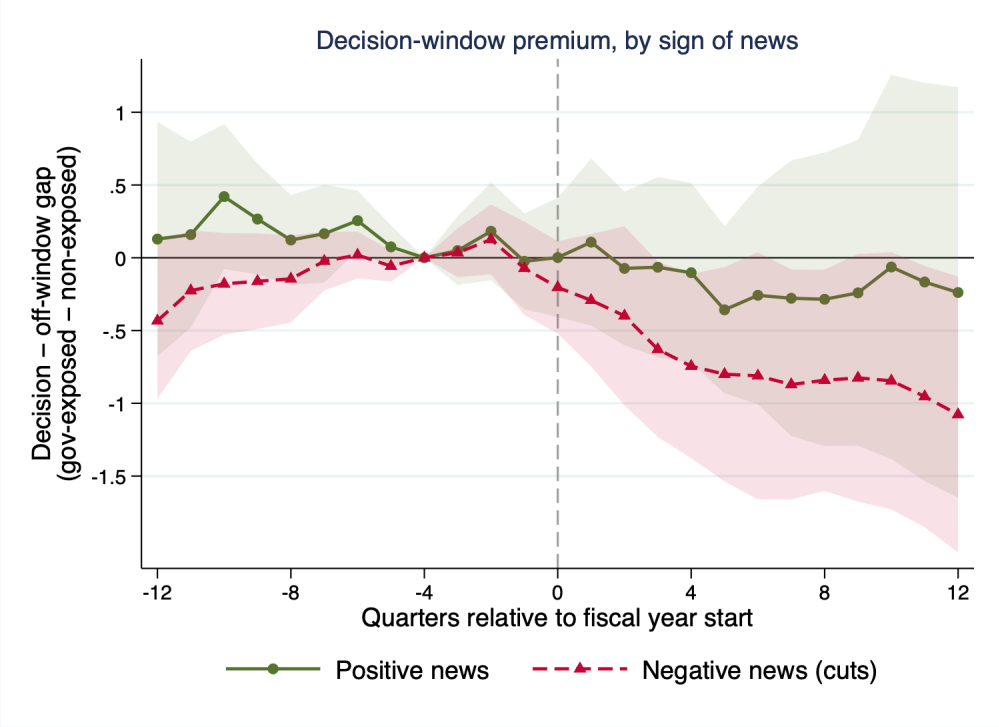


Figure 6: **Asymmetric Dynamic Employment Response to Decision-Window News.** The figure plots horizon-by-horizon local projections where the dependent variable is cumulative private employment growth from the pre-fiscal-year base to horizon h, $\Delta_h \log Y_{s,j,t} = \log Y_{s,j,t+h} - \log Y_{s,j,t-1}$. The plotted coefficients are decision-window premia: the differential response of high- versus low-government-exposure industries to decision-window news relative to off-window news. The positive series uses positive decision-window news; the negative series uses the absolute value of negative decision-window news. Government exposure is a predetermined high-exposure indicator based on the state-and-local-government purchase share. Regressions include state-by-industry, industry-by-fiscal-year, and state-by-fiscal-year fixed effects. Shaded areas are 90 percent confidence intervals based on clustered standard errors.

bination of state spending responses, the procurement channel effect, and the downstream fiscal multiplier effects. Thus, the decision window can help identify the fiscal multiplier for routine budget adjustment, which we explore next.

Rule-induced fiscal multiplier. We exploit the decision-window effect to construct a rule-induced fiscal elasticity. This differs from the standard fiscal-multiplier exercise. The usual approach exploits spending shocks that are orthogonal to local fundamentals; here, the variation comes from the systematic budget rule that arises from the government's optimization problem discussed in Section 2.

Revenue news moves spending when it is processed through the decision window, and the on- versus off-window contrast isolates the part of the real response generated by that routine fiscal adjustment. The identifying assumption is that comparable news arriving on- and off-window have similar direct effects on the private sector (see details in Appendix G). We estimate the multiplier as a Wald ratio: the reduced-form decision-window premium in the outcome divided by the first-stage decision-window premium in realized spending. The first stage, estimated directly in Table G.1 (Panel B, column 2), is 1.20: a decision-window news realization raises realized spending by 1.20 percent more than the same news off-window.[21] Dividing the aggregate real-effects decision-window premium in Table 6 by this first stage gives

$$\widehat{\mu}^{f}_{\text{Emp}} = \frac{0.41}{1.20} = 0.34, \qquad \widehat{\mu}^{f}_{\text{Wages}} = \frac{0.55}{1.20} = 0.46.$$

Thus a 1 percent increase in realized spending induced by routine budget adjustment is associated with a 0.34 percent increase in private employment and a 0.46 percent increase in private wages.

To put the employment estimate in familiar units, we benchmark to fiscal year 2023, the most recent year with near-complete coverage: realized state spending averaged roughly $33 billion and private employment 2.7 million workers per state, implying about $36,000 of induced state spending per private job-year.[22] The implied cost per job-year is the same order of magnitude as estimates from the cross-region fiscal-multiplier literature (Nakamura and Steinsson, 2014; Auerbach and Gorodnichenko, 2012; Chodorow-Reich, 2019), and comparable to the ARRA-based estimates in Chodorow-Reich et al. (2012). The comparison is only a scaling benchmark: unlike shock-based designs, this estimate uses routine forecast-mediated budget adjustment induced by the decision window.

4.2 Monetary Policy Transmission via Decision Windows

Next, we test if the decision-window effect generalizes to other types of shocks by analyzing the passthrough of monetary policy to state fiscal policy. Monetary policy shocks are

[21] This magnitude is close to the first stage implied by the forecast channel in Equation (7): the on- versus off-window revenue-forecast premium of 2.36 (Table 2), times the pass-through of 0.56 from a news-induced forecast revision to realized spending (Table 3, Column 2), gives $2.36 \times 0.56 = 1.32$.

[22] The conversion is $33 billion/$(0.34 \times 2.7$ million$) \simeq$ $36,000. Evaluated instead at the means of the full estimation sample (roughly $19.5 billion and 2.3 million private workers per state-year over 1995–2024), the same conversion gives about $25,000; we prefer the 2023 benchmark because sample means average dollars across three decades of price levels.

common across states and can affect fiscal choices through macro expectations and financing conditions. We find that in both the state-level panel and aggregate time series, tightening surprises have larger and more persistent fiscal spending effects when they arrive during the decision window.

We adapt the monthly event-time design of Section 3, replacing state-specific revenue news with a national monetary policy surprise. We compute monthly sums of shocks from four standard measures of monetary policy surprises, all constructed by Bauer and Swanson (2023): the first principal component of the high-frequency changes in one to four quarter ahead Eurodollar (SOFR in 2023) futures in the 30-minute windows around FOMC announcements, both raw and residualized by lagged macro variables, and the analogous raw high-frequency changes in 2- and 5-year U.S. Treasury yield futures. For all policy surprise measures, the identifying assumption is that asset price movements inside the tight window around FOMC announcements are responding to policy news, and are thus orthogonal to state-level economic fundamentals (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023).

We adopt the same identification as in specification (13), with the national monetary-policy surprise MPS_m in place of revenue news. Because the surprise is common to all states in a calendar month, the time fixed effects $\delta_{t,m}$ absorb its level, and identification comes only from its interaction with whether month m lies in state s's decision window for fiscal year t.

$$y_{s,t} = \theta^{y}_{\text{MPS}} \left(\text{MPS}_m \times \mathbf{1}\{m \in \mathcal{W}^{\text{Decision}}_{s,t}\} \right) + \mathbf{X}'_{s,t}\beta + \alpha_s + \delta_{t,m} + \varepsilon_{s,m,t}. \tag{15}$$

where MPS_m is in percentage-point units. As before, the sample stacks each state-year's pre-decision, decision, and deliberation months, and the outcome $y_{s,t}$ is the spending budget or realized spending relative to the prior-year original budget. The coefficient θ^{y}_{MPS} is the in-window-minus-off-window response to a tightening surprise.

We find clear evidence that the decision window mediates the pass-through of monetary policy to state fiscal policy. Using our most stringent fixed-effects specification, Table 7 shows that a 5-basis-point surprise tightening, roughly one standard deviation of the high-frequency shock, lowers budgeted and realized spending growth rate by about 0.5 percentage point when it arrives inside the decision window relative to the same surprise outside the window. The sign is negative for all four measures; the effect is statistically significant for the two futures-based surprises and the 2-year Treasury surprise, and has the same sign but is imprecise for the 5-year yield. As with revenue news, national policy news becomes state

fiscal policy mainly when it arrives during the decision window.

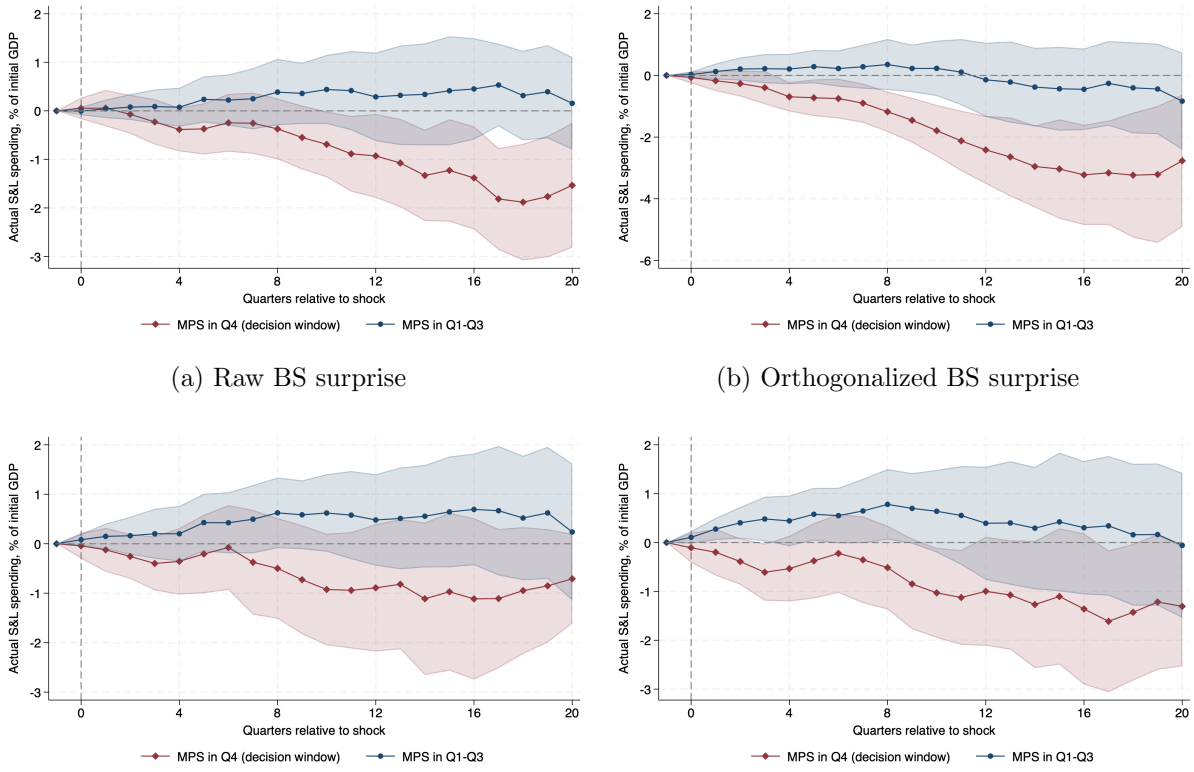
Next, we turn to aggregate time-series effects. While the monthly state-panel specification cleanly isolates the decision-window premium using cross-state variation in proposal dates, it absorbs the national time series shock and therefore cannot identify the aggregate fiscal response. Using quarterly NIPA data, we estimate local projections (Jordà, 2005) of state-and-local spending on monetary policy surprises, separating shocks that arrive in the fourth calendar quarter (Q4) from shocks that arrive in the first three calendar quarters (Q1–Q3). While states have different proposal dates, many of the decision windows overlap with calendar Q4 (see Table A.1 in Appendix A). Figure 7 plots the coefficients over the 20 quarters following each shock, the effects of Q4 shocks are in red and the effects of Q1–Q3 shocks are in blue.

A striking pattern emerges: Q4 tightenings generate persistent declines in state-and-local spending, while Q1–Q3 tightenings have negligible effects. The response to Q4 shocks is not immediate; it builds over roughly two to three years. This timing is consistent with our channel: a policy-news event affects public spending only after it has been incorporated into budgets that are subsequently enacted and executed. Across measures, the trough Q4 response is roughly 0.1-0.3 percent of initial GDP for a 10 basis point surprise. To put this scale in context, the same orthogonalized Bauer–Swanson surprise implies an aggregate GDP response of about 1-1.4 percent of initial GDP in the corresponding local projection (Appendix Figure F.4), so the direct state-and-local spending response is around 10-20 percent of the total output response.

This Q4 effect is not driven by the shocks themselves: the Q4 and Q1–Q3 surprise distributions are statistically indistinguishable in mean, variance, and shape (Appendix Figure F.3). We verify in Appendix F.4 that the Q4 difference is statistically significant. We also show in this Appendix that the Q4 effect is not driven by the December FOMC meeting alone and survives controlling for federal spending. Federal fiscal spending does not exhibit similar patterns.

Consistent with the private sector employment results, these findings show that news is incorporated into state spending during the decision window. Importantly, the differential decision-window effect appears to be persistent, likely reflecting the multipler effect of fiscal spending.

Moreover, our monetary policy results offer a complementary institutional perspective on the classic problem that monetary policy operates with long and variable lags (Friedman,



(a) Raw BS surprise

(b) Orthogonalized BS surprise

(c) 2-year Treasury surprise

(d) 5-year Treasury surprise

Figure 7: **Monetary Policy and State-and-Local Spending, by Budget Season.** Local projections of cumulative real state-and-local government spending (NIPA), scaled by lagged real GDP, on four quarterly monetary policy surprise measures: the raw and orthogonalized Bauer–Swanson surprises and the 2-year and 5-year Treasury yield surprises from Bauer and Swanson (2023). Monthly surprises are summed within each quarter and kept in their native percentage-point units. Each panel splits shocks by whether they arrive in the calendar Q4, when many states draft the upcoming budget or outside it, Q1–Q3. The local projections include four lags of state-and-local spending growth; four lags of each season-specific shock (the in-window and off-window components entered separately); and four predetermined quarterly lags of real GDP growth, payroll-employment growth, S&P 500 returns, the change in the 5-year Treasury yield, and CPI inflation. The sample period is 1988–2024, the span of the Bauer–Swanson surprises. The vertical axis is the cumulative spending response in percent of initial GDP. Shaded areas are 90 percent confidence intervals based on horizon-dependent Newey–West standard errors.

1961; Christiano et al., 1999). Standard accounts emphasize nominal rigidities, information stickiness, and financial frictions (Christiano et al., 2005; Coibion and Gorodnichenko, 2015; Bernanke et al., 1996). The evidence here points to an additional public-sector budget calendar friction: because budgets fix commitments at discrete calendar dates, the same monetary policy shock has a larger and more persistent effect if it arrives when many states

are in their decision window.

5 Conclusion

This paper studies how economic shocks that could affect revenue expectations is translated into state fiscal policy and real activity. Using hand-collected budget documents for 47 U.S. states from 1995 to 2024, we show that the original enacted budget is the key commitment margin for spending and that revenue news affects this commitment only when it arrives during the budget decision window. Comparable off-window news remains informative about realized revenue, but it is not processed into forecasts, spending plans, or realized spending in the same way. This timing difference carries through to the private economy, especially among industries and firms exposed to state and local government demand.

The main implication is that routine budget adjustment is a quantitatively important propagation mechanism for ordinary economic shocks. The real effects estimates imply that news arriving during the decision window affectsprivate employment outcomes by up to four times as much as similar news that arrives outside the decision window. Using this variation, our paper can estimate the real consequences of the endogenous budget response to economic conditions. Our finding that monetary policy transmission is stronger during a state's decision window offers an institutional lens for the long and variable lags of policy transmission to state spending.

The results also change how one should think about fiscal stabilization. If budgets respond only when information arrives inside a narrow budget decision window, then the timing of forecasts, budget revisions, and intergovernmental aid can matter as much as their total size. Aid that arrives after a budget has been committed may prevent later cuts, but it is less likely to change the same-year spending plan that governs immediate procurement and employment. Conversely, negative news that arrives while the budget is being written can be converted quickly into spending cuts and private-sector contractions. The sign-specific results suggest that this propagation is especially important in downturns, where firms and workers may have less capacity to absorb fiscal contractions.

More broadly, the paper shows that fiscal institutions shape shock transmission through timing: budget calendars determine when information becomes policy. This insight extends beyond state public finance. Many large organizations, such as government agencies, corporations, and institutional investors, receive information continuously but commit to actions

only periodically. As data on expectations and plans become increasingly available, our simple decision-window framework offers a lens for understanding how information is incorporated into decisions. Applying this framework more widely can illuminate how shocks, such as monetary or federal fiscal policy, transmit to the real economy.

References

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., and Thesmar, D. (2023). Overreaction in expectations: Evidence and theory. *The Quarterly Journal of Economics*, 138(3):1713–1764.

Andersen, A. L., Lassen, D. D., and Nielsen, L. H. W. (2012). Late budgets. *American Economic Journal: Economic Policy*, 4(4):1–40.

Anderson, T. W. and Rubin, H. (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics*, 20(1):46–63.

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Auerbach, A. J., Gorodnichenko, Y., and Murphy, D. (2020). Local fiscal multipliers and fiscal spillovers in the USA. *IMF Economic Review*, 68(1):195–229.

Azeredo da Silveira, R., Sung, Y., and Woodford, M. (2024). Optimally imprecise memory and biased forecasts. *American Economic Review*, 114(10):3075–3118.

Barrot, J.-N. and Nanda, R. (2020). The employment effects of faster payment: Evidence from the federal quickpay reform. *The Journal of Finance*, 75(6):3139–3173.

Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155.

Bernanke, B., Gertler, M., and Gilchrist, S. (1996). The financial accelerator and the flight to quality. *The Review of Economics and Statistics*, 78(1):1–15.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *Journal of Finance*, 73(1):199–227.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Bouscasse, P. and Hong, S. (2026). Monetary-fiscal interactions in the united states.

Chodorow-Reich, G. (2019). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, 11(2):1–34.

Chodorow-Reich, G., Feiveson, L., Liscow, Z., and Woolston, W. G. (2012). Does state fiscal relief during recessions increase employment? Evidence from the American Recovery and Reinvestment Act. *American Economic Journal: Economic Policy*, 4(3):118–145.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

Christiano, L. J., Eichenbaum, M. S., and Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of Macroeconomics*, 1(B):65–148.

Clemens, J. and Miran, S. (2012). Fiscal policy multipliers on subnational government spending. *American Economic Journal: Economic Policy*, 4(2):46–68.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Cox, L., Müller, G. J., Pasten, E., Schoenle, R., and Weber, M. (2024). Big g. *Journal of Political Economy*, 132(10):3260–3297.

Dessein, W., Galeotti, A., and Santos, T. (2016). Rational inattention and organizational focus. *American Economic Review*, 106(6):1522–1536.

Dessein, W. and Santos, T. (2006). Adaptive organizations. *Journal of Political Economy*, 114(5):956–985.

Feenberg, D. R., Gentry, W., Gilroy, D., and Rosen, H. S. (1989). Testing the rationality of state revenue forecasts. *The Review of Economics and Statistics*, pages 300–308.

Ferraz, C., Finan, F., and Szerman, D. (2015). Procuring firm growth: the effects of government purchases on firm dynamics. Technical report, National Bureau of Economic Research.

Frankel, J. (2011). Over-optimism in forecasts by official budget agencies and its implications. *Oxford Review of Economic Policy*, 27(4):536–562.

Frankel, J. and Schreger, J. (2013). Over-optimistic official forecasts and fiscal rules in the eurozone. *Review of world economics*, 149(2):247–272.

Friedman, M. (1961). The lag in effect of monetary policy. *Journal of Political Economy*, 69(5):447–466.

Gennaioli, N., Ma, Y., and Shleifer, A. (2016). Expectations and investment. *NBER Macroeconomics Annual*, 30(1):379–431.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Grenadier, S. R., Malenko, A., and Malenko, N. (2016). Timing decisions in organizations: Communication and authority in a dynamic environment. *American Economic Review*, 106(9):2552–2581.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Leeper, E. M., Walker, T. B., and Yang, S.-C. S. (2013). Fiscal foresight and information flows. *Econometrica*, 81(3):1115–1145.

Liebman, J. B. and Mahoney, N. (2017). Do expiring budgets lead to wasteful year-end spending? Evidence from Federal Procurement. *American Economic Review*, 107(11):3510–3549.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Malenko, N. (2024). Information flows, organizational structure, and corporate governance. In *Handbook of corporate finance*, pages 511–546. Edward Elgar Publishing.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117(4):1295–1328.

Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5):1329–1376.

Mikesell, J. L. and Ross, J. M. (2014). State revenue forecasts and political acceptance: The value of consensus forecasting in the budget process. *Public Administration Review*, 74(2):188–203.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107.

Nagel, S. and Xu, Z. (2022). Asset pricing with fading memory. *The Review of Financial Studies*, 35(5):2190–2245.

Nakamura, E. and Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from US regions. *American Economic Review*, 104(3):753–792.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, 133(3):1283–1330.

National Association of State Budget Officers (2021). Budget processes in the states. Report, National Association of State Budget Officers, Washington, DC.

Nekarda, C. J. and Ramey, V. A. (2011). Industry evidence on the effects of government spending. *American Economic Journal: Macroeconomics*, 3(1):36–59.

Olivei, G. and Tenreyro, S. (2007). The timing of monetary policy shocks. *American Economic Review*, 97(3):636–663.

Poterba, J. M. (1994). State responses to fiscal crises: The effects of budgetary institutions and politics. *Journal of Political Economy*, 102(4):799–821.

Poterba, J. M. (1996). Budget institutions and fiscal policy in the us states. *The American Economic Review*, 86(2):395.

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics*, 126(1):1–50.

Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy*, 126(2):850–901.

Reis, R. (2006). Inattentive producers. *Review of Economic Studies*, 73(3):793–821.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.

Selgrad, J. and Siani, K. (2025). Monetary policy and investment plans. *Available at SSRN*.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Suárez Serrato, J. C. and Wingender, P. (2016). Estimating local fiscal multipliers. Working Paper 22425, National Bureau of Economic Research.

Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.

Wiegand, C. (2025). The effect of fiscal policy shocks on asset prices. *Available at SSRN 5269517*.

Wilson, D. J. (2012). Fiscal spending jobs multipliers: Evidence from the 2009 American Recovery and Reinvestment Act. *American Economic Journal: Economic Policy*, 4(3):251–282.

Table 1: **Summary Statistics**

Panel A summarizes total real expenditures and total real revenues at the state-fiscal-year level over fiscal years 1995–2024 across 47 annual-budget states. "Actual" is total realized real spending or revenue, in billions of dollars. "Budget" is the log change in the original budget across fiscal years, $\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$ for spending and the analogous object for revenue. "Execution" is the within-fiscal-year gap between the realized outcome and the original budget, $\log(\text{Act}_{s,t}/\text{Orig}_{s,t})$. Variables are winsorized at the 0.5/99.5 percentiles. Panel B summarizes the state-fiscal-year-level shift-share news measure aggregated across proposal-relative windows. Letting ℓ be the calendar-month distance from state s's FY-t proposal date, the windows are pre-planning ($\ell \in [-11, -6]$), planning ($\ell \in [-5, -1]$), deliberation ($\ell \geq 0$ and $m <$ FY start), and execution (FY start $\leq m \leq$ FY end). Panel B statistics are multiplied by 100.

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
Budget-$\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$	0.023	0.082	-0.089	-0.015	0.019	0.052	0.148
Execution-$\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t})$	-0.075	0.104	-0.288	-0.098	-0.055	-0.027	0.048
Revenues							
Actual total revenue (Bn)	20.688	25.328	2.040	5.116	10.999	27.664	68.421
Budget-$\log(\text{Orig}^R_{s,t}/\text{Orig}^R_{s,t-1})$	0.022	0.089	-0.096	-0.010	0.020	0.050	0.146
Execution-$\log(\text{Act}^R_{s,t}/\text{Orig}^R_{s,t})$	-0.026	0.119	-0.172	-0.054	-0.018	0.017	0.111

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

	Mean	SD	P5	P25	Median	P75	P95
News Pre-Decision.	-0.06	0.51	-0.45	-0.14	-0.06	0.03	0.31
News Decision	-0.07	0.27	-0.47	-0.14	-0.07	0.03	0.24
News Deliberation	-0.06	0.54	-0.51	-0.14	-0.06	0.03	0.31
News Execution	-0.07	0.55	-0.55	-0.14	-0.06	0.03	0.30

Table 2: **Timing of News and Original Budget Formation**

Monthly-panel regressions of original-budget forecast revisions, forecast errors, and realized revenue growth on state-specific shift-share news by proposal-relative timing. Panel A reports separate coefficients for pre-decision news ($\text{News}^{\text{Pre}}_{s,t,m}$, $\ell \in [-11, -6]$), decision-window news ($\text{News}^{\text{Decision}}_{s,t,m}$, $\ell \in [-5, -1]$), and deliberation news ($\text{News}^{\text{Delib}}_{s,t,m}$, $\ell \geq 0$ before fiscal-year start). Panel B reports the in/out specification, where the coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}^{\text{Decision}}_{s,t}\}$ is the decision-window premium relative to pooled off-window news. All columns share a common sample (state-fiscal-years with both the forecast revision and the forecast error observed), so realized revenue growth equals the forecast revision plus the forecast error and the coefficients satisfy the adding-up identity column by column. The "one-sided p-value" row in Panel B tests the framework's directional predictions—a positive premium for the forecast and a negative premium for the forecast error; the framework makes no directional prediction for realized revenue, so that entry is left blank. The sample is a stacked monthly panel restricted to these three pre-fiscal-year windows, 1995–2024. All specifications include state and fiscal-year-by-calendar-month fixed effects and predetermined controls: four lagged quarterly windows of CFNAI, payroll employment, the S&P 500, the five-year Treasury yield, and CPI, with the most recent ending one year before budget formation; the Borusyak et al. (2022) aggregate-share control; and two-year-lagged state rainy-day-fund balances and GDP growth. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

Panel A: Window-specific responses

	Forecast $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$	Forecast error $\log(\text{Act}^R_t/\text{Orig}^R_t)$	Realized $\log(\text{Act}^R_t/\text{Orig}^R_{t-1})$
	(1)	(2)	(3)
$\text{News}^{\text{Pre}}_{s,t,m}$	0.02	1.10***	1.12
	(0.46)	(0.39)	(0.75)
$\text{News}^{\text{Decision}}_{s,t,m}$	2.71***	-0.16	2.55***
	(0.61)	(0.57)	(0.74)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.55	0.62*	1.17*
	(0.47)	(0.34)	(0.58)
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402
Adj. R^2	0.121	0.575	0.446

Panel B: Decision-window premium

	Forecast $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$	Forecast error $\log(\text{Act}^R_t/\text{Orig}^R_t)$	Realized $\log(\text{Act}^R_t/\text{Orig}^R_{t-1})$
	(1)	(2)	(3)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}^{\text{Decision}}_{s,t}\}$	2.36***	-0.96	1.40*
	(0.50)	(0.64)	(0.72)
One-sided p-value	0.000	0.071	
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402
Adj. R^2	0.121	0.575	0.446

Table 3: **Pass-Through: Revenue Forecast Updates to Spending**

The table reports 2SLS estimates of the pass-through from original-budget revenue forecast updates to spending. Panel A uses $\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$ as the dependent variable; Panel B uses $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t-1})$. The endogenous regressor is $\log(\text{Orig}^R_{s,t}/\text{Orig}^R_{s,t-1})$, instrumented separately across columns with news arriving in the pre-decision ($\ell \in [-11, -6]$), decision ($\ell \in [-5, -1]$), or deliberation ($\ell \geq 0$ before fiscal-year start) window; the other two window-news terms enter as controls. The sample is the stacked monthly panel restricted to these three pre-fiscal-year windows. All specifications include state and fiscal-year-by-calendar-month fixed effects and the predetermined controls from Table 2. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%. Placebo instruments are reported as dashes because their 2SLS estimates are not interpretable. "Effective F" is the first-stage Kleibergen–Paap cluster-robust Wald statistic. "AR 95% CI" reports weak-instrument-robust Anderson and Rubin (1949) confidence sets; these are bounded for the decision-window instrument and unbounded for the placebo instruments.

Panel A: Original spending plan

	Spending budget: $\log(\text{Orig}^E_t/\text{Orig}^E_{t-1})$		
	(1)	(2)	(3)
Revenue $\widehat{\text{budget}}$ change	–	0.37**	–
		(0.16)	
Instrument window	Pre-decision	Decision	Deliberation
Effective F	0.0	22.0	1.3
AR 95% CI	Unbounded	[0.08, 0.78]	Unbounded
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402

Panel B: Actual spending growth

	Actual spending growth: $\log(\text{Act}^E_t/\text{Orig}^E_{t-1})$		
	(1)	(2)	(3)
Revenue $\widehat{\text{budget}}$ change	–	0.56***	–
		(0.18)	
Instrument window	Pre-decision	Decision	Deliberation
Effective F	0.0	22.0	1.3
AR 95% CI	Unbounded	[0.19, 0.96]	Unbounded
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402

Table 4: **Real Effects of Decision-Window**

The table reports triple-difference regressions of one-year log changes in private employment and private wages on the state-specific shift-share news measure, $\text{News}_{s,t,m}$, interacted with an indicator for months in state s's budget decision window for fiscal year t, $\mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$, and a measure of industry exposure to state and local government demand, built from the average five-year-trailing, two-year-lagged BEA government-purchase share of each industry. Panel A uses the binary indicator GovShare^{high}, which equals one for state–industry cells with above-median exposure; Panel B uses the continuous measure GovShare^{sd}, the same share standardized to mean zero and unit standard deviation across cells. The sample is a stacked state-industry-fiscal-year-month panel covering two-digit NAICS industries from 1995 to 2024; for each state-fiscal-year, the included months fall before the fiscal-year start. All specifications include the relevant lower-order interactions. Standard errors are clustered two ways by state and calendar-month-year. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Above-Median Exposure Indicator

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$		$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$	
	(1)	(2)	(3)	(4)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}_s^{high}$	0.22*	0.26**	0.33*	0.37**
	(0.11)	(0.12)	(0.17)	(0.18)
Two-way Interactions	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓
Industry × Year FE	✓	✓	✓	✓
State × Industry FE		✓		✓
FYear × Year-Month FE		✓		✓
Observations	566,079	566,079	566,079	566,079
Adj.R^2	0.55	0.57	0.78	0.78

Panel B: Continuous Exposure Measure

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$		$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$	
	(1)	(2)	(3)	(4)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}_s^{sd}$	0.15**	0.16**	0.16*	0.18*
	(0.06)	(0.07)	(0.09)	(0.09)
Two-way Interactions	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓
Industry × Year FE	✓	✓	✓	✓
State × Industry FE		✓		✓
FYear × Year-Month FE		✓		✓
Observations	566,079	566,079	566,079	566,079
Adj.R^2	0.55	0.57	0.78	0.78

Table 5: **Where Off-Window News Goes: Budget Trajectory by Window of Arrival**

Monthly-panel regressions tracing where news arriving in each proposal-relative window appears in the subsequent budget trajectory. All outcomes are benchmarked to the current original revenue budget, Orig_t^R. Columns (1)–(2) ask whether later original budgets revise toward the news, $\log(\text{Orig}_{t+k}^R/\text{Orig}_t^R)$ for $k = 1, 2$ (the $k = 0$ revision is zero by construction); columns (3)–(5) ask whether the news appears as realized revenue surprises relative to the plan, $\log(\text{Act}_{t+k}^R/\text{Orig}_t^R)$ for $k = 0, 1, 2$, where the $k = 0$ column is the contemporaneous forecast error. As in Panel A of Table 2, the regressors are the window-specific news measures: pre-decision news ($\text{News}_{s,t,m}^{\text{Pre}}$, $\ell \in [-11, -6]$), decision-window news ($\text{News}_{s,t,m}^{\text{Decision}}$, $\ell \in [-5, -1]$), and deliberation news ($\text{News}_{s,t,m}^{\text{Delib}}$, $\ell \geq 0$ before fiscal-year start). The sample is the stacked monthly panel restricted to these three pre-fiscal-year windows, 1995–2024; the number of observations varies across columns with the availability of each outcome horizon. All specifications include state and fiscal-year-by-calendar-month fixed effects and the predetermined controls described in Table 2. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

	Budget revision $\log(\text{Orig}_{t+k}^R/\text{Orig}_t^R)$		Forecast error $\log(\text{Act}_{t+k}^R/\text{Orig}_t^R)$		
	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)
$\text{News}_{s,t,m}^{\text{Pre}}$	-0.07	1.61**	1.13***	1.04	3.38***
	(0.31)	(0.63)	(0.39)	(0.74)	(0.62)
$\text{News}_{s,t,m}^{\text{Decision}}$	-1.40*	-3.83***	-0.21	-0.70	-4.10**
	(0.72)	(1.06)	(0.58)	(1.22)	(1.70)
$\text{News}_{s,t,m}^{\text{Delib}}$	0.03	-0.52	0.70*	0.54	-0.11
	(0.42)	(0.47)	(0.36)	(0.80)	(0.60)
Observations	15,638	14,938	16,998	15,638	14,938
Adj. R^2	0.123	0.139	0.579	0.447	0.392

Table 6: **Aggregate State-Level Real Effects of Decision-Window News**

The table reports monthly-panel regressions of state-level one-year log changes in private employment and private wages on the state-specific shift-share news measure, $\text{News}_{s,t,m}$. The specification is an in/out parametrization: $\text{News}_{s,t,m}$ enters as a pooled off-window level and is interacted with an indicator for months in state s's budget decision window for fiscal year t, $\mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$. The coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$ is therefore the decision-window premium relative to off-window news, while the coefficient on $\text{News}_{s,t,m}$ is the off-window effect. The off-window category pools the pre-decision and deliberation months before the fiscal year starts. The sample is a stacked monthly panel restricted to one budget round, excluding the earliest before-pre-decision months and execution months. All regressions absorb state and fiscal-year-by-calendar-month fixed effects and include predetermined controls: four lagged quarterly windows of CFNAI, payroll employment, the S&P 500, the five-year Treasury yield, and CPI, with the most recent ending one year before budget formation; the Borusyak et al. (2022) aggregate-share control; and two-year-lagged state rainy-day-fund balances and GDP growth. Standard errors are two-way clustered by state and calendar month-year. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$	$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$
	(1)	(2)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$	0.409***	0.553**
	(0.083)	(0.210)
$\text{News}_{s,t,m}$	0.110***	0.193*
	(0.034)	(0.109)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Macro + State Controls	✓	✓
Observations	19,724	19,724
Adj. R^2	0.84	0.75

Table 7: **Monetary Policy Surprises and the Budget Decision Window**

Monthly-panel regressions in the spirit of Table 2, replacing the state-specific revenue news with a national monetary policy surprise (MPS). The sample is a stacked state–fiscal-year–calendar-month panel covering the pre-decision, decision, and deliberation windows for each state-year. The key regressor is $\text{MPS}_m \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$, where m indexes calendar months and $\mathcal{W}_{s,t}^{\text{Decision}}$ is the set of calendar months in state s's budget decision window for fiscal year t. Columns vary the surprise measure: Bauer–Swanson raw (BS Raw) and orthogonalized (BS Orth.) surprises and the Bauer–Swanson 2- and 5-year Treasury responses (Bauer and Swanson, 2023). Panel A is the spending budget and Panel B is realized spending, each measured relative to the original budget in $t-1$. All specifications include state and fiscal-year-by-calendar-month fixed effects and predetermined controls: four lagged quarterly windows of CFNAI, payroll employment, the S&P 500, the five-year Treasury yield, and CPI, with the most recent ending one year before budget formation; and two-year-lagged state rainy-day-fund balances and GDP growth. Standard errors are clustered two ways by state and calendar month-year. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Spending budget, $\log(Orig_t^E / Orig_{t-1}^E)$

	(1) BS Raw	(2) BS Orth.	(3) UST 2y	(4) UST 5y
$MPS_m \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$	-0.10**	-0.14***	-0.08**	-0.05
	(0.04)	(0.05)	(0.04)	(0.03)
State FE	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓
Macro + State Controls	✓	✓	✓	✓
Observations	18,444	18,444	18,444	18,444
Adj. R^2	0.132	0.132	0.132	0.132

Panel B: Realized spending, $\log(Act_t^E / Orig_{t-1}^E)$

	(1) BS Raw	(2) BS Orth.	(3) UST 2y	(4) UST 5y
$MPS_m \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$	-0.12**	-0.17***	-0.09*	-0.04
	(0.06)	(0.05)	(0.05)	(0.04)
State FE	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓
Macro + State Controls	✓	✓	✓	✓
Observations	18,444	18,444	18,444	18,444
Adj. R^2	0.541	0.541	0.541	0.541

Appendix Contents

A Institutional Setting and Data **50**

A.1 State Budget Calendar 50

A.2 Primary Data Source: Annual Comprehensive Financial Reports 51

A.3 Comparison with the NASBO Fiscal Survey 51

B Stylized Facts and the Commitment Margin **54**

B.1 Anchoring of the Initial Budget 56

B.2 Decomposing Budget-Cycle Variation 56

B.3 Revenue Pass-Through to Spending 59

B.4 Rainy Day Funds: Realized Surpluses vs Budget Surprises 61

C Framework Predictions **63**

D Shift-Share Measure: Construction Details **64**

D.1 Revenue Categories and National Sources 65

D.2 Exposure Weights and Elasticity Scaling 67

E Shift-Share Results: Diagnostics and Robustness **71**

E.1 Design Diagnostics 71

E.2 Distinguishing the Decision Window from Alternative Models 73

E.3 Placebo and Inference 76

E.4 Calendar and Sample Robustness 79

E.5 Specification and Measurement Robustness 90

F Real Effects: Additional Results and Robustness **95**

F.1 Budget-Side Sign Split 95

F.2 Robustness of the Industry-Level Results 95

F.3 Public-Payroll Diagnostics 100

F.4 Monetary-Policy Diagnostics 100

G Rule-Induced Multipliers **105**

A Institutional Setting and Data

This appendix provides additional detail on the state budget calendar and the primary data source underlying the analysis in Section 1.

A.1 State Budget Calendar

Table A.1 reports the budget calendar for each state in our sample. Columns include the fiscal-year start date and the typical month in which the governor submits the initial budget proposal. These dates define the planning, deliberation, and execution 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.

The NASBO snapshots are available for 1987, 1989, 1992, 1997, 1999, 2002, 2008, 2015, and 2021. Our baseline uses a piecewise-constant calendar: for each fiscal year, we assign the most recent NASBO calendar observed at or before that year and carry it forward until a newer snapshot is available. This avoids using future calendar information while preserving low-frequency institutional changes in budget timing. One exception is Arkansas: the 1987–2002 NASBO waves report the governor's submission as "Sept./Dec.," whereas the 2008–2021 waves report November for the same institution. We code Arkansas as a November proposer throughout, so its budget calendar does not shift artificially when NASBO rewords the entry.

Beyond timing, the table reports two cycle codings from the same NASBO source, taken as of the 2021 survey wave: the "budget cycle", whether the state enacts a budget every year or a single two-year budget (e.g., Texas and Indiana), and the "legislative cycle", whether the legislature holds a regular session every year or every other year (e.g., Montana and Nevada, whose legislatures convene only in odd-numbered years). Twenty states enact a single two-year budget, but most of these have an annual legislative cycle and thus revise the budget annually through supplemental appropriations. For example, Connecticut and Hawaii consider their budget cycles effectively annual in practice, despite officially having biennial budgeting (National Association of State Budget Officers, 2021). In fact, only four states—Montana, Nevada, North Dakota, and Texas—have both a biennial budget and a legislature that convenes only every other year, so that no official annual budget decision occurs in the second year of the biennium. Our baseline analysis therefore keeps all state-

years, using each state's imputed calendar as constructed; Section 3 and Appendix E.4 show that the decision-window results are essentially unchanged when second-year cohorts are excluded, re-dated, or interacted separately. Note that under GASB 34, amendments authorized before a fiscal year begins are counted in that year's original budget, so even the second year of a biennium can include relevant revisions in its "Original Budget."

Separately, four states—Wyoming, Washington, North Dakota, and Oregon—are outside the sample entirely: their ACFR budgetary comparison schedules report the two-year biennium as a whole (Washington's fiscal 2021 report, for example, presents a single "Original Budget 2019–21 Biennium" column), so the annual budget objects our outcomes require do not exist for them.

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 A.1 reproduces a representative budgetary comparison statement. This statement reports the original budget (first enacted appropriation, including pre-fiscal-year amendments), the final revised budget (original budget plus all mid-year amendments through the close of the fiscal year), and actual amounts for each major revenue and expenditure category. GASB Statement No. 34 enforced this structure beginning in fiscal years 2001-2003, where smaller states were allowed later adoption of the standard. We hand-collected and digitized analogous tables from 47 states over fiscal years 1995–2024. As discussed in Section 1, our analysis focuses on the original budget ("initial budget") and actual outcomes; the final revised budget is reported but not used as a primary analytical object because it accumulates amendments of heterogeneous institutional origin throughout the fiscal year.

A.3 Comparison with the NASBO Fiscal Survey

Much of the state-fiscal literature measures budgets with the NASBO Fiscal Survey of States, a semiannual survey of state budget offices reporting general-fund revenues and expenditures. The survey is concucted in common calendar windows: roughly February–May for the spring

Table A.1: **State Budget Calendar Timing (2021).** This table summarizes each state's fiscal-year start date, the typical range of months when initial and final budget proposals are submitted, and the legislative and budget cycles (A = annual, B = biennial). All columns reflect the 2021 NASBO snapshot, the most recent available.

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

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.

edition and August–October for the fall edition, rather than at each state's proposal, fiscal-year-close date. As a result, the budget and actual figures reported in a given survey wave can correspond to different budget vintages across states, depending on where each state is in its budget calendar. We instead build fiscal amounts from ACFR budgetary-comparison schedules. This choice is important for two reasons. First, the ACFR reports the original enacted budget, the final amended budget, and actual outcomes on a common audited budgetary basis, which is the decomposition our design requires. Second, the two sources measure different fiscal aggregates: NASBO reports state-defined general funds, while ACFR budgetary comparisons follow the appropriated budgetary basis used in each state's audited financial statements.

Table A.2 and Figure A.2 compare the two for the totals both report, over the estimation sample. The audited ACFR totals and the NASBO general-fund totals are strongly correlated in logs (about 0.93), but they do not match: only about half of state-years agree within ten percent, and year-over-year growth rates correlate only 0.38 for expenditure and 0.51 for revenue. The gaps are structural, reflecting differences in fund coverage across states: Michigan, for example, routes most spending through the School Aid Fund and other restricted funds, so its NASBO general fund ($9 billion in FY2015) is a fraction of the appropriated budget the ACFR reports ($33 billion), while the ratio runs the other way for states such as Vermont and New York.

We therefore use each source where it is authoritative: NASBO for the budget calendar (proposal and enactment dates and annual-versus-biennial cycles, as in Appendix A.1), and the hand-collected ACFR schedules for the fiscal amounts.

B Stylized Facts and the Commitment Margin

This appendix follows the order of Section 1.3. It first documents the budget benchmark used in the decomposition, then reports the distribution and variance-decomposition evidence, and finally shows how revenue surprises map into aggregate spending, spending categories, and rainy-day funds.

Table A.2: **ACFR versus NASBO: Budget-Data Comparison**

The table compares realized general-fund totals from the audited ACFR budgetary-comparison schedules (this paper) with the NASBO Fiscal Survey of States, for the totals both sources report, over fiscal years 1995–2024. Amounts are placed in common units ($ millions; ACFR is reported in thousands, NASBO in millions), and the sample is state-years with a positive actual value in both sources. "Log-level correlation" is the correlation of log actual amounts; the median ratio and percentiles refer to the ACFR/NASBO ratio of actual amounts; "Within ±10% of parity" is the share of state-years with that ratio in $[0.9, 1.1]$; "Growth-rate correlation" is the correlation of year-over-year log changes, excluding changes exceeding one log point (coverage breaks rather than fiscal news).

	Expenditure	Revenue
Matched state-years	1,119	1,034
States	46	45
Log-level correlation	0.93	0.94
Median ratio (ACFR/NASBO)	1.03	1.04
Ratio, 10th–90th pctile	0.90–2.46	0.96–2.30
Within ±10% of parity	46%	49%
Growth-rate correlation	0.38	0.51

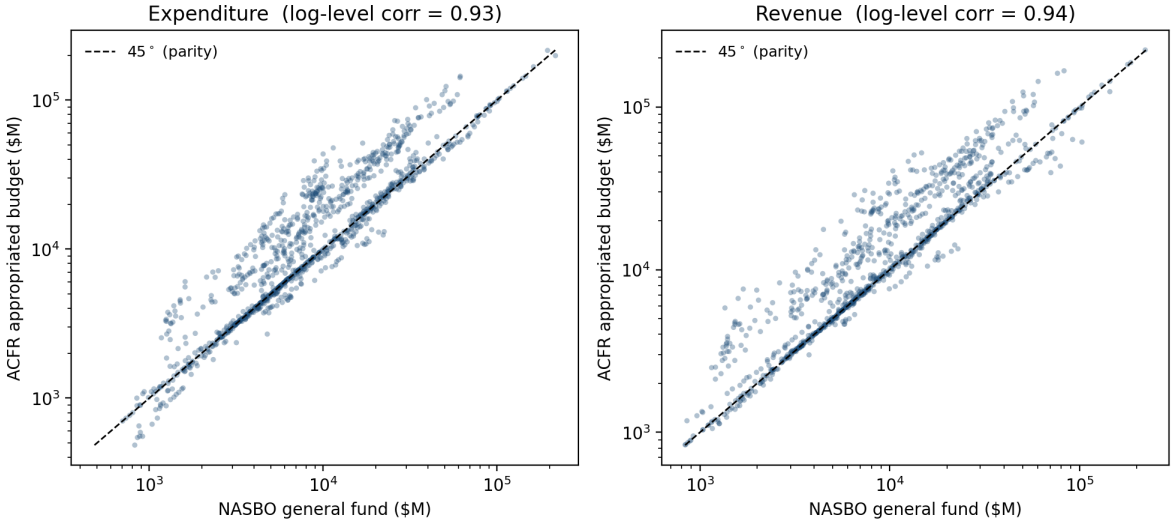


Figure A.2: **ACFR versus NASBO Actual Totals.** Each point is a state-fiscal-year, plotting the realized general-fund total from the NASBO Fiscal Survey (horizontal axis) against the audited ACFR appropriated-budget total (vertical axis), both in $ millions on a log scale, over fiscal years 1995–2024. The dashed line is parity. Left: expenditure; right: revenue. The two are strongly correlated but diverge in level for states whose general fund covers a different share of the appropriated budget.

B.1 Anchoring of the Initial Budget

A key assumption underlying our decomposition is that the prior year's initial budget, $\text{Orig}_{s,t-1}$, is the relevant reference point when the initial budget is drafted. Table B.1 provides direct evidence. We regress the log initial budget on $\log(\text{Orig}_{t-1})$ and $\log(\text{Act}_{t-2})$, the two most recent available benchmarks. This supports the use of Orig_{t-1} as the baseline in the planning component: conditional on the prior initial budget, the two-year-lagged actual adds little explanatory power.

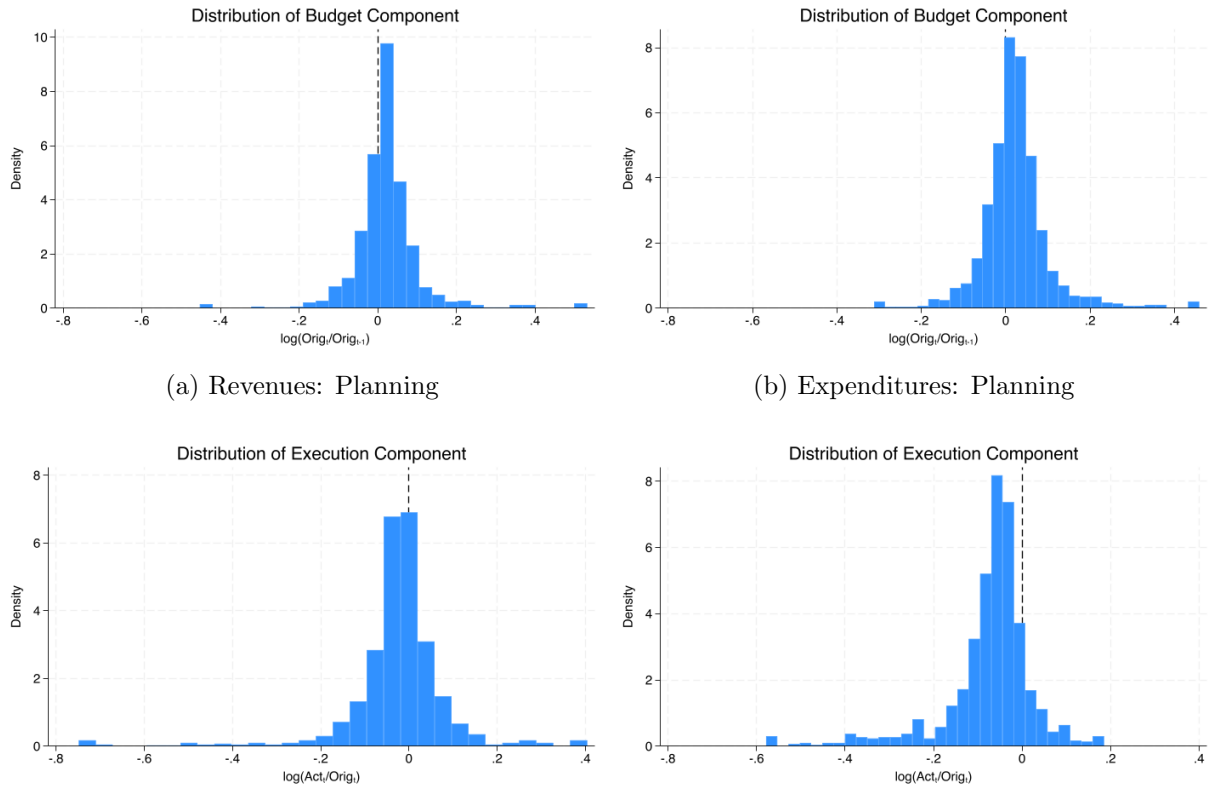
Table B.1: **Anchoring of the Initial Budget**

The table reports OLS regressions of the log initial-budget revenue level, $\log(\text{Orig}^R_{s,t})$, on lagged budget benchmarks: the prior year's initial budget, $\log(\text{Orig}^R_{s,t-1})$, and the two-year-lagged actual realization, $\log(\text{Act}^R_{s,t-2})$. The sample is an annual state-fiscal-year panel covering 47 annual-budget states and fiscal years 1995–2024. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue: $\log(\text{Orig}^R_t)$		
	(1)	(2)	(3)
$\log(\text{Orig}_{t-1})$	0.70***		0.56***
	(0.09)		(0.17)
$\log(\text{Act}_{t-2})$		0.60***	0.20
		(0.09)	(0.15)
State + FYear FE	✓	✓	✓
Observations	990	958	954
Adj.R^2	0.990	0.988	0.990

B.2 Decomposing Budget-Cycle Variation

Distribution of budget-cycle adjustments. Figure B.1 displays the empirical distributions of revenue and expenditure changes across the planning and execution phases of the budget cycle defined in Equation (1). 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 initial budget. Second, the execution-phase distributions are left-skewed: initial budgeted expenditures exceed actual outlays in nearly all observations. This systematic overestimation reflects the conservatism inherent in balanced-budget requirements and executive spending controls.



(a) Revenues: Planning

(b) Expenditures: Planning

(c) Revenues: Execution

(d) Expenditures: Execution

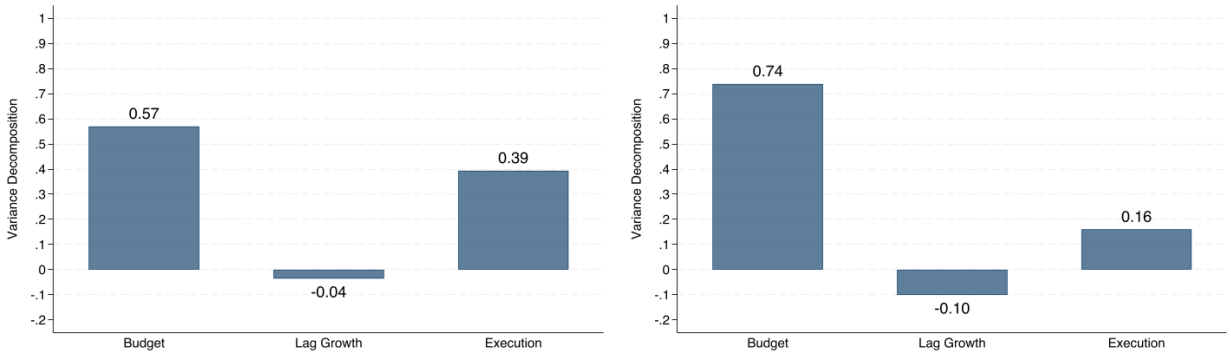
Figure B.1: **Distribution of State Revenue and Expenditure Changes over the Budget Cycle.** The figure shows the distributions of the planning and execution components in Equation (1) for both revenues and expenditures.

Variance decomposition. To quantify the contribution of each component, we decompose the variance of annual changes in actual outcomes using Equation (1). We first demean each variable within state: $\widetilde{X}_{s,t} \equiv X_{s,t} - \bar{X}_s$, then decompose the variance in actual spending growth as follows:

$$1 = \frac{\operatorname{cov}\left(\widetilde{\log(\text{Orig}_{s,t}/\text{Orig}_{s,t-1})}, \widetilde{\Delta \log(\text{Act}_{s,t})}\right)}{\operatorname{var}\left(\widetilde{\Delta \log(\text{Act}_{s,t})}\right)} + \frac{\operatorname{cov}\left(\widetilde{\log(\text{Act}_{s,t}/\text{Orig}_{s,t})}, \widetilde{\Delta \log(\text{Act}_{s,t})}\right)}{\operatorname{var}\left(\widetilde{\Delta \log(\text{Act}_{s,t})}\right)} - \frac{\operatorname{cov}\left(\widetilde{\log(\text{Act}_{s,t-1}/\text{Orig}_{s,t-1})}, \widetilde{\Delta \log(\text{Act}_{s,t})}\right)}{\operatorname{var}\left(\widetilde{\Delta \log(\text{Act}_{s,t})}\right)}. \tag{B.1}$$

Equivalently, each share is the OLS slope from a regression of the corresponding component on $\Delta \log(\text{Act}_{s,t})$ with state fixed effects absorbed.

Alternative benchmark. The main-text variance decomposition in Figure 3 benchmarks the planning component to the prior year's initial budget, $\log(\text{Orig}_{s,t}/\text{Orig}_{s,t-1})$. An alternative is to benchmark to two-year-lagged actual outcomes, $\log(\text{Orig}_{s,t}/\text{Act}_{s,t-2})$, which avoids any mechanical dependence on the prior year's initial budget. Figure B.2 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 B.2: **Decomposing the Sources of Variation in Actual Growth (Alternative Benchmark).** This figure replicates Figure 3 using two-year-lagged actual outcomes, $\log(\text{Orig}_{s,t}/\text{Act}_{s,t-2})$, as the planning-stage benchmark. Components are: Planning $\log(\text{Orig}_{s,t}/\text{Act}_{s,t-2})$, Execution $\log(\text{Act}_{s,t}/\text{Orig}_{s,t})$, and Lag Growth $\log(\text{Act}_{s,t-1}/\text{Act}_{s,t-2})$.

B.3 Revenue Pass-Through to Spending

Revenue constraints and expenditure responses. Table B.2 quantifies the relationship between revenue and expenditure changes across the budget phases. Revenue changes explain a substantial share of the variation in expenditure adjustments, with particularly strong co-movement during the planning phase. 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 B.2: **Revenue Constraints and Expenditure Responses Across Budget Phases**

The table reports OLS regressions of phase-by-phase expenditure adjustments on the corresponding phase-by-phase revenue adjustment from the decomposition in Equation (1). The dependent variable in columns (1)–(2) is the budget component for total spending, $\log(\text{Init}^E_{s,t}/\text{Init}^E_{s,t-1})$; in columns (3)–(4) it is the execution component, $\log(\text{Act}^E_{s,t}/\text{Init}^E_{s,t})$. The regressor in each column is the corresponding revenue component of Equation (1). The sample is an annual state-fiscal-year panel covering 47 annual-budget states and fiscal years 1995–2024. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Total Spending			
	Budget $\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$		Execution $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t})$	
	(1)	(2)	(3)	(4)
Rev (Budget)	0.58***	0.56***		
	(0.06)	(0.07)		
Rev (Execution)			0.49***	0.28***
			(0.09)	(0.07)
State + FYear FE		✓		✓
Obs.	990	990	1,026	1,026
Adj.R^2	0.43	0.43	0.36	0.71

Asymmetric spending adjustment. Table B.3 documents asymmetric execution-stage spending adjustment to revenue surprises. In columns (1)–(2), revenue shortfalls are associated with substantial spending cuts, while windfalls have little effect on aggregate spending, consistent with balanced-budget enforcement and executive impoundment authority. With fiscal-year fixed effects in column (3), pass-through is moderate and statistically similar across signs: the windfall slope is 0.21, and the shortfall differential is 0.08 and insignificant. Thus, the asymmetry arises in within-state changes across fiscal years and is absorbed by fiscal-year-common components rather than within-year cross-state variation.

Table B.3: **Asymmetric Spending Adjustment to Revenue Shortfalls and Windfalls**

The table reports OLS regressions of the execution-phase expenditure adjustment, $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t})$, on the execution-phase revenue adjustment, $\log(\text{Act}^R_{s,t}/\text{Orig}^R_{s,t})$, an indicator for a within-year revenue shortfall, $\mathbf{1}\{\text{Rev shortfall}\} \equiv \mathbf{1}\{\log(\text{Act}^R_{s,t}/\text{Orig}^R_{s,t}) < 0\}$, and their interaction. Column (1) includes no fixed effects; column (2) adds state fixed effects; column (3) adds state and fiscal-year fixed effects, so its coefficients are identified from within-year cross-state variation only. The sample is an annual state-fiscal-year panel covering 47 annual-budget states and fiscal years 1995–2024. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

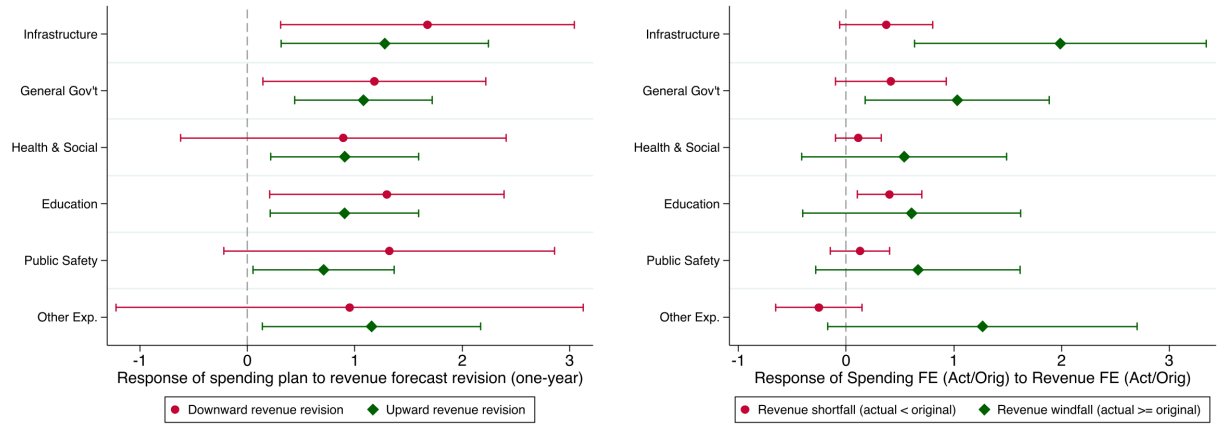
	Execution $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t})$		
	(1)	(2)	(3)
$\log(\text{Act}^R_{s,t}/\text{Orig}^R_{s,t}) \times 1(\text{Rev shortfall})$	0.54***	0.24**	0.08
	(0.15)	(0.11)	(0.10)
$\log(\text{Act}^R_{s,t}/\text{Orig}^R_{s,t})$	0.07	0.08	0.21*
	(0.15)	(0.10)	(0.11)
1(Rev shortfall)	-0.01	-0.01	-0.01
	(0.01)	(0.01)	(0.01)
State FE		✓	✓
FYear FE			✓
Obs.	1,026	1,026	1,026
Adj.R^2	0.41	0.67	0.71

Spending category adjustments. Our hand-collected schedules provide line-item spending detail, allowing us to trace budget and execution adjustments across categories. Panel (a) of Figure B.3 relates each category's original-budget growth, $\log(\text{Orig}^E_{\ell,t}/\text{Orig}^E_{\ell,t-1})$, to the revenue-forecast revision by sign. Forecast increases and decreases pass through similarly across categories; infrastructure is somewhat more responsive, but the difference is insignificant.

Panel (b) relates each category's execution deviation from budget to the revenue surprise, measured as realized revenue relative to the original forecast. Shortfalls are absorbed primarily through education cuts, whereas windfalls are concentrated in infrastructure and general government. Although the category-level windfall responses are larger in percentage terms, these categories account for only about 5 and 13 percent of total spending. Education is the largest category, so the smaller percentage shortfall response there has a larger effect on aggregate spending. The large windfall coefficients in Panel (b) therefore do not imply a larger aggregate spending response to windfalls.

Figure B.3: **Revenue Passes Through to Spending Plans Across Categories**

Each marker is the coefficient from a separate regression of one spending category's outcome on the corresponding total-revenue object, with the revenue term split by sign and 90% confidence intervals shown. Panel (a) regresses the one-year revision in a category's original spending plan, $\log(\text{Orig}^E_{\ell,t}/\text{Orig}^E_{\ell,t-1})$, on the one-year revision in the original total-revenue forecast, $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$, entered separately when that revision is downward (cranberry) and upward (green diamonds). Panel (b) regresses the within-year execution error in category spending, $\log(\text{Act}^E_{\ell,t}/\text{Orig}^E_{\ell,t})$, on the within-year total-revenue execution error, $\log(\text{Act}^R_t/\text{Orig}^R_t)$, entered separately for revenue shortfalls (cranberry) and windfalls (green). All regressions absorb state and fiscal-year fixed effects, cluster standard errors by state, and winsorize both variables at the 0.5/99.5 percentiles. The six categories are Infrastructure, General Government, Health & Social, Education, Public Safety, and Other. Sample: fiscal years ≥ 1995.



(a) Planning: spending-plan revision vs. revenue-forecast revision

(b) Execution: spending error vs. revenue error

B.4 Rainy Day Funds: Realized Surpluses vs Budget Surprises

Rainy-day-fund (RDF) balances could respond to two distinct objects, which the preceding tables make it important to separate. The first is the state's realized fiscal position: realized revenues minus realized expenditures, $\text{Act}^R_{s,t} - \text{Act}^E_{s,t}$, which can be in surplus or deficit however accurate the budget was. The second is budget surprise: the deviation of that realized surplus from the surplus written into the original budget: the same execution-stage margin studied in Table B.3. Table B.4 separates the two.

Reserves respond to positive realized surpluses. In columns (2)–(3), an additional dollar of realized surplus is associated with roughly three cents of rainy-day-fund accumulation.

The estimated deficit response is imprecise. Columns (4)–(6) show no statistically precise relationship between fund changes and surplus forecast errors. Thus, the evidence is consistent with rainy-day funds responding more to realized fiscal positions than to within-year deviations from the original budget, although the table does not separately identify the role of the planned surplus.

Table B.4: **Rainy Day Fund Response to Budget Surplus and Forecast Errors**

The table reports OLS regressions of the change in the state's rainy-day-fund balance scaled by lagged realized expenditure, $\Delta\text{RDF}_{s,t}/\text{Act}^{E}_{s,t-1}$, where $\Delta\text{RDF}_{s,t} \equiv \text{RDF}_{s,t} - \text{RDF}_{s,t-1}$. Columns (1)–(3) use the realized surplus, $\text{Surplus}_{s,t} \equiv (\text{Act}^{R}_{s,t} - \text{Act}^{E}_{s,t})/\text{Act}^{E}_{s,t-1}$; columns (4)–(6) use the surplus forecast error, the realized minus the original-budget surplus, $\text{FE}^{S}_{s,t} \equiv \text{Surplus}_{s,t} - \text{Surplus}^{\text{Orig}}_{s,t}$, where $\text{Surplus}^{\text{Orig}}_{s,t} \equiv (\text{Orig}^{R}_{s,t} - \text{Orig}^{E}_{s,t})/\text{Act}^{E}_{s,t-1}$. Columns (2) and (5) allow separate slopes by the sign of the regressor; columns (3) and (6) add state GDP growth as a control. All specifications include state and fiscal-year fixed effects. The outcome and all regressors are winsorized at the 1st and 99th percentiles. Unlike a log-change outcome, the scaled change is defined for state-years with zero fund balances, so these are retained in the sample. The sample is an annual state-fiscal-year panel covering 47 annual-budget states and fiscal years 1995–2024. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Dependent variable: $\Delta\text{RDF}_{s,t}/\text{Act}^{E}_{s,t-1}$					
	Realized surplus $(\text{Act}^{R}_{s,t}-\text{Act}^{E}_{s,t})/\text{Act}^{E}_{s,t-1}$			Surplus forecast error $\text{Surplus}_{s,t}-\text{Surplus}^{\text{Orig}}_{s,t}$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Surplus}_{s,t}$	0.062					
	(0.037)					
$\text{Surplus}_{s,t} \times 1(\text{Surplus}_{s,t} > 0)$		0.030**	0.029*			
		(0.015)	(0.015)			
$\text{Surplus}_{s,t} \times 1(\text{Surplus}_{s,t} < 0)$		0.126	0.119			
		(0.099)	(0.093)			
$\text{FE}^{S}_{s,t}$				0.015		
				(0.014)		
$\text{FE}^{S}_{s,t} \times 1(\text{FE}^{S}_{s,t} > 0)$					0.014	0.013
					(0.017)	(0.017)
$\text{FE}^{S}_{s,t} \times 1(\text{FE}^{S}_{s,t} < 0)$					0.025	0.019
					(0.063)	(0.058)
State GDP growth			0.131**			0.169*
			(0.056)			(0.093)
State + FYear FE	✓	✓	✓	✓	✓	✓
GDP control			✓			✓
Obs.	1,000	1,000	1,000	1,000	1,000	1,000
Adj.R^2	0.117	0.132	0.139	0.079	0.079	0.092

C Framework Predictions

This appendix records the forecast-incorporation predictions used in Section 2. It follows the framework section in the main text and precedes the empirical construction of the shift-share revenue-news measure.

Table C.1: **Forecast-Incorporation Patterns Under Alternative Models**

The table conditions on incrementally informative news, $\beta_w^{\text{actual}} \neq 0$. Budget loadings refer to the response of the original-budget revenue forecast; forecast-error loadings refer to the response of $R^{\text{actual}} - \widehat{R}^{\text{budget}}$. "Off-window" refers to pre-decision and deliberation windows outside the budget decision window.

Model	Budget loading	Forecast-error loading
Full information	$\beta_w^{\text{budget}} \simeq \beta_w^{\text{actual}}$	$\beta_w^{\text{fe}} \simeq 0$
Sticky information	Pre-decision news caught up at update	Pre-decision: $\beta_w^{\text{fe}} \simeq 0$
Rational inattention	$\lvert\beta_w^{\text{budget}}\rvert < \lvert\beta_w^{\text{actual}}\rvert$ for costly signals	Same direction as realized revenue
Overreaction	$\lvert\beta_w^{\text{budget}}\rvert > \lvert\beta_w^{\text{actual}}\rvert$ for recent/salient news	Opposite sign from realized revenue
Decision-window	High in decision window, low off-window	Low in decision window, positive off-window

Fiscal feedback to the revenue base. The framework outlined in the main text takes revenue expectations as a function of news, and not the government's actions. A natural concern is that realized revenue may itself respond to fiscal policy: if decision-window news raises spending, and spending raises the tax base, then realized revenue is not a pure measure of the news' fundamental revenue content. This appendix shows why that concern affects the interpretation of realized-revenue loadings but not the core forecast-error test.

Let realized revenue contain an autonomous component and a fiscal-feedback component,

$$R_{s,t}^{\text{actual}} = R_{s,t}^{0} + \kappa G_{s,t}^{\text{actual}}, \qquad \kappa \in [0, 1),$$

where $R_{s,t}^{0}$ is the revenue base absent same-year state-spending feedback and κ is the within-year own-revenue recovered per dollar of state spending. Let $\widehat{R}_{s,t}^{0} \equiv \mathbb{E}[R_{s,t}^{0} \mid \mathcal{A}_{s,t}^{\text{budget}}]$. If the original-budget revenue forecast internalizes the effects of the enacted spending plan, then

$$R_{s,t}^{\text{budget}} = \widehat{R}_{s,t}^{0} + \kappa G_{s,t}^{\text{budget}}.$$

In the benchmark case in which the enacted spending plan is realized, $G_{s,t}^{\text{actual}} = G_{s,t}^{\text{budget}}$, the feedback term cancels from the forecast error:

$$R_{s,t}^{\text{actual}} - R_{s,t}^{\text{budget}} = R_{s,t}^{0} - \widehat{R}_{s,t}^{0}.$$

Thus forecast-error predictability continues to measure information omitted from the budget forecast.

Fiscal feedback does change the interpretation of realized-revenue loadings. For news in window w,

$$\frac{\partial R_{s,t}^{\text{actual}}}{\partial Z_{s,t,w}} = \frac{\partial R_{s,t}^{0}}{\partial Z_{s,t,w}} + \kappa \frac{\partial G_{s,t}^{\text{actual}}}{\partial Z_{s,t,w}}.$$

The realized-revenue response therefore combines fundamental revenue content with any fiscal response generated by the budget. A larger realized-revenue loading for decision-window news need not mean that decision-window news is intrinsically more informative; it may partly reflect the fact that only decision-window news changes current spending. This is why the empirical test emphasizes the decomposition between the original-budget forecast and the forecast error, rather than equality of realized-revenue coefficients across windows.

If official forecasts do not incorporate the feedback from the spending plan, then the cancellation is no longer exact. In that case, decision-window spending raises realized revenue without being fully embedded in the forecast, causing decision-window news to predict positive forecast errors. This attenuates the decision-window forecast-error contrast toward zero and would work against us.

D Shift-Share Measure: Construction Details

This appendix provides the implementation details needed to reproduce the shift-share news measure: the national source series paired to each revenue category, the construction of monthly innovations, the predetermined state exposure weights, and the scaling that puts category shocks into common revenue-relevance units. We avoid restating the identification argument here; the purpose of this appendix is measurement.

D.1 Revenue Categories and National Sources

Table D.1 reports the exact national source series used in the baseline shift-share measure. The five categories are the own-source revenue bases that both vary substantially across states and have clear national tax-base proxies. Intergovernmental transfers are excluded because they are driven primarily by federal policy rather than by state tax-base exposure; property taxes are excluded because the national house-price proxy has little aggregate predictive content for state property-tax revenues.

Table D.1: **Baseline Revenue Categories and National Source Series**

Revenue category	National source series	FRED / NIPA series
Personal income tax	Personal income	PI
General sales tax	Personal consumption expenditures	PCE
Charges and misc. revenue	PCE services	PCES
Corporate income tax	Industrial production	INDPRO
Severance tax	WTI crude oil price	WTISPLC

The table lists the five own-source revenue categories used to construct the baseline state-specific shift-share news measure $\text{News}_{s,t,m}$ together with the national source series paired to each category. Source series are pulled from FRED / NIPA at monthly frequency where available and converted to monthly log changes, then residualized on a series-specific AR(4) and standardized as described in Section 3.2.

Let $x_{k,m}$ denote that source series in month m. We construct the category-level innovation directly from this paired source: convert the series to monthly frequency where needed, take monthly log changes, residualize those changes on four monthly lags of the same series, and standardize the residuals to unit variance. Denoting the raw monthly change by $\Delta x_{k,m}$, the monthly innovation for revenue category k is

$$z_{k,m} = \frac{\hat{u}_{k,m} - \overline{\hat{u}}_k}{\text{sd}(\hat{u}_k)}, \qquad \hat{u}_{k,m} \text{ from } \Delta x_{k,m} = \alpha_k + \sum_{r=1}^{4} \rho_{kr} \Delta x_{k,m-r} + \hat{u}_{k,m}, \tag{D.1}$$

Thus the baseline innovation captures news relative to recent national dynamics, not the predictable AR component of the national tax base. We do not partial out calendar-month dummies in the production shock. Instead, because proposal dates cluster in a few calendar months, Appendix Table E.12 directly tests whether exposure-specific calendar seasonality explains the decision-window premium. The filtered series are shown in Figure D.1.

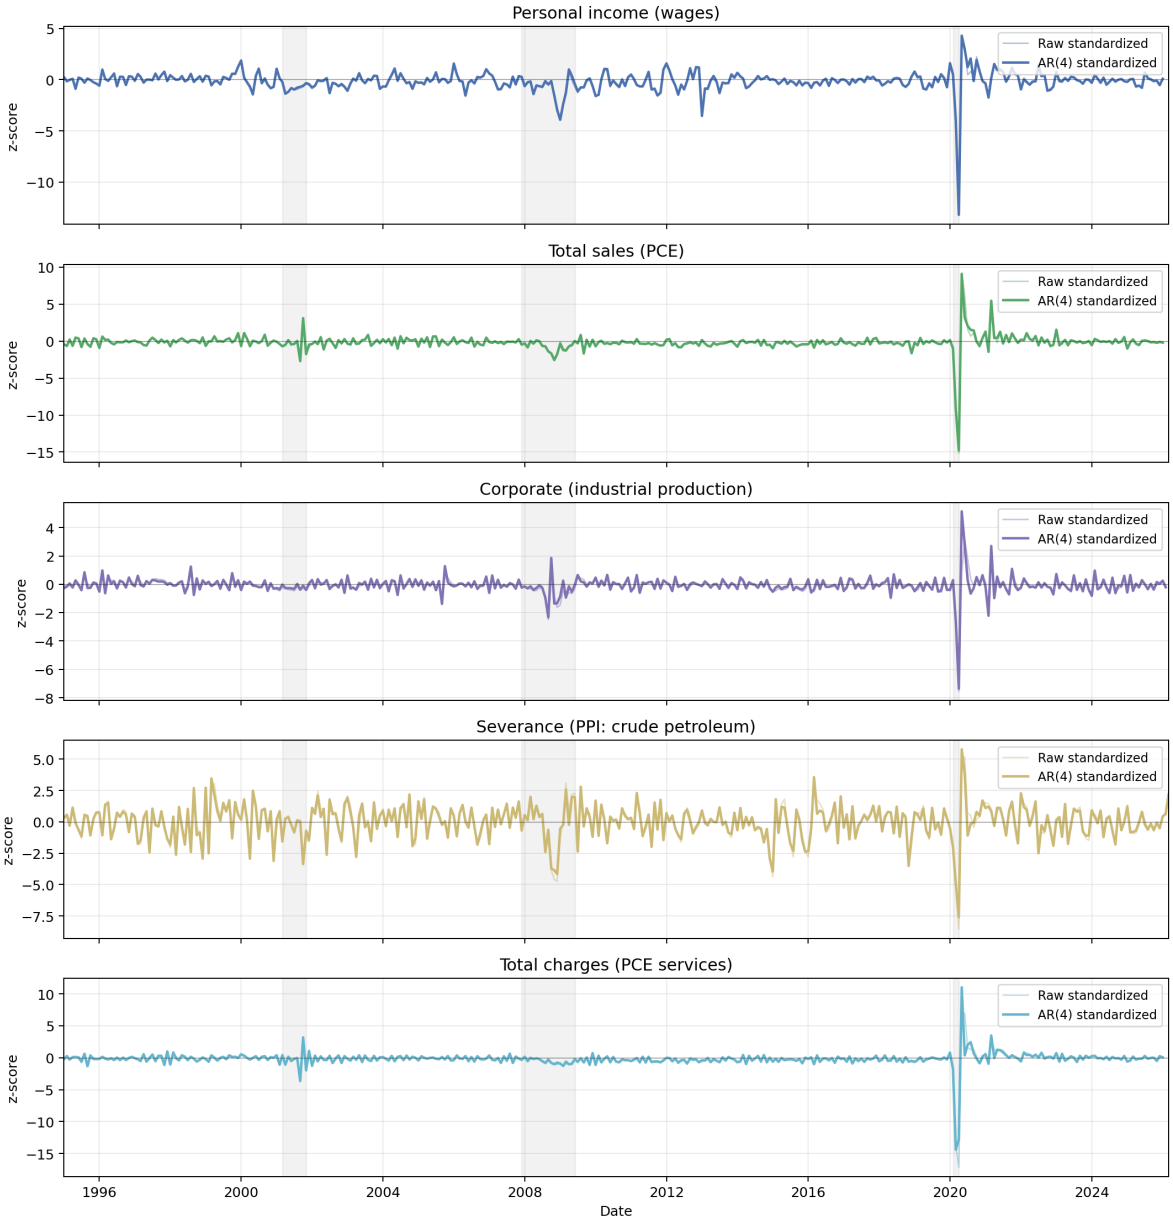


Figure D.1: **Standardized Monthly Source Innovations.** Each panel plots the raw monthly source change and the corresponding standardized innovation residualized on four own lags. The figure illustrates the source-level normalization step that makes innovations comparable in dispersion before category aggregation.

D.2 Exposure Weights and Elasticity Scaling

State exposure to revenue category k is measured using a five-year trailing average ending in fiscal year $t-2$:

$$\omega_{s,k,t} = \frac{1}{5} \sum_{\tau=t-6}^{t-2} \frac{\text{Revenue}_{s,k,\tau}}{\text{General Revenue}_{s,\tau}}. \tag{D.2}$$

Three implementation choices are important. First, the denominator is total general revenue, not the sum of the included categories. Thus the included exposure weights are not renormalized to sum to one; states with large transfer shares have mechanically lower exposure to own-source revenue news. Second, the $t-2$ endpoint makes the weights predetermined with respect to the monthly news entering the fiscal-year-t budget cycle. Third, cross-state heterogeneity in the shift-share measure comes from these exposure shares, while the monthly source innovations are common national shocks. Figure D.3 shows the broad-category revenue composition underlying this variation.

The standardized source innovations have common variance but different fiscal content. To put categories into common predicted-revenue units, we first aggregate the category-level monthly innovations into fiscal-year category shocks,

$$\text{Shock}^{US}_{k,\tau} = \sum_{m \in FY_\tau} z_{k,m}. \tag{D.3}$$

We then estimate a category-specific scaling coefficient from aggregate revenue data over the full sample of fiscal years:

$$\Delta \log \text{Rev}^{US}_{k,\tau} = \eta_k \, \text{Shock}^{US}_{k,\tau} + u_{k,\tau}. \tag{D.4}$$

The fitted slope $\hat{\eta}_k$ is a single national, category-level constant, applied uniformly across states and years. The scaling is estimated on aggregate category-specific revenue growth intentionally: state-level heterogeneity is already captured by $\omega_{s,k,t}$, while $\hat{\eta}_k$ only changes the relative units of the national source innovations before they are interacted with state exposure.

Figure D.4 reports the fitted coefficients $\hat{\eta}_k$ and the explanatory power of the category-level scaling regressions. All five coefficients are positive; severance and corporate income carry the largest $\hat{\eta}_k$ (their revenues are the most sensitive per unit shock), while sales, severance, and charges are the most predictable in R^2 terms. Personal income and corporate income have lower R^2s, reflecting the roles of discretionary tax-policy changes and the

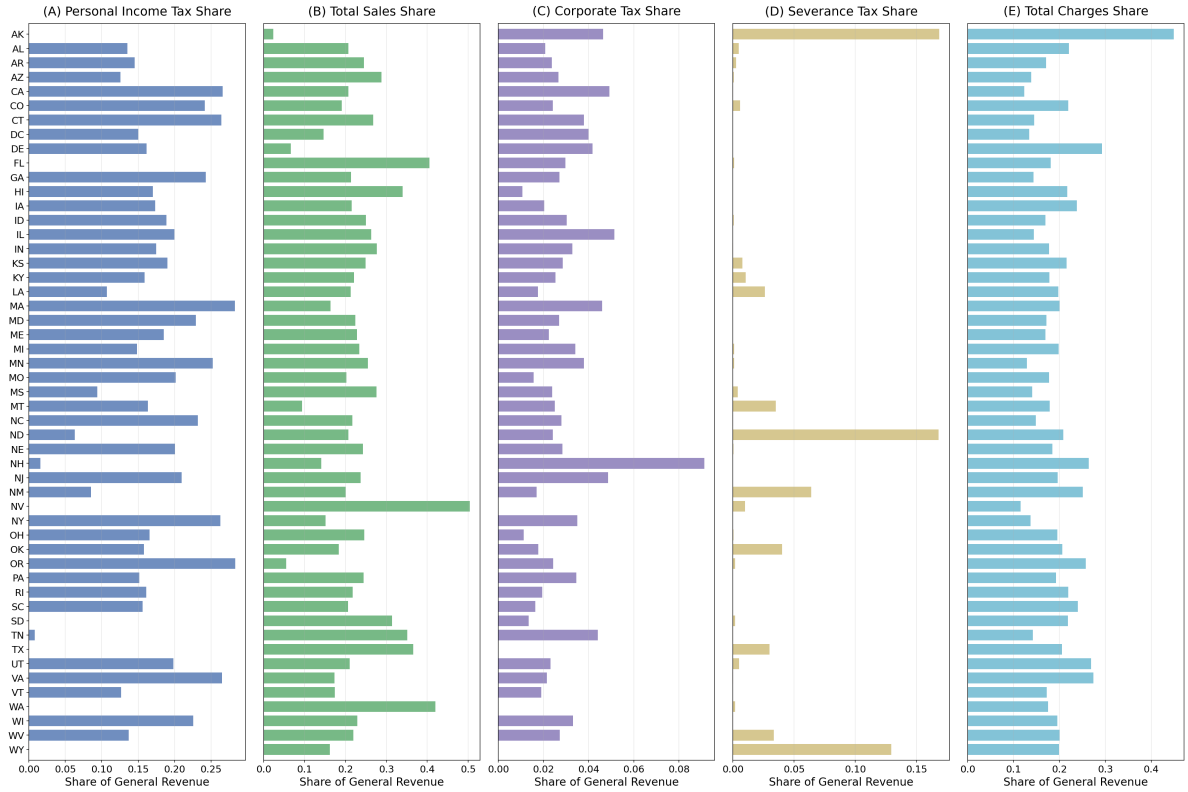


Figure D.2: **Cross-State Variation in Revenue Composition.** Each panel displays the average share of a revenue source in total general revenue for each state. Panel (A): individual income tax; Panel (B): general sales tax; Panel (C): charges and miscellaneous revenue; Panel (D): corporate income tax; Panel (E): severance tax. States are sorted alphabetically in all panels. Data source: Census Bureau Annual Surveys of State and Local Government Finance.

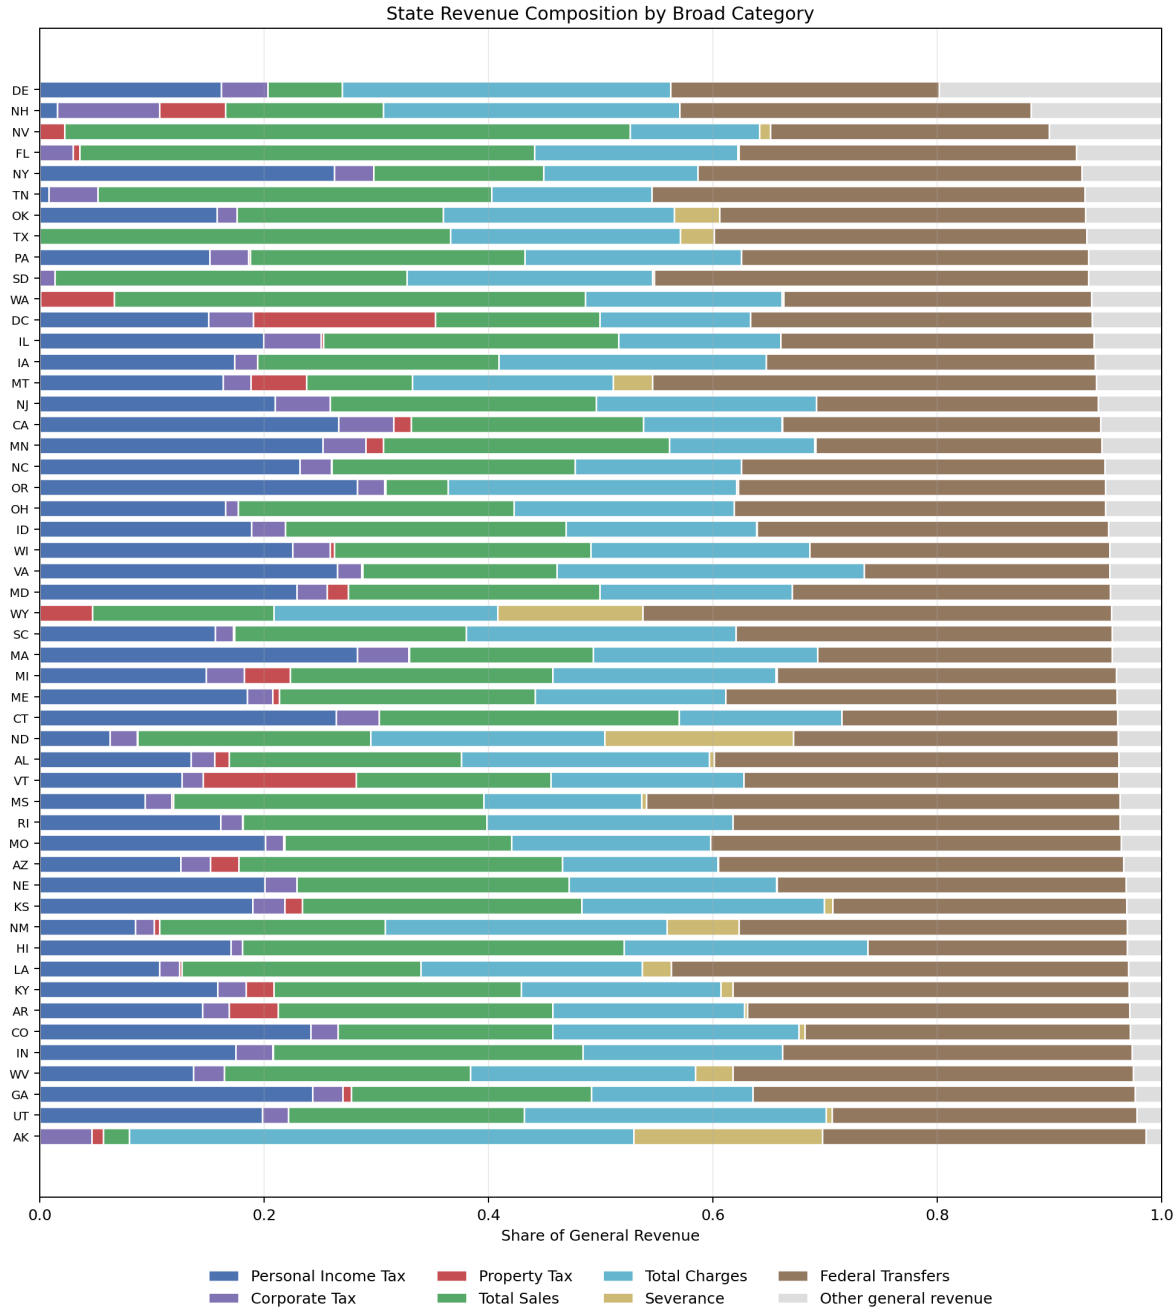


Figure D.3: **State Revenue Composition by Broad Category.** Each row shows one state's average share of general revenue across broad revenue categories over fiscal years 1995–2024. States are sorted by the residual "other general revenue" share. The five own-source categories entering the baseline shift-share (personal income, sales, corporate, severance, and total charges and miscellaneous) together average roughly 62% of general revenue. Property tax is displayed for descriptive completeness but is not used in the Bartik instrument because its aggregate R^2 against the Case-Shiller national HPI is near zero. Data source: Census Bureau Annual Surveys of State and Local Government Finance.

volatility of profits in those bases.

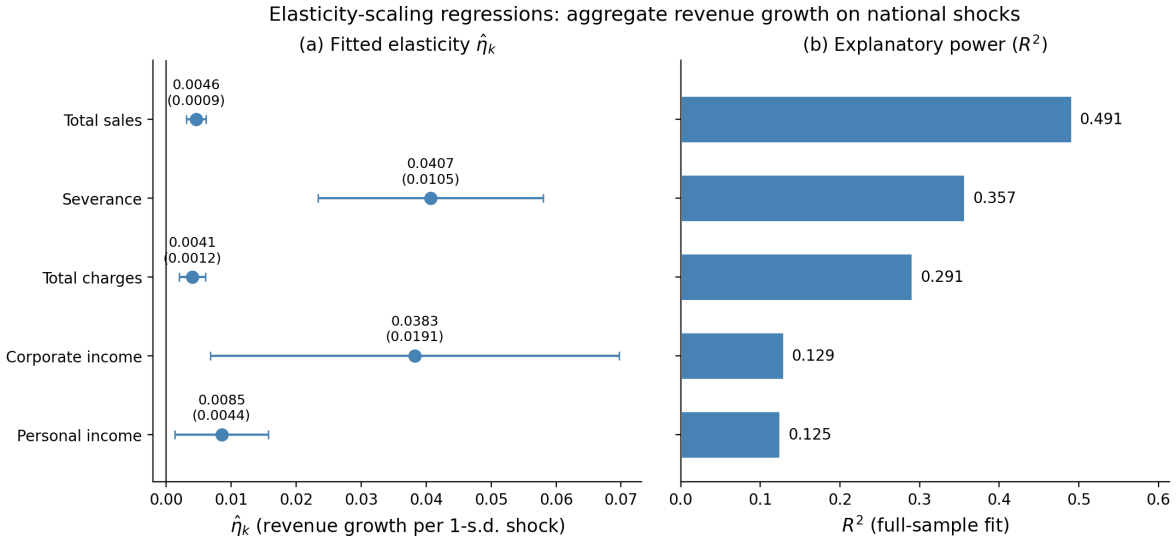


Figure D.4: **Elasticity-Scaling Regressions: Coefficients and Explanatory Power.** Both panels summarize the univariate category-level regressions of aggregate national revenue growth on the corresponding fiscal-year national tax-base shock (equation 10), estimated over the full sample of fiscal years. Panel (a) reports the fitted scaling coefficient $\hat{\eta}_k$ for each revenue category—the predicted category-revenue growth per one-standard-deviation shock—with 90% confidence intervals (horizontal whiskers) and the point estimate and its standard error printed above each marker. All five intervals exclude zero. Panel (b) reports the R^2 of the same regressions. Categories are ordered by R^2 in both panels. Shocks are AR(4)-residualized and z-scored at monthly frequency, then summed to fiscal-year frequency before estimation.

Combining standardized monthly innovations, predetermined state exposure weights, and category-specific scaling coefficients gives the monthly state-specific news measure used in Section 3.2:

$$\text{News}_{s,t,m} = \sum_{k=1}^{K} \hat{\eta}_k \, \omega_{s,k,t} \, z_{k,m}. \tag{D.5}$$

This is the same object as equation (8) in the main text. The proposal-relative windows used in the regressions do not create new shock objects; they classify this same monthly news series by each month's distance from the state's proposal date.

Table D.2: **Monthly Shift-Share News by Budget Window, Excluding COVID**

The table reproduces Table 1, Panel B, the distribution of the monthly state-specific news $\text{News}_{s,t,m}$ ($\times 100$) by budget window, excluding calendar years 2020–2021. The exclusion is by calendar rather than fiscal year because the stacked monthly panel assigns a single calendar-month shock to several fiscal-year cohorts at different event times, so dropping COVID fiscal years would leave the spring-2020 innovations in adjacent cohorts' pre-decision and deliberation windows. Sample and construction are otherwise identical to Table 1.

	Mean	SD	P5	P25	Median	P75	P95
News Pre-Decision.	-0.06	0.23	-0.31	-0.14	-0.06	0.02	0.22
News Decision	-0.08	0.26	-0.50	-0.14	-0.07	0.02	0.22
News Deliberation	-0.06	0.26	-0.34	-0.14	-0.06	0.02	0.25
News Execution	-0.06	0.26	-0.35	-0.14	-0.06	0.03	0.25

E Shift-Share Results: Diagnostics and Robustness

This appendix collects the results-side material for the shift-share analysis: diagnostics supporting the design and robustness of the decision-window premium. Construction details are in Appendix D.

E.1 Design Diagnostics

Here we document that the measure predicts realized state revenue growth, explain why the fixed-effect-rich validation is less precise, and report diagnostics supporting the predetermined-exposure interpretation.

Relevance and variance decomposition. Table E.1 regresses realized fiscal-year revenue growth, $\log(\text{Act}^R_{s,t}/\text{Act}^R_{s,t-1})$, on the fiscal-year news measure $\text{News}^{FY}_{s,t} \equiv \sum_{m \in FY_t} \text{News}_{s,t,m}$. Column (1) is the raw correlation; column (2) adds state fixed effects, lagged macro controls, and the Borusyak et al. (2022) aggregate-share control; column (3) further adds fiscal-year fixed effects. The point estimate is positive, large, and stable across columns: $\hat{\beta} = 1.69$ in column (2) and $\hat{\beta} = 1.90$ in column (3). The slope is at least as large under the more demanding identification; what changes is precision. The column (3) standard error is roughly three times larger, reducing the t-statistic from 3.5 to 1.2.

The mechanical reason is that the underlying Bartik shifters are national shocks, so most of the news measure's variation is common across states. A variance decomposition of $\text{News}^{FY}_{s,t}$ on the estimation sample confirms this: state fixed effects explain only 1.2% of its

Table E.1: **Validation: Shift-Share Measure Predicts Realized State Revenues**

The table reports OLS regressions of realized fiscal-year revenue growth, $\log(\text{Act}^R_{s,t}/\text{Act}^R_{s,t-1})$, on the fiscal-year-aggregated shift-share news measure $\text{News}^{FY}_{s,t}$, defined as the sum of the monthly state-specific news $\text{News}_{s,t,m}$ over the fiscal year. The sample is an annual state-fiscal-year panel covering 47 annual-budget states and fiscal years 1995–2024. $\text{News}_{s,t,m}$ is constructed in Section 3.2. Macro controls are four quarterly lag windows of CFNAI, payroll employment, S&P 500, five-year Treasury, and CPI. The aggregate-share control of Borusyak et al. (2022) is the unit-level sum of exposure shares across the five tax categories. Standard errors are clustered by state. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Realized revenue growth: $\log(\text{Act}^R_{s,t}/\text{Act}^R_{s,t-1})$		
	(1)	(2)	(3)
$\text{News}^{FY}_{s,t}$	1.48***	1.69***	1.90
	(0.42)	(0.48)	(1.59)
State FE		✓	✓
Fiscal-Year FE			✓
Macro Controls		✓	✓
Aggregate-share control (BHJ)		✓	✓
Observations	1,061	1,061	1,061
Adj. R^2	0.051	0.059	0.069

variance, fiscal-year fixed effects explain 82.1% (the common national component), and the remaining 17.3% is the within-state-year residual. The between-state component is absorbed by state fixed effects in both columns (2) and (3); the common-time component is absorbed by fiscal-year fixed effects in column (3) but used to identify column (2). Column (3) is therefore identified from the 17.3% within-state-year residual under the shock-exogeneity assumption of Borusyak et al. (2022). This residual is the same variation that identifies window-specific effects in the monthly panel of Section 3, where calendar-month-by-year fixed effects play the role of fiscal-year fixed effects.

Predetermined weights: tax policy does not respond to news. The predetermined interpretation of $\omega_{s,k,t}$ also requires that states do not respond to the same news by changing tax policy. We test this directly using Tax Foundation data on state individual income tax top marginal rates over 2000–2014. For each state-year, we construct indicators for any rate change, a rate increase, and a rate decrease, and regress each outcome on pre-decision, decision-window, and deliberation news with state and fiscal-year fixed effects, lagged macro controls, and state-clustered standard errors.

Table E.2 reports the results. Neither decision-window, pre-decision, nor deliberation news predicts tax-policy changes across the three outcomes. The coefficients are small,

statistically insignificant, and economically negligible, supporting the interpretation that the exposure weights are predetermined in practice as well as by construction.

Table E.2: **Bartik News and State Tax Policy Changes**

The dependent variable is an indicator equal to one if the state's individual income tax top marginal rate changed (column 1), increased (column 2), or decreased (column 3) in fiscal year t. The regressors are the fiscal-year aggregates of the monthly news over the decision window, $\text{News}_{s,t}^{\text{Decision}}$, pre-decision window, $\text{News}_{s,t}^{\text{Pre}}$, and deliberation window, $\text{News}_{s,t}^{\text{Delib}}$. Specifications include state and fiscal-year fixed effects and standard macro controls (four quarterly lag windows of CFNAI, payroll employment, the S&P 500, the five-year Treasury yield, and CPI). The sample is restricted to state-years with Tax Foundation top-rate data, 2000–2014. Standard errors are clustered by state. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	(1) Any change	(2) Increase	(3) Decrease
$\text{News}_{s,t}^{\text{Decision}}$	3.961	-0.127	4.088
	(3.922)	(0.821)	(4.044)
$\text{News}_{s,t}^{\text{Pre}}$	-0.343	-0.274	-0.068
	(2.553)	(1.163)	(2.309)
$\text{News}_{s,t}^{\text{Delib}}$	6.342	0.185	6.156
	(4.588)	(1.903)	(4.568)
State FE	✓	✓	✓
Fiscal-year FE	✓	✓	✓
Macro controls	✓	✓	✓
Observations	627	627	627
Adj. R^2	0.154	0.067	0.124

E.2 Distinguishing the Decision Window from Alternative Models

Under-incorporation of off-window news could in principle reflect information frictions rather than budget timing. Appendix Table C.1 summarizes the main alternatives: sticky information (Mankiw and Reis, 2002), rational or endogenous attention (Sims, 2003; Maćkowiak and Wiederholt, 2009), and overreaction or diagnostic expectations (Bordalo et al., 2018). These models tie incorporation to news age, processing cost, or salience; the decision-window hypothesis instead predicts a discrete jump at the budget calendar.

The window-level results in Table 2 already reject the simplest versions of these alternatives. Pre-decision news is available before the forecast is updated but is not caught up, contrary to standard sticky information. Deliberation news is the newest pre-fiscal-year information but is not overweighted, contrary to simple recency or diagnostic expectations.

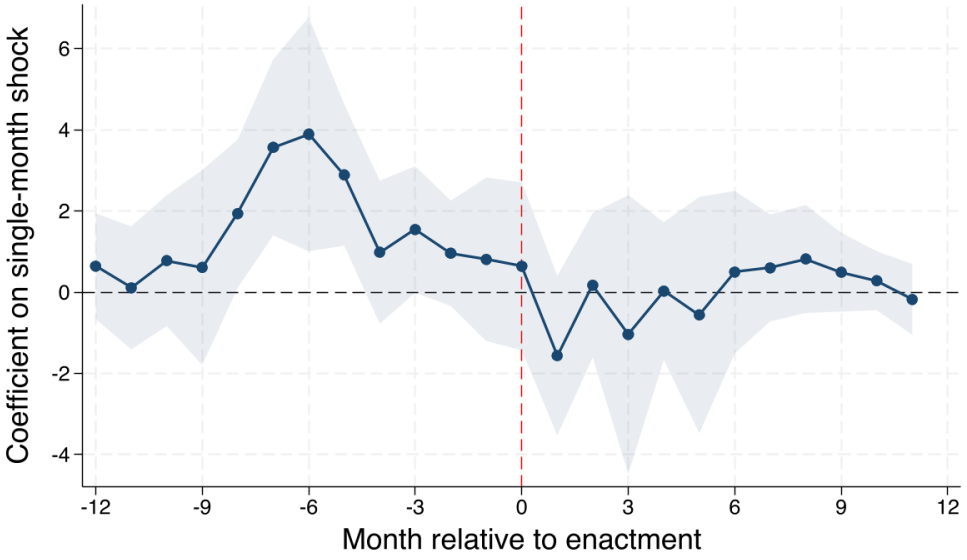


Figure E.1: **Monthly Profile of Revenue Forecast Updating Around the Enactment Date.** The figure plots coefficients from the same stacked monthly specification as Figure 4, but indexes event time relative to the legislative enactment date rather than the proposal date. The dependent variable is one-year original-budget revenue forecast growth. The specification includes state fixed effects, calendar-month-by-year fixed effects, lagged macro controls, and the aggregate-share control. Shaded areas denote 95 percent confidence intervals based on state-clustered standard errors.

The remaining concern is endogenous attention: officials might process off-window news only when it is unusually salient. The two tests below therefore ask whether incorporation varies with news age *within* the decision window, and whether unusually large shocks receive special attention outside it.

Decision window splits. Table E.3 splits the baseline window into early ($\ell \in [-5, -3]$) and late ($\ell \in [-2, -1]$) halves. The two forecast loadings are statistically indistinguishable ($p = 0.90$), so incorporation is flat within the window and drops at both boundaries—a step function in the budget calendar rather than a gradient in news age.

Table E.3: **Early versus Late Decision-Window News**

The table re-estimates Table 2, Panel A, splitting the decision window into an early half ($\ell \in [-5, -3]$) and a late half ($\ell \in [-2, -1]$), with the pre-decision ($\ell \in [-11, -6]$) and deliberation ($\ell \geq 0$ before fiscal-year start) terms unchanged. The "p: early = late" row reports the p-value of the Wald test that the two half-window loadings are equal. Outcomes, sample, fixed effects, controls, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	Forecast $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$	Forecast error $\log(\text{Act}_t^R/\text{Orig}_t^R)$	Realized $\log(\text{Act}_t^R/\text{Orig}_{t-1}^R)$
	(1)	(2)	(3)
$\text{News}_{s,t,m}^{\text{Pre}}$	0.03	1.08***	1.10
	(0.46)	(0.39)	(0.75)
$\text{News}_{s,t,m}^{\text{Decision},[-5,-3]}$	2.81***	-0.71	2.10**
	(0.74)	(0.65)	(0.95)
$\text{News}_{s,t,m}^{\text{Decision},[-2,-1]}$	2.63**	0.26	2.89**
	(1.05)	(0.64)	(1.32)
$\text{News}_{s,t,m}^{\text{Delib}}$	0.54	0.65*	1.20*
	(0.49)	(0.33)	(0.60)
p: early = late	0.896	0.153	0.660
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402
Adj. R^2	0.121	0.575	0.446

Large shocks and endogenous attention. An endogenous-attention model could also generate calendar-like incorporation if officials allocate attention around dated budget decisions. Such a model carries an additional prediction: unusually large pre-decision shocks should be more likely to attract attention and enter the forecast. Table E.4 tests this by

interacting each window's news with the absolute monthly news magnitude, $|\text{News}_{s,t,m}|$, standardized to unit standard deviation and centered at its sample mean. The pre-decision interaction is small and insignificant for both the forecast (−0.12, standard error 0.08) and the forecast error (−0.10, standard error 0.17): larger pre-decision shocks are incorporated no more than smaller ones. The evidence therefore points to calendar position rather than shock salience as the relevant margin of incorporation.

Table E.4: **Incorporation and Shock Size**

The table re-estimates Table 2, Panel A, adding for each window an interaction of the news term with the absolute monthly news magnitude $|\text{News}_{s,t,m}|$, standardized to unit standard deviation and centered at its mean over the estimation sample. Each window's base loading is thus the response at an average-magnitude shock, and the "× |News| (SD)" row is the change in that loading per one-standard-deviation-larger shock; a coefficient near zero means incorporation does not vary with shock size. Column 1 is the original-budget revenue forecast, column 2 the forecast error. Sample, fixed effects, controls, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1) Forecast	(2) Forecast error
$\text{News}^{\text{Pre}}_{s,t,m}$	0.83	1.84
	(0.81)	(1.20)
× |News| (SD)	-0.12	-0.10
	(0.08)	(0.17)
$\text{News}^{\text{Decision}}_{s,t,m}$	1.93*	1.38
	(1.10)	(1.41)
× |News| (SD)	0.28	-0.56*
	(0.26)	(0.33)
$\text{News}^{\text{Delib}}_{s,t,m}$	-0.12	0.66
	(0.84)	(0.88)
× |News| (SD)	0.09	-0.00
	(0.15)	(0.11)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Controls	✓	✓
Observations	16,402	16,402
Adj. R^2	0.121	0.575

E.3 Placebo and Inference

Pseudo-window placebo. The decision-window premium in Table 2 is identified from where a national innovation falls in the state's budget calendar. A residual concern is that the premium is a mechanical feature of the stacked monthly panel—that any five-month

block would appear to carry a premium regardless of its position. We rule this out with a pseudo-window placebo. Holding the estimation sample and the news series fixed, we relocate each state–fiscal-year's five-month "decision window" to a random block drawn uniformly from positions that are *disjoint* from the true window $\ell \in [-5, -1]$: a pre-decision block starting at $\ell = -11$ or -10 ($[-11, -7]$ or $[-10, -6]$), or a deliberation block starting at $\ell = 0, 1, 2$ ($[0, 4]$, $[1, 5]$, or $[2, 6]$). Draws are independent across state–fiscal-years. For each of 500 draws we re-estimate the *identical* in/out specification of Table 2—the news level plus its interaction with the relocated window—and record the estimated premium. Because the true decision window now falls inside the pooled news level, a correctly specified design should return a placebo premium near zero.

We relocate the window rather than permute proposal dates across states. Proposal *months* are tightly clustered (November–March), and the news is classified by calendar-month distance to the proposal; permuting proposal dates therefore leaves each state's implied window almost unchanged and still overlapping the true window, so a date-permutation placebo mechanically inherits the true signal. Relocating to a disjoint block instead places the window in the pre-decision or deliberation region, where off-window news does not enter the forecast (Table 2)—the relevant null.

Figure E.2 plots the distribution of the 500 placebo premiums. They are centered at -0.06 with a standard deviation of 0.25, and none reaches the true premium of 2.36; the permutation p-value is below $1/500 < 0.002$. The slight negative centering is expected—with the true decision window absorbed into the pooled level, an off-window block looks marginally weaker than the sample average. The true premium lies far outside the placebo distribution, so the decision-window result reflects the true budget calendar rather than a mechanical property of the stacked panel.

Shift-share inference with common national shifters. All cross-state variation in the news measure comes from predetermined exposure shares multiplied by five common national innovations. Residuals may therefore be correlated across states with similar revenue compositions, so conventional or two-way clustered standard errors can over-reject (Adão et al., 2019; Borusyak et al., 2022). Table E.5 reports inference diagnostics for the observed decision-window coefficient. The first rows report the baseline coefficient and two-way clustered standard error; the remaining rows summarize two alternative reference distributions.

The first reference distribution is a wild cluster bootstrap by calendar month-year, the

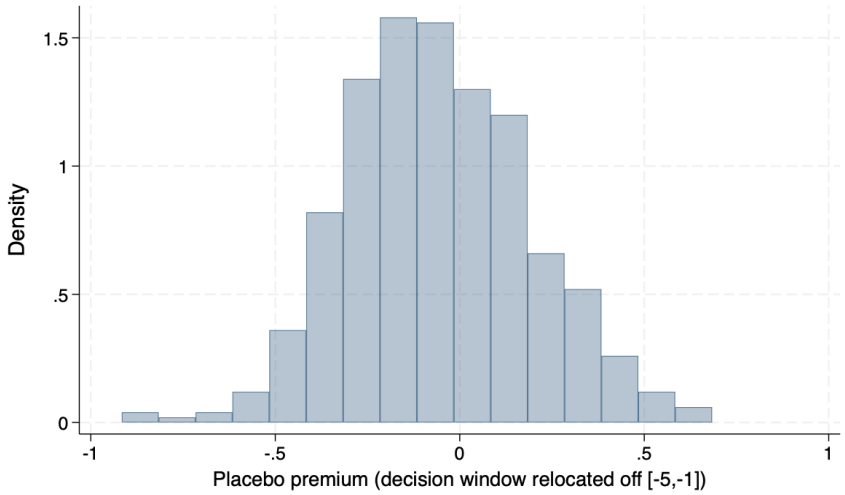


Figure E.2: **Pseudo-Window Placebo: Relocated Decision Windows Carry No Premium.** Distribution of the estimated decision-window premium across 500 placebo draws that relocate each state–fiscal-year's five-month window to a random block disjoint from the true window $\ell \in [-5, -1]$ (a pre-decision or deliberation block), holding the estimation sample and news series fixed and re-estimating the same in/out specification as Table 2. The red line marks the true premium (2.36). Placebo premiums center near zero and never reach the true value (permutation $p < 0.002$), indicating the premium reflects the true budget calendar rather than a mechanical feature of the stacked panel.

dimension on which the common shifters vary. The second is an exposure-vector permutation: we reassign the predetermined exposure vectors one-to-one across states, rebuild the news measure using the actual national innovations, and re-estimate the specification 500 times. This permutation holds each state's budget calendar fixed, unlike the pseudo-window placebo, and asks whether the observed premium is unusually large relative to random exposure-calendar alignment.

The forecast premium survives both checks. The wild-bootstrap p-value is 0.006; the exposure-permutation p-value is 0.014, with permuted coefficients centered at -0.25 and a 5th–95th percentile range of $[-1.8, 1.6]$, against the observed coefficient 2.71. The forecast-error coefficient remains a null under both diagnostics ($p = 0.91$ and 0.83). A shock-level correction in the spirit of Borusyak et al. (2022) is not one-to-one in this interacted-window design because exposure to each category-month shock depends on state-specific window timing; the exposure-vector permutation is therefore our design-specific finite-sample check for common-shifter leverage.

Table E.5: **Shift-Share Inference: Wild Bootstrap and Exposure Permutation**

The table reports inference diagnostics for the decision-window coefficient from the baseline three-window specification of Table 2, Panel A, for the forecast and forecast-error outcomes. The first two rows report the observed coefficient and its conventional standard error, two-way clustered by state and calendar month-year. "p: wild cluster bootstrap" is the wild cluster bootstrap p-value (999 replications) for the decision-window coefficient, clustered by calendar month-year and computed from the algebraically equivalent regression with explicit state and fiscal-year-by-calendar-month indicators. The permutation rows reassign entire predetermined exposure vectors one-to-one across the sample states, rebuild the news measure from the permuted exposures and the actual national innovations, and re-estimate the specification 500 times; $p = (1 + \#\{|\hat{\beta}^{perm}| \geq |\hat{\beta}|\})/(1 + 500)$, and the mean and 5th–95th percentile rows summarize the permutation distribution of the decision-window coefficient. Sample and specification are identical to Table 2.

	Forecast $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$	Forecast error $\log(\text{Act}_t^R/\text{Orig}_t^R)$
Observed decision-window coefficient	2.71	-0.16
Two-way clustered s.e.	(0.61)	(0.57)
p: wild cluster bootstrap (month-year)	0.006	0.911
p: exposure permutation	0.014	0.826
Permutation mean	-0.25	0.27
Permutation 5th–95th percentile	[-1.78, 1.58]	[-0.83, 1.35]
Permutation replications	500	500
Observations	16,402	16,402

Rotemberg weights and leverage. Table E.6 then asks which revenue categories drive the decision-window coefficient, following Goldsmith-Pinkham et al. (2020). The pooled coefficient decomposes as $\hat{\beta} = \sum_k \alpha_k \hat{\beta}_k$, where α_k is category k's share of the identifying variation in residualized decision-window news and $\hat{\beta}_k$ is the just-identified coefficient using category k's contribution alone. Severance accounts for the largest share of the weight, $\alpha = 0.65$, with $\hat{\beta}_k = 4.01$, close to the severance-only estimate in Table E.15. Personal income accounts for $\alpha = 0.23$, with $\hat{\beta}_k = 1.31$; the remaining categories each carry about four percent. The result therefore shows that identification leans heavily on the most exposure-differentiated base. For this reason, the severance-only and leave-one-out exercises reported below are important leverage checks.

E.4 Calendar and Sample Robustness

Fixed budget calendars. The baseline assigns each fiscal year the most recent NASBO calendar observed at or before that year, so low-frequency institutional changes in budget timing enter as they are reported. Reported proposal months are highly persistent within

Table E.6: **Rotemberg Weights: Which Revenue Categories Drive Identification**

The table decomposes the pooled decision-window coefficient of Table 2, Panel A, into revenue-category components following Goldsmith-Pinkham et al. (2020): $\hat{\beta} = \sum_k \alpha_k \hat{\beta}_k$, where $\alpha_k = \text{cov}(\tilde{c}_k, \tilde{N})/\text{var}(\tilde{N})$ is category k's Rotemberg weight, $\hat{\beta}_k = \text{cov}(\tilde{c}_k, \tilde{y})/\text{cov}(\tilde{c}_k, \tilde{N})$ is the just-identified coefficient using only category k's contribution, $\tilde{c}_k$ is category k's decision-window news contribution ($\hat{\eta}_k \omega_{s,k,t} z_{k,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$), $\tilde{N}$ is the composite decision-window news, and $\tilde{y}$ is the outcome, all residualized on the other window terms, controls, and fixed effects of the baseline specification. Weights sum to one by construction, and the final row verifies that the weighted components reproduce the pooled coefficients. Standard errors in parentheses are from the just-identified IV regression of the residualized outcome on the residualized composite decision-window news, instrumented by category k's contribution, two-way clustered by state and calendar month-year; $\hat{\beta}_k$ for low-weight categories rests on little identifying variation and should be read with the corresponding standard errors. Sample and specification are identical to Table 2.

Revenue category	Rotemberg weight α_k	β_k: Forecast	β_k: Forecast error
Personal income	0.227	1.31 (1.14)	1.79 (1.01)
General sales	0.043	-1.81 (3.59)	-4.05 (4.31)
Corporate income	0.043	-0.57 (2.69)	0.51 (3.50)
Severance	0.652	4.01 (0.53)	-0.31 (0.76)
Charges & misc.	0.035	-2.82 (0.66)	-6.12 (3.48)
$\sum_k \alpha_k \beta_k$ (pooled coefficient)	1.000	2.71	-0.16

state: 40 of the 51 states never change their reported proposal month across the NASBO waves covering the sample period, state fixed effects absorb 93 percent of the panel variation in proposal months, and Table E.7 lists the eleven states whose reported months change, typically by one adjacent month at a single survey wave. Table E.8 instead holds every state's proposal month fixed at its most recent (2021) NASBO wave for *all* sample years and re-dates the event window accordingly. The fixed calendar re-dates only seven states (Florida, Illinois, Montana, New Jersey, North Carolina, Rhode Island, and Tennessee), typically by one month. The decision-window forecast loading is 2.68 ($t = 4.4$), the pre-decision and deliberation loadings remain small, and the forecast-error mirror is unchanged: the results do not depend on the vintage of the calendar information.

Biennial budgets and legislative calendars. Not every sample state runs a fully annual budget process: fifteen states enact a single two-year budget, and for them the second year of each biennium has no budget proposal of its own. Appendix Table A.1 reports each state's budget and legislative cycle, and Appendix A details the institutions; the baseline keeps all state-years on their imputed calendars. This subsection shows that the results do not depend on that choice.

Table E.7: **Changes in Reported Proposal Months Across NASBO Waves**

The table lists every state whose reported proposal month changes across the NASBO survey waves assigned to fiscal years 1995–2024 under the baseline piecewise-constant calendar (each fiscal year uses the most recent wave observed at or before that year); the remaining 40 states keep the same reported month throughout. Each entry shows the assigned proposal month and the first fiscal year to which it applies.

State	Reported proposal month (first fiscal year assigned)
CO	Jan (FY1995) → Oct (FY1997)
FL	Dec (FY1995) → Jan (FY1997) → Feb (FY2008)
IL	Mar (FY1995) → Feb (FY1997)
MT	Jan (FY1995) → Nov (FY2008)
NC	Feb (FY1995) → Mar (FY2015)
NJ	Jan (FY1995) → Feb (FY2008)
OR	Dec (FY1995) → Jan (FY1997) → Nov (FY2008)
RI	Feb (FY1995) → Jan (FY2008)
TN	Jan (FY1995) → Feb (FY1999) → Jan (FY2008)
VA	Jan (FY1995) → Dec (FY1997)
WY	Dec (FY1995) → Nov (FY2015)

For biennial-budget states, a single enacted budget covers two fiscal years. The proposal and enactment occur before the first year, so the second year has no separate budget decision. Any decision window assigned to the second year is therefore imputed rather than institutional: the relevant biennium was proposed and enacted roughly twelve months earlier. Table E.9 shows that this mis-dating does not drive the baseline estimate, and that it is itself informative about the mechanism. Column (1) reproduces the baseline decision-window loading. Column (2) interacts every window-news term with a second-year indicator: the loading for cohorts with an actual decision in the window rises to 2.83, the second-year interaction is −3.40 (significant at 1%), and the implied second-year loading is −0.57—news assigned to a window containing no decision is not incorporated. The forecast-error column provides the mirror image: the same second-year news predicts forecast errors (implied loading 3.11, significant at 5%), so the information arrives and matters for revenue, but with no budget decision to absorb it, it surfaces as error instead. Column (3) drops second-year cohorts from the sample, and column (4) instead corrects the coding, relabeling all second-year pre-fiscal-year months as post-proposal (deliberation) months that serve as off-window controls. The decision-window loading is stable at 2.7–2.8 under every treatment, so we retain the baseline calendar coding as the paper's method and report these variants as robustness.

Table E.10 instead restricts the sample. Columns (1)–(2) drop the *strictly biennial* state-years, coded biennial on both margins: no budget document is written for the second year and no regular annual session exists to write one (although even here some second-

Table E.8: **Fixed 2021 NASBO Budget Calendar**

The table re-estimates the forecast and forecast-error columns of Table 2, Panel A, holding each state's proposal month fixed at its most recent (2021) NASBO survey wave for all sample years, instead of the baseline's piecewise-constant assignment of the most recent wave observed at or before each fiscal year. Relative event time and the three windows are re-dated from the fixed proposal month using the same conventions as the baseline construction (proposal months September–December are assigned to calendar year $t-2$, other months to $t-1$; decision window $[-5,-1]$, pre-decision $[-11,-6]$, deliberation from the proposal month to the fiscal-year start). The fixed calendar re-dates seven states (FL, IL, MT, NC, NJ, RI, TN), typically by one month; monthly observations required by a shifted window but outside the stored event-window panel are dropped (14 monthly rows). The news measure is unaffected by the calendar. Outcomes, fixed effects, controls, common-sample restriction, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1) Forecast	(2) Forecast error
$\text{News}^{\text{Pre}}_{s,t,m}$	0.01	1.07***
	(0.45)	(0.39)
$\text{News}^{\text{Decision}}_{s,t,m}$	2.68***	-0.15
	(0.62)	(0.56)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.56	0.64*
	(0.47)	(0.33)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Controls	✓	✓
Observations	16,393	16,393
Adj. R^2	0.121	0.575

year adjustment occurs—North Carolina's even-year "short session," for example, formally amends the enacted biennium). This drops Maine, Montana, Nevada, North Carolina, Texas, and Wisconsin in all sample years, Kentucky through fiscal year 2001 (its 2000 constitutional amendment added annual sessions), Ohio from fiscal year 2021, and Arkansas throughout (Arkansas does not enter the estimation sample in any case because its ACFRs do not report budgeted revenue). Columns (3)–(4) keep only states whose *legislative* cycle is coded annual in every NASBO wave, dropping the thirteen states ever coded with a biennial legislative cycle (CA, KS, KY, MA, ME, MT, NC, NV, OH, TX, VT, WI, and AR). Columns (5)–(6) keep only states whose *budget* cycle is coded annual in every wave, dropping the fifteen states coded biennial throughout (CT, HI, IN, KY, ME, MN, MT, NC, NE, NH, NV, OH, TX, VA, WI) plus Arkansas; NASBO's mixed annual/biennial code ("A,B", reported for Kansas and Missouri in 1997–2014) is counted as annual, consistent with NASBO's own tallies. The decision-window forecast loading is 2.6–2.9 and significant at 1% under all three restrictions, and no restriction produces a decision-window forecast-error loading distinguishable from zero.

Post-2002 sample. To address potential changes in ACFR budgetary reporting around the implementation of GASB Statement No. 34, Table E.11 repeats the baseline specification using fiscal years 2002 onward. The decision-window forecast loading is 2.63 (s.e. 0.60), nearly identical to the full-sample estimate of 2.71 (s.e. 0.61). Pre-decision news remains absent from the forecast and predicts forecast errors, whereas decision-window news does not predict forecast errors. The realized-revenue column retains positive off-window estimates, although the pre-decision estimate is imprecise. The decision-window pattern therefore does not depend on the earlier ACFR records.

Leave-one-out. Figures E.3 and E.4 re-estimate the decision-window premium dropping one fiscal-year cohort or one state at a time. Dropping any single fiscal year leaves the premium between 1.92 (dropping FY2010) and 2.68, always significant at the one percent level. Dropping any single state leaves the premium between 1.40 and 2.54, always positive; the smallest value comes from dropping Alaska, and once the FY2010 crisis cohort is excluded the leave-one-state-out premia tighten to 1.66–2.12 with a minimum t-statistic of 2.3. Figures E.5 and E.6 repeat the exercise for the window-level forecast-error loadings: the pre-decision loading is positive in every leave-one-out sample, and the deliberation loading in nearly all. No single state or cohort drives the result.

Table E.9: **Biennial Budgets: Second-Year Cohorts and the Decision Window**

The table re-estimates the forecast and forecast-error columns of Table 2, Panel A, under four treatments of biennial-budget states' second-year cohorts, whose imputed decision window contains no budget decision (the biennium covering the year was proposed and enacted roughly twelve months earlier). Biennial-budget state-years are NASBO budget-cycle code "B": CT, HI, IN, KY, ME, MN, MT, NC, NE, NH, NV, OH, TX, VA, and WI in the estimation sample. The first fiscal year of each biennium is even for odd-year enactors, which is every biennial sample state except Kentucky and Virginia (which enact in even years); the second-year indicator follows this parity. Columns (1)–(2) reproduce the baseline. Columns (3)–(4) add interactions of each window-news term with the second-year indicator, together with the indicator and its interactions with the three window indicators, so the news interactions are not driven by level differences. Columns (5)–(6) drop second-year cohorts from the sample. Columns (7)–(8) relabel all second-year pre-fiscal-year months as deliberation (post-proposal) months and re-estimate the baseline specification on the full sample. Cycle codings are assigned to fiscal years from the most recent NASBO survey wave at or before each year. Outcomes, window definitions, fixed effects, controls, common-sample restriction, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	Baseline		Off-window interaction		Excl. yr. 2		Yr. 2 as controls	
	(1) Forecast	(2) Forecast error	(3) Forecast	(4) Forecast error	(5) Forecast	(6) Forecast error	(7) Forecast	(8) Forecast error
$\text{News}^{\text{Pre}}_{s,t,m}$	0.02	1.10***	0.03	1.12**	-0.00	1.12**	0.13	1.08***
	(0.46)	(0.39)	(0.43)	(0.47)	(0.36)	(0.51)	(0.23)	(0.32)
× Biennial, yr2			-0.22	-0.04				
			(0.35)	(0.41)				
$\text{News}^{\text{Decision}}_{s,t,m}$	2.71***	-0.16	2.83***	-0.29	2.77***	-0.12	2.65***	-0.35
	(0.61)	(0.57)	(0.55)	(0.65)	(0.54)	(0.59)	(0.63)	(0.64)
× Biennial, yr2			-3.40***	3.41**				
			(1.25)	(1.59)				
$\text{News}^{\text{Delib}}_{s,t,m}$	0.55	0.62*	0.58	0.65*	0.59*	0.79**	0.48	0.71***
	(0.47)	(0.34)	(0.46)	(0.34)	(0.30)	(0.35)	(0.48)	(0.26)
× Biennial, yr2			-0.29	-0.34				
			(0.22)	(0.79)				
State FE	✓	✓	✓	✓	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Observations	16,402	16,402	16,402	16,402	13,222	13,222	16,402	16,402
Adj. R^2	0.121	0.575	0.124	0.576	0.127	0.602	0.121	0.575

Table E.10: **Biennial Budgets: Sample Restrictions**

The table re-estimates the forecast and forecast-error columns of Table 2, Panel A, on three restricted samples. Columns (1)–(2) drop the state-years NASBO codes as biennial on both margins—a biennial budget cycle and a biennial legislative cycle—so that no budget document is written for the second year and no regular annual session exists to write one (Maine, Montana, Nevada, North Carolina, Texas, and Wisconsin in all sample years; Kentucky through fiscal year 2001; Ohio from fiscal year 2021; Arkansas throughout). Columns (3)–(4) keep only states whose legislative cycle is coded annual in every NASBO wave, dropping the thirteen states ever coded with a biennial legislative cycle (CA, KS, KY, MA, ME, MT, NC, NV, OH, TX, VT, WI, and AR). Columns (5)–(6) keep only states whose budget cycle is coded annual in every wave, dropping the fifteen states coded biennial throughout (CT, HI, IN, KY, ME, MN, MT, NC, NE, NH, NV, OH, TX, VA, WI) plus Arkansas; NASBO's mixed annual/biennial code ("A,B", reported for Kansas and Missouri in 1997–2014) is counted as annual, consistent with NASBO's own tallies. Cycle codings are assigned to fiscal years from the most recent NASBO survey wave at or before each year. Outcomes, window definitions, fixed effects, controls, common-sample restriction, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	Excl. strictly biennial		Annual legislature		Annual budget cycle	
	(1) Forecast	(2) Forecast error	(3) Forecast	(4) Forecast error	(5) Forecast	(6) Forecast error
$\text{News}^{\text{Pre}}_{s,t,m}$	0.31	1.19***	0.14	1.15***	0.23	1.24
	(0.51)	(0.42)	(0.54)	(0.39)	(0.49)	(1.04)
$\text{News}^{\text{Decision}}_{s,t,m}$	2.64***	-0.11	2.94***	-0.12	2.67***	-0.24
	(0.58)	(0.65)	(0.54)	(0.68)	(0.56)	(0.99)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.40	0.68**	0.44	0.55*	0.52	0.78
	(0.39)	(0.27)	(0.47)	(0.30)	(0.39)	(0.59)
State FE	✓	✓	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
Observations	13,805	13,805	11,502	11,502	9,989	9,989
Adj. R^2	0.111	0.591	0.107	0.600	0.123	0.632

Table E.11: **Forecast Incorporation in the FY2002+ Sample**

The table repeats the baseline three-window specification from Panel A of Table 2, restricting the sample to fiscal years 2002–2024. The outcomes are the original-budget revenue forecast revision, the within-year forecast error, and realized revenue growth, as labeled. Window definitions, fixed effects, predetermined controls, and two-way clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	Forecast $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$	Forecast error $\log(\text{Act}_t^R/\text{Orig}_t^R)$	Realized $\log(\text{Act}_t^R/\text{Orig}_{t-1}^R)$
	(1)	(2)	(3)
$\text{News}_{s,t,m}^{\text{Pre}}$	0.01	1.15***	1.16
	(0.48)	(0.32)	(0.72)
$\text{News}_{s,t,m}^{\text{Decision}}$	2.63***	-0.34	2.29***
	(0.60)	(0.55)	(0.68)
$\text{News}_{s,t,m}^{\text{Delib}}$	0.44	0.43	0.87*
	(0.50)	(0.28)	(0.49)
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	14,520	14,520	14,520
Adj. R^2	0.129	0.586	0.466

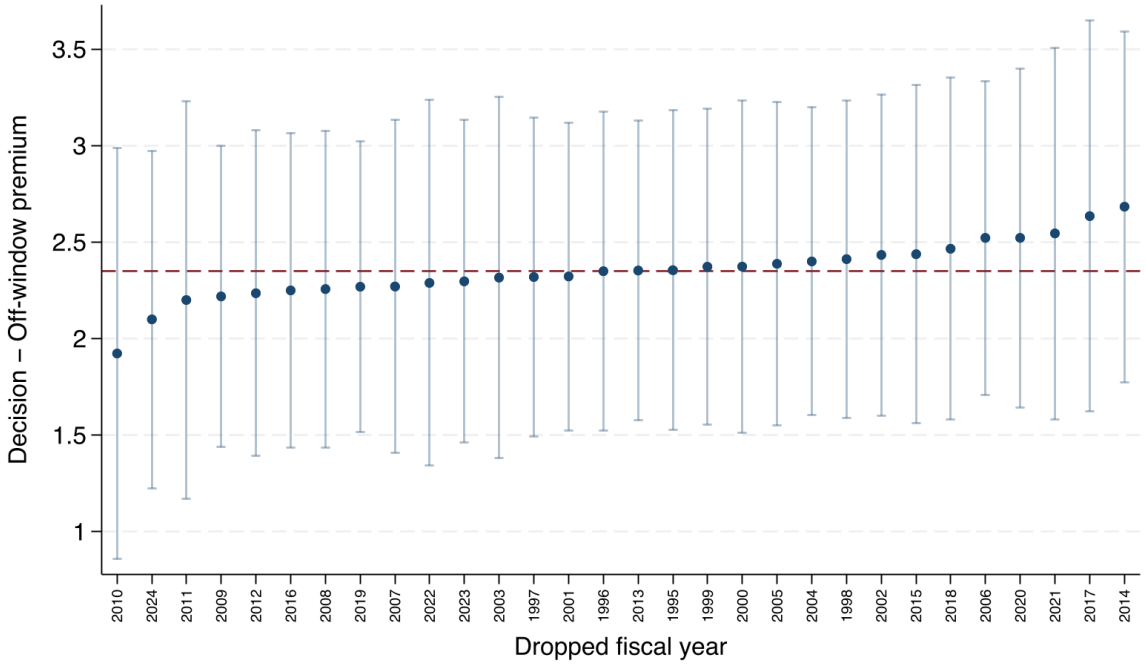
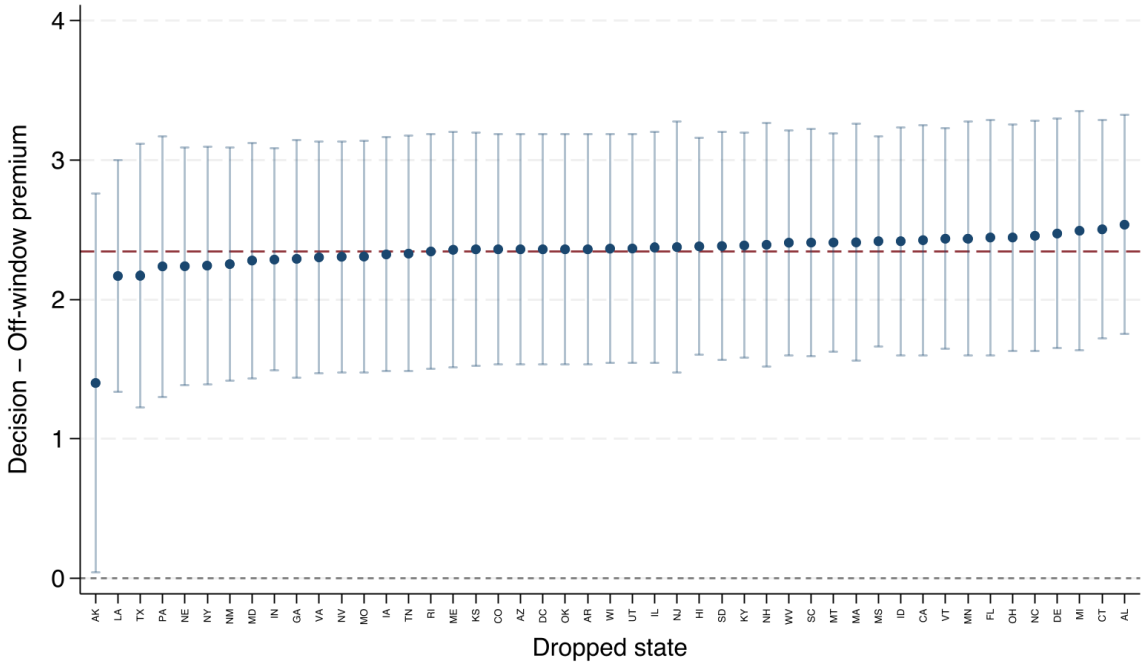
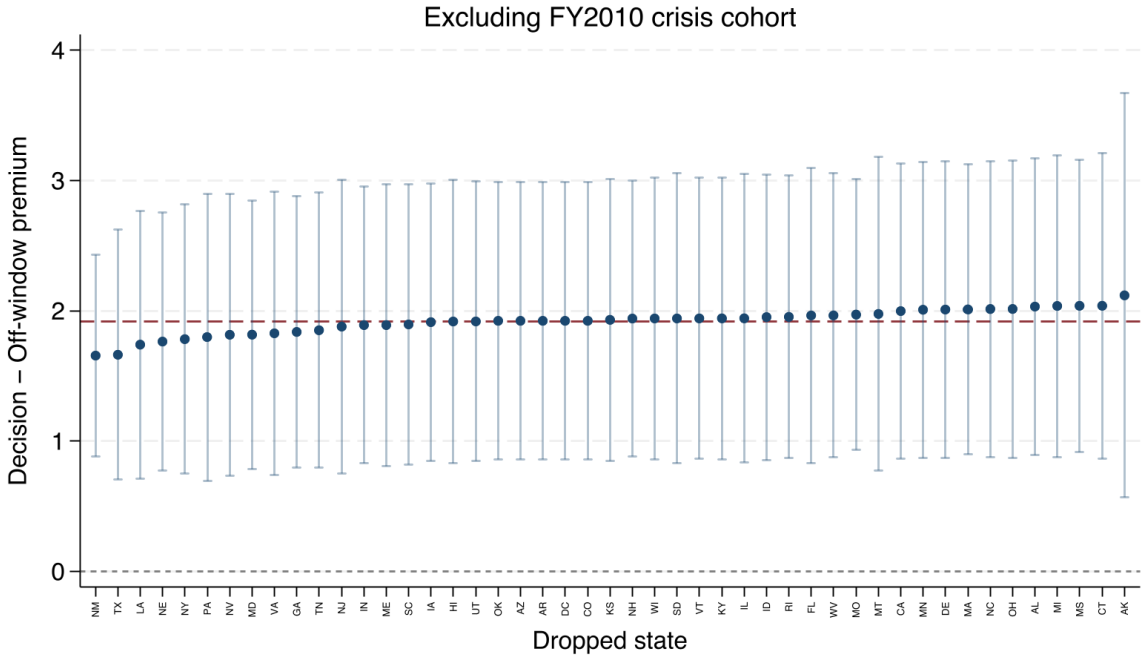


Figure E.3: **Leave-One-Fiscal-Year-Out Estimates of the Decision-Window Premium.** Each marker re-estimates the in/out decision-window premium of Table 2, Panel B, column (1) after dropping the indicated fiscal-year cohort. Cohorts are ordered by the estimated premium. Vertical bars are 90 percent confidence intervals based on standard errors two-way clustered by state and calendar month-year; the dashed horizontal line marks the mean across leave-one-out estimates.

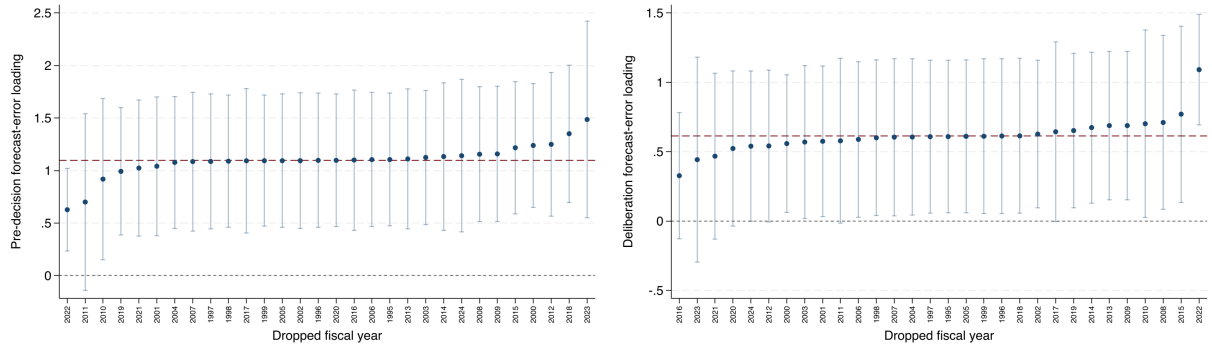


(a) Baseline sample



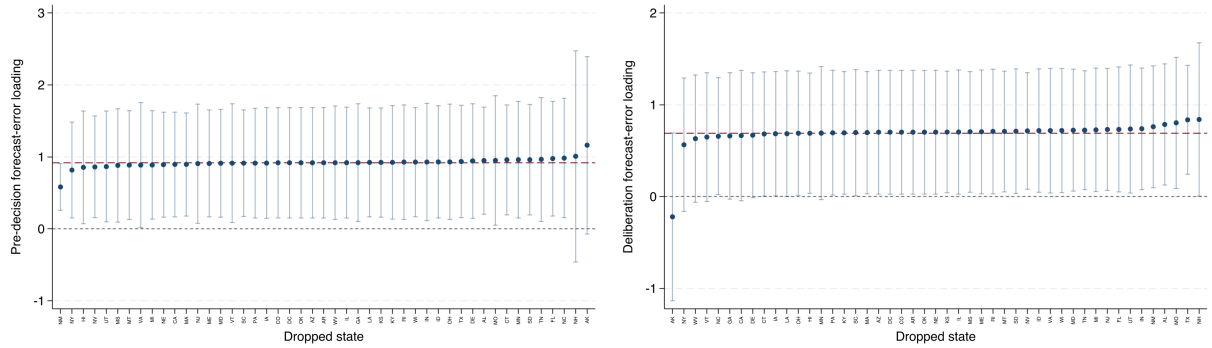
(b) Excluding the FY2010 crisis cohort

Figure E.4: **Leave-One-State-Out Estimates of the Decision-Window Premium.** Each marker re-estimates the in/out decision-window premium of Table 2, Panel B, column (1)—the coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$ for the original-budget revenue forecast—after dropping the indicated state. Panel (b) additionally excludes the FY2010 cohort from every regression. States are ordered by the estimated premium. Vertical bars are 90 percent confidence intervals based on standard errors two-way clustered by state and calendar month-year; the dashed horizontal line marks the mean across leave-one-out estimates.



(a) Pre-decision loading $\beta^{\text{fe}}_{\text{Pre}}$ (b) Deliberation loading $\beta^{\text{fe}}_{\text{Delib}}$

Figure E.5: **Leave-One-Fiscal-Year-Out Estimates of the Forecast-Error Loadings.** Each marker re-estimates the window-level forecast-error regression of Table 2, Panel A, column (2) on the common sample after dropping the indicated fiscal-year cohort, and plots the loading on pre-decision news (panel a) or deliberation news (panel b). Cohorts are ordered by the estimated loading. Vertical bars are 90 percent confidence intervals based on standard errors two-way clustered by state and calendar month-year; the dashed horizontal line marks the mean across leave-one-out estimates.



(a) Pre-decision loading $\beta^{\text{fe}}_{\text{Pre}}$ (b) Deliberation loading $\beta^{\text{fe}}_{\text{Delib}}$

Figure E.6: **Leave-One-State-Out Estimates of the Forecast-Error Loadings, Excluding FY2010.** Each marker re-estimates the window-level forecast-error regression of Table 2, Panel A, column (2) on the common sample after dropping the indicated state, with the FY2010 cohort additionally excluded from every regression, and plots the loading on pre-decision news (panel a) or deliberation news (panel b). States are ordered by the estimated loading. Vertical bars are 90 percent confidence intervals based on standard errors two-way clustered by state and calendar month-year; the dashed horizontal line marks the mean across leave-one-out estimates.

Why is Alaska influential? The Alaska sensitivity reflects a single state-cohort rather than a generic dependence on one state. Alaska combines the sample's most extreme severance exposure with a December proposal date, so its state-specific news realizations are the largest in absolute value, and the late-2008 collapse in oil prices fell inside its FY2010 decision window—the months just before its December 2008 proposal for fiscal year 2010. The Alaska–FY2010 cell is therefore the single largest decision-window shock in the sample, paired with a large contemporaneous forecast revision. Two patterns in the leave-one-out estimates confirm that the influence runs through this one cell. First, dropping the FY2010 cohort alone moves the premium only to 1.92, still significant at the one percent level, so Alaska's other twenty-nine budget rounds are not what the full-sample estimate leans on. Second, once FY2010 is excluded, dropping Alaska no longer lowers the premium at all: the Alaska-dropped estimate (2.12) is the *largest* of the forty-seven leave-one-state-out estimates in the no-FY2010 sample. Outside the crisis cohort, Alaska if anything dampens the premium. The influential observation is thus the intersection of the most exposed state with the one national shock that landed squarely in its decision window—exactly the variation the design is built to exploit—rather than an anomaly of Alaskan budgeting.

Exposure-specific calendar seasonality. Proposal months cluster in the late fall and winter, so the window indicators are close to deterministic functions of state and calendar month. The fixed effects absorb common calendar-month levels, but states with different revenue compositions could also have different seasonal gradients in forecasts or outcomes, which the design would then attribute to the window. Table E.12 addresses this with period-interacted share controls (Goldsmith-Pinkham et al., 2020): calendar-month-specific slopes on the five predetermined exposure shares, using each cohort's own weights and, more conservatively, weights fixed at their 1995 values. The decision-window loading is essentially unchanged (2.71 at baseline, 2.68 and 2.80 with the controls), and the forecast-error coefficient remains indistinguishable from zero.

E.5 Specification and Measurement Robustness

Specification build-up. Table E.13 traces the forecast and forecast-error loadings across three specifications, holding the estimation sample fixed at the common sample of Table 2, Panel A: state fixed effects only, adding fiscal-year-by-calendar-month fixed effects, and adding the predetermined controls (the main-text specification). The decision-window forecast loading falls from 4.10 to 2.99 once the common national innovation is absorbed and is es-

Table E.12: **Exposure-Specific Calendar Seasonality**

The table re-estimates Table 2, Panel A, adding calendar-month-specific slopes on the five predetermined revenue-exposure shares, the standard period-interacted share controls for shift-share designs. For each outcome, the *Base* column reproduces the baseline; *Shares* × *month* interacts calendar-month indicators with each cohort's own predetermined exposure shares $\omega_{s,t}$ (the five-year trailing averages lagged two years that also enter the news measure); *Shares* × *month (1995)* instead interacts the month indicators with the shares fixed at their initial 1995 values, $\omega_{s,1995}$, a conservative variant immune to any drift in the weights. The share-by-month terms are level controls, so they absorb exposure-specific seasonal gradients in the outcomes while leaving the window-timing variation intact. Outcomes, window definitions, fixed effects, controls, common-sample restriction, and clustering are identical to Table 2. *, **, and *** denote significance at 10%, 5%, and 1%.

	Forecast			Forecast error		
	(1) Base	(2) Shares × month	(3) Shares × month (1995)	(4) Base	(5) Shares × month	(6) Shares × month (1995)
$\text{News}^{\text{Pre}}_{s,t,m}$	0.02	0.08	0.03	1.10***	1.06*	1.12**
	(0.46)	(0.53)	(0.48)	(0.39)	(0.59)	(0.52)
$\text{News}^{\text{Decision}}_{s,t,m}$	2.71***	2.68***	2.80***	-0.16	-0.35	-0.16
	(0.61)	(0.77)	(0.72)	(0.57)	(0.70)	(0.76)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.55	0.55	0.57	0.62*	0.54	0.64
	(0.47)	(0.52)	(0.50)	(0.34)	(0.50)	(0.48)
$\omega_{s,t}$× Cal. month		✓			✓	
$\omega_{s,1995}$× Cal. month			✓			✓
State FE	✓	✓	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
Observations	16,402	16,402	16,402	16,402	16,402	16,402
Adj. R^2	0.121	0.121	0.118	0.575	0.585	0.574

sentially unchanged by the controls (2.71); the off-window loadings are near zero throughout, and the forecast-error mirror—pre-decision news predicting errors, decision-window news leaving none—appears once the common shock is absorbed.

Table E.13: **Specification Build-Up for the Window-Specific Loadings**

The table reports the window-specific loadings of Table 2, Panel A, across three specifications. Columns (1)–(3) use the original-budget revenue forecast revision, $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$, as the dependent variable; columns (4)–(6) use the within-year forecast error, $\log(\text{Act}_t^R/\text{Orig}_t^R)$. Columns (1) and (4) include state fixed effects only; columns (2) and (5) add fiscal-year-by-calendar-month fixed effects; columns (3) and (6) add the predetermined controls of Table 2 and reproduce its Panel A columns. All columns are restricted to the common estimation sample of Table 2, Panel A (the full-specification sample), so observations are identical across columns and coefficient changes reflect specification rather than sample composition. The sample is the stacked monthly panel restricted to the pre-decision, decision, and deliberation windows, 1995–2024. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

	Forecast $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$			Forecast error $\log(\text{Act}_t^R/\text{Orig}_t^R)$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{News}_{s,t,m}^{\text{Pre}}$	-0.02	0.12	0.02	0.16	1.12**	1.10***
	(0.71)	(0.51)	(0.46)	(0.34)	(0.45)	(0.39)
$\text{News}_{s,t,m}^{\text{Decision}}$	4.10***	2.99**	2.71***	1.15	0.05	-0.16
	(0.84)	(1.16)	(0.61)	(0.73)	(0.76)	(0.57)
$\text{News}_{s,t,m}^{\text{Delib}}$	0.57	0.78	0.55	0.21	0.80*	0.62*
	(0.39)	(0.53)	(0.47)	(0.74)	(0.44)	(0.34)
State FE	✓	✓	✓	✓	✓	✓
FYear × Year-Month FE		✓	✓		✓	✓
Controls			✓			✓
Observations	16,402	16,402	16,402	16,402	16,402	16,402
Adj. R^2	0.027	0.088	0.121	0.511	0.548	0.575

Alternative decision-window boundaries. Table E.14 re-estimates the window-specific loadings of Table 2, Panel A, under four alternative decision-window definitions: $[-4,-1]$, $[-6,-1]$, $[-5,0]$ (adding the proposal month itself), and $[-5,-2]$ (dropping the final pre-proposal month). In every variant the pre-decision window begins at $\ell = -11$ and ends one month before the decision window, and deliberation begins one month after it, so the three windows always partition the sample months. The forecast loading is significant under every definition and largest for $[-4,-1]$; dropping the final pre-proposal month ($[-5,-2]$) barely moves it, so the result is not driven by last-minute news, while adding the proposal month ($[-5,0]$) dilutes it, consistent with the budget being largely locked by the submission

month. The forecast-error columns show the mirror in every variant: the decision-window error loading is never distinguishable from zero, while off-window news predicts errors.

Table E.14: **Alternative Decision-Window Boundaries**

The table re-estimates the window-specific loadings of Table 2, Panel A, with the decision window redefined as $\ell \in [-4,-1]$, $[-6,-1]$, $[-5,0]$, or $[-5,-2]$ (columns as labeled; the baseline $[-5,-1]$ is Table 2). Columns (1)–(4) use the original-budget revenue forecast revision, $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$, as the dependent variable; columns (5)–(8) use the within-year forecast error, $\log(\text{Act}_t^R/\text{Orig}_t^R)$. In each variant the pre-decision window is $[-11, \ell_{\text{lo}}-1]$ and deliberation runs from $\ell_{\text{hi}}+1$ to the fiscal-year start, so the three windows partition the sample months. The sample is the common estimation sample of Table 2 (stacked monthly panel, 1995–2024). All specifications include state and fiscal-year-by-calendar-month fixed effects and the predetermined controls of Table 2. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

	Forecast $\log(\text{Orig}_t^R/\text{Orig}_{t-1}^R)$				Forecast error $\log(\text{Act}_t^R/\text{Orig}_t^R)$			
	(1) $[-4,-1]$	(2) $[-6,-1]$	(3) $[-5,0]$	(4) $[-5,-2]$	(5) $[-4,-1]$	(6) $[-6,-1]$	(7) $[-5,0]$	(8) $[-5,-2]$
$\text{News}_{s,t,m}^{\text{Pre}}$	0.09	0.18	-0.00	0.03	0.82***	0.96	1.10**	1.09***
	(0.38)	(0.51)	(0.46)	(0.46)	(0.30)	(0.64)	(0.42)	(0.39)
$\text{News}_{s,t,m}^{\text{Decision}}$	3.02***	2.17***	2.13***	2.96***	0.30	0.20	-0.13	-0.29
	(0.87)	(0.76)	(0.75)	(0.70)	(0.53)	(0.62)	(0.47)	(0.60)
$\text{News}_{s,t,m}^{\text{Delib}}$	0.58	0.50	0.44	0.68	0.66**	0.65*	0.78	0.58*
	(0.48)	(0.51)	(0.58)	(0.45)	(0.31)	(0.33)	(0.48)	(0.34)
State FE	✓	✓	✓	✓	✓	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Observations	16,402	16,402	16,402	16,402	16,402	16,402	16,402	16,402
Adj. R^2	0.121	0.121	0.121	0.121	0.575	0.575	0.575	0.575

Severance-only news. Table E.15 re-estimates the forecast-incorporation decomposition of Table 2, Panel A, replacing the five-category news measure with its severance component alone, $\hat{\eta}_{\text{sev}}\,\omega_{s,\text{sev},t}\,z_{\text{oil},m}$, so the only national shock is the oil-price innovation and the only exposure variation is predetermined severance reliance. The decision-window pattern is reproduced within this single transparent shock family: decision-window oil news moves the original-budget revenue forecast (3.92, s.e. 0.54), leaves no predictable forecast error (-0.29, s.e. 0.76), and carries strong realized-revenue content (3.63, s.e. 0.80). Pre-decision oil news does not enter the forecast (0.15, s.e. 0.53); its realized-revenue content is imprecisely estimated in this single-shock design (1.82, s.e. 1.21) and sits in the forecast error rather than the forecast.

Table E.15: **Decision-Window Estimates Using Severance (Oil) News Only**

The table re-estimates the window-specific specification of Table 2, Panel A, replacing the five-category shift-share news measure with its severance component only: $\text{News}^{\text{sev}}_{s,t,m} = \hat{\eta}_{\text{sev}}\, \omega_{s,\text{sev},t}\, z_{\text{oil},m}$, where $z_{\text{oil},m}$ is the standardized AR(4) innovation in the WTI crude-oil price and $\omega_{s,\text{sev},t}$ is the predetermined severance revenue share. Dependent variables are the original-budget revenue forecast revision $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$ (column 1), the within-year revenue forecast error $\log(\text{Act}^R_t/\text{Orig}^R_t)$ (column 2), and realized revenue growth $\log(\text{Act}^R_t/\text{Orig}^R_{t-1})$ (column 3). All columns share the common sample of Table 2, Panel A, so the realized-revenue loading equals the sum of the forecast and forecast-error loadings coefficient by coefficient. The sample is the stacked monthly panel restricted to the pre-decision, decision, and deliberation windows, 1995–2024. All specifications include state and fiscal-year-by-calendar-month fixed effects and the predetermined controls of Table 2. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

	Forecast $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$	Forecast error $\log(\text{Act}^R_t/\text{Orig}^R_t)$	Realized $\log(\text{Act}^R_t/\text{Orig}^R_{t-1})$
	(1)	(2)	(3)
$\text{News}^{\text{Pre}}_{s,t,m}$	0.15	1.67	1.82
	(0.53)	(1.17)	(1.21)
$\text{News}^{\text{Decision}}_{s,t,m}$	3.92***	-0.29	3.63***
	(0.54)	(0.76)	(0.80)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.95*	1.05	2.00
	(0.47)	(0.70)	(1.20)
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	16,402	16,402
Adj. R^2	0.122	0.575	0.447

F Real Effects: Additional Results and Robustness

This appendix collects supporting material for Section 4. It reports additional robustness checks for the industry-level real-effects results and diagnostics for the monetary-policy exercise.

F.1 Budget-Side Sign Split

The budget rule in equation (3) allows pass-through to differ by the sign of the news: when forecast revenue falls short of baseline needs, the balanced-budget constraint binds and forecast updates pass through one-for-one, whereas good news can flow partly into reserves. If this margin is active at the budget stage, negative decision-window news should pass through more strongly than positive news. Table F.1 tests this by splitting decision-window news into a positive component and the absolute value of the negative component, for the revenue forecast, the spending plan, and realized spending.

The revenue forecast responds with the expected sign in both directions: positive decision-window news raises the forecast and negative news lowers it. The sign-specific spending responses point the same way: positive news raises the spending budget and realized spending, while the negative-news responses are smaller and at most marginally significant. The asymmetry test does not reject symmetry at conventional levels: the sums $\beta_+ + \beta_-$ are positive rather than negative in all three columns and statistically indistinguishable from zero ($p = 0.30$, 0.15, and 0.12).

If anything, the point estimates suggest slightly stronger pass-through of positive news, though never significantly so. Any downside concentration in the real-effects estimates of Section 4 is accordingly not a mechanical inheritance of budget-stage asymmetry.

F.2 Robustness of the Industry-Level Results

Excluding the financial-crisis budget round. The financial crisis is a useful stress test because FY2010 contains the September 2008 shock for many state budget cycles. Table F.2 repeats the main industry specification excluding FY2010. The coefficients are similar to the baseline estimates in Table 4, indicating that the industry response is not driven by the crisis budget round alone.

Table F.1: **Asymmetric Pass-Through Through the Decision Window**

The table reports monthly-panel regressions of revenue-budget growth, $\log(\text{Orig}^R_{s,t}/\text{Orig}^R_{s,t-1})$ (column 1), spending-budget growth, $\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$ (column 2), and realized expenditure growth benchmarked to the lagged spending plan, $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t-1})$ (column 3), on sign-split budget-decision-window news. Budget-decision-window news, $\text{News}^{\text{Decision}}_{s,t,m}$, is split into a positive component $\text{News}^{\text{Decision},+}_{s,t,m}$ and an absolute-value-of-negative component $|\text{News}^{\text{Decision},-}_{s,t,m}|$, where the decision window is $\ell \in [-5,-1]$ and ℓ is the calendar-month distance from state s's FY-t proposal date. News in the other windows enters as signed single coefficients: pre-decision ($\ell \in [-11,-6]$), deliberation ($\ell \geq 0$ and $m <$ FY start), and after FY start ($m \geq$ FY start); the omitted reference is the before-pre-decision window ($\ell < -11$), captured by the level coefficient "News (baseline)." The row "$\beta_+ + \beta_-$ (asymmetry)" reports the sum of the positive and absolute-value-of-negative decision-window coefficients, with the p-value of the test of symmetry below it. The sample is a stacked monthly panel: for each state-fiscal-year, a 30-month window running from July of calendar year $t-3$ to December of calendar year $t-1$. All specifications include predetermined controls: four lagged quarterly windows of CFNAI, payroll employment, the S&P 500, the five-year Treasury yield, and CPI, with the most recent ending one year before budget formation; the Borusyak et al. (2022) aggregate-share control; and two-year-lagged state rainy-day-fund balances and GDP growth. Standard errors are two-way clustered by state and calendar month-year. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Revenue budget $\log(\text{Orig}^R_t/\text{Orig}^R_{t-1})$	Spending budget $\log(\text{Orig}^E_t/\text{Orig}^E_{t-1})$	Spending growth $\log(\text{Act}^E_t/\text{Orig}^E_{t-1})$
	(1)	(2)	(3)
$\text{News}^{\text{Pre}}_{s,t,m}$	0.02	0.15	0.97***
	(0.46)	(0.25)	(0.32)
$\text{News}^{\text{Decision},+}_{s,t,m}$	3.50***	2.06***	2.69***
	(0.95)	(0.58)	(0.68)
$\lvert\text{News}^{\text{Decision},-}_{s,t,m}\rvert$	-2.22**	-0.70	-1.29*
	(0.85)	(0.65)	(0.76)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.56	0.18	0.41
	(0.46)	(0.43)	(0.25)
$\beta_+ + \beta_-$ (asymmetry)	1.28 (1.22)	1.37 (0.93)	1.40 (0.88)
p(symmetry)	0.300	0.150	0.120
State FE	✓	✓	✓
FYear × Year-Month FE	✓	✓	✓
Controls	✓	✓	✓
Observations	16,402	18,787	18,787
Adj. R^2	0.121	0.126	0.535

Table F.2: **Real Effects of Decision-Window News, Excluding FY2010**

The table repeats the industry-level triple-difference regressions from Table 4, excluding fiscal year 2010. The dependent variables are one-year log changes in private employment and private wages. The key regressor is the state-specific shift-share news measure, $\text{News}_{s,t,m}$, interacted with an indicator for months in state s's budget decision window for fiscal year t, $\mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$, and a measure of industry exposure to state and local government demand. Panel A uses GovShare^{high}, an indicator for above-median exposure. Panel B uses GovShare^{sd}, the continuous exposure share standardized to mean zero and unit variance. The sample is a stacked state–industry–fiscal-year–month panel covering two-digit NAICS industries; for each state-fiscal-year, included months fall before the fiscal-year start. All specifications include the relevant lower-order interactions. Standard errors are clustered two ways by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Above-Median Exposure Indicator

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$		$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$	
	(1)	(2)	(3)	(4)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}_s^{high}$	0.29*	0.31*	0.43*	0.46*
	(0.15)	(0.17)	(0.22)	(0.24)
Two-way Interactions	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓
Industry × Year FE	✓	✓	✓	✓
State × Industry FE		✓		✓
FYear × Year-Month FE		✓		✓
Observations	549,395	549,395	549,395	549,395
Adj.R^2	0.55	0.57	0.78	0.78

Panel B: Continuous Exposure Measure

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$		$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$	
	(1)	(2)	(3)	(4)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}_s^{sd}$	0.19**	0.20**	0.23**	0.25**
	(0.08)	(0.09)	(0.11)	(0.11)
Two-way Interactions	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓
Industry × Year FE	✓	✓	✓	✓
State × Industry FE		✓		✓
FYear × Year-Month FE		✓		✓
Observations	549,395	549,395	549,395	549,395
Adj.R^2	0.55	0.57	0.78	0.78

Table F.3: **Budget Trajectory by Window of Arrival, Excluding FY2010**

The table repeats Table 5, excluding fiscal year 2010, whose decision window contains the September 2008 shock for many states. The columns follow later revenue budgets and realized revenue surprises relative to the current original revenue budget, Orig_t^R. $\text{News}_{s,t,m}^{\text{Pre}}$, $\text{News}_{s,t,m}^{\text{Decision}}$, and $\text{News}_{s,t,m}^{\text{Delib}}$ are the window-specific news measures for the same proposal-relative windows used in Table 5. All specifications include the fixed effects and controls from the main table. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	Budget revision $\log(\text{Orig}_{t+k}^R/\text{Orig}_t^R)$		Forecast error $\log(\text{Act}_{t+k}^R/\text{Orig}_t^R)$		
	t+1	t+2	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)
$\text{News}_{s,t,m}^{\text{Pre}}$	-0.15	1.37**	0.94**	0.73	2.59***
	(0.33)	(0.56)	(0.44)	(0.69)	(0.63)
$\text{News}_{s,t,m}^{\text{Decision}}$	-0.21	-2.19	1.16	1.63	-0.50
	(0.71)	(1.31)	(0.78)	(1.40)	(2.14)
$\text{News}_{s,t,m}^{\text{Delib}}$	-0.03	-0.71*	0.78*	0.60	-0.30
	(0.44)	(0.39)	(0.44)	(0.82)	(0.56)
Observations	15,009	14,309	16,369	15,009	14,309
Adj. R^2	0.131	0.158	0.594	0.463	0.434

Table F.4: **Aggregate State-Level Real Effects, Excluding FY2010**

The table repeats the aggregate state-level regressions from Table 6, excluding fiscal year 2010. The dependent variables are one-year log changes in private employment and private wages. The specification is the same in/out parametrization as in the main table: $\text{News}_{s,t,m}$ is the pooled off-window response and $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$ is the decision-window premium. The sample is a stacked monthly panel restricted to one budget round, excluding the earliest before-pre-decision months and execution months. All regressions absorb state and fiscal-year-by-calendar-month fixed effects and include the predetermined controls used in Table 6. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	$\log(\text{Emp}_{s,t}/\text{Emp}_{s,t-1})$	$\log(\text{Wages}_{s,t}/\text{Wages}_{s,t-1})$
	(1)	(2)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$	0.486***	0.666***
	(0.098)	(0.242)
$\text{News}_{s,t,m}$	0.120***	0.213*
	(0.040)	(0.121)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Macro + State Controls	✓	✓
Exclude FY 2010	✓	✓
Observations	18,990	18,990
Adj. R^2	0.84	0.72

Leave-One-State-Out Robustness. Figure F.1 re-estimates the main industry specification after omitting one state at a time. The plotted coefficients are the decision-window triple interaction for private employment and private wages. The dashed horizontal line is the full-sample estimate. The exercise checks whether the main industry result is driven by a single large state or a single unusual budget calendar.

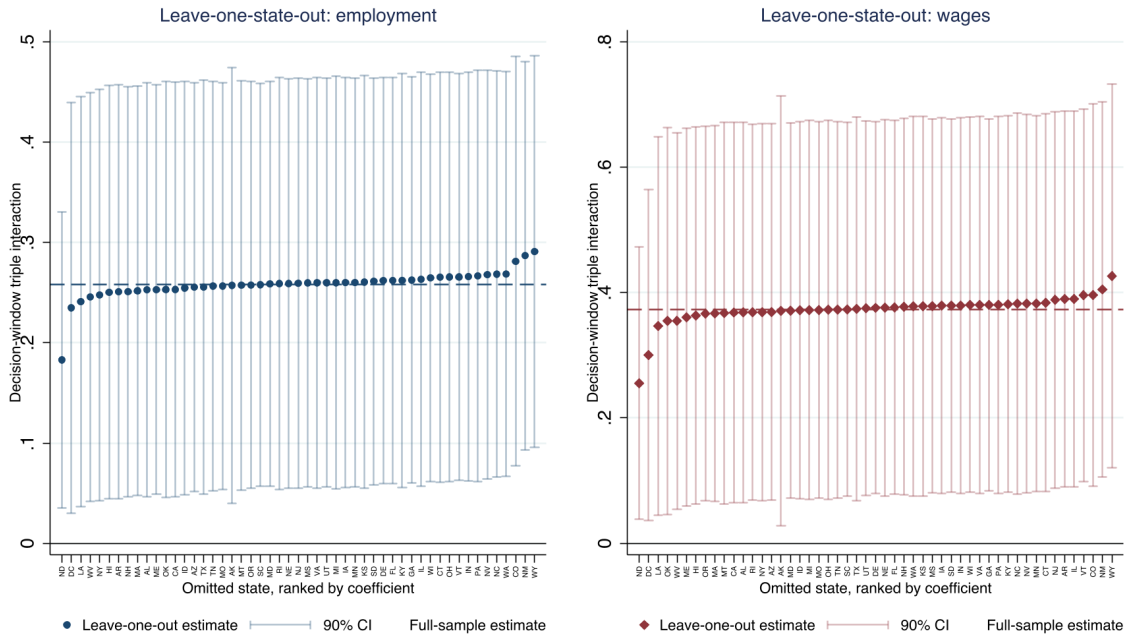


Figure F.1: **Leave-One-State-Out Robustness for Government-Exposed Industries.** The figure re-estimates the main industry-level specification after omitting one state at a time. Each point is the coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}^{high}$, with 90 percent confidence intervals. The dashed line is the full-sample estimate. The x-axis labels indicate the omitted state, ranked by the leave-one-out coefficient.

Industry-Exposure Randomization Inference. Because government exposure is measured at the two-digit industry level, Table F.5 reports a randomization check that permutes the high-exposure indicator across industries and re-estimates the baseline specification. The observed employment coefficient, 0.258, lies near the upper tail of the placebo distribution: the permutation mean is -0.023, the 5th–95th percentile range is $[-0.214, 0.186]$, and the two-sided permutation p-value is 0.059. The response is therefore concentrated in the industries actually exposed to government demand, rather than reflecting a generic industry-level pattern.

Table F.5: **Government-Exposure Permutation Inference**

The table reports randomization inference for the employment coefficient in the main industry-level specification. The observed coefficient is the coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\} \times \text{GovShare}^{high}$. The placebo distribution permutes the above-median government-exposure indicator across two-digit industries, re-estimates the baseline specification, and compares the observed coefficient with the resulting placebo coefficients. The specification, fixed effects, sample, and clustering match Table 4.

Statistic	Value
Observed coefficient	0.258
Observed standard error	0.123
Permutation p-value	0.059
Permutation mean	-0.023
Permutation 5th–95th pctile	[-0.214, 0.186]
Permutations	100
Observations	566,079

Dynamic response with continuous exposure. The main text reports the dynamic response using the high-exposure indicator. As an additional check, the production build estimates the same local projection using a continuous government-exposure measure, see Figure F.2. This exercise confirms that the delayed and persistent response in Figure 5 is not tied to the binary exposure definition.

F.3 Public-Payroll Diagnostics

Table F.6 asks whether decision-window news directly moves public-sector payroll. We construct public employment and wages from the annual QCEW panel as total employment and wages minus private employment and wages, aggregate to the state-year level, and estimate the same decision-window news specification. The coefficients are negative and statistically insignificant for both public employment and public wages. Thus the private-sector response in Table 4 is not simply a direct public-payroll response. It is more consistent with adjustment through procurement, grants, contracts, or other non-payroll spending margins.

F.4 Monetary-Policy Diagnostics

Season balance of the surprises. Section 4.2 uses monetary policy surprises as an external test of the budget-timing mechanism. The aggregate exercise compares shocks arriving in Q4, when many states are writing budgets, to shocks arriving in Q1–Q3. A natural con-

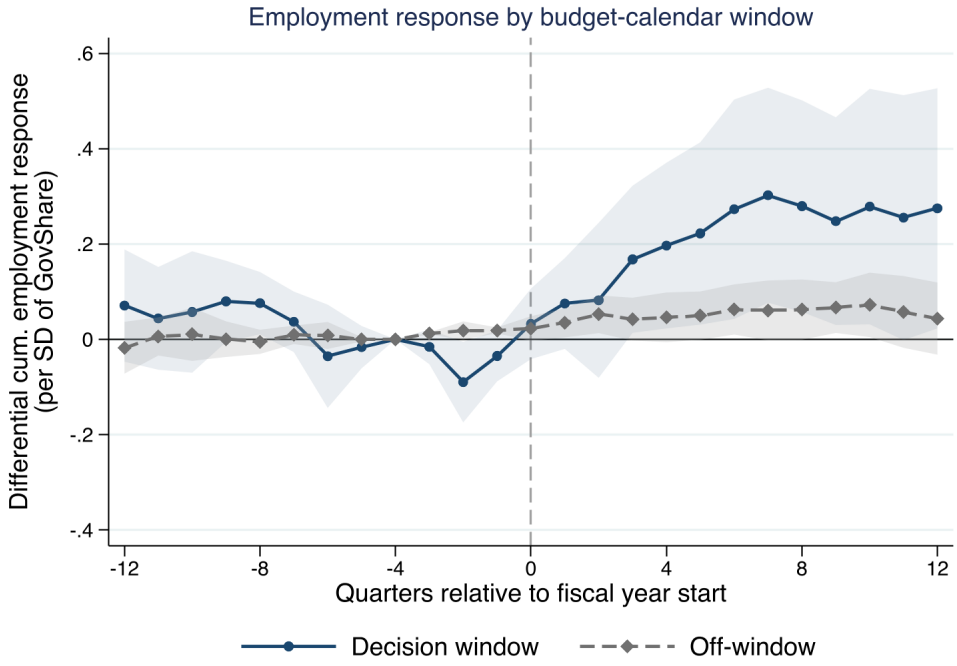


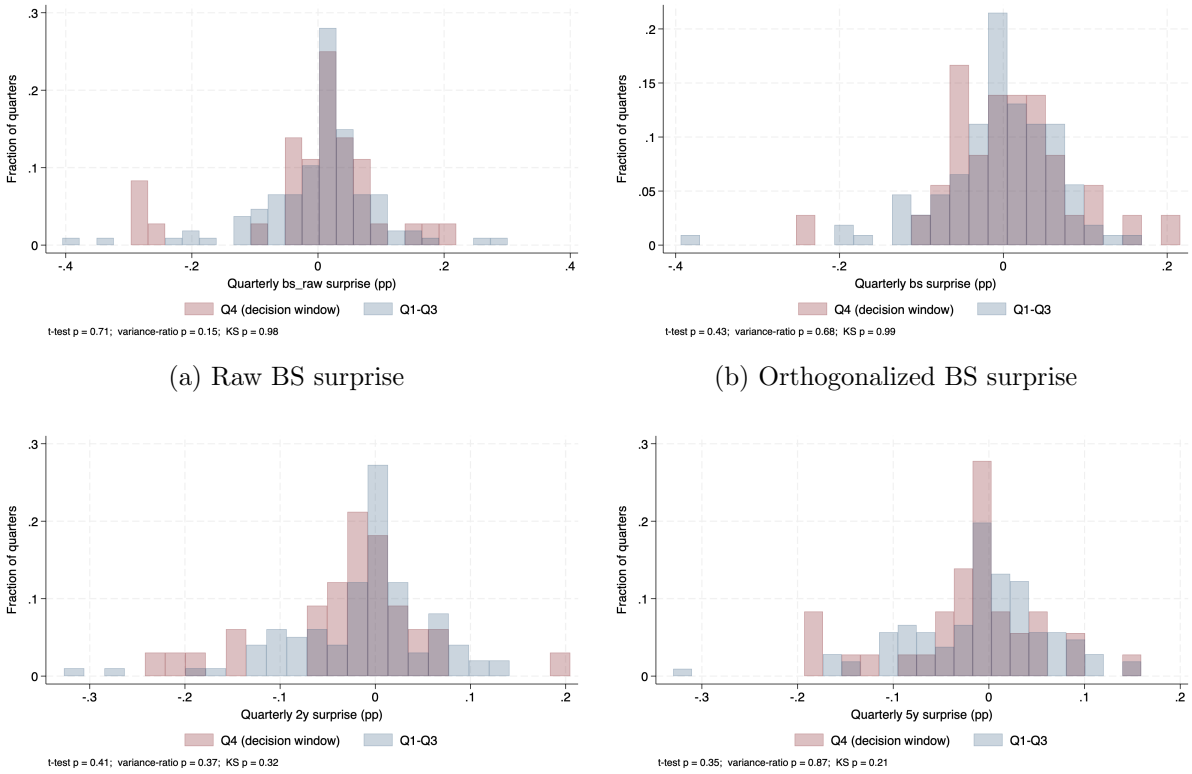
Figure F.2: **Dynamic Employment Response: Continuous Exposure.** The figure repeats the dynamic local projections in Figure 5 using the standardized continuous government-exposure measure. The dependent variable is cumulative private employment growth from the pre-fiscal-year base to horizon h. Regressions include state-by-industry, industry-by-fiscal-year, and state-by-fiscal-year fixed effects. Shaded areas are confidence intervals based on clustered standard errors.

Table F.6: **Decision-Window News and Public-Sector Payroll**

The table estimates state-level regressions of public-sector employment and wages on decision-window news. Public employment is $\text{Emp}^{\text{total}} - \text{Emp}^{\text{private}}$, and public wages are $\text{Wages}^{\text{total}} - \text{Wages}^{\text{private}}$, both constructed from the annual QCEW state-industry panel and aggregated to the state-year level. The dependent variables are one-year log changes. All specifications include state and fiscal-year-by-calendar-month fixed effects; standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	Public employment	Public wages
	(1)	(2)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}_{s,t}^{\text{Decision}}\}$	-1.050	-0.174
	(0.728)	(0.419)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Observations	25,444	25,444
Adj. R^2	0.043	0.083

cern is that the monetary surprises themselves differ by season. Figure F.3 compares the distributions of the four surprise measures used in Figure 7. The balance diagnostics do not suggest that Q4 surprises are systematically larger or differently distributed than surprises in the rest of the year.



(a) Raw BS surprise

(b) Orthogonalized BS surprise

(c) 2-year Treasury surprise

(d) 5-year Treasury surprise

Figure F.3: **Are Monetary Policy Surprises Different by Budget Season?** The figure plots the distribution of each quarterly monetary policy surprise, split by whether it arrives in Q4, when many states draft the upcoming budget, or outside that season. Histograms use common bins and are normalized by the number of quarters within each season. Each panel reports p-values from a two-sample t-test of equal means, a variance-ratio test of equal dispersion, and a two-sample Kolmogorov–Smirnov test of equal distributions.

Magnitude benchmark. A second input to the aggregate exercise is a magnitude benchmark. Figure F.4 reports the total response of real GDP to the orthogonalized Bauer–Swanson surprise from the same local-projection design, in the same units as the spending responses (percent of initial GDP), so the two are directly comparable.

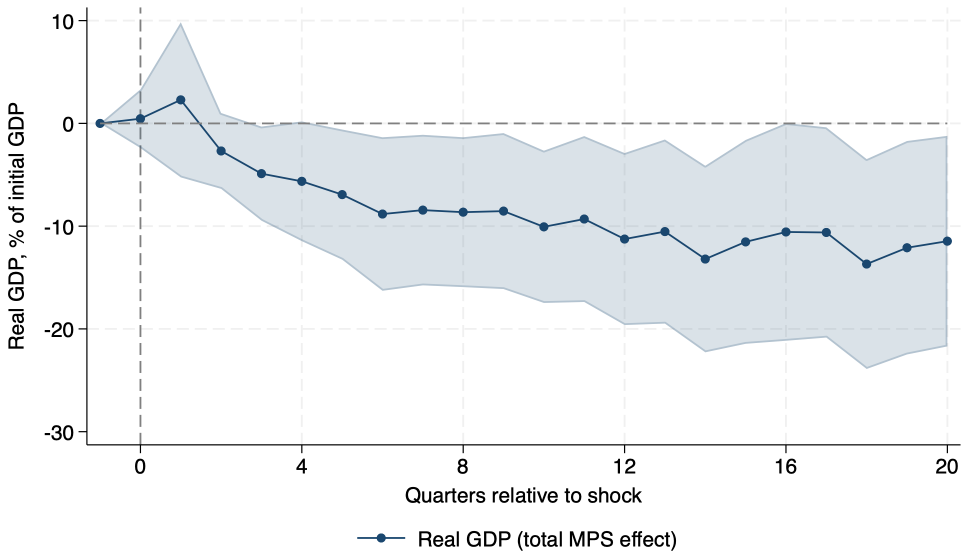


Figure F.4: **Magnitude Benchmark: Total Response of Real GDP to the Monetary Policy Surprise.** The figure plots local-projection estimates of the cumulative change in real GDP from quarter $t-1$ to $t+h$, scaled by real GDP in quarter $t-1$ and multiplied by 100 (percent of initial GDP—the same units as the spending responses in Figure 7), on the quarterly sum of orthogonalized Bauer–Swanson monetary policy surprises (Bauer and Swanson, 2023). The regression uses the total shock, not the Q4/Q1–Q3 split, since monetary transmission to aggregate output is not seasonal. Controls follow the spending local projections: four lags of the shock, four lags of state-and-local spending growth, and four quarterly lags (quarters $t-2$ through $t-5$) of log changes in real GDP, payroll employment, the S&P 500, the five-year Treasury yield, and the CPI; the contemporaneous quarter is excluded because asset prices respond to the shock within the quarter. Quarterly NIPA data, 1988–2023. Shaded areas are 90 percent confidence intervals based on Newey–West standard errors with lag length $\max(h+1, 4)$.

Distinguishing Q4 timing from Q4 confounds. Unlike the state-panel specification in Table 7, whose fiscal-year-by-calendar-month fixed effects absorb every national factor common to a calendar quarter, the aggregate Q4-versus-Q1–Q3 local projection of Section 4.2 has no such absorption. We therefore report a formal test that the Q4 and Q1–Q3 responses differ and a set of checks for the confounds specific to Q4.

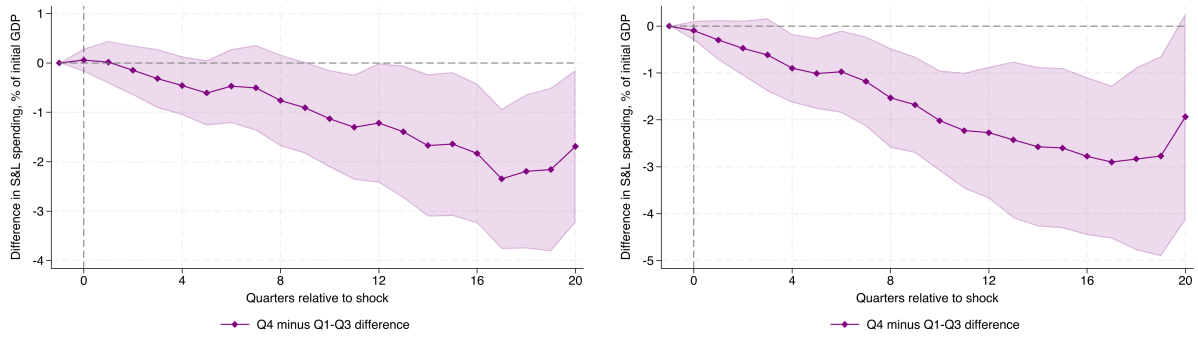
Table F.7 tests the difference directly. Reparametrizing the local projection so that the coefficient on the Q4 interaction is the Q4-minus-Q1–Q3 difference in the spending response at each horizon, the two impulse responses are statistically distinct for the futures-based Bauer–Swanson surprises: the difference is significant horizon by horizon, and the moving-block-bootstrap joint p-value over the two-to-four-year window is 0.055 for the orthogonalized surprise and 0.080 for the raw surprise. The yield-based surprises share the same negative sign but are not jointly significant. Figure F.5 plots the difference responses for both Bauer–Swanson surprises.

Table F.7: **Q4-versus-Q1–Q3 Difference in the Aggregate Spending Response**

The table reports the decision-window difference in the aggregate state-and-local spending response, $\Delta(h)$ = (Q4 response) − (Q1–Q3 response), from the local projection of Figure 7 reparametrized so that the coefficient on the Q4 interaction is the difference at horizon h. Columns report Δ at horizons $h = 8, 12, 16$ quarters, with horizon-dependent Newey–West standard errors in parentheses. "Joint p" is a two-sided moving-block bootstrap p-value (block length 12 quarters, 400 replications) for the null that the window-averaged difference over $h \in [8, 16]$ is zero. The orthogonalized Bauer–Swanson surprise is the primary measure; the four surprises are correlated, so they are not four independent tests. The specification, controls, and sample match Figure 7. *, **, and *** denote significance at 10%, 5%, and 1%.

Surprise	Δ at h=8	Δ at h=12	Δ at h=16	Joint p
BS raw	-0.759	-1.217*	-1.833**	0.080
	(0.562)	(0.735)	(0.858)	
BS orthogonalized	-1.533**	-2.275***	-2.775***	0.055
	(0.645)	(0.853)	(1.022)	
2-year Treasury	-1.123	-1.372	-1.809*	0.460
	(0.806)	(1.026)	(1.094)	
5-year Treasury	-1.295*	-1.393	-1.663	0.295
	(0.760)	(1.114)	(1.284)	

Table F.8 addresses the confounds specific to calendar Q4, which is also the start of the *federal* fiscal year and contains the year-end FOMC meeting. The premium is not a federal-fiscal-year artifact: federal current expenditures show no significant Q4 response (a clean, if imprecise, null), and the state-and-local premium is essentially unchanged when federal-spending lags are added as controls. It is not driven by the December FOMC meeting: rebuilding the Q4 surprise from October and November only leaves the premium intact.



(a) Raw BS surprise

(b) Orthogonalized BS surprise

Figure F.5: **Q4-minus-Q1–Q3 Difference in the Spending Response.** The figure plots the difference impulse response $\Delta(h)$ from Table F.7 for the two Bauer–Swanson surprises, with 90 percent confidence intervals based on horizon-dependent Newey–West standard errors.

Figure F.6 plots the federal-spending difference impulse response for both surprises.

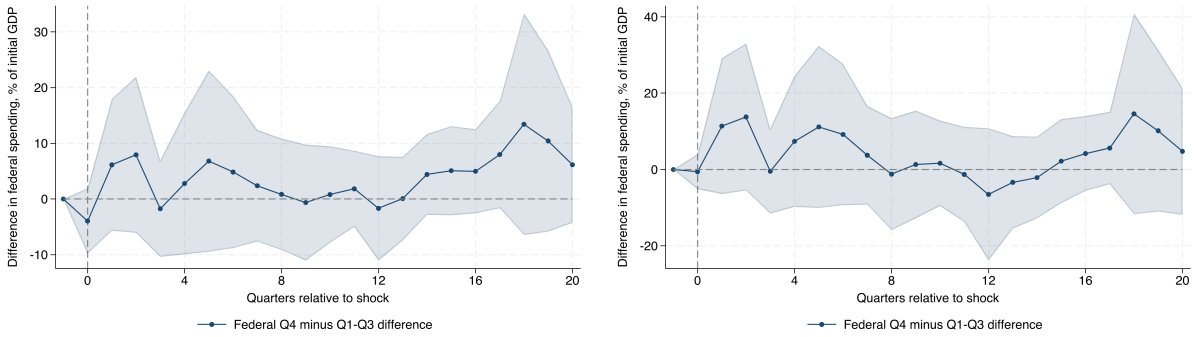
Table F.8: **Q4 Confound Checks**

The table reports the Q4-minus-Q1–Q3 difference in the cumulative state-and-local spending response at the three-year horizon ($h = 12$) under each confound check, for the two Bauer–Swanson surprises (orthogonalized and raw), with horizon-dependent Newey–West standard errors in parentheses. "Federal placebo" replaces the outcome with federal current expenditures (FRED FGEXPND), deflated and scaled identically to the state-and-local series; "S&L, control federal spending" adds four predetermined quarterly lags of federal-spending growth to the state-and-local regression; and "No-December" rebuilds the Q4 surprise from October and November only, dropping December's monthly contribution before quarterly aggregation. The specification, controls, and sample otherwise match Figure 7. *, **, and *** denote significance at 10%, 5%, and 1%.

Confound check (h=12)	BS raw	BS orthogonalized
Federal placebo (federal spending)	1.829 (4.096)	-1.291 (7.418)
S&L, control federal spending	-1.663*** (0.575)	-1.909*** (0.687)
No-December (Oct+Nov Q4)	-2.005*** (0.568)	-2.403*** (0.751)

G Rule-Induced Multipliers

This appendix formalizes the scaling exercise in Section 4.1.3. The object is a rule-induced fiscal elasticity: the response of aggregate outcomes to realized spending variation generated by routine forecast-mediated budget adjustment. This differs from a conventional fiscal multiplier identified from residual or discretionary spending shocks.



(a) Raw BS surprise

(b) Orthogonalized BS surprise

Figure F.6: **Federal-Spending Placebo.** Q4-minus-Q1–Q3 difference in the cumulative *federal* spending response (FRED federal current expenditures), from the same local-projection design as the state-and-local difference in Figure F.5, for the two Bauer–Swanson surprises. A band covering zero indicates no federal-fiscal-year artifact. Shaded areas are 90 percent confidence intervals from horizon-dependent Newey–West standard errors.

Let $Z_{s,t,w}$ denote revenue news arriving in window w, with D denoting the budget decision window and O the pooled off-window comparison. Define the reduced-form responses of the outcome and of realized spending to window-w news,

$$\gamma_w^Y \equiv \frac{\partial \log Y_{s,t}}{\partial Z_{s,t,w}}, \qquad \Gamma_w \equiv \frac{\partial \log G_{s,t}^{\text{act}}}{\partial Z_{s,t,w}},$$

together with the fiscal elasticity $\mu_Y^f \equiv \partial \log Y_{s,t} / \partial \log G_{s,t}^{\text{act}}$.

Revenue news can affect Y through a private channel and a fiscal channel:

$$\gamma_w^Y = \underbrace{\eta_w^Y}_{\text{private channel}} + \underbrace{\mu_Y^f \Gamma_w}_{\text{fiscal channel}},$$

where η_w^Y is the non-fiscal response to the news. The fiscal channel is the multiplier times Γ_w, and news enters this channel only through the budget it funds.

The identifying condition is that decision-window and off-window news have the same private-channel effect:

$$\eta_D^Y = \eta_O^Y.$$

This condition allows revenue news to affect private activity directly; it requires only that the direct effect does not depend on whether the news arrives inside or outside the budget decision window. This is plausible in our setting because on- and off-window news are the

same underlying innovations; what varies across states and years is only where a given realization falls in the statutory budget calendar. That calendar governs whether the news is incorporated into the current budget, not how firms and households respond to its economic content, so the direct private-sector effect should not shift discretely at the window boundary. Under this condition, differencing the decision-window and off-window reduced forms removes the private channel:

$$\gamma_D^Y - \gamma_O^Y = \mu_Y^f \left(\Gamma_D - \Gamma_O\right).$$

The differential spending response $\Gamma_D - \Gamma_O$ is the decision-window premium in realized spending, which we estimate directly (Table G.1, Panel B, column 2). Then, the estimated multiplier is

$$\widehat{\mu}_Y^f = \frac{\gamma_D^Y - \gamma_O^Y}{\Gamma_D - \Gamma_O}.$$

Table G.1: **Spending Response to Decision-Window Revenue News**

Monthly-panel regressions of the enacted spending budget and realized spending on state-specific shift-share news by proposal-relative timing. Column (1) is enacted spending-budget growth, $\log(\text{Orig}^E_{s,t}/\text{Orig}^E_{s,t-1})$; column (2) is realized spending benchmarked to the prior enacted budget, $\log(\text{Act}^E_{s,t}/\text{Orig}^E_{s,t-1})$. Panel A reports separate coefficients for pre-decision news ($\text{News}^{\text{Pre}}_{s,t,m}$, $\ell \in [-11,-6]$), decision-window news ($\text{News}^{\text{Decision}}_{s,t,m}$, $\ell \in [-5,-1]$), and deliberation news ($\text{News}^{\text{Delib}}_{s,t,m}$, $\ell \geq 0$ before fiscal-year start). Panel B reports the in/out specification, where the coefficient on $\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}^{\text{Decision}}_{s,t}\}$ is the decision-window premium relative to pooled off-window news. The sample is a stacked monthly panel restricted to these three pre-fiscal-year windows, 1995–2024. All specifications include state and fiscal-year-by-calendar-month fixed effects and the predetermined controls from Table 2. Standard errors are two-way clustered by state and calendar month-year. *, **, and *** denote significance at 10%, 5%, and 1%.

Panel A: Window-specific responses

	Spending budget: $\log(\text{Orig}^E_t/\text{Orig}^E_{t-1})$	Spending growth: $\log(\text{Act}^E_t/\text{Orig}^E_{t-1})$
	(1)	(2)
$\text{News}^{\text{Pre}}_{s,t,m}$	0.14	0.97***
	(0.25)	(0.32)
$\text{News}^{\text{Decision}}_{s,t,m}$	1.23***	1.83***
	(0.43)	(0.56)
$\text{News}^{\text{Delib}}_{s,t,m}$	0.17	0.39
	(0.45)	(0.27)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Controls	✓	✓
Observations	18,787	18,787
Adj. R^2	0.126	0.535

Panel B: Decision-window premium

	Spending budget: $\log(\text{Orig}^E_t/\text{Orig}^E_{t-1})$	Spending growth: $\log(\text{Act}^E_t/\text{Orig}^E_{t-1})$
	(1)	(2)
$\text{News}_{s,t,m} \times \mathbf{1}\{m \in \mathcal{W}^{\text{Decision}}_{s,t}\}$	1.07**	1.20**
	(0.51)	(0.49)
State FE	✓	✓
FYear × Year-Month FE	✓	✓
Controls	✓	✓
Observations	18,787	18,787
Adj. R^2	0.126	0.535