Supply-Side Beliefs: Interpretation, Aggregates, and Monetary Policy*

Rupal Kamdar
Indiana University, Bloomington

Walker Ray
Federal Reserve Bank of Chicago
and CEPR

October 2026 (*latest version*)

Abstract

We study supply-side beliefs: the tendency of households to interpret economic news as informative about supply. In a New Keynesian model, agents observe an ambiguous signal about future supply and demand shocks. News which raises inflation forecasts raises realized inflation and, when beliefs are sufficiently supply-tilted, lowers realized output, regardless of the true mix of supply and demand shocks. We test this prediction using news-driven revisions in household inflation expectations, measured from the daily timing of survey interviews around CPI releases. Following a release which raises inflation expectations, both inflation and unemployment rise over the following two years, whereas windows without a release deliver the opposite comovement. Beliefs tilt toward supply because of the same asymmetry preventing monetary policy from neutralizing supply shocks. The stagflationary response to news presents the central bank with additional trade-offs, but also introduces a belief-shaping motive: when attention to news is costly enough, optimal policy fully anchors inflation expectations without fully stabilizing inflation, by lowering the value of learning about supply shocks.

Keywords: inflation expectations, information, supply shocks

JEL Classification: D83, D84, E31, E32, E52

*rkamdar@iu.edu, walkerdray@gmail.com. We thank numerous seminar participants at the St. Louis Fed, the Federal Reserve Board, the Bank of England, the San Francisco Fed, Ohio University, the Midwest Macroeconomics Meeting and the Bologna Macro Meeting for excellent comments. The views expressed in this document are those of the authors and not those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

"In one sense, we are all Keynesians now; in another, nobody is."—Milton Friedman, 1966

1 Introduction

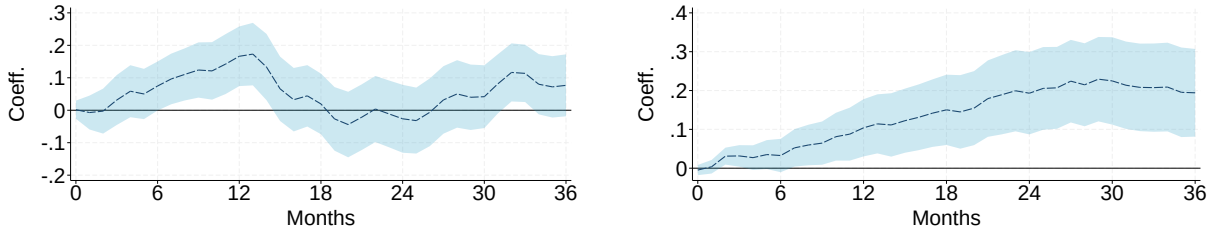
The conventional wisdom among monetary policymakers holds that most short-run fluctuations are demand-driven: inflation and output move together, and policy stabilizes both. Households, in contrast, appear to be supply-siders. When asked to explain inflation, households cite costs and markups rather than demand; consumer forecasts of inflation and unemployment are consistently *positively* correlated; and, relative to experts, households' stated reasons regarding the propagation of shocks typically mimic supply-side mechanisms.[1] This paper studies why household expectations depart from full-information rational expectations in this direction, whether the departure has aggregate consequences, and what it implies for the conduct of monetary policy.

We argue that these *supply-side beliefs* arise because supply shocks are costlier than demand shocks, and that the same asymmetry makes information events themselves stagflationary. We develop a New Keynesian model in which agents observe an ambiguous signal about next period's supply and demand shocks. Because beliefs act as anticipated fundamentals, news which raises inflation forecasts raises realized inflation and, when beliefs are supply-tilted, *lowers* realized output. We test this prediction with news-driven revisions in household inflation expectations, measured from the daily timing of survey interviews around CPI releases. Following a release which raises inflation expectations, both inflation and unemployment rise over the following two years (Figure 1); windows without a release deliver the opposite comovement.

Our results imply that news acts as a cost-push shock under any interest rate rule, even when the news is pure noise, and thus complicates the policy choices of the central bank. Moreover, monetary policy itself may reinforce supply-side beliefs. A rule which stabilizes the output gap offsets demand shocks as a by-product but only redistributes the damage of supply shocks between inflation and output, tilting beliefs further toward supply. An optimizing central bank therefore takes into account how beliefs respond to its plan. This *belief-shaping motive* calls for more aggressive inflation stabilization than expectations management alone: when inflation responds less to supply shocks, agents have less incentive to learn about future shocks. Optimal policy seeks to anchor inflation expectations, but does so without fully stabilizing realized inflation: full anchoring requires only that forecasting inflation is not worth the cost.

Our framework is based on a standard New Keynesian model with one change to the information structure. Agents observe a public signal about next period's shocks: a linear combination of the future supply shock, the future demand shock, and noise. The

[1] Shiller (1996) documents the narratives; Kamdar (2019) and Candia et al. (2020) the forecast correlation; Andre et al. (2022) and Andre et al. (2026) the supply-side reasoning relative to experts.



(a) Inflation Response **(b)** Unemployment Response

Figure 1: Aggregate Responses to an Inflation News Shock

Notes: Local projections of CPI inflation (panel a) and the unemployment rate (panel b) on the news shock: the difference in average one-year inflation expectations between Michigan Survey respondents interviewed in the five days after a CPI release and those interviewed in the five days before. Coefficients are per standard deviation of the shock (approximately one percentage point), with 90% confidence intervals. Section 3 describes the data and shock construction, and Section 4 describes the specification, controls, and sample.

interpretation of the signal is summarized by two Kalman gains: how much of a given realization agents attribute to each shock. We first study the positive implications of the model, which hold for arbitrary gains. We then explore the drivers of belief formation by endogenizing the information problem, and study the implications for policy and welfare.

The first key observation is that the contemporaneous reaction of the aggregate economy to news occurs purely through revisions in subjective beliefs. Regardless of whether the signal truly reflects fundamental demand or supply shocks in the next period (or is simply noise), the response of the economy today is driven by how agents interpret the signal. We show that news which raises subjective inflation forecasts always raises realized inflation today, but the response of output depends on the ratio of the supply gain to the demand gain. When beliefs are sufficiently *supply-tilted* (when the gain ratio exceeds a threshold), inflationary news *lowers* output today. The supply tilt does not need to be extreme: the threshold is below one because forecasts incorporate the anticipated policy response. Under an active rule, agents expect the central bank to tighten when next period's supply shock arrives, so an anticipated supply shock moves output forecasts by more than an anticipated demand shock.

We show that a sufficient condition for the stagflationary response to news can be recovered from household surveys. If household forecasts of inflation and activity comove negatively, inflationary news lowers realized output. Together with the positive comovement of realized inflation and output, the same survey moment bounds the supply tilt in beliefs from below. Both conditions hold in the data. The same logic implies that Phillips curve regressions which impose full-information rational expectations are biased toward flat or incorrectly signed slopes, whereas survey-augmented estimation is consistent. Further, a Phillips curve estimated from household forecasts alone has a negative slope, which we show is a robust feature of the data.

We test the prediction by measuring news shocks directly, from the daily timing of

survey interviews. The shock in a given month is the difference in average one-year inflation expectations between Michigan Survey respondents interviewed in the five days after the CPI release and those interviewed in the five days before it. Interview timing is set by the survey's fieldwork, so the two groups are comparable, and our identifying assumption is that the release is the only systematic aggregate event in the window which causes households to update. The shock is unpredictable from current and lagged inflation and unemployment, and essentially uncorrelated with the reaction of ten-year Treasury yields and oil prices to the same releases. A shock which raises inflation expectations also raises the share of households who expect unemployment to rise and lowers consumer sentiment. This conditional comovement of forecasts is the survey moment above, and it places the economy in the region in which news is stagflationary in realized outcomes.

Inflationary news is indeed stagflationary. After a one standard deviation shock, which raises expected inflation by roughly one percentage point, realized inflation rises by about 0.17 percentage points at its peak a year later. Unemployment rises gradually, by 0.10 percentage points after one year and 0.19 after two; industrial production falls over the same horizons. Standard inference based on our local projection specification shows these responses are highly significant. Additionally, because the measured shock likely contains survey sampling error, these estimates understate the response to a one percentage point revision in mean beliefs. We also assess statistical significance on the basis of more conservative placebo tests, which show that the pattern is specific to CPI release windows. The randomization p-value of the unemployment response over pseudo-release dates is 0.001 at two years; further, the placebo medians have the signs of demand-driven fluctuations.

A quantitative version of the model (with persistent shocks, habits, interest-rate smoothing, and more forward-looking news) reproduces the dynamics found in Figure 1 quantitatively. Moreover, the estimated model implies that fundamentals are demand-dominated: supply innovations account for only a small fraction of the innovation variance, and yet beliefs still tilt heavily towards supply. A supply-side reading of news thus coexists with an economy driven mostly by demand.

Why do agents interpret news as driven by supply? Under rational inattention and under robustness, the answer lies in the same asymmetry which prevents monetary policy from neutralizing supply shocks. Rationally inattentive agents "pay for information" and thus direct their limited attention toward the shock about which forecast errors are costlier. Robust agents distrust the composition of shocks and guard against the "worst case" where costlier shocks are more prevalent. Under an active policy rule, demand shocks are offset as a by-product of stabilizing inflation, whereas the costs of a supply shock can only be distributed between inflation and output. Mistakes about supply are therefore the costly ones, and both kinds of agents read news as supply news.

Because monetary policy partly sets the relative cost of the two shocks, the policy regime also shapes the interpretation of news: a rule which responds more strongly to

the output gap offsets more of a demand shock but only redistributes a supply shock, and tilts beliefs further toward supply. Monetary policy is thus partly responsible for the stagflationary impact of news, and that impact in turn generates further policy trade-offs for the central bank.

We show that no simple Taylor-type rule offsets news shocks, but the plan the central bank commits to determines the benefits of learning about future supply shocks. Under any rule, a response to news can neutralize the output response but not the inflation response, because the revision in expected inflation enters the Phillips curve directly as a cost-push term. Any central bank would therefore prefer that subjective inflation forecasts never moved; but we show that the degree to which a planner anchors inflation expectations depends on the degree of commitment and which mechanisms of the model are internalized by the planner.

We characterize optimal policy in the case of rationally inattentive private agents and study a ladder of planners who differ in what they internalize. Every planner neutralizes demand shocks, so supply is the only news worth learning about. The discretionary planner, however, can do nothing about forecasts, which depend on private expectations of the central bank's actions next period. With commitment, the central bank has an incentive to manage expectations. A planner who takes attention as given therefore promises to deliver a smaller inflation reaction to supply shocks next period, but never actively removes the private incentive to collect information. A planner who also internalizes that information collection is endogenous recognizes that more aggressive stabilization of inflation anchors expectations further because agents optimally choose to learn less about supply shocks. When attention is costly enough, the optimal policy is therefore to fully anchor inflation expectations. Anchoring does not require a commitment to full inflation stabilization, which is never optimal because of the usual inflation-output trade-off; it requires only that forecasting inflation is not worth the cost of attention. Thus, whereas a policy rule tilts beliefs away from demand by removing the value of learning about demand shocks, optimal policy anchors beliefs by lowering the value of learning about supply shocks as well.

Related Literature. Our paper builds on the literature which uses survey data to document departures from full-information rational expectations and embeds them in general equilibrium models.[2] Closest to our paper are the evidence that households reason about the economy in supply-side terms (Kamdar, 2019; Candia et al., 2020; Andre et al., 2022) and the work of Bhandari et al. (2022), who estimate household belief distortions

[2]Models of sticky information, noisy information, rational inattention, diagnostic expectations, and learning have found support in surveys (e.g., Carroll 2003; Mankiw et al. 2003; Coibion and Gorodnichenko 2012; Bordalo et al. 2020); see Coibion et al. (2018) for a review. General equilibrium models with deviations from FIRE include Mankiw and Reis (2007); Maćkowiak and Wiederholt (2015); Woodford (2013); Melosi (2014).

from survey data in a business cycle model in which the distortions act as wedges; they model the distortions with the same robustness device as our second foundation. Relative to these papers, we derive the direction of the belief distortion from the relative costs of the two shocks, which the policy regime partly determines; we test the causal transmission of the distortion with an identified news shock; and we study its consequences for the design of monetary policy. Kamdar and Ray (2025) study rationally inattentive consumers with a richer set of shocks; here we focus on the interpretation of a single ambiguous signal, and on how the policy rule shapes that interpretation.

Second, we contribute to the literature on news- and noise-driven fluctuations. In Beaudry and Portier (2004) and the work summarized in Beaudry and Portier (2014), agents receive news about future productivity; in Lorenzoni (2009) and Blanchard et al. (2013), noise about future productivity acts as a demand shock and generates a boom with higher inflation. In our model the signal is ambiguous about the *type* of shock, and a noise shock which is read as supply news is stagflationary. Empirically, that literature identifies news shocks with vector autoregressions under long-run or medium-run restrictions (e.g., Beaudry and Portier 2006; Barsky and Sims 2011, 2012), and Chahrour and Jurado (2018) clarify the relation between news and noise representations. We instead measure the response of beliefs to news directly, from the high-frequency response of household inflation expectations to CPI releases.

Third, we contribute to work on the high-frequency response of household expectations to announcements and data releases. Binder et al. (2022) and York (2023) show that household inflation expectations respond to CPI and unemployment releases and to some monetary policy announcements; Coibion et al. (2022) show in a randomized experiment that the form of central bank communication determines whether households update. Related evidence covers announcements in other countries and heterogeneity across households (Lamla and Vinogradov, 2019; Rast, 2021; Binder et al., 2024; De Fiore et al., 2022), and Weber et al. (2025) document that attention to inflation rises with the level of inflation. Our approach is complementary. We interpret the high-frequency response of inflation expectations to CPI releases as a news shock, and we show that news shocks have macroeconomic effects. Our anchoring result also offers a policy interpretation of the finding that attention rises with the level of inflation: agents attend to inflation when forecasting it is worth the cost, and the conduct of policy determines when that is the case.

Fourth, our normative results relate to work on optimal monetary policy with endogenous information. Paciello and Wiederholt (2014) show that a policy which stabilizes prices can remove firms' incentive to acquire information about some shocks, and Angeletos et al. (2016) study the social value of information in New Keynesian economies. Maćkowiak and Wiederholt (2009, 2015) develop the rational-inattention business cycle model on which our first foundation draws, and Afrouzi (2024) shows that firms' attention

to aggregate conditions is strategic and depends on market structure, and studies how attention responds to the monetary policy rule. Relative to these papers, the attention margin in our model is the interpretation of a signal about two shocks. The policy rule determines the direction of the interpretation, and the plan the central bank commits to determines the value of attending to news at all. As a result, the value of commitment includes a belief-shaping component in addition to the classical gain from managing expectations. Our anchoring result relates to models in which anchoring is the outcome of learning (Carvalho et al., 2023). In our model, expectations are anchored when the value of learning about future supply shocks falls below the cost of attention, and the central bank's plan determines that value.

2 Model

This section develops the model and derives our key predictions. We first characterize the equilibrium of the model and describe the transmission of anticipated future shocks for arbitrary subjective beliefs. We then introduce the information structure and characterize when agents' interpretation of news is *supply-tilted*. Finally, we derive testable predictions for both beliefs and aggregate dynamics which can be taken to the data.

2.1 Setup

Our setting is a textbook New Keynesian model (Galí, 2015): the economy is populated by a representative household, a continuum of monopolistically competitive firms $i \in [0,1]$ subject to Calvo pricing frictions, and a monetary authority. Our key departure from the textbook full-information rational expectations (FIRE) benchmark is how private agents form *forecasts* in response to new information regarding *future* innovations. We denote by $\tilde{E}_t$ the *subjective* expectation operator. We return to the full characterization of how agents form beliefs in Section 2.3, but we first characterize the equilibrium of the model using only a few simplifying assumptions regarding subjective expectations: the information set is common across agents and includes the history of shocks, and beliefs are internally consistent (agents understand the structural model, beliefs satisfy the law of iterated expectations, and agents do not expect their own future forecasts to be revised in a predictable direction).

Households. The representative household chooses consumption C_t, labor N_t, and nominal bonds B_t to maximize

$$\tilde{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\varsigma} - 1}{1-\varsigma} - \frac{N_t^{1+\varphi}}{1+\varphi} \right] \quad \text{s.t.} \quad P_t C_t + Q_t e^{-\zeta_t} B_t = B_{t-1} + W_t N_t + P_t T_t, \tag{1}$$

where ς is the inverse intertemporal elasticity of substitution, φ is the inverse Frisch labor elasticity, $Q_t = (1+i_t)^{-1}$ is the price of a one-period nominal bond, W_t is the nominal wage, and T_t collects lump-sum transfers and profits. The term ζ_t is an exogenous shock which raises households' effective return (and which can be interpreted equivalently as exogenous risk premia, convenience yields, credit spreads, flight-to-safety shocks, and so on). The precise microfoundation of this shock is unimportant for our purposes; what matters is only that ζ_t generates an "aggregate demand" shock in the model through the dynamic IS curve.

Firms. Firm i produces with labor $Y_t(i) = N_t(i)$ and chooses prices $P_t(i)$ subject to a Calvo friction: firms can only reset prices with probability $1-\theta$. The firm faces the demand curve

$$Y_t(i) = \left(P_t(i)/P_t\right)^{-\epsilon_t} Y_t,$$

where Y_t is aggregate output and P_t is the usual aggregate price index. Exogenous shocks to the elasticity of substitution $\epsilon_t > 1$ move desired markups $\mathcal{M}e^{\mu_t} \equiv \epsilon_t/(\epsilon_t - 1)$, where $\mathcal{M}$ is the steady-state desired markup. As with the demand shock, the purpose of this source of exogenous variation is to generate an "aggregate supply" shock of the model through the Phillips curve.

Monetary Policy. The monetary authority sets the nominal interest rate i_t. In the baseline, we assume the simple Taylor rule $i_t = \phi_\pi \pi_t$ (in deviations from steady state). We consider general rules and optimal policy in Section 5.

Log-Linear System. Log-linearizing the optimality conditions around the zero-inflation steady state delivers the standard system (see Appendix A for all details),

$$x_t = \tilde{E}_t x_{t+1} - \varsigma^{-1}\left(i_t - \tilde{E}_t \pi_{t+1} - v_t\right), \tag{2}$$

$$\pi_t = \beta \tilde{E}_t \pi_{t+1} + \kappa x_t + u_t, \tag{3}$$

where x_t is the output gap and π_t is inflation. The Phillips curve parameters depend on preferences and nominal rigidities in the standard way: $\kappa \equiv \lambda(\varsigma+\varphi), \lambda \equiv (1-\theta)(1-\beta\theta)/\theta$.

The aggregate demand and supply shocks are $v_t \equiv -\zeta_t, u_t \equiv \lambda\mu_t$. Formally, these terms arise from exogenous variation in the risk-premium wedge and the demand elasticity. However, any alternative choice delivering the log-linearized system (2)–(3) does not change any of our results, up to the weights in the loss function (17) below. This includes an alternative specification with flexible prices but nominal wage rigidities (and re-labeling inflation as wage inflation).

2.2 Transmission of Anticipated Shocks

We simplify to the case of exogenous, mutually independent, iid Gaussian shocks: $u_t \sim N(0, \sigma_u^2)$ and $v_t \sim N(0, \sigma_v^2)$. Simplifying to the iid case allows us to isolate the role of anticipated shocks and derive closed-form results; the quantitative model of Section 4.3 has persistent shocks and delivers the same signs of the responses to news. We write $\tilde{E}_t u_{t+1}$ and $\tilde{E}_t v_{t+1}$ for the subjective forecast of the supply and demand innovations next period. For now we treat these expectations as given and study how the economy responds. Note that this setting also nests the standard news-shock case (where agents perfectly anticipate next period's shocks: $\tilde{E}_t u_{t+1} = u_{t+1}, \tilde{E}_t v_{t+1} = v_{t+1}$). Section 2.3 returns to our microfoundation for how these forecasts are formed.

With iid shocks and the simple monetary policy rule, the model can be solved period by period given the current shocks and subjective forecasts of inflation $\tilde{E}_t \pi_{t+1}$ and the output gap $\tilde{E}_t x_{t+1}$. First consider the case with no anticipated shocks $\tilde{E}_t u_{t+1} = \tilde{E}_t v_{t+1} = 0$. Then $\tilde{E}_t \pi_{t+1} = \tilde{E}_t x_{t+1} = 0$ and the solution is simply

$$x_t = \chi \varsigma^{-1}(v_t - \phi_\pi u_t), \quad \pi_t = \chi(u_t + \kappa \varsigma^{-1} v_t), \quad \chi \equiv 1/(1 + \kappa \varsigma^{-1} \phi_\pi) \in (0, 1). \tag{4}$$

Note that given our sign normalizations, an increase in u_t generates a supply-driven recession and an increase in v_t generates a demand-driven expansion; that is, both shocks generate an increase in inflation, but differ in their transmission to the output gap.

Next, because agents understand the structural model, forecasts of endogenous outcomes are formed by applying (4) to their expectations of the shocks, which gives

$$\tilde{E}_t x_{t+1} = \chi \varsigma^{-1} \left(\tilde{E}_t v_{t+1} - \phi_\pi \tilde{E}_t u_{t+1} \right), \quad \tilde{E}_t \pi_{t+1} = \chi \left(\tilde{E}_t u_{t+1} + \kappa \varsigma^{-1} \tilde{E}_t v_{t+1} \right). \tag{5}$$

Finally, substituting (5) into (2)–(3) and solving for current outcomes, we find

$$x_t = \chi \varsigma^{-1} (v_t - \phi_\pi u_t) + b_{xu} \tilde{E}_t u_{t+1} + b_{xv} \tilde{E}_t v_{t+1}, \tag{6}$$

$$\pi_t = \chi \left(u_t + \kappa \varsigma^{-1} v_t \right) + b_{\pi u} \tilde{E}_t u_{t+1} + b_{\pi v} \tilde{E}_t v_{t+1}, \tag{7}$$

where the loadings are the sensitivities of current outcomes to anticipated future shocks:

$$b_{xu} = \chi^2 \varsigma^{-1} \left[1 - \phi_\pi (1 + \beta) \right], \quad b_{xv} = \chi^2 \varsigma^{-1} \left[1 + \kappa \varsigma^{-1} (1 - \beta \phi_\pi) \right], \tag{8}$$

$$b_{\pi u} = \chi^2 \left[\beta - \kappa \varsigma^{-1} (\phi_\pi - 1) \right], \quad b_{\pi v} = \chi^2 \kappa \varsigma^{-1} \left[1 + \beta + \kappa \varsigma^{-1} \right]. \tag{9}$$

Thus, comparing with (4), we see that contemporaneous shocks transmit exactly as in a textbook model, and anticipated shocks enter only through the forecasts $\tilde{E}_t u_{t+1}$ and $\tilde{E}_t v_{t+1}$. To simplify the characterization of news transmission, we assume the following

parametric conditions:

$$\phi_\pi > \frac{1}{\beta}, \tag{10}$$

$$\phi_\pi < 1 + \frac{\beta}{\kappa\varsigma^{-1}}. \tag{11}$$

Condition (10) is a slightly strengthened Taylor principle, and places a lower bound on the aggressiveness of policy. Condition (11) is an *upper* bound; this condition holds for typical Taylor rule coefficients when the composite slope $\kappa\varsigma^{-1}$ is small, as it is with a flat Phillips curve or with risk-averse households. These two conditions imply that an anticipated shock "acts like" the same shock today.

To see this, consider agents who expect a supply shock next period: by (5), $\tilde{E}_t\pi_{t+1} > 0$ and $\tilde{E}_t x_{t+1} < 0$. Households smooth consumption by saving more today, reducing current demand. The policy rule adds downward pressure on demand, both through the response to realized inflation today and through the anticipated response at t+1. Hence $b_{xu} < 0$ whenever the rule is active enough; (10) is sufficient. Firms face two opposing forces. First, weak current demand puts downward pressure on inflation (the term κx_t). However, higher expected inflation (the term $\beta\tilde{E}_t\pi_{t+1}$) puts upward pressure on inflation. Condition (11) is exactly the condition which guarantees the forward-looking force dominates. Thus $b_{\pi u} > 0$, and inflation rises while output falls.

Now consider agents who expect a demand shock: by (5), $\tilde{E}_t\pi_{t+1} > 0$ and $\tilde{E}_t x_{t+1} > 0$. Households reduce savings today in expectation of higher consumption tomorrow; additionally, higher expected inflation also puts upward pressure on demand. Firms increase prices in anticipation of higher inflation, but the increase in inflation implies that the policy rate increases as well, which puts downward pressure on output. Thus $b_{\pi v} > 0$, and $b_{xv} > 0$ if policy is not too aggressive, which holds under (11).

Thus, given the parameter restrictions (10) and (11), an anticipated supply shock is stagflationary and an anticipated demand shock is inflationary and expansionary.

2.3 News and Belief Formation

We now specify what generates movements in subjective forecasts. Our key departure from a standard treatment of news is that information about next period's shocks is not available independently for each shock. Instead, agents observe news which is informative about both shocks, and so beliefs about the individual shocks depend on how the news is interpreted. The following assumption formally states the information structure.

Assumption 1 (Information and belief formation)**.**

(i) Full current information. *All agents observe the complete history of endogenous variables and shocks through t.*

(ii) News. *At t, agents observe public news z_t about next period's innovations, independent of the current and past state. Expectations of next period's innovations are linear in z_t.*

(iii) Internal consistency. *Agents know the model and the equilibrium mapping; future news is subjectively mean zero; and agents do not anticipate future deviations of their own beliefs from the rule they apply at t.*

Assumption 1 does not specify what the news contains or how agents interpret news. By (ii), expectations are linear in the news, which defines the gains $(\tilde{K}_u, \tilde{K}_v)$:

$$\tilde{E}_t u_{t+1} \equiv \tilde{K}_u z_t, \quad \tilde{E}_t v_{t+1} \equiv \tilde{K}_v z_t. \tag{12}$$

The gains are the interpretation of the news: how much of a given realization agents attribute to each shock. All of our positive results are stated for arbitrary gains, and thus in this section we take as given $\tilde{K}_u$ and $\tilde{K}_v$; Section 2.4 endogenizes the gains.

We normalize the news so that a positive realization raises inflation expectations, so we call a positive z_t an *inflationary news shock*. Our baseline case considers a scalar signal, though our general results do not require this restriction. Suppose that news is

$$z_t = \alpha_u u_{t+1} + \alpha_v v_{t+1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2) \text{ iid}, \tag{13}$$

independent of all other disturbances, where η_t is pure noise. Agents update beliefs by Bayes' rule, under a possibly distorted prior $(\tilde{\sigma}_u^2, \tilde{\sigma}_v^2)$ regarding the innovation variances. Then the gains are the Kalman gains under the prior,

$$\tilde{K}_u = \frac{\alpha_u \tilde{\sigma}_u^2}{\alpha_u^2 \tilde{\sigma}_u^2 + \alpha_v^2 \tilde{\sigma}_v^2 + \sigma_\eta^2}, \quad \tilde{K}_v = \frac{\alpha_v \tilde{\sigma}_v^2}{\alpha_u^2 \tilde{\sigma}_u^2 + \alpha_v^2 \tilde{\sigma}_v^2 + \sigma_\eta^2}. \tag{14}$$

Note that under rational expectations, $(\tilde{\sigma}_u^2, \tilde{\sigma}_v^2) = (\sigma_u^2, \sigma_v^2)$ and (14) are the objective Bayesian gains. Under rational inattention, the signal loadings and noise $(\alpha_u, \alpha_v, \sigma_\eta)$ are chosen, and updating occurs under the true prior; under robustness, the loadings are taken as given and agents update under a distorted prior. Both are special cases of (14). The following proposition characterizes the transmission of news for arbitrary gains.

Proposition 1 (Transmission of news)**.** *Suppose Assumption 1 holds, $i_t = \phi_\pi \pi_t$, and $\tilde{K}_u, \tilde{K}_v \geq 0$, not both zero. Equilibrium outcomes are (6)–(7) with the expectations (12). News moves subjective inflation and output forecasts in opposite directions if and only if*

$$\frac{\tilde{K}_u}{\tilde{K}_v} > \frac{1}{\phi_\pi}. \tag{15}$$

The aggregate response to news is

$$\frac{\partial x_t}{\partial z_t} = b_{xu}\tilde{K}_u + b_{xv}\tilde{K}_v, \quad \frac{\partial \pi_t}{\partial z_t} = b_{\pi u}\tilde{K}_u + b_{\pi v}\tilde{K}_v,$$

with the loadings (8)–(9). *Under* (10)–(11), $b_{xu} < 0, b_{xv} > 0, b_{\pi u} > 0, b_{\pi v} > 0$. *Hence, news moves realized inflation and realized output in opposite directions if and only if*

$$\frac{\tilde{K}_u}{\tilde{K}_v} > T_x \equiv \frac{b_{xv}}{|b_{xu}|} = \frac{1 + \kappa\varsigma^{-1}(1 - \beta\phi_\pi)}{\phi_\pi(1+\beta) - 1}, \quad T_x < \frac{1}{\phi_\pi} < 1. \tag{16}$$

Proposition 1 has three implications. First, beliefs enter the equilibrium exactly like anticipated fundamentals. The response to news depends on the gains, not on whether the realization of z_t reflects u_{t+1}, v_{t+1}, or η_t; news about true future shocks and purely noise-driven beliefs are nested in a single expression.

Second, a news shock which increases inflation forecasts always raises *realized* inflation, because under (10)–(11) both an anticipated demand shock and an anticipated supply shock raise realized inflation. Therefore, news lowers realized output if and only if the interpretation loads on supply strongly enough relative to demand. Condition (16) states when the mixture of the stagflationary response to anticipated supply shocks and the expansionary response to anticipated demand shocks is still stagflationary.

Third, the ordering in (16) shows that news may decrease realized output even if output forecasts do not fall. To see why, consider the boundary case $\tilde{K}_u/\tilde{K}_v = 1/\phi_\pi$; by (5), agents then interpret news as neither demand nor supply from the perspective of output. Thus, $\tilde{E}_t x_{t+1}$ does not respond to z_t, and only $\tilde{E}_t \pi_{t+1}$ rises. Such a "pure" expected-inflation shock is still contractionary. Through the Phillips curve, expected inflation raises current inflation by β, and through the policy response the real rate rises when $\beta\phi_\pi > 1$, which is exactly condition (10), and output falls.

Supply-Side Beliefs. The transmission of news thus depends on whether belief formation is "supply-tilted," and to what degree: the predictions of the model change qualitatively when the gain $\tilde{K}_u$ on supply is sufficiently large relative to the gain $\tilde{K}_v$ on demand. We therefore say that agents hold *supply-side beliefs* when the ratio $|\tilde{K}_u|/|\tilde{K}_v|$ is large relative to a neutral benchmark.

We consider three different definitions of a "neutral" benchmark, all relative to the Bayesian gains ratio for different signals. The *content-neutral* benchmark corresponds to a signal loading equally on both shocks; the *standardized* benchmark corresponds to equal loadings on the shocks normalized by their objective volatility; and the *strict* benchmark corresponds to a signal which generates a gain ratio of 1.[3] When supply shocks are less volatile than demand shocks, these thresholds are nested: $\sigma_u^2/\sigma_v^2 < \sigma_u/\sigma_v < 1$ so that

[3]The three benchmarks correspond to $\alpha_u = \alpha_v$, $\alpha_j \propto 1/\sigma_j$, and $\alpha_j \propto 1/\sigma_j^2$.

the "strict" supply tilt implies the "standardized" supply tilt, which in turn implies the "content-neutral" supply tilt. Note further that strict over-weighting is sufficient to satisfy both conditions (15) and (16) of Proposition 1; but strict over-weighting is not necessary for our results.

2.4 The Origins of Supply-Side Beliefs

So far we have treated the gains $\tilde{K}_u$ and $\tilde{K}_v$ as unrestricted parameters. We now derive the conditions under which beliefs are "supply-tilted" under two standard foundations of belief formation: rational inattention and robustness. Throughout, we measure the welfare cost of aggregate fluctuations by the standard second-order approximation to household utility,

$$L_t = \frac{1}{2}\left(\varpi_\pi \pi_t^2 + \varpi_x x_t^2\right), \quad \varpi_\pi = \bar{\epsilon}/\lambda, \quad \varpi_x = \varsigma + \varphi, \tag{17}$$

where $\bar{\epsilon}$ is the steady-state elasticity of substitution.[4]

Rational Inattention. Under *rational inattention*, a single information collector chooses what to learn about $\varepsilon_{t+1} \equiv (u_{t+1}, v_{t+1})$, subject to a cost of information. The collector minimizes the standard objective

$$\mathbb{E}\left[e^\top Q\, e\right] + \mathcal{C}, \quad Q \equiv \varpi_\pi\, b_\pi\, b_\pi^\top + \varpi_x\, b_x\, b_x^\top, \tag{18}$$

where $e \equiv \tilde{E}_t \varepsilon_{t+1} - \varepsilon_{t+1}$ is the vector of posterior-mean errors, and $\mathcal{C}$ represents the costs of collecting information, which are proportional to the Shannon mutual information between the posterior beliefs and the realized shocks. The vectors $b_\pi \equiv (b_{\pi u}, b_{\pi v})^\top$ and $b_x \equiv (b_{xu}, b_{xv})^\top$ collect the sensitivities in (8)–(9). Attention is allocated to shocks according to their volatility and the degree to which misperceptions of the shocks are costly. The optimal signal takes the form (13) with *endogenous* signal loadings and noise $(\alpha_u, \alpha_v, \sigma_\eta)$ and an undistorted prior. Additionally, rational inattention implies a scalar signal whenever information costs are not too small (Appendix B1). The sensitivities

[4] We assume that information is acquired collectively, by a single collector acting on behalf of all agents, and that the costs of imperfect information are evaluated according to this social welfare function; under collective information acquisition, this is the appropriate loss function. Collective information acquisition keeps the analysis transparent and removes the fixed point which arises with higher-order beliefs. Dispersed private signals would introduce the higher-order uncertainty studied by Morris and Shin (2002), Woodford (2003), and Angeletos and Lian (2018), which Kamdar et al. (2025) document in household inflation expectations.

Additionally, if instead each agent privately acquires information, beliefs become heterogeneous, and in particular will differ on average across firms and households. This complicates the analysis but does not generally overturn our main qualitative findings. In particular, in our setting the only forward-looking variable relevant to any price or wage setter is expected inflation; when the Phillips curve is flat, the conditions under which beliefs tilt towards supply are therefore generally *weaker*. This may not hold for pure consumers (whose only forward-looking choice is consumption); for the full characterization of consumers' private information collection under rational inattention, see Kamdar and Ray (2025) who characterize the conditions under which supply over-weighting occurs with a richer set of shocks.

which direct attention are equilibrium coefficients which do not depend on the gains, so belief formation requires no fixed point between beliefs and aggregates (Appendix B3).

Robustness. Under *robustness*, agents know the total variance of innovations, $S = \sigma_u^2 + \sigma_v^2$, but are not certain of how much to attribute to supply relative to demand shocks. Agents form beliefs against the worst case: nature distributes S between the two shocks so as to maximize the expected loss caused by the innovations themselves when they are realized, and agents update by Bayes' rule but under the implied prior. Like the sensitivities which direct attention, nature's objective depends only on equilibrium coefficients and not on the gains (Appendix B2). As in the case of rational inattention, we assume that worst-case costs are evaluated according to the social loss (17); robustness then implies all agents hold identical beliefs.

Rational inattention is the minimal departure from FIRE which delivers our predictions; robustness is further from FIRE and delivers the tilt in its extreme form, but is useful for simple closed-form derivations. The following proposition states the interpretation under the two foundations.

Proposition 2 (Endogenous supply-side beliefs)**.** *Suppose* (10)–(11) *hold.*

(i) Under rational inattention, suppose that the off-diagonal element of Q in (18) *is positive (which holds if and only if condition* (B3) *holds). Then there is a strictly increasing threshold function r^* such that, for any cutoff $c > 0$,*

$$\frac{\tilde{K}_u}{\tilde{K}_v} > c \iff \frac{\sigma_u^2}{\sigma_v^2} > r^*(c).$$

The thresholds on the fundamental variance ratio for a stagflationary news response, for negative forecast comovement, and for the strict supply-tilt are ordered:

$$\underbrace{r^*(T_x)}_{\frac{\partial x_t}{\partial z_t}<0} < \underbrace{r^*(1/\phi_\pi)}_{\mathrm{Cov}(\tilde{E}_t\pi_{t+1},\tilde{E}_t x_{t+1})<0} < \underbrace{r^*(1)}_{|\tilde{K}_u|>|\tilde{K}_v|}. \tag{19}$$

Moreover, $r^*(1) < 1$ if and only if $Q_{uu} > Q_{vv}$.

(ii) Under robustness, for any signal which loads on supply ($\alpha_u \neq 0$), $\tilde{K}_v = 0 < \tilde{K}_u$ whenever $\varpi_\pi(\varsigma^2 - \kappa^2) + \varpi_x(\phi_\pi^2 - 1) > 0$; under (10)*, a sufficient condition is $\kappa\varsigma^{-1} \leq 1$. The supply tilt thus exceeds every cutoff in* (19)*.*

The intuition for part (i) follows from the general principle of rational inattention: learn more about the shock which is more costly to misperceive. The optimal signal is the combination of the two future shocks which, for a given amount of information,

minimizes the welfare cost of the resulting forecast errors.[5] All else equal, more attention is paid to shocks which are more volatile, or for which misperception generates larger welfare losses. The threshold function $r^*(c)$ is the relative volatility of supply shocks at which $\tilde{K}_u/\tilde{K}_v = c$. Since r^* is increasing, the ordering in (19) follows from Proposition 1. Importantly, unless the Phillips curve is steep, a unit forecast error about supply is costlier than a unit forecast error about demand (the condition $Q_{uu} > Q_{vv}$), so $r^*(1) < 1$: beliefs tilt toward supply even when supply shocks are less volatile.[6]

The results regarding robustness follow the same intuition but in an even more extreme form. From a welfare perspective, a unit supply shock is more costly than a unit demand shock unless the Phillips curve is very steep (the condition in part (ii)). Robust agents thus act as if news is only informative about supply, regardless of the physical volatilities of demand and supply shocks. In this case, robustness delivers an infinite gain ratio.

Supply Tilt and Divine Coincidence. The exact parametric conditions under which each foundation tilts beliefs towards supply differ. However, both foundations tilt toward supply for the same underlying reason, which is also the reason behind divine coincidence: monetary policy can neutralize demand shocks but not supply shocks. Per unit of variance, supply shocks therefore generate larger welfare losses.

Appendix B3 formalizes this insight across a wider class of *cost-ranked* belief-formation models: interpretation of ambiguous news tilts toward the shock for which misperceptions are costlier under the agent's own subjective criterion. Both rational inattention and robustness are specific instances of this general class. The direction of belief tilts is therefore a property of the environment; agents would interpret news as informative about demand if the economy were such that demand shocks were more costly. The class is also restrictive; for instance, representativeness-based updating is excluded (where beliefs tilt towards the shock which the news "resembles"); and more generally, purely behavioral "hard-wired" supply heuristics are excluded, as these models do not allow for the gain ratio to change if the cost ranking changes.

2.5 Testable Predictions

Whether news transmits as supply or as demand depends on structural parameters, the policy rule, and the equilibrium interpretation through the gains. Rather than taking a stand on each of these, we now ask which observable moments reveal the interpretation of news and imply that news is stagflationary.

A natural candidate is the unconditional comovement of realized inflation and output,

[5] If information costs are sufficiently low, agents acquire a second signal and learn about the two shocks separately. Appendix B1 discusses the conditions on information costs.

[6] The condition holds with a wide margin at standard parameter values: $\beta = 0.99$, $\varsigma = \varphi = 1$, $\theta = 0.75$, $\bar{\epsilon} = 6$, and $\phi_\pi = 1.5$ imply that the three thresholds in (19) are 0.19, 0.26, and 0.37.

but on its own this moment is uninformative. $\text{Cov}(\pi_t, x_t)$ is the sum of a fundamental term (which depends on the volatilities of both shocks) and a news term (which depends on the overall variance of news $\sigma_z^2 \equiv \alpha_u^2\sigma_u^2 + \alpha_v^2\sigma_v^2 + \sigma_\eta^2$); see (A5) in Appendix A4. Both terms scale with σ_v^2 (whenever the signal places any weight on demand); however, whenever news is stagflationary (Proposition 1), the news contribution to $\text{Cov}(\pi_t, x_t)$ is negative. Thus, even with a "demand-dominated" economy, inflation and output may not be positively correlated. When the news interpretation is supply-tilted, even taking $\sigma_u^2/\sigma_v^2 \to 0$ we may still have $\text{Cov}(\pi_t, x_t) < 0$: the news channel can generate negative realized comovement in an economy with no supply shocks at all. While such an extreme news-driven economy is of course unrealistic, this thought experiment demonstrates clearly that the unconditional correlation of inflation and output cannot identify either the mix of fundamental shocks or the interpretation of news.

The unconditional comovement of *forecasts*, in contrast, is informative. The following proposition shows that negative comovement of forecasts implies that news is stagflationary, and that together with $\text{Cov}(\pi_t, x_t)$ it bounds the belief tilt from below.

Proposition 3 (Forecast comovement)**.** *Suppose $\tilde{K}_u, \tilde{K}_v \geq 0$, not both zero, and (10)–(11) hold.*

(i) $\text{Cov}(\tilde{E}_t\pi_{t+1}, \tilde{E}_t x_{t+1}) < 0$ *implies that news shocks raise realized inflation and lower realized output. Negatively correlated forecasts are not necessary: for $\tilde{K}_u/\tilde{K}_v \in (T_x, 1/\phi_\pi)$, news is stagflationary while forecasts comove positively.*

(ii) If $\text{Cov}(\tilde{E}_t\pi_{t+1}, \tilde{E}_t x_{t+1}) < 0$ *and* $\text{Cov}(\pi_t, x_t) > 0$*, then*

$$\frac{\sigma_u^2}{\sigma_v^2} < \frac{\kappa\varsigma^{-1}}{\phi_\pi}, \quad \frac{\tilde{K}_u}{\tilde{K}_v} > \frac{\varsigma}{\kappa}\frac{\sigma_u^2}{\sigma_v^2}. \tag{20}$$

Thus, beliefs are supply-tilted relative to the content-neutral benchmark if $\kappa\varsigma^{-1} \leq 1$, and relative to the standardized benchmark if $\kappa\varsigma^{-1} \leq 1/\phi_\pi$.

Proposition 1 shows that the sign of the forecast correlation determines whether the gain ratio exceeds $1/\phi_\pi$. Part (i) of Proposition 3 adds that the forecast correlation also implies the causal prediction for realized outcomes in response to news. Part (ii) shows that the forecast correlation combined with the correlation of actual inflation and output is sufficient to bound the supply tilt itself from below (along with the mild parameter restrictions).[7]

Both moments are needed to bound the belief tilt for the following reason. When news is tilted towards supply, inflation and output are less correlated than in a FIRE "demand-driven" economy (or even negatively correlated). Hence, when the economy faces stagflationary news shocks, a positive correlation of inflation and output is only

[7] Under the same illustrative values of the previous footnote, $\kappa\varsigma^{-1} = 0.17$ and $\kappa\varsigma^{-1}\phi_\pi = 0.26$.

possible when demand shocks are more volatile than supply shocks by at least the factor $\phi_\pi/(\kappa\varsigma^{-1})$: the first inequality of (20) is a necessary (but not sufficient) condition for positively correlated aggregates. With a relatively flat Phillips curve, the actual ratio of supply to demand volatility must therefore be small. A negative forecast correlation, in turn, requires $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$ regardless of the volatilities of the shocks. Together, the two moments therefore imply that beliefs are tilted towards supply relative to the fundamental variance ratio by at least the factor ς/κ.

Preliminary Empirical Results. By Proposition 3, supply-tilted beliefs are a testable prediction of the model. As is well-known, in survey data the forecast comovement is negative. Household inflation expectations and unemployment expectations are positively correlated, or equivalently inflation expectations and expectations of economic activity are negatively correlated (Kamdar, 2019; Candia et al., 2020). Table 1 verifies this fact for the sample and the measures used throughout the paper.[8]

Table 1: Unconditional Forecast Comovement

	$\tilde{E}_{i,t}U^+_{t+1}$	$\tilde{E}_{i,t}U^-_{t+1}$	$\tilde{E}_{i,t}s_{t+1}$
$\tilde{E}_{i,t}\pi_{t+1}$	1.002***	-0.575***	-0.0350***
	(0.021)	(0.016)	(0.0004)
Time FE	Yes	Yes	Yes
Obs.	224,236	224,236	225,840
R^2	0.065	0.047	0.137

Notes: Pooled regressions across respondents of indicators (times 100) for expecting unemployment to rise ($\tilde{E}_{i,t}U^+_{t+1}$) and to fall ($\tilde{E}_{i,t}U^-_{t+1}$) over the next year, and of the sentiment index ($\tilde{E}_{i,t}s_{t+1}$, the first component of the categorical forward- and backward-looking questions, Section 3.4), on the one-year inflation forecast ($\tilde{E}_{i,t}\pi_{t+1}$), with time fixed effects (one for each survey month). HC1 standard errors in parentheses. Michigan Survey of Consumers, January 1982 to February 2020.

The same results also have implications for the estimation of the Phillips curve. When news drives part of the variation in the output gap, estimates of the slope which are based on FIRE are biased downward (and can even turn negative when the news share is large enough), because episodes driven by news trace out combinations of high inflation and low output, as supply shocks would; replacing expected inflation with realized future inflation does not restore consistency, because the expectation error is predictable from the news shock. Survey-augmented estimation is consistent, since the Phillips curve (3) holds exactly with the subjective expectation on the right-hand side. Appendix A5 states both results in closed form.

Our model also implies that the "fully subjective" Phillips curve regression delivers a negative slope when beliefs are tilted towards supply. Consider a regression of inflation

[8] We use the Michigan Survey of Consumers; Section 3.1 describes the features of the survey used to construct the shock series.

forecasts on forecasts of activity and of inflation further ahead (that is, $\tilde{E}_t\pi_{t+1}$ regressed on $\tilde{E}_t\pi_{t+2}, \tilde{E}_t x_{t+1}$). Since $\tilde{E}_t\pi_{t+2} = 0$ and both remaining variables are proportional to z_t, the estimated slope converges to

$$\hat{\kappa}^{SUBJ} \to \frac{\varsigma\tilde{K}_u + \kappa\tilde{K}_v}{\tilde{K}_v - \phi_\pi\tilde{K}_u}, \quad \hat{\kappa}^{SUBJ} < 0 \text{ whenever } \tilde{K}_u/\tilde{K}_v > 1/\phi_\pi. \tag{21}$$

Although agents' forecasts are internally consistent, so that agents understand the Phillips curve, the fully subjective estimation of the Phillips curve fails, and in fact delivers a negative estimate whenever beliefs are sufficiently supply-tilted. The reason is an omitted variable. Subjective forecasts do satisfy the subjective Phillips curve, $\tilde{E}_t\pi_{t+1} = \beta\tilde{E}_t\pi_{t+2} + \kappa\tilde{E}_t x_{t+1} + \tilde{E}_t u_{t+1}$, but the expected supply shock $\tilde{E}_t u_{t+1} = \tilde{K}_u z_t$ is not observed and enters the regression error. Because expected supply is driven by the same news as $\tilde{E}_t x_{t+1}$, and a supply-tilted interpretation lowers expected output when it raises the expected supply shock, the error is negatively correlated with the regressor.

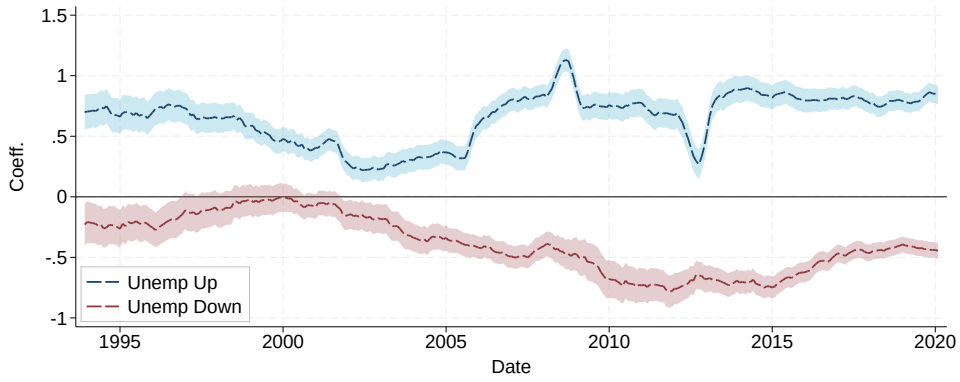


Figure 2: Rolling Subjective Phillips Curve Estimates

Notes: Estimates of $\tilde{E}_{i,t}\pi_{t+1} = \hat{\alpha} + \hat{\beta}^{SUBJ}\tilde{E}_{i,t}\pi_{t+5} + \hat{\kappa}^{SUBJ+}\tilde{E}_{i,t}U^{+}_{t+1} + \hat{\kappa}^{SUBJ-}\tilde{E}_{i,t}U^{-}_{t+1} + \varepsilon^{SUBJ}_{i,t+1}$: subjective one-year-ahead inflation expectations on five-year-ahead expectations and indicators for expecting unemployment to increase or decrease. MSC micro data; four-year rolling windows ending between December 1993 and February 2020 (the five-year question is asked from 1990), pooled across households, HC1 standard errors. 90% confidence intervals. The model counterpart is the negative subjective slope (21).

Figure 2 reports the empirical counterpart in MSC micro data: four-year rolling regressions of one-year inflation expectations on five-year inflation expectations and on indicators for expecting unemployment to rise and to fall (the same survey responses $\tilde{E}_{i,t}U^{+}_{t+1}$ and $\tilde{E}_{i,t}U^{-}_{t+1}$ from Table 1). The coefficient on $\tilde{E}_{i,t}U^{+}_{t+1}$ is positive in every window, and the coefficient on $\tilde{E}_{i,t}U^{-}_{t+1}$ is negative in every window, in line with the model prediction (21), since expecting higher unemployment means expecting lower activity.

3 Measuring News Shocks

This section constructs the empirical counterpart of an inflationary news shock. Testing the model requires an unanticipated revision in household beliefs which is driven by information rather than by contemporaneous fundamentals. We utilize the daily timing of survey interviews immediately before and after scheduled CPI releases to construct an inflation expectations shock. We show that the shock is unpredictable from current and lagged macroeconomic conditions and from the reaction of financial markets to the same releases, and then confirm the model's first prediction: an information event which raises inflation expectations lowers expectations of economic activity.

3.1 Data

We measure household expectations using the Michigan Survey of Consumers (MSC). We obtain the exact interview date of each respondent from the Inter-University Consortium for Political and Social Research (ICPSR) for January 1982 through December 2018, and directly from the MSC for January 2019 through April 2023. The MSC collects approximately 500 responses each month, with interviews spread roughly uniformly over the month. The MSC elicits a numerical forecast of inflation over the next year; most of its other questions are categorical. For instance, respondents are asked whether unemployment will increase, stay the same, or decrease over the next year, and analogous questions cover business conditions and household income. We use the one-year inflation forecast to construct the shock, and the categorical forward-looking questions to measure the comovement of forecasts.

We combine the survey with the dates of CPI releases from the release calendar of the Federal Reserve Bank of St. Louis (ALFRED, Consumer Price Index release), keeping one scheduled release per month.[9] Our estimation sample runs from January 1982 to February 2020; our baseline results exclude the COVID period, but results are similar when the sample is extended through 2023 (Appendix Figure E6). We also use standard monthly macroeconomic series: CPI inflation, the unemployment rate, industrial production, the federal funds rate, and the price of West Texas Intermediate crude oil. We use daily ten-year Treasury yields and daily oil prices to compare household belief revisions with the reaction of financial markets to the same releases. All macroeconomic and financial series are from FRED.

[9]The annual seasonal-factor updates and the September-2000 correction are not releases and are excluded, and the December-1995 CPI, delayed by the government shutdown to February 1, 1996, is treated as the January-1996 release.

3.2 News Shock Construction

Our news shock is constructed using variation in survey-based measures of household inflation expectations around CPI releases. We focus on CPI releases for three reasons. First, releases are recurrent and scheduled: the dates are set in advance, so the timing of the event is exogenous to economic conditions in the window. Second, households pay attention to these releases, at least indirectly through news coverage (see also Binder et al. 2022). Third, the information is directly about inflation, a forecast which the survey elicits numerically.

The mapping from CPI releases to the model requires some care. CPI releases are informative about *realized* prices, while the signal in the model concerns future shocks (the model of course also assumes all date-t variables are perfectly observed, but this is for tractability). We therefore interpret learning from CPI releases broadly. The release and the surrounding news coverage shift household beliefs about the path of the economy, and the signal z_t in the model represents that composite information event. Our empirical design only requires that the belief revision within the window is driven by information.

We construct our shocks as follows. Let t index CPI release months. Around each release we form two groups of respondents: those interviewed in the w days before the release and those interviewed in the w days after it. Let $\tilde{E}_t^{\text{pre}}\pi_{t+1}$ and $\tilde{E}_t^{\text{post}}\pi_{t+1}$ denote the average one-year inflation forecast in each group. Our empirical *news shock* is the difference,

$$\Delta\tilde{E}_t\pi_{t+1} \equiv \tilde{E}_t^{\text{post}}\pi_{t+1} - \tilde{E}_t^{\text{pre}}\pi_{t+1}. \tag{22}$$

The baseline window is $w = 5$ days. The release day itself is excluded from both groups, each interview is assigned to its nearest release, and means are unweighted. With $w = 5$, the median number of respondents in each group is 63.

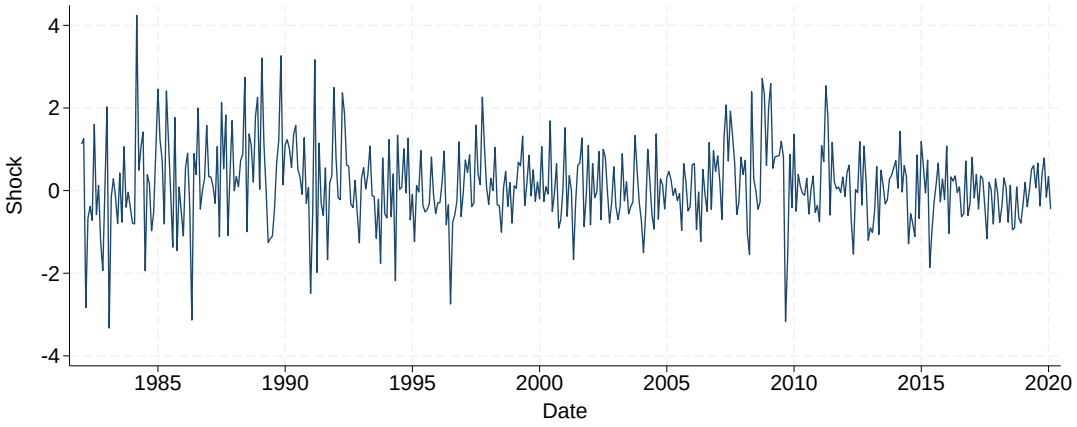


Figure 3: News Shock Time Series

Notes: News shock time series calculated by taking the difference in average expected inflation in the 5 days after and the 5 days before CPI releases.

Figure 3 plots the series from January 1982 to February 2020. The mean and median are near zero (0.14 and 0.08 percentage points, respectively), and the standard deviation is roughly one percentage point (0.99). The volatility of the shock varies over the sample: it is higher in the 1980s (a standard deviation of 1.36) than from 1990 onward (0.87 through 2007 and 0.88 thereafter).

3.3 Identification

Interview timing within the month is set by the survey's fieldwork rather than by respondents' beliefs, so we treat respondents interviewed just before and just after the release as comparable. Our key identifying assumption is then that, within the event window, the CPI release is the only systematic aggregate event which causes households to update beliefs. Under this assumption, the difference (22) isolates the revision in mean household beliefs caused by the information revealed at the release. The shock is therefore purged of everything known before the release, and of any monthly variation common to beliefs before and after the release. Two further assumptions are not required to test the model's predictions. First, the CPI release itself does not need to be an exogenous shock; indeed the CPI release reports past prices, which themselves are equilibrium objects. The only variation we require is the revision in beliefs. Second, households do not need to read the CPI release directly; transmission through the media is part of the information event. Under this identifying assumption, our shock series can be used as the regressor in a local projections (Jordà, 2005) specification.

However, the identifying assumption would fail if belief revisions in event windows reflected omitted contemporaneous shocks which independently move future economic activity. We first assess this concern directly. Table 2 shows that the shock is unpredictable from contemporaneous and lagged inflation and unemployment, and is essentially uncorrelated with daily changes in ten-year Treasury yields and oil prices around the same releases. Across specifications the R^2 is at most 0.016, and we cannot reject that all coefficients are zero (p-values between 0.46 and 0.90).

Of course, the results of Table 2 are consistent with our identifying assumption, but we cannot categorically rule out all potential confounding factors which might simultaneously impact inflation forecast revisions and economic activity. Thus, our inference on the aggregate responses is based on placebo randomization: our estimates are compared to estimates based on shocks constructed according to (22) but around randomly selected placebo dates. If the source of variation driving our forecast revisions reflects confounding factors correlated with economic activity, we would expect such a relationship to hold at any given date. That is, confounders which are not specific to CPI release windows would produce the same relationship around placebo dates, and placebo estimates would then be indistinguishable from our baseline estimates. We discuss our randomization inference

Table 2: News Shock (Un)predictability

	(1)	(2)	(3)	(4)	(5)
$\Delta y_t^{(10)}$	0.396		0.176		0.327
	(0.575)		(0.539)		(0.522)
$\Delta p_t^{(OIL)}$		0.010	0.009		0.015
		(0.020)	(0.020)		(0.021)
U_t				0.037	0.359
				(0.343)	(0.362)
π_t				-0.053	-0.071
				(0.142)	(0.147)
U_{t-1}				-0.051	-0.338
				(0.336)	(0.355)
π_{t-1}				0.085	0.139
				(0.147)	(0.152)
Obs.	435	396	394	452	394
R^2	0.002	0.002	0.002	0.003	0.016
p-value	0.491	0.609	0.852	0.899	0.458

Notes: The news shock (difference in average expected inflation in the 5 days after and before CPI releases) is regressed on the change in ten-year Treasury yields between the day after and the day before the CPI release ($\Delta y_t^{(10)}$), West Texas Intermediate oil price inflation between the day after and the day before the CPI release ($\Delta p_t^{(OIL)}$), contemporaneous-month unemployment (U_t), one-month lagged unemployment (U_{t-1}), contemporaneous-month inflation (π_t), and/or lagged-month inflation (π_{t-1}). The final row reports the p-value of the test that all slope coefficients are zero.

procedure in detail in Section 4.

3.4 Validation: Conditional Forecast Responses

The model predicts that when beliefs are sufficiently supply-tilted ($\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$), a news shock which raises inflation expectations *lowers* expectations of economic activity (Proposition 1). Table 3 tests this prediction using the change in other MSC forecasts across the same event windows. Following a one percentage point shock to inflation expectations, the share of households who expect unemployment to rise increases by 1.6 percentage points, and the share who expect unemployment to fall decreases by 0.8 percentage points. Consumer sentiment, measured as the first component of the categorical forward- and backward-looking MSC questions excluding inflation (Kamdar, 2019), declines. The p-values of the three responses are <0.001, 0.033, and <0.001.

The results of Table 3 are the *conditional* analogue of the unconditional forecast comovement of Section 2.5. This is a stronger test than Table 1, since it isolates the variation caused by an identified information event, rather than pooling all sources of variation in forecasts. The same pattern holds across the other forward-looking questions of the MSC; we report the sentiment index rather than each question separately.

By Propositions 1 and 3(i), and since in the model forecasts move only with news, the

Table 3: News Shocks and Expectations

	$\Delta\tilde{E}_t U^+_{t+1}$	$\Delta\tilde{E}_t U^-_{t+1}$	$\Delta\tilde{E}_t s_{t+1}$
$\Delta\tilde{E}_t \pi_{t+1}$	1.564***	-0.836**	-0.044***
	(0.410)	(0.391)	(0.008)
Obs.	452	452	452
R^2	0.036	0.014	0.071

Notes: The changes in the percent of households expecting unemployment to rise ($\Delta\tilde{E}_t U^+_{t+1}$) and fall ($\Delta\tilde{E}_t U^-_{t+1}$) across the event window are regressed on the news shock ($\Delta\tilde{E}_t \pi_{t+1}$) in columns (1) and (2). In column (3) the change in average sentiment ($\Delta\tilde{E}_t s_{t+1}$) is regressed on the news shock. Sentiment is calculated as the fitted first component of the categorical forward- and backward-looking questions excluding inflation; respondents who skip some questions are scored on the ones they answer. HC1 standard errors in parentheses. Sample: release months from January 1982 to February 2020.

negative conditional comovement of forecasts implies $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi > T_x$, so the survey moment alone places the economy in the region in which news shocks are stagflationary in *realized* outcomes. The next section tests this implication directly.

4 The Aggregate Effects of News Shocks

This section tests the model's central prediction for aggregate dynamics: news which raises inflation expectations increases realized inflation and lowers activity. We estimate the responses of inflation, unemployment, and industrial production to the news shock by local projection. Inflation and unemployment both rise and industrial production falls, and shocks constructed around placebo dates show that this pattern is specific to CPI release windows. Finally, a quantitative version of the model with additional sources of persistence reproduces the magnitudes and the dynamics of the responses.

4.1 Local Projections and Randomization Inference

Under the identifying assumption of Section 3.3, our series (22) is an identified news shock, so we estimate the aggregate responses by local projection (Jordà, 2005):

$$y_{t+h} = \alpha_h + \beta_h \Delta\tilde{E}_t \pi_{t+1} + \gamma_h^\top X_t + \varepsilon_{t+h}, \tag{23}$$

where y_{t+h} is the realization of an aggregate variable h months after the shock (either year-over-year CPI inflation, the unemployment rate, or log industrial production). In our baseline regression, X_t contains four lags of inflation, unemployment, the federal funds rate, oil price inflation, industrial production, and the shock itself. Results are robust to alternative controls (including no controls).

We estimate (23) for horizons up to 36 months. The estimation sample runs from January 1982 to February 2020; our main sample drops COVID and the ensuing high

inflation period, but we explore alternative sample periods as well. For the specifications which include lags of the shock among the controls, we report heteroskedasticity-robust standard errors (Montiel Olea and Plagborg-Møller, 2021); for specifications with no lags of the shock as controls, we report Newey–West standard errors with truncation lag equal to $h + 1$.

The coefficient β_h is the response to a shock which moves average one-year inflation expectations within the release window. We report all responses per standard deviation of the shock (but note that the standard deviation of the shock is almost precisely one percentage point, so the reported magnitudes correspond closely to the non-normalized coefficients). With roughly 63 respondents on each side of the window, our shock series contains finite-sample noise. Since this error is orthogonal to the true revision, it attenuates β_h toward zero; our estimates are therefore conservative.

We assess whether the responses are specific to CPI release windows via placebo analysis: randomization inference over pseudo-release dates. For each release month we draw one pseudo-event date uniformly from the admissible days of that calendar month, reconstruct the shock around the pseudo dates exactly as in Section 3.2, and re-estimate (23). A day d is admissible if $|d - r| \geq w + 1$ for every CPI release date r. The minimum gap between CPI releases in our sample is 16 days, so every month retains admissible days. We repeat the full construction for 1000 draws.

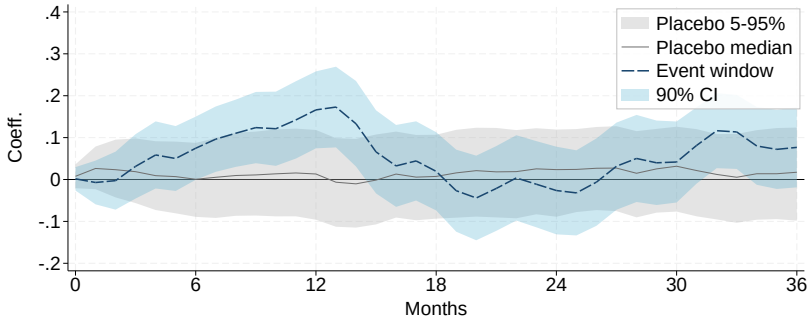
We exclude only CPI releases from the pseudo dates. Pseudo dates are "generic" days on which no CPI release occurs, but may include the release of other macroeconomic news. The placebo distribution therefore reflects the "typical" mix of shocks facing households. This approach is also conservative: pseudo windows which contain other economic releases may contain genuine belief revisions, which widens the placebo distribution against which we compare our estimates.

4.2 Results

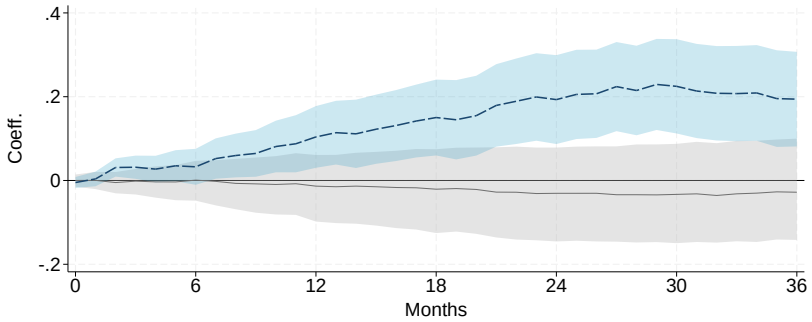
Figure 4 reports the responses of inflation, unemployment, and industrial production, each against the distribution of placebo estimates; panels (a) and (b) repeat Figure 1 with the placebo bands added. Following a news shock, realized inflation increases over the course of a year and reaches a peak effect of roughly 0.17 percentage points at horizon 13; the response then declines and becomes insignificant. Realized unemployment rises steadily: after one year it is roughly 0.10 percentage points higher, and after two years roughly 0.19 percentage points higher, where it remains through the three-year horizon; the 90 percent confidence interval excludes zero from month 7 onward. Industrial production declines gradually, by 0.67 percent after two years. The maximal response occurs at horizon 26; confidence intervals exclude zero from month 15 to month 33. The timing of the two activity measures agrees: activity is unchanged for roughly half a year after the shock

and deteriorates thereafter.

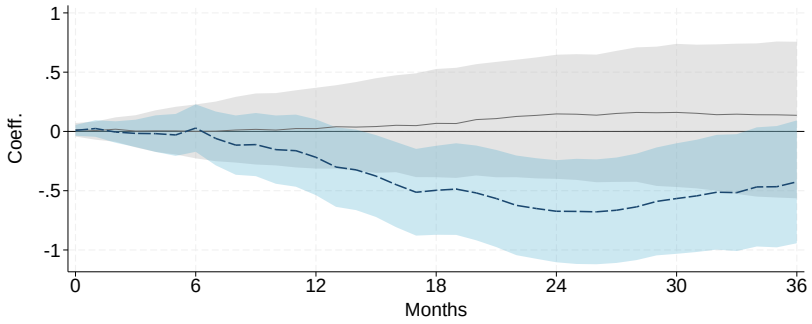
The placebo results show that the pattern is specific to CPI release windows: placebo bands include zero at every horizon for every outcome, and placebo medians are positive for inflation and industrial production but negative for unemployment, the pattern of demand-driven fluctuations. The placebo medians are also much smaller than the baseline responses at the headline horizons.



(a) Inflation Response



(b) Unemployment Response



(c) Industrial Production Response

Figure 4: Aggregate Responses and Randomization Inference

Notes: Local projections of year-over-year CPI inflation (panel a), the unemployment rate (panel b), and log industrial production, times 100 and linearly detrended over the sample (panel c), on the news shock (dashed, with 90% bands), against the 5–95% range and median of the same estimates across 1,000 placebo replications with pseudo release dates (shaded, gray). Pseudo dates are drawn uniformly over the admissible days of each release month (at least $w + 1$ days from every true CPI release); the full shock construction is rerun on each draw. Coefficients per standard deviation of the shock. Four lags of inflation, unemployment, the fed funds rate, oil price inflation, industrial production, and the shock are included as controls.

We formally design our placebo-based statistical test taking into account that our

shock is an *increase* in inflation expectations; that is, the shock is signed, so the tests are one-sided (upper tail for inflation and unemployment, lower tail for industrial production). By Proposition 1, any news which raises inflation forecasts raises realized inflation. Hence, the unemployment response is necessary to distinguish supply from demand interpretations. Thus, our placebo-based test is the marginal randomization p-value of the unemployment response, which is $p = 0.001$ at $h = 24$ and $p = 0.010$ familywise across all horizons. Similarly, the decline in industrial production gives $p = 0.006$ at $h = 24$ and $p = 0.041$ familywise across all horizons.

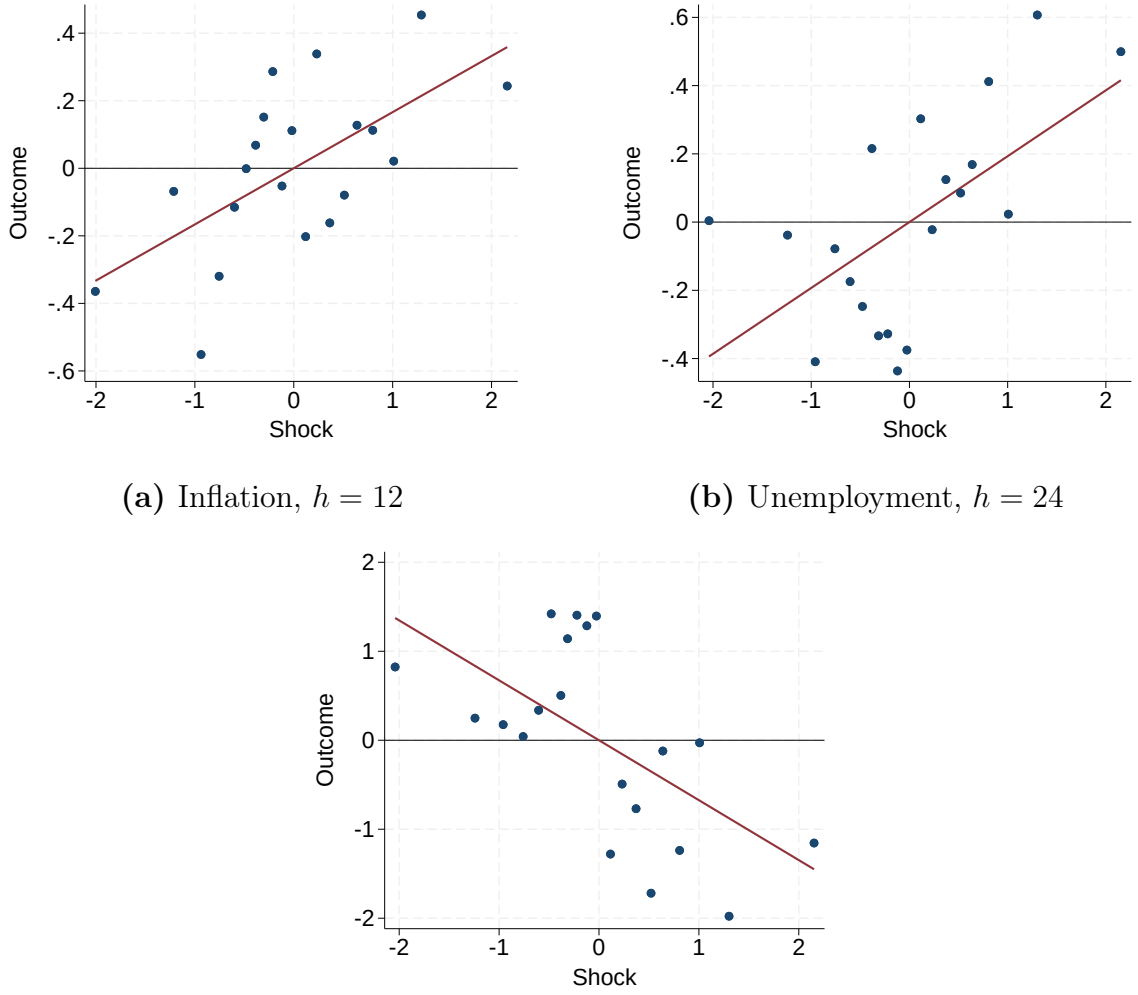
We also find that the inflation response at $h = 12$ is larger than found in the placebo sample ($p = 0.010$). Considering inflation and unemployment jointly, only about 20 percent of pseudo draws produce positive responses of both at a given horizon, and no draws exceed the true pair of responses at $h = 12$ and $h = 24$.

Our results are therefore consistent with the pattern in the model of a signal interpreted as supply news: realized inflation and activity move in opposite directions. Alternative explanations would imply comovements which the data reject. For instance, if the revision in expectations reflected news about future demand under FIRE, unemployment would fall. If higher expected inflation simply reflected "pessimism" about future economic activity driven largely by demand shocks, realized inflation would fall. Our model's supply-side interpretation predicts the joint rise of inflation and unemployment after an identified news shock which increases inflation forecasts.

The placebo distribution is itself informative about the model. Across placebo draws the estimated inflation and unemployment responses are mostly negatively correlated (the correlation across draws ranges from -0.27 to 0.07 over horizons), and the placebo medians are slightly positive for inflation (0.013 at $h = 12$) and slightly negative for unemployment (-0.031 at $h = 24$). This pattern is expected if "generic" event windows are (weakly) demand-driven. CPI release windows deliver the opposite comovement, not merely larger coefficients, which further reinforces the interpretation of our results through the lens of the model.

Robustness. Our results are robust along several dimensions, reported in Appendix E. First, we consider alternative event window sizes: shocks based on three-day and seven-day windows give similar responses (Figures E2 and E3). Second, the responses do not depend on the control set: the specification without controls delivers similar responses, and placebo-based randomization inference delivers the same conclusion (Figure E4). Third, controlling for the content of the CPI release itself does not change our findings: additionally including lagged core, energy, and food CPI inflation leaves the responses intact (Figure E1). Finally, the responses are not driven by the high-inflation years at the start of the sample (Figure E5), and including the COVID period does not meaningfully change our estimates (Figure E6). More generally, the responses reflect a broad pattern rather

than a few extreme release months. Figure 5 reports binned scatter plots for each outcome at the relevant headline horizon against the news shock, both residualized on the baseline controls. The binned scatter plots show that our baseline findings hold throughout the distribution of the shock, and not only at extreme values.



(a) Inflation, $h = 12$

(b) Unemployment, $h = 24$

(c) Industrial Production, $h = 24$

Figure 5: Binned Scatter Plots of the Baseline Responses

Notes: The local projections of Figure 4 at horizon h, shown as binned scatter plots. The horizontal axis is the news shock and the vertical axis the outcome h months later (year-over-year CPI inflation and the unemployment rate in percentage points; industrial production in 100 times log deviations from a linear trend), each residualized on the baseline controls over the estimation sample, January 1982 to February 2020. The shock is in units of its standard deviation. Dots are means within 20 equal-sized bins of the residualized shock; the line has the slope of the local-projection coefficient per standard deviation of the shock.

4.3 Quantitative Model Comparison

The analytical model of Section 2 delivers the signs of the responses to news. However, our stylized setting thus far implies that the reaction to news occurs on impact. Our empirical results show that the responses build over one to two years. We now ask

whether a more realistic quantitative version of the model delivers the magnitudes and the dynamics. We extend the model along four dimensions; Appendix D contains the full model and the details of the estimation. First, the shocks follow AR(1) processes. Second, households have external habits in consumption. Third, the policy rule features interest-rate smoothing. Fourth, news arrives twelve months ahead, the horizon of the measured inflation expectations. The longer anticipation window, along with the additional sources of persistence, generates hump-shaped responses to news without assuming that the news itself follows such a pattern.

Beliefs are Bayesian with respect to one public signal per month about the innovations twelve months ahead. We impose pure supply news: the signal loads only on the supply innovation, so $\tilde{K}_v = 0$, as at the robust corner of Section 2.4. With the signal's loading on the supply innovation normalized, the only free belief parameter is the noise in the signal, which we estimate.

Table 4: Calibration and Estimation of the Quantitative Model

Parameter	Value	Target/Source
β	0.9975	3% annual rate
ς, φ	1, 1	log utility, unit Frisch elasticity
$\bar{\epsilon}$	10	11% markup
b	0.9	habit, 0.7 per quarter
θ	0.985	Phillips-curve slope 0.003 per month
ϕ_π, ϕ_x	2, 0	inflation-only rule
ρ_i	0.95	rate smoothing
h	12	one-year forecast horizon
$\tilde{K}_v, \tilde{K}_u$	0, 0.00107	pure supply news
σ_η^2	0.585	inflation LP, $h = 12$
ρ_u	0.945	unemployment LP, $h = 24$
ρ_v	0.994	s.d., unemployment gap
σ_u, σ_v	0.00169, 0.0141	s.d., yoy inflation; corr., inflation and gap

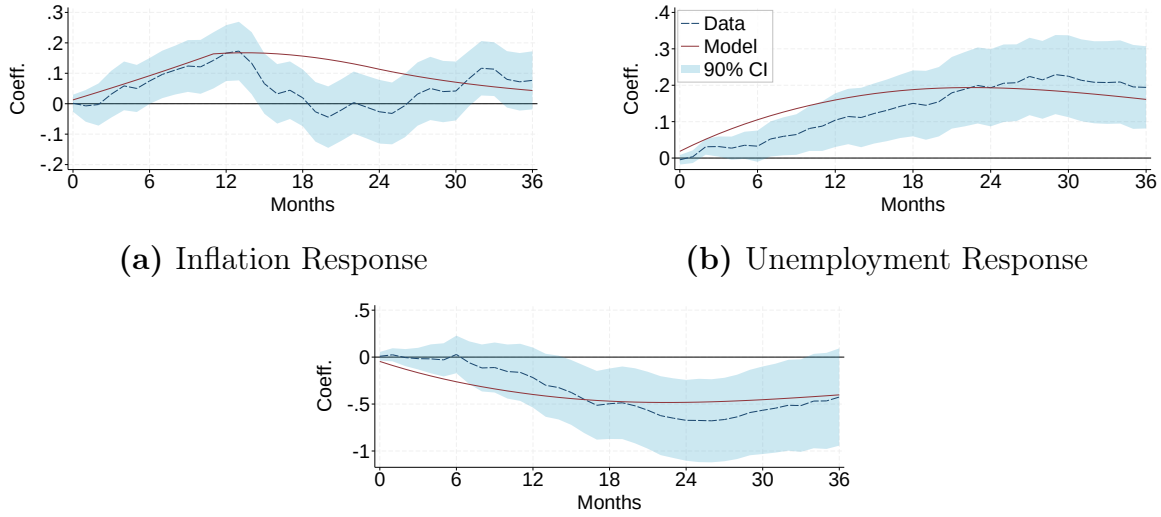
Notes: Monthly model. The first nine rows are fixed at standard values; the gains impose pure supply news. The final rows show the five estimated parameters, exactly identified by the five moments in the third column: the responses of year-over-year inflation at $h = 12$ and of unemployment at $h = 24$ to the news shock, per standard deviation of the shock (Figure 1), and three unconditional moments over 1982m1–2020m2, the standard deviation of year-over-year CPI inflation (1.38), the standard deviation of the unemployment gap relative to the CBO natural rate (1.61), and the correlation of year-over-year inflation with the unemployment gap (−0.11). Unemployment is mapped from the output gap with an Okun coefficient of −0.5.

Table 4 reports the parameters. All but five take standard monthly values, with a Phillips-curve slope of 0.009 per quarter (Hazell et al., 2022). The remaining parameters are estimated: the volatilities and persistences of the supply and demand shocks, and the noise in the signal. They are exactly identified by five moments: the responses of inflation at one year and of unemployment at two years to the news shock, and the standard deviations of year-over-year inflation and of the unemployment gap and their correlation over the sample. We map the model's output gap into unemployment with an Okun

coefficient of -0.5 and into industrial production using the empirically estimated slope of the industrial-production gap on the unemployment gap (-2.5); industrial production is not targeted in the estimation.

The model counterparts of the responses are population projections on the news-driven revision of year-ahead inflation forecasts. To ensure comparable interpretation of the model-implied and empirical estimates, we normalize the belief revision by its standard deviation. The two conditional responses pin down the size of the news-driven revision and the persistence of the supply factor; the unconditional moments pin down the other parameters. The estimates imply demand-dominated fundamentals, with variance of supply innovations only 1.4 percent of the total. The supply-side reading of news thus coexists with demand-dominated fundamentals.

As a check, the model's own rational-inattention solution at the estimated economy devotes 66 percent of its attention to supply news. The relative gain is 1.4 per standard deviation of each innovation, and thus supply-tilted even under the less extreme assumption of rational inattention compared to robustness (Appendix D4).



(a) Inflation Response

(b) Unemployment Response

(c) Industrial Production Response

Figure 6: News Shock Responses: Model and Data

Notes: Empirical local projections with 90% confidence intervals (Figure 4) against the model-implied responses to a news shock at the estimates of Table 4. Coefficients are per standard deviation of the shock in the data and per standard deviation of the news-driven revision of year-ahead inflation forecasts in the model.

Figure 6 compares the model-implied responses with the data at every horizon. The estimation targets only two points, the inflation response at one year and the unemployment response at two years, and leaves the shapes of the responses and the entire response of industrial production free. The model reproduces the main features of our baseline empirical results. Inflation rises over the first year, and activity deteriorates gradually: unemployment builds over the first two years and remains elevated into the third (similarly, industrial production falls and then begins to recover). From the end of

the first year onward, the model's unemployment and industrial production responses lie within the confidence bands of the data. The model's activity responses begin somewhat earlier than in the data, and the inflation response declines more slowly after the first year, since the anticipated shock implies a persistent path of the supply factor once the innovation actually arrives.

5 Implications for Monetary Policy

This section now turns to the design of monetary policy. Supply-side beliefs change the problem of the central bank in two ways. First, news acts as a cost-push shock: no simple Taylor rule stabilizes both inflation and output in response to news (even when the signal is pure noise). Second, the policy rule determines private agents' value of learning about future innovations because it influences the aggregate effects of supply and demand shocks. When information collection is endogenous, optimal policy therefore has a belief-shaping motive in addition to the usual motive to manage expectations. By credibly ensuring that inflation is sufficiently stabilized in response to both demand and supply shocks, the central bank reduces the incentive to collect information and so anchors expectations.

Because policy changes the value of information to private agents, the analysis requires a model of how agents form beliefs. We proceed in two steps. We first take the gains as given, and ask how news affects outcomes under alternative policy rules and what the news is worth to rationally inattentive agents (Section 2.4) under each rule. We then let the gains respond to the rule, first for simple rules and then for an optimizing central bank which takes the response of attention into account.

5.1 Policy Rules and the Interpretation of News

We first consider a more general set of simple policy rules of the form

$$i_t = \phi_\pi \pi_t + \phi_x x_t + \phi_z z_t. \tag{24}$$

Relative to the baseline rule, we allow the central bank to respond to the output gap and directly to news. The following proposition shows that no rule of this type with finite coefficients removes the trade-off created by news; that is, news itself leads to a failure of divine coincidence. Appendix C contains the derivations.

Proposition 4 (Divine coincidence failure with news)**.** *Suppose $\tilde{K}_u, \tilde{K}_v \geq 0$, not both zero, and consider any rule (24) with finite coefficients $\phi_\pi, \phi_x \geq 0$. The impact responses to news satisfy*

$$\frac{\partial \pi_t}{\partial z_t} - \kappa \frac{\partial x_t}{\partial z_t} = \beta \frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t}, \quad \frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} = \chi_g \left[(1 + \varsigma^{-1} \phi_x) \tilde{K}_u + \kappa \varsigma^{-1} \tilde{K}_v \right] > 0,$$

where $\chi_g \equiv [1 + \varsigma^{-1}(\phi_x + \kappa\phi_\pi)]^{-1}$*, and the revision in expected inflation does not depend on* ϕ_z*. Hence news implies a strict trade-off between inflation and output: any choice of* ϕ_z *which achieves* $\frac{\partial x_t}{\partial z_t} = 0$ *leaves* $\frac{\partial \pi_t}{\partial z_t} = \beta \frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} > 0$*. This holds even when* $z_t = \eta_t$ *is pure noise, and even when fundamentals are purely demand-driven.*

Intuitively, the belief revision enters the Phillips curve directly through expected inflation, and no contemporaneous interest rate response offsets the revision. A direct response to z_t can neutralize the output response, but the revision in expected inflation does not depend on ϕ_z and passes into current inflation. Fundamentally, shifts in subjective inflation forecasts present the policymaker with the same trade-off as a cost-push shock, with one difference: the size of this cost-push shock depends on policy. This provides an additional motive for anchoring inflation expectations.

Policy Mandates. Taking the interpretation of news as given, no rule of the form (24) with finite coefficients fully anchors inflation expectations. Studying two strict mandates allows for insights into the trade-offs the central bank faces in the presence of news. Under strict inflation targeting the central bank keeps $\pi_t = 0$ in every period (the limit $\phi_\pi \to \infty$); under strict output gap targeting the central bank keeps $x_t = 0$ in every period (the limit $\phi_x \to \infty$). Proposition C1 in Appendix C1 characterizes outcomes and the value of news under each mandate. The value of news changes because each mandate is credible: the targeted variable is not only kept at its steady state value today, but agents *expect* the targeted variable to remain at its steady state value next period. That is, strict inflation targeting implies $\pi_t = \tilde{E}_t \pi_{t+1} = 0$, and strict gap targeting implies $x_t = \tilde{E}_t x_{t+1} = 0$.

Strict inflation targeting therefore implies that the Phillips curve pins down the output gap today: $x_t = -u_t/\kappa$. Demand shocks are neutralized, but supply shocks pass fully into output. Therefore, gap forecasts respond to news (whenever $\tilde{K}_u \neq 0$): $\tilde{E}_t x_{t+1} = -\tilde{K}_u z_t/\kappa$. However, note that while beliefs still move, under strict inflation targeting *realized* outcomes today do not respond to news. The consumption-smoothing motive generated by fluctuations in output forecasts is fully offset by monetary policy.

In contrast, strict gap targeting implies $\pi_t = u_t + \beta \tilde{E}_t \pi_{t+1}$. Contemporaneous demand shocks are once again neutralized, but inflation reacts to both contemporaneous supply shocks and to subjective inflation forecasts. Unlike inflation targeting, both beliefs and outcomes respond to news (whenever $\tilde{K}_u \neq 0$): $\tilde{E}_t \pi_{t+1} = \tilde{K}_u z_t$ and $\pi_t = u_t + \beta \tilde{K}_u z_t$. Unlike the consumption-smoothing motive under strict inflation targeting, monetary policy does not offset the forward-looking price-setting motive.

The two mandates of course differ in how outcomes respond to contemporaneous shocks. But more importantly, the mandates have different implications for the private value of information about future shocks. The reason lies in what the two mandates imply for forward-looking decisions. Strict gap targeting naturally neutralizes the consumption-smoothing motive, but price setters still react to their beliefs regarding inflation tomor-

row. Strict inflation targeting inherently removes all forward-looking motives of price setters. But the monetary policy rule which implements strict inflation targeting simultaneously neutralizes the consumption-smoothing motive as well. Thus, strict inflation targeting shuts down *all* forward-looking decisions: since no information about future shocks changes any agent's action today, the value of information (the attention problem (18)) is zero. Strict gap targeting also removes the value of information regarding future demand shocks. But under strict gap targeting, information about future supply remains valuable since inflation reacts today to beliefs about future supply.

Policy Shapes the Interpretation of News. The same logic applies to rules with finite coefficients. Section 2.4 derives the gains under a fixed rule, but agents form beliefs based on the equilibrium consequences of shocks (under rational inattention, robustness, or any cost-ranked belief problem of Appendix B3), and the rule in part determines those consequences. Consider the response to the output gap. A rule with finite ϕ_x only partly neutralizes demand shocks, so beliefs about demand retain some value. As ϕ_x rises, however, demand shocks are increasingly offset while the consequences of supply shocks are only redistributed between inflation and output, so on both margins of the attention problem beliefs about supply gain value relative to beliefs about demand. Proposition C2 in Appendix C2 states the consequences under rational inattention. News is read as stagflationary if and only if the responses of inflation and output to anticipated shocks are negatively correlated under the prior, $b_\pi^\top \Sigma b_x < 0$, and a stronger response to the output gap can only make this correlation more negative. The set of economies in which news is stagflationary therefore expands with ϕ_x; a stagflationary reading is never undone by a stronger response to the gap; and for any composition of shocks, news becomes stagflationary once ϕ_x is large enough. Strict gap targeting (Proposition C1(ii)) is the limit of this process. Policy also lowers the target itself: holding the interpretation fixed, the stagflation threshold T_x falls in both ϕ_x and ϕ_π, so any given reading of news is stagflationary under a wider range of rules.[10] Under robustness, the relative damage of supply shocks rises with both coefficients, so a more aggressive response to either inflation or the output gap can switch the robust worst case from the demand corner to the supply corner, but never the reverse.

Therefore, active policy reinforces supply-side beliefs, and with them the stagflationary response to news. Policy is unable to fully inoculate the economy against news; and policy itself is partially responsible for the source of news-based stagflationary shocks. What a rule changes is the direction of attention. Lowering the stakes of attending to news at all is the margin on which the optimal plan of Section 5.2 operates.

[10] The response to inflation is the mixed case. A more hawkish rule shifts the consequences of supply shocks from inflation into output, so it tilts the output margin of the attention problem toward supply but the inflation margin away from it, and its net effect on the interpretation has no general sign (Appendix C2).

5.2 Optimal Policy and the Belief-Shaping Motive

We now replace the rule with a planner. The planner's problem is the standard linear-quadratic problem with one additional constraint: the attention best response of private agents. We focus on the case of rational inattention with an information cost of ϑ. We evaluate every plan according to the welfare loss: the unconditional expected loss (17) plus the cost of the information which private agents acquire,

$$\mathcal{W} \equiv \mathbb{E}\left[\varpi_\pi \pi_t^2 + \varpi_x x_t^2\right] + \vartheta \mathcal{I}, \tag{25}$$

where $\mathcal{I}$ is the mutual information between the signal and next period's innovations. The first term is twice the expected period loss from macroeconomic volatility, which we call the *macro loss* and write $2L$. In order to understand the trade-offs faced by the central bank in our model, we consider a hierarchy of different planner types who differ only in what they internalize when choosing policy.

Let $c_t = u_t + \beta \tilde{E}_t \pi_{t+1}$ denote the contemporaneous *effective cost-push* term, so that $x_t = (\pi_t - c_t)/\kappa$ by the Phillips curve. We can then write

$$\varpi_\pi \pi_t^2 + \varpi_x x_t^2 = \underbrace{\varpi_\pi \Theta c_t^2}_{\text{cost-push loss}} + \underbrace{\frac{\varpi_x}{\kappa^2 \Theta}\left(\pi_t - \Theta c_t\right)^2}_{\text{targeting-rule deviations}}, \tag{26}$$

where $\Theta \equiv \varpi_x/(\varpi_x + \kappa^2 \varpi_\pi) \in (0,1)$ is the static trade-off weight: the share of an effective cost-push shock which the targeting rule passes into inflation. The second term is proportional to the squared residual of the familiar targeting rule $\kappa \varpi_\pi \pi_t + \varpi_x x_t = 0$. For each planner, the optimality conditions below determine when a deviation from the targeting rule is worthwhile. The answer depends on the planner's degree of commitment, and in particular on whether the planner's conduct affects only agents' expectations or also how much attention agents pay to news.

Policy is a stationary linear plan which maps the public state, consisting of the current shocks $\varepsilon_t \equiv (u_t, v_t)^\top$ and the common posterior $\nu_t \equiv (\tilde{E}_t u_{t+1}, \tilde{E}_t v_{t+1})^\top$ about next period's shocks, into outcomes: $\pi_t = a^\top \varepsilon_t + a_\nu^\top \nu_t$, where $a \equiv (a_u, a_v)^\top$, and the interest rate is determined residually by the IS curve. Agents forecast tomorrow's inflation by applying the plan to their beliefs, $\tilde{E}_t \pi_{t+1} = a^\top \nu_t$. The restriction to plans without memory is not innocuous, as it precludes promises about tomorrow's response to today's shock. We relax this assumption below, but our main analytical results are for the memoryless case.

All the planners we consider share two properties (Appendix C3, Lemma C1). First, demand shocks are fully neutralized, $a_v = 0$. Second, given a, every planner applies the targeting rule to the belief-driven cost-push term, and so responds to beliefs in the same way: $a_\nu = \Theta \beta a$. This follows because neither expectations nor attention can be shaped through a_ν: inflation forecasts depend on a alone, and by (26) any departure from the

targeting rule in the response to beliefs adds the same deviation term to the realized loss and to the private value of information. Thus, the realized loss depends on a single coefficient a_u, the anticipated response of inflation to supply shocks. This coefficient matters for two reasons: it converts today's beliefs about supply into today's expected inflation, and its square scales the value to private agents of attending to news. Hence there is an *attention threshold* $\hat{a}$, increasing in the cost of information ϑ: for $a_u \leq \hat{a}$, rationally inattentive agents acquire no information, and the belief channel is inactive.

The planner hierarchy can be understood intuitively through the lens of anchoring inflation expectations. Once demand shocks are neutralized and the response to beliefs follows the targeting rule, the macro loss is the static loss from fundamentals plus a belief block proportional to the variance of subjective inflation forecasts, $\varpi_\pi \Theta \beta^2 \operatorname{Var}(\tilde{E}_t \pi_{t+1})$ (Appendix C3, Lemma C1). Every planner would therefore prefer that subjective inflation forecasts never moved: by Proposition 4, the forecast enters today's effective cost-push term, and the trade-off it creates cannot be undone by the contemporaneous instrument. Since $\tilde{E}_t \pi_{t+1} = a_u \tilde{E}_t u_{t+1}$, forecasts can be anchored through either the promised response of inflation to supply shocks, or the forecast of the supply shock itself. The planners differ in which of the two they recognize as under their control.

We consider four planners who differ only in what they internalize. Under *discretion*, the planner re-optimizes each period and takes both expectations and attention as given. Under *expectations-managing commitment*, the planner commits to a plan and internalizes how expectations respond, but treats the interpretation of news as fixed. Under *belief-shaping commitment*, the planner additionally internalizes that the interpretation of news responds to the plan, but minimizes the macro loss component of (25) only. Finally, under *cost-internalizing belief shaping*, the planner internalizes the same response but minimizes the welfare loss (25), counting the cost of attention as well.[11]

Each planner thus internalizes one more channel than the previous one; we label the planners D, N, S, and C, respectively.[12] The following proposition characterizes and orders the four plans.

Proposition 5 (The planner ladder)**.** *Define* $\psi \equiv \beta^2 \Theta(1-\Theta)$ *and* $a_1 \equiv \Theta/(1+\psi)$*, and let* $a_u^D, a_u^N, a_u^S, a_u^C$ *denote the supply-response coefficients chosen under discretion, expectations management, belief shaping, and cost-internalizing, respectively. Then:*

(i) Discretion: $a_u^D = \Theta$. *The discretionary planner applies the targeting rule to every impulse and treats belief revisions as cost-push shocks. Because the planner neutralizes demand shocks, all remaining attention is directed at supply.*

[11] We do not need cost-internalizing versions of discretion and expectations management. Both planners take the signal design as given, so the cost of information is a constant in their objectives.

[12] The labels are mnemonics for what each planner internalizes, and follow Kamdar and Ray (2026). D is discretion. N is for Nash: the plan and the attention choice are mutual best responses, each taken as given by the other. S is for Stackelberg: the planner moves first and internalizes the response of attention to the plan. C is for cost-internalizing.

(ii) Expectations management: *if* $\hat{a} \geq \Theta$*, agents pay no attention to news and* $a_u^N = \Theta$*. Otherwise* $a_u^N \in (\max(a_1, \hat{a}), \Theta)$ *is the unique fixed point of the planner's best response to given attention and the attention best response to the plan.*

(iii) Belief shaping: $a_u^S = \min\{\Theta, \max(a_1, \hat{a})\}$*, with three regimes. If* $\hat{a} \geq \Theta$ *("self-anchoring"), attention is not worth its cost even under discretion and all planners coincide. If* $a_1 \leq \hat{a} < \Theta$ *("anchoring"), the plan sets* $a_u^S = \hat{a}$*, which shuts down the news channel, noise included. If* $\hat{a} < a_1$ *("partial anchoring"), anchoring would require too large a deviation from the targeting rule, and the planner sets the response at its interior optimum* a_1*.*

(iv) Cost-internalizing belief shaping: $a_u^C = \Theta$ *if* $\hat{a} \geq \Theta$*;* $a_u^C = \hat{a}$ *if* $\Theta/(1+2\psi) \leq \hat{a} < \Theta$*; and otherwise* a_u^C *is the larger root of* $(1+\psi)a^2 - \Theta a + \psi\hat{a}^2 = 0$*, which lies in* $(\Theta/(1+2\psi), a_1)$*. The anchoring regime is wider than under belief shaping, and in the partial-anchoring regime the plan lowers the response below* a_1*.*

(v) Ordering: $a_u^C \leq a_u^S \leq a_u^N \leq a_u^D = \Theta$ *and* $\mathcal{W}^C \leq \mathcal{W}^S \leq \mathcal{W}^N \leq \mathcal{W}^D$*, and the value of commitment decomposes as*

$$\mathcal{W}^D - \mathcal{W}^C = \underbrace{\mathcal{W}^D - \mathcal{W}^N}_{\textit{expectations management}} + \underbrace{\mathcal{W}^N - \mathcal{W}^S}_{\textit{belief shaping}} + \underbrace{\mathcal{W}^S - \mathcal{W}^C}_{\textit{info-cost reduction}}. \tag{27}$$

Part (i) of Proposition 5 is the optimal-policy counterpart of Proposition 4: under optimal discretion, agents continue to read news as supply news. Because tomorrow's discretionary planner neutralizes demand shocks, beliefs about future demand do not move any forecast of outcomes, and the only information worth collecting is about supply. The discretionary planner can do nothing to anchor forecasts: today's subjective inflation forecast depends on what tomorrow's planner is expected to do, and tomorrow's planner applies the targeting rule. Part (ii) is the typical improvement from commitment. The expectations-managing planner can lower subjective inflation forecasts, but only directly by promising that inflation will react less strongly to future supply shocks. Relative to discretion, this lowers the pass-through of supply beliefs to current inflation. But because attention is taken as given, the forecast of the supply shock is, from this planner's perspective, an additional exogenous shock, and the plan moves only the forecasts of endogenous outcomes. This is the problem a planner faces when private information is exogenous, as in the literature on optimal policy with news; the belief revisions $\tilde{E}_t u_{t+1}$ and $\tilde{E}_t v_{t+1}$ are then additional shocks to be accommodated rather than objects of policy. Seen this way, full anchoring requires that inflation not respond to supply shocks at all: $a_u = 0$, which is strict inflation targeting. The planner stops short of strict inflation targeting because the welfare cost of loading the entire supply shock onto output is too high. Expectations management alone therefore never lowers the response of inflation to

supply shocks solely to anchor expectations.

Parts (iii) and (iv) contain the belief-shaping motive, which is specific to endogenous attention. The belief-shaping planners recognize that the forecast of the supply shock is itself a response to the plan, and so can anchor inflation forecasts through the forecast of the supply shock itself as well: a smaller promised response lowers not only the pass-through of a given forecast, but also the attention which produces it. Part (iii) shows that, in an intermediate range of information costs, the belief-shaping planner deliberately anchors expectations. Recall that strict inflation targeting guaranteed the value of news is zero. Our result shows that anchoring is possible without strict inflation targeting (which corresponds to $a_u = 0$). Expectations are anchored when the value of forecasting inflation is less than the cost, which requires only that the promised response to supply shocks be small enough that the value of the best available signal falls below the cost of attention. Only a belief-shaping planner sets the supply response exactly at the attention threshold; no planner who treats attention as given chooses such a plan. Part (iv) shows that counting the cost of information strengthens the belief-shaping motive. The cost-internalizing planner sets the supply response at the attention threshold $\hat{a}$ over a wider range of attention costs, since it also counts what agents spend on information. For low enough information costs, the cost-internalizing planner also anchors expectations only partially, but lowers the response further than the planner of part (iii), for the same reason. In both the anchoring and the partial-anchoring regimes, the belief-shaping planners are more hawkish than the targeting rule: they accept larger output fluctuations in response to supply shocks in exchange for less attention to news.

Part (v) separates the three channels: the usual welfare gain from managing expectations, the gain from shaping attention when only the volatility generated by attention is counted, and the additional gain from counting the cost of attention.

All else equal, the belief-shaping planners decrease the incentive for agents to learn about future shocks because, in the iid economy, anticipating a shock has no allocative value. Agents in general have a private incentive to acquire information about future shocks, but in equilibrium this induces additional macroeconomic volatility. Rationally inattentive agents reduce their forecast errors, but the better agents forecast future supply shocks, the more today's inflation responds to news, so more accurate private forecasts raise the planner's loss (Appendix C3). Through additional inflation stabilization, optimal policy therefore seeks to anchor expectations by reducing the incentive to collect information.

The same logic explains why strict inflation targeting is not optimal, even though this removes the incentive to collect information for any nonzero information cost. To see why, suppose we start from the strict inflation target, where the private benefit of collecting information is strictly less than the cost. Tolerating a small positive response of inflation to supply shocks improves the static trade-off to first order, but the value of collecting

information remains below the cost. The planner therefore always moves away from the strict target.

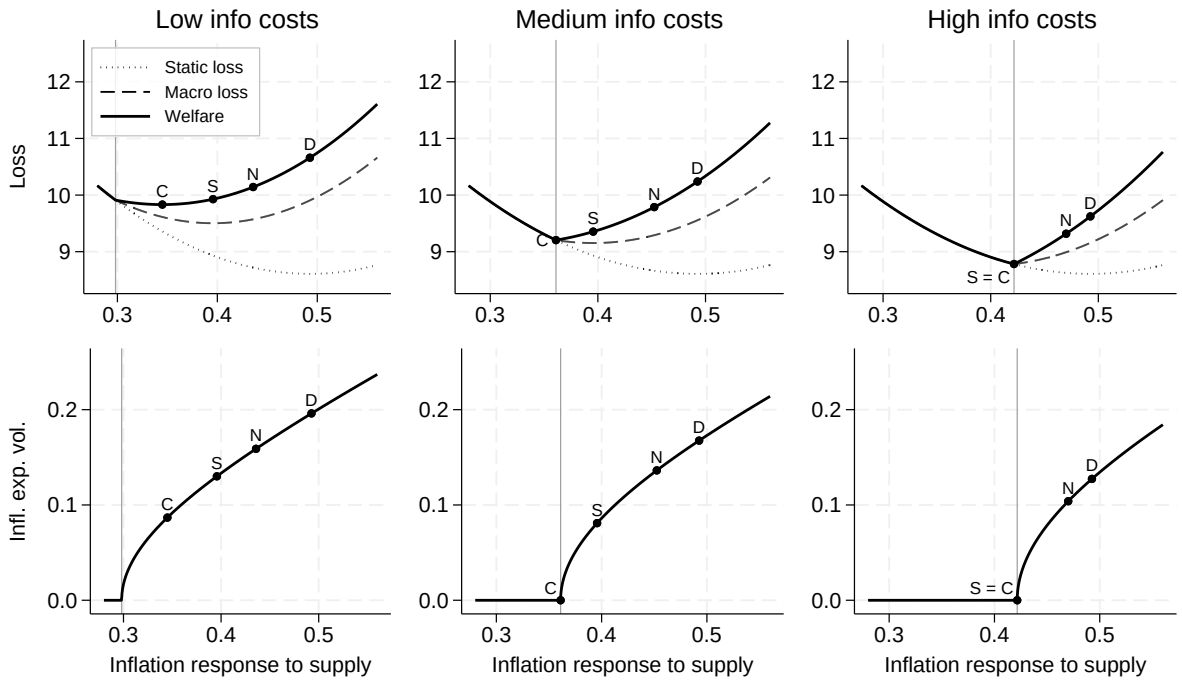


Figure 7: The Planner Ladder

Notes: Parameter values are in Table C1. The x-axis is the response of inflation to supply shocks a_u. Columns correspond to information costs: $\vartheta = 1.5$ (left), 2.2 (middle), and 3 (right). Top row: the dotted line is the static block $S(a_u)$, the dashed line the macro loss $2L(a_u) = S(a_u) + N(a_u)$ of (C16), and the solid line welfare $\mathcal{W}(a_u)$ of (25), which adds the cost of information. The vertical line marks the attention threshold $\hat{a}$, at which the macro-loss and welfare curves have a kink. D, N, S, and C mark the four memoryless planners. Bottom row: the standard deviation of the subjective inflation forecast, $\text{SD}(\tilde{E}_t \pi_{t+1})$.

Figure 7 illustrates the four plans at the parameter values of Section 2.4, for three costs of attention. The top row plots the static block, the macro loss, and welfare against the promised response a_u; the bottom row plots the standard deviation of subjective inflation forecasts, whose square is the belief block of the macro loss up to the constant $\varpi_\pi \Theta \beta^2$. The macro-loss and welfare curves have a kink at the attention threshold $\hat{a}$, and the four planners sit on both curves. The bottom row shows anchoring directly: forecasts do not move under any plan with $a_u \leq \hat{a}$, their dispersion rises with the promised response beyond the threshold, and the discretionary and expectations-managing planners leave forecasts unanchored at every cost of attention. As the cost of attention rises but still remains moderate, belief shaping accounts for a larger share of the value of commitment. When attention is cheap ($\vartheta = 1.5$), both belief-shaping planners anchor expectations only partially, with $a_u^S = 0.396$ and $a_u^C = 0.345$. The value of commitment is measured moving from discretion to the cost-internalizing planner: expectations management accounts for about 63 percent, belief shaping for about 26 percent, and counting the cost of information for about 12 percent. At an intermediate cost ($\vartheta = 2.2$), the cost-internalizing planner

anchors at $\hat{a} = 0.361$ while the belief-shaping planner still anchors only partially, and the three shares are about 44, 42, and 14 percent. With higher costs ($\vartheta = 3$), both belief-shaping planners set $a_u = \hat{a} = 0.422$ so that agents do not collect information under either plan; the information-cost term in (27) is then zero, and the other two shares are about 36 and 64 percent. As $\vartheta \to \infty$, eventually all planners coincide since the inflation response to supply shocks is sufficient even under discretion to remove the incentive to collect information.

The plans above are memoryless, which shuts down the textbook commitment channel of promising tomorrow's response to today's shock. The last two columns of Table C1 restore that channel: the plan may respond to the lagged supply shock, and the optimal promise partially reverses today's inflation (Appendix C3, Proposition C3). Note that the lagged shock is known at t and therefore does not affect the value of news; memory lowers the loss of both belief-shaping planners for every cost of attention through the static block, and leaves the attention threshold unchanged. Memory also widens the anchoring windows (at $\vartheta = 2.2$ both planners with memory anchor).

6 Conclusion

This paper studies the causes and consequences of supply-side beliefs: the tendency of households to interpret economic news as news about supply. In a New Keynesian model in which agents observe an ambiguous signal about next period's shocks, beliefs enter the equilibrium as anticipated fundamentals would, so the response of the economy to news is driven by how agents interpret the signal rather than by what the signal contains. News which raises inflation forecasts raises realized inflation and, when beliefs are sufficiently supply-tilted, lowers realized output. Using the daily timing of household survey interviews around CPI releases to measure news-driven revisions in inflation expectations, we find that inflationary news is indeed stagflationary, and a quantitative version of the model reproduces the dynamics in an economy whose fundamentals are dominated by demand shocks.

Agents interpret news as driven by supply because mistakes about supply are costlier than mistakes about demand: under an active policy rule, demand shocks are offset as a by-product of stabilizing inflation, whereas the costs of a supply shock can only be distributed between inflation and output. Because monetary policy partly sets the relative cost of the two shocks, the policy regime shapes the interpretation of news, and an optimizing central bank acquires a belief-shaping motive. A central bank which internalizes that attention responds to its plan can anchor inflation expectations by lowering the value of learning about supply shocks, and when attention is costly enough it fully anchors expectations without fully stabilizing inflation. Whereas a policy rule tilts beliefs away from demand by removing the value of learning about demand shocks, optimal policy anchors

beliefs by lowering the value of learning about supply shocks as well.

References

Afrouzi, H. (2024). Strategic Inattention, Inflation Dynamics, and the Nonneutrality of Money. *Journal of Political Economy*, 132(10):3378–3420.

Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2026). Narratives about the Macroeconomy. *Review of Economic Studies*. Forthcoming.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples. *The Review of Economic Studies*, 89(6):2958–2991.

Angeletos, G.-M., Iovino, L., and La'O, J. (2016). Real Rigidity, Nominal Rigidity, and the Social Value of Information. *American Economic Review*, 106(1):200–227.

Angeletos, G.-M. and Lian, C. (2018). Forward Guidance without Common Knowledge. *American Economic Review*, 108(9):2477–2512.

Barsky, R. B. and Sims, E. R. (2011). News Shocks and Business Cycles. *Journal of Monetary Economics*, 58(3):273–289.

Barsky, R. B. and Sims, E. R. (2012). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review*, 102(4):1343–1377.

Beaudry, P. and Portier, F. (2004). An Exploration into Pigou's Theory of Cycles. *Journal of Monetary Economics*, 51(6):1183–1216.

Beaudry, P. and Portier, F. (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review*, 96(4):1293–1307.

Beaudry, P. and Portier, F. (2014). News-Driven Business Cycles: Insights and Challenges. *Journal of Economic Literature*, 52(4):993–1074.

Bhandari, A., Borovička, J., and Ho, P. (2022). Survey Data and Subjective Beliefs in Business Cycle Models. *Available at SSRN 2763942*.

Binder, C., Kamdar, R., and Ryngaert, J. M. (2024). Partisan Expectations and COVID-Era Inflation. NBER Working Paper 32650, National Bureau of Economic Research.

Binder, C. C., Campbell, J., and Ryngaert, J. (2022). Consumer Inflation Expectations: Daily Dynamics. MPRA Paper 117628, University Library of Munich, Germany.

Blanchard, O. J., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, Noise, and Fluctuations: An Empirical Exploration. *American Economic Review*, 103(7):3045–3070.

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

Candia, B., Coibion, O., and Gorodnichenko, Y. (2020). Communication and the Beliefs of Economic Agents. NBER Working Papers 27800, National Bureau of Economic Research, Inc.

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics*, 118(1):269–298.

Carvalho, C., Eusepi, S., Moench, E., and Preston, B. (2023). Anchored Inflation Expectations. *American Economic Journal: Macroeconomics*, 15(1):1–47.

Chahrour, R. and Jurado, K. (2018). News or Noise? The Missing Link. *American Economic Review*, 108(7):1702–1736.

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). Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). The Formation of Expectations, Inflation, and the Phillips Curve. *Journal of Economic Literature*, 56(4):1447–1491.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary Policy Communications and Their Effects on Household Inflation Expectations. *Journal of Political Economy*, 130(6):1537–1584.

De Fiore, F., Lombardi, M., and Schuffels, J. (2022). Are Households Indifferent to Monetary Policy Announcements? CEPR Discussion Papers 17041, C.E.P.R. Discussion Papers.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.

Hazell, J., Herreño, J., Nakamura, E., and Steinsson, J. (2022). The Slope of the Phillips Curve: Evidence from U.S. States. *Quarterly Journal of Economics*, 137(3):1299–1344.

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

Kamdar, R. (2019). The Inattentive Consumer: Sentiment and Expectations. Technical report, Center for Applied Economics and Policy Research.

Kamdar, R. and Ray, W. (2025). Attention-Driven Sentiment and the Business Cycle. Technical report, Center for Applied Economics and Policy Research.

Kamdar, R. and Ray, W. (2026). Optimal Unconventional Policy in a New Keynesian Preferred Habitat Model. CEPR Discussion Paper 22023, CEPR Press, Paris & London. JEL codes: E52, E44, E43, D14.

Kamdar, R., Ray, W., and Ulate, M. (2025). We Think That They Think: Political Affiliation and Higher-Order Beliefs. *Journal of Political Economy Macroeconomics*. Forthcoming.

Lamla, M. J. and Vinogradov, D. V. (2019). Central Bank Announcements: Big News for Little People? *Journal of Monetary Economics*, 108(C):21–38.

Lorenzoni, G. (2009). A Theory of Demand Shocks. *American Economic Review*, 99(5):2050–2084.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.

Mankiw, N. G. and Reis, R. (2007). Sticky Information in General Equilibrium. *Journal of the European Economic Association*, 5(2-3):603–613.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 18:209–248.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature*, 52(1):124–188.

Melosi, L. (2014). Estimating Models with Dispersed Information. *American Economic Journal: Macroeconomics*, 6(1):1–31.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823.

Morris, S. and Shin, H. S. (2002). Social Value of Public Information. *American Economic Review*, 92(5):1521–1534.

Paciello, L. and Wiederholt, M. (2014). Exogenous Information, Endogenous Information, and Optimal Monetary Policy. *Review of Economic Studies*, 81(1):356–388.

Rast, S. (2021). Central Bank Communication with the General Public: Survey Evidence from Germany. Manuscript.

Shiller, R. J. (1996). Why Do People Dislike Inflation? Cowles Foundation Discussion Papers 1115, Cowles Foundation for Research in Economics, Yale University.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., Kenny, G., and Ponce, J. (2025). Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings. *Econometrica*, 93(1):229–264.

Woodford, M. (2003). Imperfect Common Knowledge and the Effects of Monetary Policy. In Aghion, P., Frydman, R., Stiglitz, J. E., and Woodford, M., editors, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, chapter 1, pages 25–58. Princeton University Press, Princeton.

Woodford, M. (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Review of Economics*, 5(1):303–346.

York, J. (2023). Do Household Inflation Expectations Respond to Macroeconomic Data Releases? Manuscript.

Appendix A Baseline Model: Derivations and Proofs

A1 Household and Firm Optimality

With multiplier $\beta^t \lambda_t$ on the budget constraint in (1), the first-order conditions for C_t, N_t, B_t are $C_t^{-\varsigma} = \lambda_t P_t$, $N_t^{\varphi} = \lambda_t W_t$, and $\lambda_t Q_t e^{-\zeta_t} = \beta \tilde{E}_t \lambda_{t+1}$. These give labor supply $W_t/P_t = C_t^{\varsigma} N_t^{\varphi}$ and the Euler equation

$$1 = \beta(1+i_t) e^{\zeta_t}\, \tilde{E}_t \left[(C_{t+1}/C_t)^{-\varsigma} P_t/P_{t+1} \right]. \tag{A1}$$

Log-linearizing around the zero-inflation steady state and imposing $c_t = y_t = x_t$ yields (2) with $v_t = -\zeta_t$. The natural real rate is $r_t^n = v_t$ and the natural level of output is constant.

A firm which resets its price at t chooses P^* to maximize

$$\tilde{E}_t \sum_{k \geq 0} \theta^k Q_{t,t+k} \left[(P^*/P_{t+k})^{1-\epsilon_{t+k}} Y_{t+k} - mc_{t+k} (P^*/P_{t+k})^{-\epsilon_{t+k}} Y_{t+k} \right],$$

with $Q_{t,t+k} = \beta^k (C_{t+k}/C_t)^{-\varsigma}$ and $mc_t = W_t/P_t$. The first-order condition log-linearizes to

$$p_t^* = (1-\beta\theta)(p_t + \widehat{mc}_t + \mu_t) + \beta\theta\, \tilde{E}_t p_{t+1}^*, \quad \mu_t = \log \frac{\epsilon_t}{\epsilon_t - 1} - \log \mathcal{M}, \tag{A2}$$

where variation in the elasticity survives to first order only through μ_{t+k}. The price index gives $\pi_t = (1-\theta)(p_t^* - p_{t-1})$; combining with (A2) and $\widehat{mc}_t = (\varsigma + \varphi) x_t$ delivers (3) with $\lambda = (1-\theta)(1-\beta\theta)/\theta$, $\kappa = \lambda(\varsigma + \varphi)$, and $u_t = \lambda \mu_t$. The derivation of the wage-markup reading of Section 2.1 is in the Online Appendix.

A2 Equilibrium

Substituting the rule and (3) into (2) and collecting terms, using $1 - \kappa\chi\varsigma^{-1}\phi_\pi = \chi$:

$$x_t = \chi \left[\varsigma^{-1}(1-\beta\phi_\pi) \tilde{E}_t \pi_{t+1} + \tilde{E}_t x_{t+1} + \varsigma^{-1} v_t - \varsigma^{-1} \phi_\pi u_t \right], \quad \pi_t = \beta \tilde{E}_t \pi_{t+1} + \kappa x_t + u_t. \tag{A3}$$

With iid shocks, agents who know the model forecast outcomes at $t+1$ by applying the no-news solution (4) to their expectations $(\tilde{E}_t u_{t+1}, \tilde{E}_t v_{t+1})$, which gives (5). Substituting (5) into (A3) yields (6)–(7) with the loadings (8)–(9). The loadings are the responses to perfectly anticipated shocks because the derivation does not use the relation between the expectations and the realized innovations: setting $(\tilde{E}_t u_{t+1}, \tilde{E}_t v_{t+1}) = (u_{t+1}, 0)$ or $(0, v_{t+1})$ gives the response to a perfectly anticipated u_{t+1} or v_{t+1}.

A3 Signs and Thresholds: Proofs of Propositions 1 and 3

Write $a \equiv \kappa\varsigma^{-1}$. Under (10)–(11): $b_{\pi v} > 0$ always; $b_{xu} < 0$ for $\phi_\pi > 1/(1+\beta)$, which is implied by (10); $b_{\pi u} > 0$ is exactly (11); and $b_{xv} > 0$ follows from (11), since multiplying (11) by $a\beta$ gives $a\beta\phi_\pi < a\beta + \beta^2$, and hence $1 + a(1 - \beta\phi_\pi) > (1 - \beta^2) + a(1 - \beta) > 0$.

Proof of Proposition 1. Under Assumption 1, substituting (12) into (5) gives the response of forecasts to news,

$$\frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} = \chi(\tilde{K}_u + \kappa\varsigma^{-1}\tilde{K}_v), \quad \frac{\partial \tilde{E}_t x_{t+1}}{\partial z_t} = \chi\varsigma^{-1}(\tilde{K}_v - \phi_\pi \tilde{K}_u).$$

With $\tilde{K}_u, \tilde{K}_v \geq 0$, not both zero, the first is positive, and the second is negative if and only if $\phi_\pi \tilde{K}_u > \tilde{K}_v$, that is, $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$ (with $\tilde{K}_u/\tilde{K}_v = \infty$ when $\tilde{K}_v = 0$), which is (15). This step uses no restriction on $\phi_\pi > 0$. Since forecasts move only with z_t, with $\sigma_z^2 \equiv \mathrm{Var}(z_t)$,

$$\mathrm{Cov}(\tilde{E}_t \pi_{t+1}, \tilde{E}_t x_{t+1}) = \chi^2 \varsigma^{-1}(\tilde{K}_u + \kappa\varsigma^{-1}\tilde{K}_v)(\tilde{K}_v - \phi_\pi \tilde{K}_u)\sigma_z^2, \tag{A4}$$

so forecasts comove negatively if and only if (15) holds.

Substituting (12) into (6)–(7) gives the stated form of $\frac{\partial x_t}{\partial z_t}$ and $\frac{\partial \pi_t}{\partial z_t}$. Under (10)–(11), $\frac{\partial \pi_t}{\partial z_t} = b_{\pi u}\tilde{K}_u + b_{\pi v}\tilde{K}_v > 0$, and $\frac{\partial x_t}{\partial z_t} = b_{xu}\tilde{K}_u + b_{xv}\tilde{K}_v < 0 \iff \tilde{K}_u/\tilde{K}_v > T_x = b_{xv}/|b_{xu}|$. For the ordering, the denominator of T_x is positive under (10), and $T_x\phi_\pi < 1$ reduces, after clearing denominators, to $(\beta\phi_\pi - 1)(1 + a\phi_\pi) > 0$, which holds under (10). Hence $T_x < 1/\phi_\pi$, and $1/\phi_\pi < 1$ since $\phi_\pi > 1/\beta > 1$. □

Proof of Proposition 3. Part (i). By (A4) and Proposition 1, negative forecast comovement is equivalent to $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$. Since $T_x < 1/\phi_\pi$ (Proposition 1), this implies $\tilde{K}_u/\tilde{K}_v > T_x$, and hence the stagflationary response. For $\tilde{K}_u/\tilde{K}_v$ in the nonempty interval $(T_x, 1/\phi_\pi)$, news is stagflationary while forecasts comove positively, so the converse fails.

Part (ii). By part (i), negative forecast comovement implies $\partial\pi_t/\partial z_t > 0 > \partial x_t/\partial z_t$, so the news term in (A5) is negative, and positive realized comovement requires $\kappa\varsigma^{-1}\sigma_v^2 > \phi_\pi\sigma_u^2$, the first inequality in (20). Combining this with $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$ gives $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi = (\varsigma/\kappa)(\kappa\varsigma^{-1}/\phi_\pi) > (\varsigma/\kappa)\,\sigma_u^2/\sigma_v^2$, the second inequality. The content-neutral claim follows since $\varsigma/\kappa \geq 1$ when $\kappa\varsigma^{-1} \leq 1$; for the standardized claim, $\sigma_u/\sigma_v < (\kappa\varsigma^{-1}/\phi_\pi)^{1/2} \leq 1/\phi_\pi < \tilde{K}_u/\tilde{K}_v$ when $\kappa\varsigma^{-1} \leq 1/\phi_\pi$. □

A4 Comovements

With independent shocks, (6)–(7) give

$$\mathrm{Cov}(\pi_t, x_t) = \chi^2\varsigma^{-1}\left(\kappa\varsigma^{-1}\sigma_v^2 - \phi_\pi\sigma_u^2\right) + \frac{\partial \pi_t}{\partial z_t}\frac{\partial x_t}{\partial z_t}\sigma_z^2. \tag{A5}$$

Consider the demand-dominated limit ($\sigma_v^2 \to \infty$, with σ_u^2 and the signal loadings fixed) and suppose that the news term is stagflationary. Both terms of (A5) scale with σ_v^2 whenever $\alpha_v \neq 0$, and

$$\text{Cov}(\pi_t, x_t) > 0 \iff \chi^2 \kappa \varsigma^{-2} > \left|\frac{\partial \pi_t}{\partial z_t}\frac{\partial x_t}{\partial z_t}\right| \alpha_v^2. \tag{A6}$$

Condition (A6) is a genuine parameter restriction: when it fails, the news channel generates negative realized comovement in an economy with no supply shocks. Whether it fails depends on the belief foundation. Under FIRE, $\tilde{K}_u \to 0$ in the demand-dominated limit and the news term vanishes. Under rational inattention the endogenous loadings turn the signal toward demand in the same limit. At the robust corner, $\tilde{K}_u \to 1/\alpha_u$, and the condition depends on the composition of the signal.

A5 Implications for Phillips Curve Estimation

Section 2.5 draws three implications of Proposition 1 for the estimation of the Phillips curve. This appendix states the first two in closed form and derives all three. The results concern signs rather than magnitudes, and for transparency we state them for pure supply news, $\tilde{K}_v = 0$, in the limit without realized supply shocks, $\sigma_u^2 \to 0$, so that the interpretation of news is the only source of supply-type comovement in the data. The derivations follow the statements.

First, estimation which imposes FIRE delivers a flat or wrongly signed Phillips curve. Episodes driven by news trace out combinations of high inflation and low output, as supply shocks would, so the OLS estimate of the slope is biased downward. Let s_z denote the share of the variance of the output gap which is due to news. Then OLS of π_t on x_t converges to

$$\hat{\kappa}^{OLS} \to \kappa - s_z \frac{\beta\,(\varsigma + \kappa\phi_\pi)}{\phi_\pi(1+\beta) - 1}, \tag{A7}$$

which lies below κ for any positive news share and turns negative once the share exceeds $\kappa[\phi_\pi(1+\beta)-1]/[\beta(\varsigma+\kappa\phi_\pi)]$, about 0.27 in the example economy of Section 2.4. Replacing expected inflation with realized future inflation, the usual remedy, does not restore consistency: the expectation error $\beta(\tilde{E}_t\pi_{t+1} - \pi_{t+1})$ is predictable from the news shock, which is a component of x_t. The model thus rationalizes the documented instability of estimated Phillips curve slopes (Mavroeidis et al., 2014).

Second, survey-augmented estimation is consistent. The Phillips curve (3) holds exactly with the subjective expectation on the right-hand side, so a regression of π_t on $\tilde{E}_t\pi_{t+1}$ and x_t recovers (β, κ). This is in line with the finding that Phillips curves estimated with household survey expectations are better behaved than their FIRE counterparts (Coibion and Gorodnichenko, 2015; Coibion et al., 2018). (Note that with realized supply shocks, x_t requires an instrument which is orthogonal to u_t, and measurement error in survey

expectations may still generate bias.)

Derivations. Take $\tilde{K}_v = 0$ and $\sigma_u^2 \to 0$, as above. Then $\pi_t = \chi(\kappa\varsigma^{-1}v_t + a_z z_t)$ and $x_t = \chi\varsigma^{-1}(v_t + b_z z_t)$ with $a_z = (\beta - \kappa\varsigma^{-1}(\phi_\pi - 1))\chi\tilde{K}_u$ and $b_z = (1 - \phi_\pi(1+\beta))\chi\tilde{K}_u < 0$.

(i) FIRE-based estimation. For OLS of π_t on x_t,

$$\hat{\kappa}^{OLS} \to \kappa + \frac{b_z(\varsigma a_z - \kappa b_z)\sigma_z^2}{\sigma_v^2 + b_z^2\sigma_z^2} = \kappa - \frac{\beta\tilde{K}_u^2(\phi_\pi(1+\beta) - 1)(1 + \kappa\varsigma^{-1}\phi_\pi)\sigma_z^2}{\varsigma^{-1}\left[\tilde{K}_u^2(\phi_\pi(1+\beta) - 1)^2\sigma_z^2 + (1 + \kappa\varsigma^{-1}\phi_\pi)^2\sigma_v^2\right]} < \kappa, \tag{A8}$$

where the second equality uses $\varsigma a_z - \kappa b_z = \beta\varsigma\tilde{K}_u$, which follows from $\chi(1 + \kappa\varsigma^{-1}\phi_\pi) = 1$. Since the news share of the variance of x_t is $s_z = b_z^2\sigma_z^2/(\sigma_v^2 + b_z^2\sigma_z^2)$ and $\tilde{K}_u/|b_z| = 1/[\chi(\phi_\pi(1+\beta) - 1)]$, (A8) is (A7). Replacing expected with realized future inflation does not restore consistency. The regression error is $\beta(\tilde{E}_t\pi_{t+1} - \pi_{t+1})$, and $\mathbb{E}[x_t(\tilde{E}_t\pi_{t+1} - \pi_{t+1})] = \chi^2\varsigma^{-1}b_z(\tilde{K}_u\sigma_z^2 - \kappa\varsigma^{-1}\alpha_v\sigma_v^2) \neq 0$: expectation errors are predictable from the news shock, which is a component of x_t.

(ii) Survey-augmented estimation. The structural equation $\pi_t = \beta\tilde{E}_t\pi_{t+1} + \kappa x_t + u_t$ holds exactly with the subjective expectation on the right-hand side. With $\sigma_u^2 \to 0$ it is an identity in the regressors, so OLS recovers (β, κ). With $\sigma_u^2 > 0$, a valid instrument for x_t must be orthogonal to u_t; variation which runs through z_t is collinear with $\tilde{E}_t\pi_{t+1}$ and therefore fails relevance. With classical measurement error in the survey expectation, the news shock, which is observable at high frequency, is the natural instrument for the measured expectation.

(iii) Subjective Phillips curve. In a regression of $\tilde{E}_t\pi_{t+1}$ on $(\tilde{E}_t\pi_{t+2}, \tilde{E}_t x_{t+1})$, $\tilde{E}_t\pi_{t+2} = 0$ and both remaining variables are proportional to z_t by (5) and (12), so the slope converges to the ratio of the two coefficients in (5), which is (21); the denominator is negative exactly when forecasts comove negatively, by (A4).

Appendix B Belief Formation

B1 Rational Inattention

A single collector chooses what to learn about $\varepsilon_{t+1} \equiv (u_{t+1}, v_{t+1})^\top$ on behalf of all agents. All agents share the posterior, so outcomes are given by Proposition 1 applied to the posterior means ν. Under certainty equivalence the loss from imperfect information is quadratic in the posterior-mean errors $e \equiv \nu - \varepsilon_{t+1}$, weighted by the sensitivities of the loss-relevant outcomes to beliefs: this is the tracking problem (18) with $b_\pi = (b_{\pi u}, b_{\pi v})^\top$ and $b_x = (b_{xu}, b_{xv})^\top$. With Gaussian priors $\Sigma = \mathrm{diag}(\sigma_u^2, \sigma_v^2)$ and information priced at ϑ per nat, the solution is reverse water-filling on $\Omega \equiv \Sigma^{1/2}Q\Sigma^{1/2}$, with eigenvalues $\omega_1 \geq \omega_2$: direction i is learned if and only if $\omega_i > \vartheta/2$. A single signal is optimal if and only if

$\omega_2 \leq \vartheta/2 < \omega_1$; it is then of the form (13) with

$$\alpha^* = \Sigma^{-1/2} b^*, \quad (\tilde{K}_u, \tilde{K}_v)^\top = \frac{\Sigma \alpha^*}{1+\sigma_\eta^2}, \quad \frac{\tilde{K}_u}{\tilde{K}_v} = \frac{\alpha_u^* \sigma_u^2}{\alpha_v^* \sigma_v^2}, \tag{B1}$$

where b^* is the top eigenvector of Ω and the normalization is $\alpha^\top \Sigma \alpha = 1$. The one-signal condition is far from binding: in the example economy of Section 2.4, ω_2/ω_1 attains a maximum of 0.156 over the entire range of variance ratios, against the bound $1/2$ at a signal-to-noise ratio of one.

Lemma B1 (Positive gains)**.** *Impose, in addition to* (10)–(11)*,*

$$\varpi_\pi \, b_{\pi u} \, b_{\pi v} > \varpi_x \, |b_{xu}| \, b_{xv}, \tag{B2}$$

that is, the off-diagonal element of Q is positive: the inflation margin, on which the two shocks agree, outweighs the output margin, on which they disagree. Then $\tilde{K}_u, \tilde{K}_v > 0$. In primitives, with $a \equiv \kappa \varsigma^{-1}$, (B2) *is equivalent to*

$$\bar{\epsilon}\varsigma \left[\beta - a(\phi_\pi - 1)\right](1 + \beta + a) > \left[\phi_\pi(1+\beta) - 1\right]\left[1 + a(1 - \beta\phi_\pi)\right], \tag{B3}$$

and (10)*,* (11)*, and* (B2) *jointly define a bounded window $\phi_\pi \in (1/\beta, \bar{\phi})$ with $\bar{\phi} < 1 + \beta/a$.*

Proof. The diagonal of $\Omega = \Sigma^{1/2} Q \Sigma^{1/2}$ is positive by construction, and (B2) makes the off-diagonal element positive, so every entry of Ω is strictly positive. By the Perron–Frobenius theorem its top eigenvector is strictly positive, and $\tilde{K}_u, \tilde{K}_v > 0$ by (B1). Writing $Q_{uv} = \varpi_\pi b_{\pi u} b_{\pi v} + \varpi_x b_{xu} b_{xv}$ in the loadings (8)–(9), substituting $\varpi_\pi = \bar{\epsilon}/\lambda$, $\varpi_x = \varsigma + \varphi$, and $\kappa = \lambda(\varsigma + \varphi)$, and canceling the common factor $\chi^4(\varsigma + \varphi)/\varsigma^2$ gives (B3). The left side of (B3) vanishes at the boundary of (11) while the right side remains positive there, which bounds the window. □

In the example economy of Section 2.4 the window is wide: $\bar{\phi} = 6.07$ against the bound of 6.77 implied by (11), and (B2) holds with a wide margin at $\phi_\pi = 1.5$. With two shocks, (B2) is *equivalent* to strict positivity of the gains, since the diagonal of Ω is automatically positive; this equivalence allows Proposition 2(i), which is stated in terms of the off-diagonal element of Q, to be read as a statement about positive gains. With more than two shocks, positive gains no longer imply an entrywise-positive Ω, and the primitive condition is the operative one.

Proof of Proposition 2(i). Under (B2) the gains are positive (Lemma B1), so each condition in the main text takes the form $\tilde{K}_u/\tilde{K}_v > c$. The condition $\tilde{K}_u = c\tilde{K}_v$ states that b^* is orthogonal to $\ell \equiv \Sigma^{1/2}(1, -c)^\top$. In two dimensions this holds if and only if $\ell^\perp = (c\sigma_v, \sigma_u)^\top$ is an eigen-direction; since $\ell^\perp$ is strictly positive for $c > 0$, it is the *top* eigen-direction

under (B2), because a positive eigenvector of an entrywise-positive matrix is associated with the Perron root. Solving $\ell^\top \Omega \ell^\perp = 0$ for the variance ratio gives the threshold

$$\frac{\sigma_u^2}{\sigma_v^2} = r^*(c) \equiv \frac{c(Q_{vv} + cQ_{uv})}{Q_{uv} + cQ_{uu}}, \quad \frac{dr^*}{dc} = \frac{Q_{uv}(Q_{vv} + 2cQ_{uv} + c^2 Q_{uu})}{(Q_{uv} + cQ_{uu})^2} > 0 \text{ under (B2).} \quad \text{(B4)}$$

Monotonicity of r^* together with the ordering $T_x < 1/\phi_\pi < 1$ (Appendix A3) orders the three thresholds. At $c = T_x$ the orthogonality condition implies $b_\pi^\top \Sigma b_x = 0$, which gives the closed form $r^*(T_x) = b_{\pi v} b_{xv} / (b_{\pi u} |b_{xu}|)$. Evaluating (B4) at $c = 1$ gives $r^*(1) = (Q_{vv} + Q_{uv})/(Q_{uu} + Q_{uv})$, whose denominator is positive under (B2), so $r^*(1) < 1$ if and only if $Q_{uu} > Q_{vv}$. Finally, the optimal direction b^*, and hence $\tilde{K}_u/\tilde{K}_v$, depends only on Ω, and not on ϑ or σ_η^2: the interpretation is invariant to the cost of information. □

B2 Robustness: The Corner

Agents know the total innovation variance $S = \sigma_u^2 + \sigma_v^2$ but distrust its composition. Nature chooses the shares $\tilde{\sigma}_u^2 + \tilde{\sigma}_v^2 = S$ to maximize the expected loss caused by the innovations themselves when they are realized,

$$\tfrac{1}{2}\left(D_u \tilde{\sigma}_u^2 + D_v \tilde{\sigma}_v^2\right), \quad D_h \equiv \varpi_\pi (\partial \pi_t / \partial h)^2 + \varpi_x (\partial x_t / \partial h)^2, \quad \text{(B5)}$$

where the derivatives are the contemporaneous responses (4), and agents update by (14) under the distorted prior. By (6)–(7), realized innovations transmit as in the textbook model whatever the beliefs, so D_u and D_v do not depend on the gains. The objective (B5) is linear in the variance shares and is therefore maximized at a corner.

Proposition B1 (Robust corner)**.** *The worst case is the supply corner, $(\tilde{\sigma}_u^2, \tilde{\sigma}_v^2) = (S, 0)$, if and only if $D_u > D_v$, that is, if and only if $\varpi_\pi(\varsigma^2 - \kappa^2) + \varpi_x(\phi_\pi^2 - 1) > 0$, for which $\phi_\pi \geq 1$ and $\kappa \leq \varsigma$ (one strict) is sufficient for any loss weights. At the supply corner, for any signal with $\alpha_u \neq 0$,*

$$\tilde{K}_v = 0, \quad \tilde{K}_u = \frac{\alpha_u S}{\alpha_u^2 S + \sigma_\eta^2}. \quad \text{(B6)}$$

Proof. From (4), $D_u = \chi^2[\varpi_\pi + \varpi_x \varsigma^{-2} \phi_\pi^2]$ and $D_v = \chi^2 \varsigma^{-2}[\varpi_\pi \kappa^2 + \varpi_x]$, so $D_u > D_v$ is the stated inequality once the common factor χ^2 is canceled and both sides are multiplied by ς^2. The objective (B5) is linear in $\tilde{\sigma}_u^2 \in [0, S]$ with slope $\frac{1}{2}(D_u - D_v)$, so its maximum is uniquely the supply corner if and only if $D_u > D_v$. At the supply corner the distorted prior is $(S, 0)$, and (14) gives (B6); with the normalization of Section 2.3, $\alpha_u > 0$ and hence $\tilde{K}_u > 0$. Under (10), $\phi_\pi > 1$, so $\kappa \leq \varsigma$ suffices, and Proposition 2(ii) follows. □

Including the Response to News. An alternative formulation lets nature also account for the loss caused at t by the economy's response to the news, evaluated under the distorted prior $\tilde{\Sigma} \equiv \mathrm{diag}(\tilde{\sigma}_u^2, \tilde{\sigma}_v^2)$. This loss occurs one period before the innovations

are realized, so the fundamental component (B5) is discounted by β. With the gains $\tilde{K} = \tilde{\Sigma}\alpha/(\alpha^\top\tilde{\Sigma}\alpha + \sigma_\eta^2)$ from (14) and the perceived variance of the news $\alpha^\top\tilde{\Sigma}\alpha + \sigma_\eta^2$, the belief-response component is

$$\tfrac{1}{2}\left[\varpi_\pi(b_\pi^\top\tilde{K})^2 + \varpi_x(b_x^\top\tilde{K})^2\right]\left(\alpha^\top\tilde{\Sigma}\alpha + \sigma_\eta^2\right) = \frac{1}{2}\,\frac{\varpi_\pi(b_\pi^\top\tilde{\Sigma}\alpha)^2 + \varpi_x(b_x^\top\tilde{\Sigma}\alpha)^2}{\alpha^\top\tilde{\Sigma}\alpha + \sigma_\eta^2}, \tag{B7}$$

with b_π and b_x as in (18). This component is nonlinear in the variances shares, but it is convex. With $s \equiv \tilde{\sigma}_u^2$, each numerator in (B7) is the square of a function which is affine in s, the denominator is affine in s and positive, and the ratio of a squared affine function to a positive affine function is convex. The full objective is therefore convex in s, and its maximum on $[0, S]$ is again attained at a corner. At corner h the prior is rank one, and the full objective equals $\frac{1}{2}S(\beta D_h + Q_{hh}\, g_h)$, where $g_h \equiv \alpha_h^2 S/(\alpha_h^2 S + \sigma_\eta^2) \in [0, 1)$ and Q_{hh} is the corresponding diagonal element of Q in (18). Since $0 \leq g_h < 1$, the supply corner remains the worst case whenever

$$\beta D_u > \beta D_v + Q_{vv}, \tag{B8}$$

a condition which does not involve the signal.

The two foundations are connected at the edge of the window of Lemma B1. At $\phi_\pi = \bar{\phi}$ the off-diagonal element of Q is zero and the top eigen-direction of the rational-inattention problem is the u axis, so the rational-inattention solution coincides with the robust corner.

B3 The Cost-Ranked Class

Definition B1 (Cost-ranked beliefs)**.** For a belief-formation model, let C_h denote the cost which the model's own criterion attaches to ignorance about innovation $h \in \{u, v\}$ under the equilibrium mapping. The model is *cost-ranked* if, in every admissible environment, its interpretation of ambiguous news tilts toward the costlier ignorance: $C_u > C_v \Rightarrow |\tilde{K}_u|\sigma_v > |\tilde{K}_v|\sigma_u$, and symmetrically.

Definition B1 measures the asymmetry of the gains against the standardized benchmark of Section 2.3, under which $|\tilde{K}_u|\sigma_v = |\tilde{K}_v|\sigma_u$ holds for an undistorted reader exactly when the signal devotes equal weight to the two shocks per unit of standard deviation.

Definition B2 (Cost of ignorance)**.** For rational inattention, $C_h^{\mathrm{RI}} = \Omega_{hh} = Q_{hh}\sigma_h^2$: the increase in the tracking loss from forgoing information about innovation h, holding the other margin fixed. For robustness, $C_h^{\mathrm{MM}} = D_h S$: the expected loss which nature's objective (B5) attaches to corner h.

Proposition B2 (Both foundations are cost-ranked)**.** *(i) Rational inattention satisfies Definition B1 in every admissible environment (under the one-signal condition), as an*

exact equivalence with equality at ties. (ii) Robustness satisfies it in every admissible environment without ties, $C_u^{\mathrm{MM}} \neq C_v^{\mathrm{MM}}$, for any signal which loads on the costlier shock.

Proof. (i) In standardized coordinates the optimal direction is the top eigenvector of Ω, whose diagonal is $(C_u^{\mathrm{RI}}, C_v^{\mathrm{RI}})$. With $\Delta \equiv \Omega_{uu} - \Omega_{vv}$ and $m \equiv \Omega_{uv} \neq 0$, the top eigenvector satisfies $b_u^*/b_v^* = m/(\omega_1 - \Omega_{uu})$ with $\omega_1 - \Omega_{uu} = \frac{1}{2}(-\Delta + \sqrt{\Delta^2 + 4m^2}) > 0$, so

$$|b_u^*| \gtrless |b_v^*| \iff 2|m| \gtrless -\Delta + \sqrt{\Delta^2 + 4m^2} \iff \Delta \gtrless 0, \tag{B9}$$

with no restriction on the sign of m; if $m = 0$ the claim is immediate. Since $|\tilde{K}_u|\sigma_v/(|\tilde{K}_v|\sigma_u) = |b_u^*|/|b_v^*|$, the chain delivers Definition B1 as an equivalence. (ii) By Proposition B1, nature selects the corner with the larger C_h^{MM}. At the supply corner $|\tilde{K}_u|\sigma_v > 0 = |\tilde{K}_v|\sigma_u$, and symmetrically at the demand corner. □

The information problem requires no fixed point between beliefs and aggregates, for three reasons: full current observability means that endogenous variables never serve as signals; common information means that the sensitivities b_π, b_x are anticipated-shock coefficients which do not contain the gains; and linearity means that belief-dependent second moments do not feed back into decision rules. Relaxing any of the three reintroduces a fixed point.

Appendix C Monetary Policy: Derivations

C1 General Rules and Strict Mandates

Under the rule family (24), define $\chi_g \equiv [1 + \varsigma^{-1}(\phi_x + \kappa\phi_\pi)]^{-1}$. Since future signals are subjectively mean zero, the ϕ_z term is expected to be inactive at $t{+}1$, and believed outcomes at $t{+}1$ are

$$\tilde{E}_t\pi_{t+1} = \chi_g\left[(1 + \varsigma^{-1}\phi_x)\tilde{K}_u + \kappa\varsigma^{-1}\tilde{K}_v\right] z_t, \quad \tilde{E}_t x_{t+1} = \chi_g\varsigma^{-1}(\tilde{K}_v - \phi_\pi\tilde{K}_u)z_t. \tag{C1}$$

The responses at t satisfy

$$\frac{\partial x_t}{\partial z_t} = \chi_g\left[\varsigma^{-1}(1 - \beta\phi_\pi)\frac{\partial \tilde{E}_t\pi_{t+1}}{\partial z_t} + \frac{\partial \tilde{E}_t x_{t+1}}{\partial z_t} - \varsigma^{-1}\phi_z\right], \quad \frac{\partial \pi_t}{\partial z_t} = \beta\frac{\partial \tilde{E}_t\pi_{t+1}}{\partial z_t} + \kappa\frac{\partial x_t}{\partial z_t}. \tag{C2}$$

Proof of Proposition 4. The identity is the second equation of (C2). With $\phi_\pi, \phi_x \geq 0$ finite, $\chi_g \in (0, 1]$ and $1 + \varsigma^{-1}\phi_x \geq 1$; with $\tilde{K}_u, \tilde{K}_v \geq 0$ not both zero, the bracket in (C1) is strictly positive, so $\frac{\partial \tilde{E}_t\pi_{t+1}}{\partial z_t} > 0$. The belief revision (C1) is invariant to ϕ_z, while $\frac{\partial x_t}{\partial z_t}$ in (C2) is affine in ϕ_z with slope $-\chi_g\varsigma^{-1} \neq 0$, so some ϕ_z achieves $\frac{\partial x_t}{\partial z_t} = 0$, and then $\frac{\partial \pi_t}{\partial z_t} = \beta\frac{\partial \tilde{E}_t\pi_{t+1}}{\partial z_t} > 0$. Thus a direct response to the information event can neutralize the output response, but no contemporaneous choice of the instrument removes the belief

revision from the Phillips curve. The argument uses neither the source of the gains (so it covers $z_t = \eta_t$ and $\tilde{K}_u = 0 < \tilde{K}_v$) nor conditions (10)–(11). □

Comparative Statics of the Trade-Off. For given gains, differentiating (C1) gives

$$\frac{\partial}{\partial \phi_\pi}\frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} = -\kappa \varsigma^{-1} \chi_g \frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} < 0, \quad \frac{\partial}{\partial \phi_x}\frac{\partial \tilde{E}_t \pi_{t+1}}{\partial z_t} = \kappa \varsigma^{-2} \chi_g^2 \left(\phi_\pi \tilde{K}_u - \tilde{K}_v\right), \tag{C3}$$

so the revision falls in ϕ_π and rises in ϕ_x if and only if $\tilde{K}_u/\tilde{K}_v > 1/\phi_\pi$, the condition under which $\frac{\partial \tilde{E}_t x_{t+1}}{\partial z_t} < 0$ in (C1). Since $\chi_g \to 0$ as $\phi_\pi \to \infty$ and $\chi_g(1+\varsigma^{-1}\phi_x) \to 1$, $\chi_g \kappa \varsigma^{-1} \to 0$ as $\phi_x \to \infty$, the revision converges to 0 and to $\tilde{K}_u$ in the two limits, which are the values under the strict mandates of Proposition C1.

Proposition C1 (Strict mandates and the value of news)**.**

(i) Under strict inflation targeting, $\pi_t = 0$ and $\tilde{E}_t \pi_{t+1} = 0$, so realized outcomes do not respond to news: $x_t = -u_t/\kappa$, and the policy rate $i_t = v_t + (u_t - \tilde{K}_u z_t)/(\kappa \varsigma^{-1})$ falls *after inflationary news. Beliefs still move, $\tilde{E}_t x_{t+1} = -\tilde{K}_u z_t/\kappa$, but no outcome depends on beliefs. The sensitivity of outcomes to beliefs is zero, $Q = 0$, and so are the stakes of news in the attention problem (18).*

(ii) Under strict gap targeting, $x_t = 0$, $\pi_t = u_t + \beta \tilde{K}_u z_t$, and $i_t = v_t + \tilde{K}_u z_t$: belief revisions pass into realized inflation at full strength. Outcomes depend on beliefs about supply alone, $Q = \varpi_\pi \beta^2 e_u e_u^\top$, so the stakes of news in the attention problem (18) are positive, $\omega_1 = \varpi_\pi \beta^2 \sigma_u^2$, and any attention attracted by the news is directed at supply: $\tilde{K}_v = 0$.

(iii) Under either mandate, demand shocks are neutralized and beliefs about future demand move no outcome, so the value of information about future demand is zero.

Proof of Proposition C1. Under strict inflation targeting, $\pi_t = 0$ in every period and agents believe this, so $\tilde{E}_t \pi_{t+1} = 0$; the Phillips curve gives $x_t = -u_t/\kappa$, hence $\tilde{E}_t x_{t+1} = -\tilde{K}_u z_t/\kappa$, and the IS curve gives

$$\frac{\partial \pi_t}{\partial z_t} = \frac{\partial x_t}{\partial z_t} = 0, \quad \tilde{E}_t x_{t+1} = -\frac{\tilde{K}_u z_t}{\kappa}, \quad i_t = v_t + \frac{u_t - \tilde{K}_u z_t}{\kappa \varsigma^{-1}}. \tag{C4}$$

Under gap targeting, $x_t = 0$ in every period; the Phillips curve gives $\pi_t = u_t + \beta \tilde{E}_t \pi_{t+1}$ with $\tilde{E}_t \pi_{t+1} = \tilde{E}_t u_{t+1} = \tilde{K}_u z_t$, since future news is subjectively mean zero, and the IS curve gives $i_t = v_t + \tilde{K}_u z_t$. Both mandates are reached as limits of rules (24): inflation targeting as $\phi_\pi \to \infty$, and gap targeting as $\phi_x \to \infty$. The value of news follows from the tracking problem (18). Under strict inflation targeting the sensitivities of outcomes to beliefs are $b_\pi = b_x = 0$, so $Q = 0$ and $\omega_1 = 0$. Under gap targeting $b_\pi = (\beta, 0)^\top$ and

$b_x = 0$, so $Q = \varpi_\pi \beta^2 e_u e_u^\top$, Ω has rank one on the supply axis with $\omega_1 = \varpi_\pi \beta^2 \sigma_u^2$, and the optimal signal loads on u_{t+1} alone, so $\tilde{K}_v = 0$. In the demand direction both weight matrices vanish, which is (iii). □

Ranking of the Mandates. Per unit of supply-innovation variance, the fundamental damages at the two poles are $D_u^{\text{gap}} = \varpi_\pi$ and $D_u^{\text{SIT}} = \varpi_x/\kappa^2$, with ratio $\varpi_\pi \kappa^2/\varpi_x = \bar{\epsilon}\kappa$; demand shocks are neutralized at both poles. The news components break the near tie: gap targeting adds $\varpi_\pi \beta^2 \tilde{K}_u^2 \sigma_z^2$ to the loss, and strict inflation targeting adds nothing. With the rational-inattention gain at the gap-targeting pole, $\tilde{K}_u = \sigma_u/(1+\sigma_\eta^2)$ and $\sigma_z^2 = 1+\sigma_\eta^2$, strict inflation targeting dominates if and only if

$$\bar{\epsilon}\kappa > \frac{1}{1+\beta^2/(1+\sigma_\eta^2)}. \tag{C5}$$

At standard parameter values $\bar{\epsilon}\kappa = 1.03$, while the right side equals 0.67 at a signal-to-noise ratio of one.

C2 Policy and the Interpretation of News

Under the rule (24) with $\phi_z = 0$, write $m \equiv 1 + \varsigma^{-1}\phi_x \geq 1$ and $\chi_g = (m + a\phi_\pi)^{-1}$ with $a \equiv \kappa\varsigma^{-1}$. Substituting (C1) into (C2), the sensitivities of current outcomes to anticipated shocks are

$$b_{xu} = \chi_g^2 \varsigma^{-1}\left[m - \phi_\pi(1+\beta m)\right], \quad b_{xv} = \chi_g^2 \varsigma^{-1}\left[1 + a(1-\beta\phi_\pi)\right], \tag{C6}$$

$$b_{\pi u} = \chi_g^2\left[m(\beta m + a) - a\phi_\pi\right], \quad b_{\pi v} = \chi_g^2 a\left[1 + a + \beta m\right], \tag{C7}$$

which reduce to (8)–(9) at $\phi_x = 0$. The signs established under (10)–(11) survive any $\phi_x \geq 0$: b_{xv} and $b_{\pi v}$ are positive as before; $b_{xu} < 0$ since $m(1-\beta\phi_\pi) < 0 < \phi_\pi$ by (10); and $b_{\pi u} \geq \chi_g^2[\beta + a - a\phi_\pi] > 0$ by (11) and $m \geq 1$. Write $b_{\pi u}/b_{\pi v}$ and $|b_{xu}|/b_{xv}$ for the supply tilts of the inflation and output margins of Q. Then

$$\begin{aligned} \frac{\partial}{\partial m}\left(\frac{b_{\pi u}}{b_{\pi v}}\right) &= \frac{\beta^2 m^2 + 2\beta m(1+a) + a(1+a) + a\beta\phi_\pi}{a(1+a+\beta m)^2} > 0, \\ \frac{\partial}{\partial m}\left(\frac{|b_{xu}|}{b_{xv}}\right) &= \frac{\beta\phi_\pi - 1}{1 + a(1-\beta\phi_\pi)} > 0, \end{aligned} \tag{C8}$$

whereas at $\phi_x = 0$,

$$\frac{d}{d\phi_\pi}\left(\frac{b_{\pi u}}{b_{\pi v}}\right) = -\frac{1}{1+\beta+a} < 0, \quad \frac{d}{d\phi_\pi}\left(\frac{|b_{xu}|}{b_{xv}}\right) = \frac{1+\beta+a}{[1+a(1-\beta\phi_\pi)]^2} > 0. \tag{C9}$$

A stronger response to the output gap tilts both margins toward supply. A more hawkish rule tilts the output margin toward supply and the inflation margin away from it. The

following proposition, referred to in Section 5.1, states the consequences for the interpretation.

Proposition C2 (Policy and the stagflationary reading of news)**.** *Under rational inattention and the rule* (24) *with* $\phi_z = 0$*, on the region where* (10)–(B2) *hold at the rule:*

(i) News is read as stagflationary, $\tilde{K}_u/\tilde{K}_v > T_x$*, if and only if the responses of inflation and output to anticipated shocks are negatively correlated under the prior,*

$$b_\pi^\top \Sigma\, b_x < 0 \iff \frac{\sigma_u^2}{\sigma_v^2} \cdot \frac{b_{\pi u}}{b_{\pi v}} \cdot \frac{|b_{xu}|}{b_{xv}} > 1. \tag{C10}$$

(ii) Both supply tilts in (C10) *are strictly increasing in* ϕ_x*. Hence the set of variance ratios* σ_u^2/σ_v^2 *for which news is stagflationary expands with* ϕ_x*; a stagflationary reading is never undone by a stronger response to the output gap; and for any* $\sigma_u^2/\sigma_v^2 > 0$*, news is stagflationary for all* ϕ_x *sufficiently large.*

(iii) The stagflation threshold T_x *is strictly decreasing in both* ϕ_x *and* ϕ_π*.*

Under robustness, the relative damage of supply shocks, which determines the worst case, is strictly increasing in both ϕ_π *and* ϕ_x*.*

Proof of Proposition C2. (i) By the Perron–Frobenius argument of Lemma B1, which uses only (B2), the gains are positive, and by (B4) r^* is strictly increasing in c, so $\tilde{K}_u/\tilde{K}_v > T_x$ if and only if $\sigma_u^2/\sigma_v^2 > r^*(T_x)$. From the proof of Proposition 2(i), $r^*(T_x) = b_{\pi v} b_{xv}/(b_{\pi u}|b_{xu}|)$, which rearranges to the second inequality in (C10); multiplying through by $\sigma_v^2 b_{\pi v} b_{xv} > 0$ and using $b_{xu} = -|b_{xu}|$ gives $\sigma_u^2 b_{\pi u} b_{xu} + \sigma_v^2 b_{\pi v} b_{xv} < 0$, which is $b_\pi^\top \Sigma b_x < 0$. (ii) Monotonicity of both tilts is (C8), so $r^*(T_x)$ is strictly decreasing in ϕ_x; the first two claims follow from (i). For the limit, from (C6)–(C7),

$$r^*(T_x) = \frac{a(1 + a + \beta m)}{m(\beta m + a) - a\phi_\pi} \cdot \frac{1 + a(1 - \beta\phi_\pi)}{\phi_\pi + m(\beta\phi_\pi - 1)},$$

and each factor is $O(1/m)$, so $r^*(T_x) \to 0$ as $\phi_x \to \infty$. Condition (B2) holds for all large ϕ_x, since $b_{\pi u} \to \beta$, $b_{\pi v} = a\beta/m + O(m^{-2})$, and $b_{xu} b_{xv} = O(m^{-3})$, so $Q_{uv} = \varpi_\pi a \beta^2/m + O(m^{-2}) > 0$. (iii) $T_x = b_{xv}/|b_{xu}| = [1 + a(1 - \beta\phi_\pi)]/[\phi_\pi + m(\beta\phi_\pi - 1)]$. The numerator is decreasing in ϕ_π and independent of m; the denominator is increasing in m by (10) and increasing in ϕ_π with derivative $1 + \beta m > 0$. For robustness, the contemporaneous responses under the general rule are $x_t = \chi_g \varsigma^{-1}(v_t - \phi_\pi u_t)$ and $\pi_t = \chi_g[(1 + \varsigma^{-1}\phi_x)u_t + \kappa\varsigma^{-1} v_t]$, so the damage ratio is

$$R \equiv \frac{D_u}{D_v} = \frac{\varpi_\pi(1 + \varsigma^{-1}\phi_x)^2 + \varpi_x \varsigma^{-2}\phi_\pi^2}{\varsigma^{-2}(\varpi_\pi \kappa^2 + \varpi_x)}, \tag{C11}$$

which is strictly increasing in both arguments. □

Remark. The response to inflation has no counterpart to (ii). By (C9), a more hawkish rule raises the output tilt and lowers the inflation tilt in (C10), so the covariance $b_\pi^\top \Sigma b_x$, and with it the stagflationary reading, need not move monotonically in ϕ_π; only the threshold result (iii) is unconditional.

Configurations Outside the Maintained Region. Define $\phi_s \equiv 1 + \beta/a$, the boundary of (11) at which $b_{\pi u}$ changes sign, and $\phi_b \equiv \beta^{-1}(1 + 1/a)$, at which b_{xv} changes sign; $\phi_b - \phi_s$ is positive and of order $1 - \beta$. Under the loss weights (17), the robust corner switches from demand to supply at $\phi_R \equiv \sqrt{1 + \bar{\epsilon}\varsigma(a - 1/a)}$ on the branch $a > 1$. Three configurations follow, corresponding to a flat, a moderately steep, and a steep Phillips curve; the Online Appendix contains the full statements. Two results are relevant for the main text. First, forecast comovement classifies the belief regime perfectly (it is positive if and only if news is read as demand news), and is invariant to every switch in the character of transmission, because expectations are formed under the believed shock's own one-step mapping while realized responses run through the policy-distorted mapping. Second, the joint observation of stagflationary realized responses to inflationary news *and* negative forecast comovement occurs only in the configuration with supply-read news and $\phi_\pi < \phi_s$. The data therefore restrict the economy to the region in which (11) holds.

C3 Optimal Policy

A policy strategy is a stationary linear map from the public state (u_t, v_t, ν_t) to outcomes,

$$\pi_t = a_u u_t + a_v v_t + a_\nu^\top \nu_t, \quad x_t = \frac{1}{\kappa}(\pi_t - \beta \tilde{E}_t \pi_{t+1} - u_t), \quad \tilde{E}_t \pi_{t+1} = a^\top \nu_t, \quad a \equiv (a_u, a_v)^\top, \tag{C12}$$

where ν_t is the common posterior about ε_{t+1} and the instrument is determined by the IS curve. Agents forecast tomorrow's inflation by applying the plan to their beliefs, $\tilde{E}_t \pi_{t+1} = a^\top \nu_t$, so today's effective cost-push term is $c_t = u_t + \beta a^\top \nu_t$ and $\kappa x_t = \pi_t - c_t$. Given a strategy σ, outcomes respond to the posterior with sensitivities $b_\pi = a_\nu$ and $b_x = (a_\nu - \beta a)/\kappa$, and the collector solves the attention problem at those sensitivities:

$$\mathcal{B}(\sigma): \min_{\text{signal}} \mathbb{E}[(\nu - \varepsilon)^\top Q(\sigma)(\nu - \varepsilon)] + \vartheta \mathcal{I}, \quad Q(\sigma) = \varpi_\pi b_\pi b_\pi^\top + \varpi_x b_x b_x^\top, \tag{C13}$$

where $\mathcal{I}$ is the mutual information between the signal and ε. The strategy is set first; attention is designed once, knowing the strategy; shocks then realize period by period. Welfare is (25). Under discretion and expectations management the signal design is taken as given, so the cost of information is a constant in the planner's objective; the belief-shaping planner minimizes the macro loss $2L$ and the cost-internalizing planner minimizes $\mathcal{W}$, each taking the collector's best response into account.

The derivations use two facts. First, the posterior ν_t is independent of the current

shocks, so the macro loss breaks into a fundamental block and a belief block, $2L = F(a) + B$, with

$$F(a) = \sigma_u^2\big[\varpi_\pi\Theta + \Gamma(a_u - \Theta)^2\big] + \Gamma\sigma_v^2 a_v^2, \quad B = \mathbb{E}[\nu_t^\top Q \nu_t], \quad \Gamma \equiv \frac{\varpi_x}{\kappa^2\Theta}, \tag{C14}$$

where F follows from the identity (26) at $(\pi, c) = (a_u, 1)$ per unit of u_t and at $(a_v, 0)$ per unit of v_t. Second, the same identity at $(\pi, c) = (a_\nu^\top y, \beta a^\top y)$ for every vector y gives

$$Q = \varpi_\pi\Theta\beta^2\, aa^\top + \Gamma\, dd^\top, \quad d \equiv a_\nu - \Theta\beta a, \tag{C15}$$

so $Q \succeq \varpi_\pi\Theta\beta^2 aa^\top$, with equality if and only if $d = 0$. The first term is the sensitivity of outcomes to beliefs created by the promised response a through the cost-push term; the second is present only when the response to beliefs departs from the targeting rule.

Lemma C1 (Common reduction)**.** *For each of the four planners: (i) the response to beliefs applies the targeting rule to the belief-driven cost-push term, $a_\nu = \Theta\beta a$, which simultaneously minimizes the belief block for given beliefs and drives $Q(\sigma)$ to its positive-semidefinite minimum $\varpi_\pi\Theta\beta^2 aa^\top$. (ii) $a_v = 0$, and attention (if any) is exclusively directed at supply: the collector observes at most the single signal $z_t = u_{t+1} + \eta_t$, and $\tilde{K}_v = 0$. (iii) The macro loss and welfare reduce to functions of a_u alone,*

$$2L(a_u) = S(a_u) + N(a_u), \quad S(a_u) = \sigma_u^2\left[\varpi_\pi\Theta + \frac{\varpi_x}{\kappa^2\Theta}(a_u - \Theta)^2\right], \quad N(a_u) = \max\big(0, \bar{\omega}(a_u) - \vartheta/2\big), \tag{C16}$$

$$\mathcal{W}(a_u) = 2L(a_u) + \vartheta \log\max(1, a_u/\hat{a}), \tag{C17}$$

with $\bar{\omega}(a_u) = \varpi_\pi\Theta\beta^2 a_u^2\sigma_u^2$, the stakes of supply news, and attention threshold

$$\hat{a} = \frac{1}{\beta}\sqrt{\frac{\vartheta}{2\varpi_\pi\Theta\sigma_u^2}}, \quad \bar{\omega}(\hat{a}) = \vartheta/2. \tag{C18}$$

Any other optimal strategy yields the same outcomes.

Proof. The collector's problem (C13) is the reverse water-filling problem of Appendix B1 on $\Omega = \Sigma^{1/2} Q \Sigma^{1/2}$, with eigenvalues $\omega_1 \geq \omega_2$: along eigenvector i the resolved share of prior variance is $r_i = \max(0, 1 - \vartheta/(2\omega_i))$, so the tracking loss is $\sum_i \min(\omega_i, \vartheta/2)$, the belief block is $B = \sum_i \max(0, \omega_i - \vartheta/2)$, and $B + \vartheta\mathcal{I} = \sum_i n_C(\omega_i)$ with $n_C(\omega) \equiv \max(0, \omega - \vartheta/2) + (\vartheta/2)\log\max(1, 2\omega/\vartheta)$. Both B and $B + \vartheta\mathcal{I}$ are nondecreasing in each eigenvalue, and the two values add up to $\operatorname{tr}\Omega$, which the strategy fixes.

(i) Under discretion the planner takes $\tilde{E}_t\pi_{t+1} = a^{+\prime}\nu_t$ as given, where a^+ is its successors' response, and minimizes the period loss realization by realization; the targeting rule $\pi_t = \Theta c_t$ gives $a_u = \Theta$, $a_v = 0$, and $a_\nu = \Theta\beta a^+$, and the Markov condition $a^+ = a$ gives $d = 0$. Under expectations management the distribution of ν_t is held fixed, a_ν

enters only $B = \varpi_\pi \Theta \beta^2 \mathbb{E}(a^\top \nu_t)^2 + \Gamma \mathbb{E}(d^\top \nu_t)^2$, and $d = 0$ attains the lower bound whatever that distribution is. Under either belief shaper the distribution of ν_t moves with a_ν through Ω. By (C15), $\Omega = \varpi_\pi \Theta \beta^2 \tilde{a}\tilde{a}^\top + \Gamma \tilde{d}\tilde{d}^\top$ with $\tilde{a} = \Sigma^{1/2} a$ and $\tilde{d} = \Sigma^{1/2} d$, so $\omega_1 \geq \omega_a \equiv \varpi_\pi \Theta \beta^2 (a_u^2 \sigma_u^2 + a_v^2 \sigma_v^2)$ for every d (evaluate the quadratic form at $\tilde{a}/|\tilde{a}|$), while $d = 0$ gives $\omega_1 = \omega_a$ and $\omega_2 = 0$. Since the belief part of each objective is nondecreasing in both eigenvalues, $d = 0$ minimizes it; leaning harder against belief-driven inflation only shifts sensitivity from inflation to output and raises the stakes. (ii) With $d = 0$, F contains the term $\Gamma \sigma_v^2 a_v^2$ and the belief part of each belief shaper's objective is nondecreasing in ω_a, which is increasing in a_v^2, so $a_v = 0$. For the expectations manager, at the collector's best response $\mathrm{Var}(\nu_t) a = r_1 \Sigma a$, so the first-order condition for a_v is $(\Gamma + \varpi_\pi \Theta \beta^2 r_1) \sigma_v^2 a_v = 0$. With $a_v = 0$, Ω is rank one on the supply axis, and the collector learns about u_{t+1} alone. (iii) With $a_v = 0$ and $d = 0$, $F = S(a_u)$, $\omega_1 = \bar{\omega}(a_u)$, the resolved share is $r(a_u) = \max(0, 1 - (\hat{a}/a_u)^2)$, $B = N(a_u)$, and $\mathcal{I} = -\frac{1}{2}\log(1 - r) = \log \max(1, a_u/\hat{a})$. Uniqueness of outcomes: equality in the bounds above requires $a_v = 0$ and either no attention or $d^\top \nu_t = 0$, in which case a_ν does not affect outcomes. □

Proof of Proposition 5. By Lemma C1, each planner chooses a single number $a = a_u \geq 0$. In units of $\Gamma \sigma_u^2$, the macro loss and welfare are

$$\ell(a) = \Theta(1-\Theta) + (a-\Theta)^2 + \psi \max(0, a^2 - \hat{a}^2), \quad \ell^{\mathcal{W}}(a) = \ell(a) + 2\psi \hat{a}^2 \log \max(1, a/\hat{a}), \tag{C19}$$

using $\varpi_\pi \Theta/\Gamma = \Theta(1-\Theta)$, $\varpi_\pi \Theta \beta^2/\Gamma = \psi$, and $\vartheta/(\Gamma \sigma_u^2) = 2\psi \hat{a}^2$. Both functions are convex: the news term is the maximum of two convex functions, and the information term $G(a) \equiv a^2 - \hat{a}^2 + 2\hat{a}^2 \log(a/\hat{a})$ on $a > \hat{a}$ has $G'' = 2 - 2\hat{a}^2/a^2 > 0$, with slope jumping from 0 to $4\hat{a}$ at the kink.

(i) Discretion. The targeting rule gives $a_u^D = \Theta$ whatever the collector does, so the fixed point with the collector is immediate: the collector attends if and only if $\hat{a} < \Theta$.

(ii) Expectations management. The planner holds the resolved share r fixed and minimizes $(a - \Theta)^2 + \psi r a^2$; the collector's share is water-filled. The equilibrium is a fixed point of

$$a(r) = \frac{\Theta}{1 + \psi r}, \quad r(a) = \max\left(0, 1 - (\hat{a}/a)^2\right), \tag{C20}$$

which is unique because $a(\cdot)$ is decreasing and $r(\cdot)$ is nondecreasing: $a_u^N = \Theta$ with no attention if $\hat{a} \geq \Theta$, and otherwise the positive root of $(1 + \psi)a^2 - \Theta a - \psi \hat{a}^2 = 0$,

$$a_u^N = \frac{\Theta + \sqrt{\Theta^2 + 4\psi(1+\psi)\hat{a}^2}}{2(1+\psi)} \in \left(\max(a_1, \hat{a}), \Theta\right), \tag{C21}$$

where the location follows from the sign of the quadratic at $\hat{a}$, a_1, and Θ. The strict inequality $a_u^N > \hat{a}$ whenever $\hat{a} < \Theta$ implies that this planner never anchors.

(iii) Belief shaping. The planner minimizes ℓ. On $a \leq \hat{a}$ the objective is the parabola

$(a-\Theta)^2$ plus a constant, decreasing for $a<\Theta$; on $a\geq\hat{a}$ its derivative is proportional to $(1+\psi)a-\Theta$, which vanishes at a_1. Hence $a_u^S=a_1$ if $\hat{a}<a_1$, $a_u^S=\hat{a}$ if $a_1\leq\hat{a}<\Theta$, and $a_u^S=\Theta$ if $\hat{a}\geq\Theta$.

(iv) Cost-internalizing belief shaping. The planner minimizes $\ell^{\mathcal{W}}$. On $a>\hat{a}$ the sign of the derivative is that of $h(a)\equiv(1+\psi)a^2-\Theta a+\psi\hat{a}^2$, and the right derivative at the kink has the sign of $(1+2\psi)\hat{a}-\Theta$. If $\hat{a}\geq\Theta$ the static minimizer lies in the no-attention region. If $\Theta/(1+2\psi)\leq\hat{a}<\Theta$ the objective is decreasing up to $\hat{a}$ and, by convexity, nondecreasing beyond it, so the plan sits at the kink. If $\hat{a}<\Theta/(1+2\psi)$, then $h(\hat{a})<0$, so $\hat{a}$ lies between the two roots of h and the minimizer is the larger root; the larger root is decreasing in $\hat{a}$, tends to a_1 as $\hat{a}\to 0$, and equals $\Theta/(1+2\psi)$ at the edge of the anchoring regime.

(v) Ordering. $a_u^S\leq a_u^N\leq\Theta$ follows from (ii) and (iii). For $a_u^C\leq a_u^S$: if $\hat{a}<\Theta/(1+2\psi)$, then $h(a_1)=\psi\hat{a}^2>0$ and $a_1>\hat{a}$, so a_1 lies above the larger root and $a_u^C<a_1=a_u^S$; if $\Theta/(1+2\psi)\leq\hat{a}<a_1$, then $a_u^C=\hat{a}<a_1=a_u^S$; otherwise the two coincide. All four planners are evaluated on the convex function $\ell^{\mathcal{W}}$, which is minimized at a_u^C and is nondecreasing to the right of its minimizer, so $\mathcal{W}^C\leq\mathcal{W}^S\leq\mathcal{W}^N\leq\mathcal{W}^D$, and (27) is an identity. □

Proposition C3 (One lag of memory)**.** *Augment the plan with the term $c\,u_{t-1}$. The new term is known at t, so b_π, b_x, Q, $\hat{a}$, and Lemma C1 are unchanged, and the two commitment margins separate exactly. The static block becomes*

$$S_1(a,c)=\Big[\varpi_\pi(a^2+c^2)+\frac{\varpi_x}{\kappa^2}\big((a-\beta c-1)^2+c^2\big)\Big]\sigma_u^2,\quad c(a)=\frac{\Theta\beta(a-1)}{1+\Theta\beta^2}<0, \tag{C22}$$

where $c(a)$ is the optimal promise given a: a partial reversal*, which lowers today's effective cost-push term to $(1+\beta c)u_t$. Concentrating out the promise, $S_1(a,c(a))/(\Gamma\sigma_u^2)$ is a parabola with minimum at a_1 and curvature $\gamma\equiv(1+\psi)/(1+\Theta\beta^2)<1$ relative to the memoryless static block. The belief-shaping optimum is $a_u^{1,S}=\min\{a_1,\max(a_2,\hat{a})\}$ with $a_2=\Theta/[1+\psi(2+\beta^2\Theta)]=\gamma a_1/(\gamma+\psi)$. The cost-internalizing optimum is $a_u^{1,C}=a_1$ if $\hat{a}\geq a_1$, $a_u^{1,C}=\hat{a}$ if $\gamma a_1/(\gamma+2\psi)\leq\hat{a}<a_1$, and otherwise the larger root of $(\gamma+\psi)a^2-\gamma a_1 a+\psi\hat{a}^2=0$. For both planners the improvement over the memoryless optimum is strict for every ϑ, the anchoring region is wider in proportional terms, and the policy is time-inconsistent in the classical sense: re-optimizing at t sets the response to the inherited shock to zero while promising $c(a)$ again for the future.*

Proof. Tomorrow's plan responds to u_t, which is known at t, so $\tilde{E}_t\pi_{t+1}=a^\top\nu_t+cu_t$ and the cost-push term becomes $c_t=(1+\beta c)u_t+\beta a^\top\nu_t$: the promise changes the supply component and leaves the belief component unchanged, so the sensitivities (C15) and Lemma C1 are unaffected, and u_{t-1} is independent of (ε_t,ν_t). Collecting the u_t and u_{t-1} terms of π_t and $\kappa x_t=\pi_t-c_t$ gives S_1. As a function of c, S_1/σ_u^2 is $Ac^2-2(\varpi_x/\kappa^2)\beta(a-1)c$

plus terms without c, with $A = \Gamma(1+\Theta\beta^2)$, whose minimizer is $c(a)$ in (C22) (using $\varpi_x/\kappa^2 = \Gamma\Theta$); the minimized value is $\sigma_u^2[\varpi_\pi a^2 + \Gamma\Theta(a-1)^2/(1+\Theta\beta^2)]$, so memory divides the weight on the output gap by $1+\Theta\beta^2$. Dividing by $\Gamma\sigma_u^2$ and using $\varpi_\pi/\Gamma = 1-\Theta$, the first-order condition without attention is $(1-\Theta)a + \Theta(a-1)/(1+\Theta\beta^2) = 0$, which gives a_1 because $(1-\Theta)(1+\Theta\beta^2) + \Theta = 1+\psi$; the curvature is $(1-\Theta) + \Theta/(1+\Theta\beta^2) = \gamma$. The concentrated objectives are therefore $\gamma(a-a_1)^2 + \psi\max(0, a^2-\hat{a}^2)$ and the same plus $2\psi\hat{a}^2\log\max(1, a/\hat{a})$, up to constants, and the three-case arguments of parts (iii) and (iv) of the proof of Proposition 5 apply with $(\Theta, 1)$ replaced by (a_1, γ): the interior solutions are $a_2 = \gamma a_1/(\gamma+\psi)$ and the larger root of $(\gamma+\psi)a^2 - \gamma a_1 a + \psi\hat{a}^2 = 0$, and the anchoring regimes are $[a_2, a_1)$ and $[\gamma a_1/(\gamma+2\psi), a_1)$. The memoryless optima satisfy $a_u < 1$, so the promise $c(a_u)$ is nonzero and $S_1(a_u, c(a_u)) < S_1(a_u, 0) = S(a_u)$ while the news and information terms are unchanged; hence memory pays strictly for every ϑ, even when attention is dead. Since ϑ is proportional to $\hat{a}^2$, the anchoring regime of S spans a factor $(a_1/a_2)^2 = [(\gamma+\psi)/\gamma]^2$ in the cost of attention, which exceeds the memoryless factor $(1+\psi)^2$ because $\gamma < 1$, and likewise for C. Finally, the benefit of the promise cu_{t-1}, a smaller cost push at $t-1$, has been collected by date t; at t the term is a pure deviation from the targeting rule, by (26) at $c = 0$, so a planner who re-optimizes sets it to zero while promising $c(a)u_t$ for $t+1$. □

Table C1: The Planner Ladder: Illustration

	Memoryless				One lag of memory	
	D	N	S	C	S	C
$\vartheta = 1.5$						
a_u	0.493	0.436	0.396	0.345	0.306	0.298
$2L$	9.906	9.572	9.503	9.614	7.187	7.190
$\mathcal{W}$	10.659	10.142	9.928	9.832	7.228	7.190
$\vartheta = 2.2$						
a_u	0.493	0.452	0.396	0.361	0.361	0.361
$2L$	9.556	9.291	9.153	9.205	6.949	6.949
$\mathcal{W}$	10.240	9.787	9.355	9.205	6.949	6.949
$\vartheta = 3$						
a_u	0.493	0.470	0.422	0.422	0.396	0.396
$2L$	9.156	8.991	8.782	8.782	6.915	6.915
$\mathcal{W}$	9.623	9.317	8.782	8.782	6.915	6.915

Notes: Illustrative parameter values ($\beta = 0.99$, $\varsigma = \varphi = 1$, $\theta = 0.75$, $\bar{\epsilon} = 6$, $\sigma_u = 0.5$; implying $\Theta = 0.493$, $a_1 = 0.396$, and $a_2 = 0.306$). ϑ is the unit cost of attention; a_u is the response of inflation to supply shocks, $2L$ twice the per-period expected macro loss, and $\mathcal{W}$ welfare (25). D: discretion; N: expectations management; S: belief shaping; C: cost-internalizing belief shaping; the last two columns add one lag of memory to S and C. The anchoring windows are $\vartheta \in [2.64, 4.09)$ for S and $[1.84, 4.09)$ for C without memory, and $[1.58, 2.64)$ and $[1.05, 2.64)$ with one lag. Computed from the closed forms of Propositions 5 and C3.

Appendix D The Quantitative Model

This appendix states the quantitative model used in Section 4.3. The full development and numerical verification are in the Online Appendix.

D1 Environment

Period utility is $(C_t - b\bar{C}_{t-1})^{1-\varsigma}/(1-\varsigma) - N_t^{1+\varphi}/(1+\varphi)$, with external habit in lagged aggregate consumption $\bar{C}_{t-1}$. The demand shock is the risk-premium shock of the baseline model. The cost-push shock is a proportional revenue tax on firms which is rebated lump-sum, $u_t \equiv \lambda\hat{\tau}_t$. This shock is observationally equivalent to the markup shock at first order, but is preferable for exact welfare accounting: both shocks are then wedges which do not affect preferences or technology. The shocks follow independent AR(1) processes with persistence ρ_u and ρ_v. The log-linear system is

$$(1+b)x_t = bx_{t-1} + \tilde{E}_t x_{t+1} - \frac{1-b}{\varsigma}(i_t - \tilde{E}_t\pi_{t+1} - v_t), \tag{D1}$$

$$\pi_t = \beta\tilde{E}_t\pi_{t+1} + \lambda\, mrs_t + u_t, \quad mrs_t = \varphi x_t + \frac{\varsigma}{1-b}(x_t - bx_{t-1}), \tag{D2}$$

with the rule $i_t = \rho_i i_{t-1} + (1-\rho_i)(\phi_\pi\pi_t + \phi_x x_t)$. Because both shocks are wedges, the efficient allocation does not respond to them and x_t is the welfare-relevant gap.

One scalar public signal per period, $z_t = \alpha^\top\varepsilon_{t+h} + \eta_t$, carries news about the innovation vector h periods ahead. Beliefs are $\tilde{E}_t[\varepsilon_{t+j}] = \tilde{K} z_{t+j-h}$, so that at each date there are h outstanding vintages of news, each of which is held fixed until it matures (Lemma D1). The state is $(v_t, u_t, x_{t-1}, m_{1,t}, \ldots, m_{h,t})$ with stored signals $m_{j,t} = z_{t+j-h}$. We solve the system by the generalized Schur method; the objective law of motion applies the same policy functions to the true innovations and signals. In the limit $h = 1$, $b = 0$, $\rho_u = \rho_v = 0$ the solution reproduces Proposition 1 exactly.

D2 Exact Welfare under Habit

Proposition D1 (Quadratic loss with external habit)**.** *Let the constant employment wedge satisfy $1 - \bar{\tau} = (\bar{\epsilon} - 1)/[\bar{\epsilon}(1 - \beta b)]$, so that the steady state is efficient inclusive of the habit externality. Then to second order, up to terms independent of policy and beliefs,*

$$\mathcal{W} = -\frac{\Lambda(1-\beta b)}{2}\,\mathbb{E}\sum_{t\geq 0}\beta^t\ell_t + \text{t.i.p.}, \quad \ell_t = \frac{\bar{\epsilon}}{\lambda}\pi_t^2 + \varphi x_t^2 + \frac{\varsigma}{(1-b)(1-\beta b)}(x_t - bx_{t-1})^2, \tag{D3}$$

with $\Lambda > 0$ constant. At $b = 0$, ℓ_t is twice the loss (17).

Sketch. Expand the utility of consumption in the habit-adjusted argument $X_t = C_t - b\bar{C}_{t-1}$. Discounting moves quasi-differences onto single dates via $\sum_t \beta^t(a_t - ba_{t-1}) =$

$(1-\beta b)\sum_t \beta^t a_t$ + t.i.p. The labor term contributes the price-dispersion sum $\sum_t \beta^t \ln \Delta_t = \frac{\bar{\epsilon}}{2\lambda}\sum_t \beta^t \pi_t^2$ + t.i.p. The linear terms cancel under the planner's steady-state condition $(1-\beta b)\bar{X}^{-\varsigma} = \bar{N}^{\varphi}$, which the stated wedge delivers. The log-linear dynamics do not depend on $\bar{\tau}$. The full derivation, which we verify against exact nonlinear welfare, is in the Online Appendix. □

D3 Belief Formation at Horizon h

Lemma D1 (Fixed vintages). *The belief about ε_{t+h} formed at t is constant until $t+h$, when the truth is observed. Innovations at other dates and the noise in later signals are independent of ε_{t+h}, and aggregates realized in the interim are functions of states and common beliefs which are already known to the collector.*

Lemma D2 (Belief linearity and superposition). *Equilibrium outcomes satisfy $(x_t, \pi_t)^\top = F_f(v_t, u_t, x_{t-1})^\top + \sum_{j=1}^{h} \Gamma_j \nu_{j,t}$, where the matrices F_f and Γ_j do not depend on $\tilde{K}$. For each vintage, the objective path equals the full-information path plus the pure-noise response path scaled by the belief error $\nu - \varepsilon_{t+h}$, and belief errors are independent across vintages.*

Proposition D2 (Tracking matrix). *The conditional loss which is attributable to belief errors equals the sum over vintages of $\mathbb{E}[e^\top Q e]$ with*

$$Q_{ij} = \sum_{t\geq 0} \beta^t \, G_i(t)^\top \, W \, G_j(t), \tag{D4}$$

where G_i is the response path to a unit signal in the model solved with $\tilde{K} = e_i$, and W collects the welfare weights of (D3). Q does not depend on $\tilde{K}$, so the attention problem requires no fixed point. In the iid limit without habit, (D4) reduces to (18).

The rational-inattention problem in the quantitative model is thus the tracking problem (18) with the matrix (D4) and with Σ defined over innovations; the robust corner is computed from the damages of unanticipated innovations under W. Proofs of Lemmas D1 and D2 and of Proposition D2 are in the Online Appendix. With habit, the endogenous state generates a covariance between the anticipation and realization phases of each vintage; this term is excluded from the collector's tracking criterion by construction, and it is identically zero at $b = 0$.

D4 Beliefs, Model Moments, and Estimation

In the baseline the signal is $z_t = \alpha^\top \varepsilon_{t+h} + \eta_t$ with the normalization $\alpha^\top \Sigma \alpha = 1$, and the gains are

$$\tilde{K} = \frac{\Sigma\alpha}{1+\sigma_\eta^2} = \Sigma^{1/2} d\,(1-v_1), \quad d \equiv \Sigma^{1/2}\alpha, \quad v_1 \equiv \frac{\sigma_\eta^2}{1+\sigma_\eta^2}, \tag{D5}$$

where d is the unit direction of the signal in standardized shock space and v_1 is the share of the prior variance in that direction which the signal leaves unresolved. Pure supply news is $d = (0, 1)^\top$, so that $\tilde{K}_v = 0$ and $\tilde{K}_u = \sigma_u(1 - v_1)$. The model counterpart of the survey forecast is the year-ahead forecast $R_t \equiv \tilde{E}_t \sum_{k=1}^{12} \pi_{t+k}$, whose revision at t is linear in the innovations, and the counterpart of the news shock is the part of the revision driven by the signal, R_t^z. Because the signal is unpredictable from past states, the model counterpart of the local-projection coefficient at horizon h is the population projection of y_{t+h} on R_t^z, per standard deviation of R_t^z, computed from the solution (D1)–(D2) under the subjective law of motion, in which the anticipated innovation realizes on average at maturity. For year-over-year inflation the monthly responses are summed over twelve months; unemployment is ωx with the Okun coefficient $\omega = -0.5$; and industrial production is $\hat{\delta}\omega x$, with $\hat{\delta} = -2.5$ the sample slope of the industrial-production gap on the unemployment gap. Under Bayesian beliefs the response of year-over-year inflation at $h = 12$ per unit of R_t^z is exactly one, since the revision is the conditional expectation of π_{t+12}^{yoy}. The unconditional moments follow from the Lyapunov equation of the state under the objective law of motion.

The five estimated parameters $(\sigma_u, \sigma_v, v_1, \rho_u, \rho_v)$ minimize the sum of squared relative deviations of the five model moments from their data counterparts, searched over the logs of the volatilities and the logits of v_1 and of the annual persistences; the system is exactly identified and the minimum is zero to numerical precision (Table 4). Two objects are computed at the estimates and are not targeted. The forecast comovement is the correlation of the revisions in the year-ahead forecasts of inflation and of the output gap over all innovations. The strict rational-inattention solution is the water-filling solution of (18) with the matrix (D4), with the information cost $\vartheta = 2\omega_1 v_1$ set so that the collector leaves the same unresolved share along its top eigenvector as the estimated signal; the two belief models then differ only in the direction of attention. At the estimates the collector devotes 66 percent of attention to supply news, with a relative gain of 1.4; the damage of a unit supply innovation under the welfare criterion of Appendix D2 exceeds that of a unit demand innovation by a factor of 30.1, so the robust corner is the supply corner as well. The forecast comovement at the estimates is -0.17, the sign of the survey moment of Table 1.

The identification is close to recursive. Because beliefs are Bayesian, realized year-over-year inflation one year ahead moves one-for-one with the revision in the year-ahead forecast, so the inflation response at $h = 12$ pins down the standard deviation of the news-driven revision (0.17 percentage points). The ratio of the unemployment response at two years to the inflation response at one year then depends, given the fixed parameters, only on the persistence of the supply wedge, which is 0.50 at an annual rate (a half-life of about one year). The three unconditional moments pin down the volatility of demand relative to supply, the overall scale, and the persistence of the demand wedge, 0.93 at an annual

rate; with a flat Phillips curve, transitory demand shocks move output much more than inflation, and persistent demand shocks move inflation through expected future gaps, which allows a demand-driven economy to match the volatility of inflation. The local projections use the unemployment rate, and the unconditional moments the gap relative to the CBO natural rate, which trends down over the sample together with inflation and has no counterpart in the stationary model; the natural rate moves smoothly at the quarterly frequency and is not expected to respond to a release-window shock. The model's correlation between monthly inflation and the output gap is 0.17, and between year-over-year inflation and the output gap 0.11, the data value.

Appendix E Additional Tables and Figures



Figure E1: Local Projections with Released-Content Controls

Notes: The baseline four-lag control set augmented with four lags each of core, energy, and food CPI inflation (year-over-year, matching the headline convention). The release in month t reports month $t-1$, so the first lag is close to the content of the release itself. Coefficients scaled per standard deviation of the shock; 90% confidence intervals.

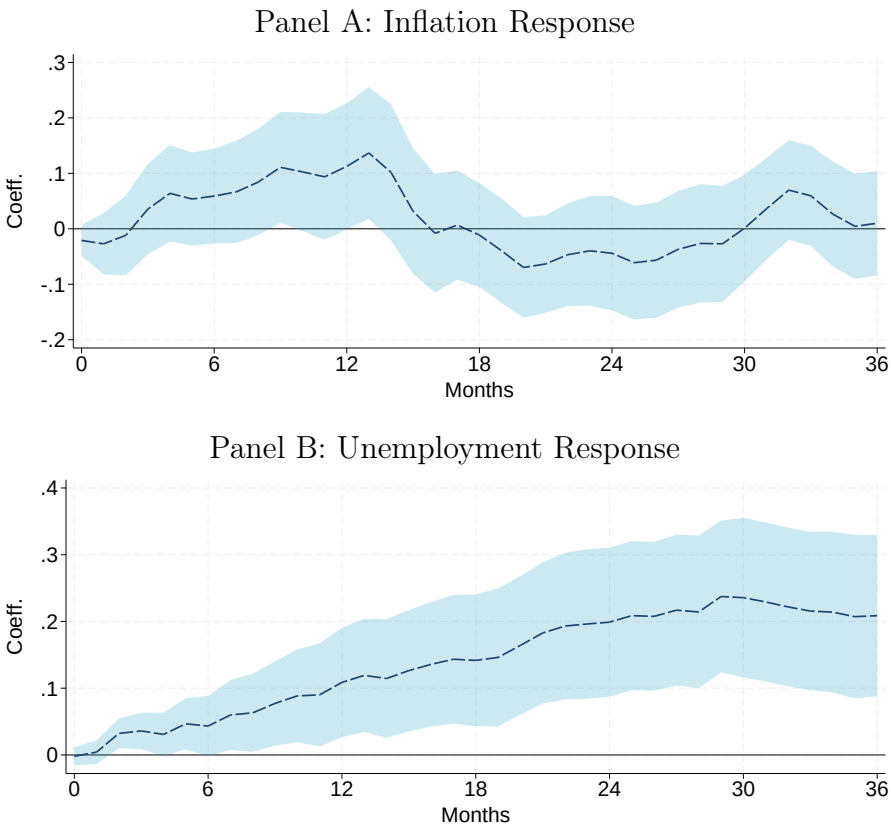


Figure E2: Response to a 3-Day News Shock

Notes: Local projections of inflation and unemployment on a news shock based on 3-day windows around CPI releases. Baseline controls. Coefficients scaled per standard deviation of the shock; 90% confidence intervals.

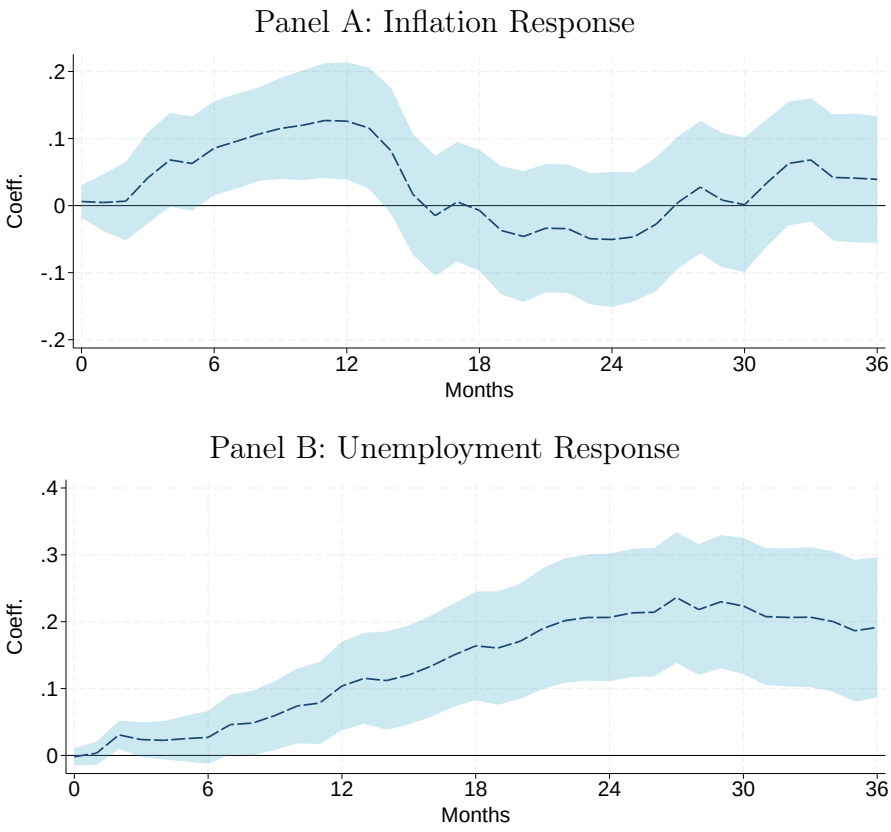


Figure E3: Response to a 7-Day News Shock

Notes: Local projections of inflation and unemployment on a news shock based on 7-day windows around CPI releases. Baseline controls. Coefficients scaled per standard deviation of the shock; 90% confidence intervals.

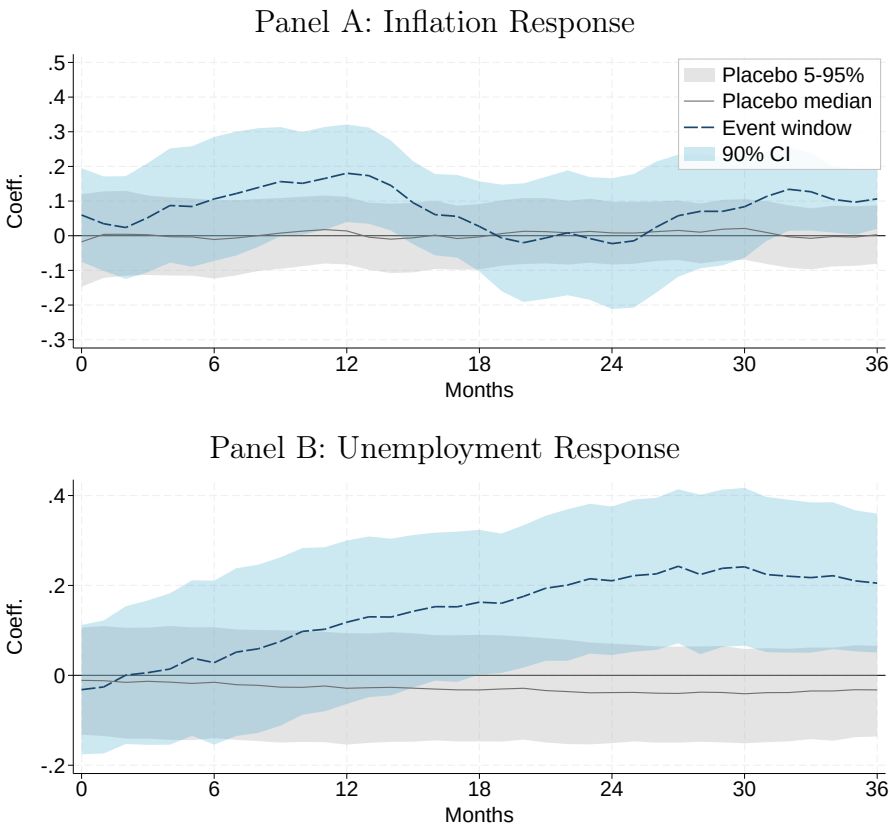


Figure E4: Randomization Inference: No-Control Specification

Notes: As Figure 4, with no controls in either the true or the placebo regressions. Coefficients scaled per standard deviation of the shock; 90% confidence intervals from Newey–West standard errors with truncation lag equal to the horizon plus one.

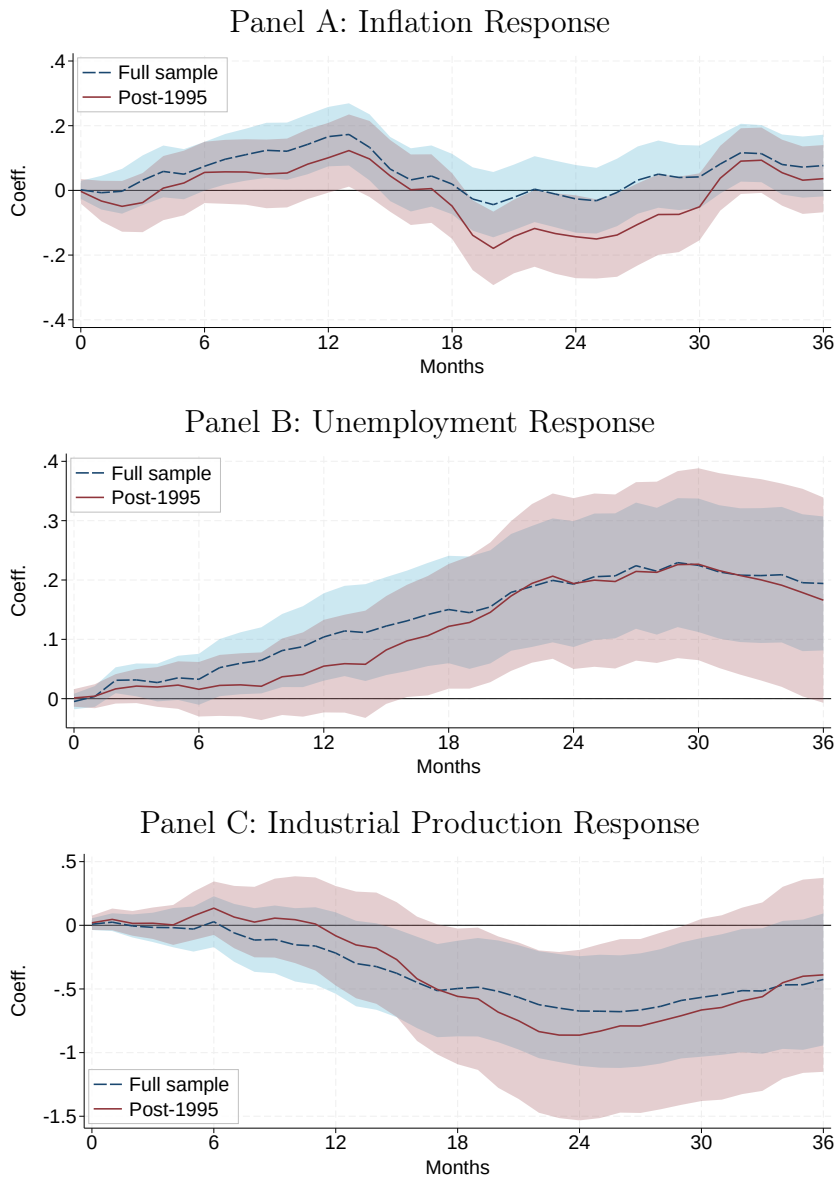


Figure E5: Response to a News Shock: Full and Post-1995 Samples

Notes: Local projections of inflation, unemployment, and industrial production on the news shock, estimated on the full sample (January 1982 to February 2020, dashed) and on release months from January 1995 to February 2020 (solid). Baseline controls. Coefficients scaled per standard deviation of the shock within each sample; shaded areas are 90% confidence intervals.

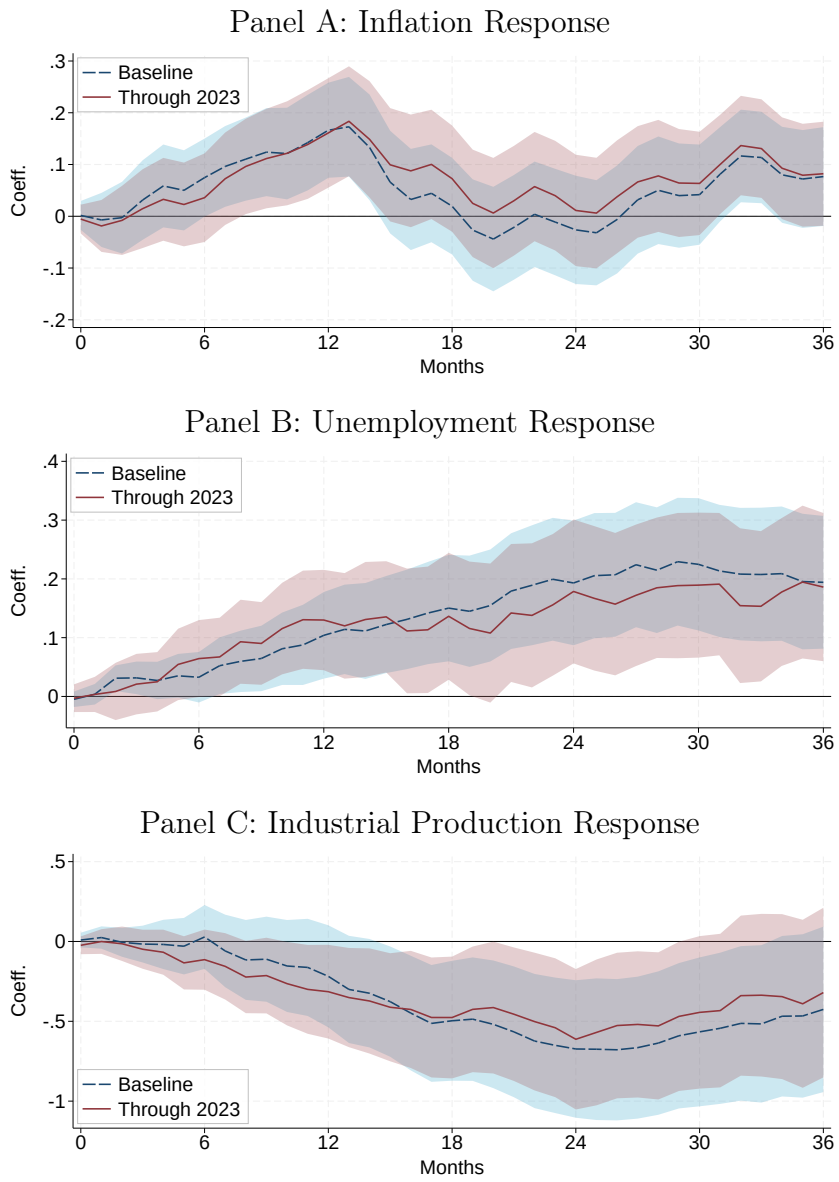


Figure E6: Response to a News Shock: Sample Through 2023

Notes: Local projections of inflation, unemployment, and industrial production on the news shock, estimated on the baseline sample (January 1982 to February 2020, dashed) and on a sample which adds the COVID period and the inflation which followed (solid): shocks through April 2023, the last month with interview dates, and macroeconomic data through September 2023. Baseline controls; the trend of industrial production is fitted over the baseline sample in both. Coefficients scaled per standard deviation of the shock within each sample; shaded areas are 90% confidence intervals.