The Effects of News Shocks and Supply-Side Beliefs*

Rupal Kamdar
Indiana University, Bloomington

Walker Ray
London School of Economics

Not for Circulation - March 2024

Abstract

This paper investigates the causes and consequences of "supply-side beliefs" —the tendency for households to overweight the importance of aggregate supply shocks relative to demand-side factors. We develop a New Keynesian model where agents receive news about future supply and demand shocks. Because supply shocks are more costly, potential model misspecification causes agents to endogenously downweight the likelihood that news is informative about demand shocks. The model rationalizes a number of empirical puzzles, such as the observed positive correlation between survey-based inflation and unemployment expectations, and the instability of estimated Phillips curve slope coefficients. Next, we empirically test a key prediction of the model: news shocks cause realized inflation and unemployment to move in the same direction. Using daily survey-based inflation expectations around CPI releases, we construct a novel measure of news-driven inflation expectation shocks. Consistent with the model, following a positive inflation expectation surprise, both realized inflation and unemployment rise. We find that a one percentage point shock to inflation expectations boosts inflation by roughly 0.1% and unemployment by 0.2% over the next 2 years.

Keywords: inflation expectations, information, supply shocks

JEL Classification: D83, D84, E31, E32

*Kamdar Email: rkamdar@iu.edu. Ray Email: w.d.ray@lse.ac.uk. We thank seminar participants at the St. Louis Fed Applied Time Series Workshop and the San Francisco Fed for excellent comments.

1 Introduction

Nearly all economic decisions are based on agents' perceptions about the current economy and expectations about future economic outcomes. Thus, it is important to understand how agents form their beliefs. The workhorse approach to modeling these beliefs has been full-information rational expectations (FIRE), which posits that agents not only understand the data-generating process but also know all relevant information, past and present. However, survey data of expectations are often at odds with the implications of FIRE.

This paper focuses on one important bias of consumer expectations: "supply-side" beliefs. A robust feature in surveys is that households consistently seem to overweight supply factors as drivers of the economy and report supply-side reasoning when asked why they hold their beliefs. For example, households often report "greed" and "big business and corporate profits" as drivers of inflation (e.g. see Shiller 1996). Supply-side reasoning is also evident in the positive correlation between households' inflation and unemployment expectations (e.g. see Kamdar 2019). Not only do households use supply-side reasoning, but they seem to do so more than experts or what rational forecasts would suggest. Relative to experts, households consistently use supply-side reasoning more and demand-side reasoning less (e.g. see Andre et al., 2023, 2022).

In this paper, we investigate the causes and consequences of supply-side beliefs. Specifically, we develop a New Keynesian model where private agents overweight the likelihood that news is informative about supply shocks. That is, they receive a signal that is a combination of next period's supply and demand shock. They interpret the information as being more informative about supply than it actually is. We show that such over-weighting can arise endogenously when households face frictions when acquiring information; or when households face concerns regarding model misspecification. In either case, the intuition is that supply shocks are particularly damaging to household utility. Therefore, households either dedicate more attention to supply shocks (under rational inattention), or act as if supply shocks are more likely (when facing model misspecification concers).

The model yields a variety of theoretical implications. For instance, in response to news shocks which raise inflation expectations, agents will decrease their output expectations as they attribute the news to an adverse supply shock. This type of behavior is readily observed in survey data. But additionally, the model predicts that *realized* inflation should also rise and that *realized* output should fall, as agents behave as though a negative supply shock will occur in the next period.

We test this novel prediction of the model using a high-frequency approach to identify shocks to household inflation expectations. We use daily survey data to estimate how inflation expectations change in the days before and after CPI releases. We interpret these high-frequency changes in inflation expectations to be the result of the news. While our shocks have weak correlation with alternative measures of inflation expectations (such as those implied by changes in fixed income yields), we show that these shocks capture important and novel variation in household beliefs which is not accounted for in financial markets.

Using our novel shock series, we trace out the macroeconomic effects. Using a local projection, a 1 percentage point shock to inflation expectations results in a 0.1 percentage point increase in inflation in the following year. Further, unemployment also increases, with a peak response of 0.2 percentage points two years after the shock. These estimates are economically meaningful and statistically strong. Additionally, our estimates are robust to the exact construction of the shock or the sample period. We also show that "placebo" shocks constructed using arbitrary event dates do not feature the same effects; in particular, the estimated inflation and unemployment responses are smaller in magnitude and statistically insignificant (and in fact, the point estimates for the response of unemployment are negative).

The theoretical model also has implications for the literature that works on estimating the Phillips curve. Specifically, FIRE-based estimations will be biased; however, estimations using a subjective, survey-based measure of inflation expectations are unbiased.

Literature Review. There is a large literature documenting empirical deviations from FIRE: Carroll (2003) Mankiw et al. (2003) Bordalo et al. (2020) Coibion and Gorodnichenko (2012) (see Coibion et al. (2018) for a review of this literature).

General equilibrium models with deviations from FIRE: Mankiw and Reis (2007), Maćkowiak and Wiederholt (2015), Bhandari et al. (2022), Woodford

(2013)

News and noise shocks: Beaudry and Portier (2014), Barsky and Sims (2011), Chahrour and Jurado (2018)

High-frequency responses to announcements: Rast (2021) Binder et al. (2022) York (2023) Lamla and Vinogradov (2019) De Fiore et al. (2022) Nagel and Yan (2022) Candia et al. (2020)

2 Model

The setup of our model is a completely standard representative agent New Keynesian (RANK) model: households supply labor and consume differentiated goods produced by firms facing nominal rigidities when choosing prices. The aggregate dynamics of the model are summarized by the familiar (log-linearized) New Keynesian Phillips curve and aggregate Euler equations:

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

$$x_t = -\sigma^{-1}(i_t - \tilde{E}_t \pi_{t+1} - v_t) + \tilde{E}_t x_{t+1}. \tag{2}$$

We denote cost-push shocks by u_t (which we refer to as "aggregate supply shocks"); and denote aggregate discount factor shocks by v_t (which refer to as "aggregate demand shocks"). We assume that both follow AR(1) processes:

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u, \tag{3}$$

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v, \tag{4}$$

where the innovations $\varepsilon_t^u \sim N(0, \sigma_u^2)$ and $\varepsilon_t^v \sim N(0, \sigma_v^2)$, and are assumed independent.

Our point of departure from standard RANK models is how agents form beliefs and collect information. We assume that agents, both consumers and firms, observe all period t variables perfectly (as well as the history of such variables). The key departure from FIRE is the incorporation of news shocks and how the news shocks are perceived. Put another way, subjective expectations denoted by $\tilde{E}$ will not be equal to FIRE expectations denoted by $\mathbb{E}$.

Agents receive news which is informative about next period's demand and supply shocks:

$$z_t = \varepsilon^u_{t+1} + \varepsilon^v_{t+1} + \eta_t. \tag{5}$$

The signal contains noise $\eta_t \sim N(0, \sigma^2_\eta)$. In this setting, fully rational Bayesian updating would imply:

$$\mathbb{E}_t \left[\varepsilon^u_{t+1} | z_t\right] = K_u z_t \quad \text{and} \quad \mathbb{E}_t \left[\varepsilon^v_{t+1} | z_t\right] = K_v z_t, \tag{6}$$

$$K_u = \frac{\sigma^2_u}{\sigma^2_u + \sigma^2_v + \sigma^2_\eta} \quad \text{and} \quad K_v = \frac{\sigma^2_v}{\sigma^2_u + \sigma^2_v + \sigma^2_\eta}, \tag{7}$$

where K_u and K_v are the typical Kalman gain terms. However, agents misinterpret the news shock as more informative about supply relative to demand, so that:

$$\tilde{E}_t \left[\varepsilon^u_{t+1} | z_t\right] = \tilde{K}_u z_t \quad \text{and} \quad \tilde{E}_t \left[\varepsilon^v_{t+1} | z_t\right] = \tilde{K}_v z_t, \tag{8}$$

$$\tilde{K}_u > K_u > 0 \quad \text{and} \quad 0 \leq \tilde{K}_v < K_v, \tag{9}$$

where the subjective Kalman gain for supply shocks ($\tilde{K}_u$) is greater than the FIRE Kalman gain for supply shocks (K_u), and the the subjective Kalman gain for demand shocks ($\tilde{K}_v$) is smaller than the FIRE Kalman gain for demand shocks (K_v).

This bias could be the result of a variety of different informational frictions. In Appendix B.2, we show that such a signal structure arises endogenously when agents are rationally inattentive (as in Sims 2003) under most standard parametric assumptions. Alternatively, in Appendix B.3, we show that this signal structure also occurs when agents are "robust optimizers" and worry about model misspecification concerns regarding the volatility of supply and demand shocks. In both cases, the intuition is similar: when supply shocks are more costly from a welfare perspective, agents learn more about such shocks (under rational inattention) or act as if such shocks are more likely (under robustness). Of course, this signal structure could also simply arise due to behavioral bias:

$$\tilde{z}_t = \tilde{\alpha}_u \varepsilon^u_{t+1} + \tilde{\alpha}_v \varepsilon^v_{t+1} + \tilde{\sigma}_\eta \eta_t, \tag{10}$$

where $\tilde{\alpha}_u > \alpha_u$ and $\tilde{\alpha}_v < \alpha_v$.

2.1 Analytical Results

In order to deliver the simplest intuitive results, we focus on the simplest version of the model: $\tilde{\alpha}_u = 1, \tilde{\alpha}_v = 0, \tilde{\sigma}_\eta = 0$, and $\alpha_u = \alpha_v = 0$ (so that z is a pure noise shock). Further, assume that supply and demand shocks are iid; that is, $\rho_u = \rho_v = 0$. Agents form expectations:

$$
\begin{aligned}
E_t[u_{t+1}|\tilde{z}_t] &= \rho_u u_t + \tilde{z}_t = \tilde{z}_t, \\
E_t[v_{t+1}|\tilde{z}_t] &= \rho_v v_t = 0.
\end{aligned}
$$

This contrasts with FIRE-based beliefs, which are given by:

$$
\begin{aligned}
\mathbb{E}_t[u_{t+1}|z_t] &= \rho_u u_t = 0, \\
\mathbb{E}_t[v_{t+1}|z_t] &= \rho_v v_t = 0.
\end{aligned}
$$

Assuming a simple Taylor rule where the monetary authority responds to inflation $i_t = \phi_\pi \pi_t$, aggregate dynamics are given by

$$
\begin{aligned}
\pi_t &= \beta E_t \pi_{t+1} + \kappa x_t + u_t \\
x_t &= -\sigma^{-1}(\phi_\pi \pi_t - E_t \pi_{t+1} - v_t) + E_t x_{t+1} \\
E_t u_{t+1} = \tilde{z}_t, &\quad E_t v_{t+1} = 0 \\
E_t u_{t+k} = 0, &\quad E_t v_{t+k} = 0 \quad k > 1
\end{aligned}
$$

Assuming determinacy conditions are met, theunique equilibrium is given by

$$
\begin{aligned}
\pi_t &= \beta \chi \tilde{z}_t + \kappa x_t + u_t, \\
x_t &= -\sigma^{-1}\phi_\pi \pi_t - \sigma^{-1}\chi(\phi_\pi - 1)\tilde{z}_t + \sigma^{-1} v_t, \\
\Longrightarrow \quad \pi_t &= \chi \left[\kappa\sigma^{-1} v_t + u_t + (\beta - \kappa\sigma^{-1}(\phi_\pi - 1))\chi \tilde{z}_t\right], \\
x_t &= \chi \sigma^{-1} \left[v_t - \phi_\pi u_t + (1 - \phi_\pi(1+\beta))\chi \tilde{z}_t\right].
\end{aligned}
$$

Note that the sign of $\frac{\partial \pi_t}{\partial \tilde{z}_t}$ is ambiguous. However, if $\kappa\sigma^{-1}(\phi_\pi - 1) < \beta$ (which will hold in realistic parameterizations; for example, when the Phillips curve is not too steep), then we have:

$$\frac{\partial \pi_t}{\partial v_t} > 0, \frac{\partial \pi_t}{\partial u_t} > 0, \frac{\partial \pi_t}{\partial \tilde{z}_t} > 0$$
$$\frac{\partial x_t}{\partial v_t} > 0, \frac{\partial x_t}{\partial u_t} < 0, \frac{\partial x_t}{\partial \tilde{z}_t} < 0$$

Hence v_t is indeed a typical demand shock, u_t is a typical supply shock, and the response to $\tilde{z}_t$ is similar to a supply shock.

2.1.1 Implications for Phillips Curve Estimation

In addition, our model has important implications for empirical estimates of the Phillips Curve.

First, consider the "expectations-augmented Phillips Curve" OLS regression:

$$\pi_t = \hat{\beta} E_t \pi_{t+1} + \hat{\kappa} x_t + \epsilon_{t+1}.$$

The OLS estimates converge to

$$\mathbf{b}^{EXP} = \begin{bmatrix} Var E_t \pi_{t+1} & Cov(E_t \pi_{t+1}, x_t) \\ Cov(E_t \pi_{t+1}, x_t) & Var x_t \end{bmatrix}^{-1} \begin{bmatrix} Cov(E_t \pi_{t+1}, \pi_t) \\ Cov(\pi_t, x_t) \end{bmatrix} \neq \begin{bmatrix} \beta \\ \kappa \end{bmatrix}.$$

That is, this regression is biased. However, suppose that supply shock volatility $\sigma_u \to 0$. Then we find

$$\mathbf{b}^{EXP} \to \begin{bmatrix} \chi(\beta + 2\beta\pi_\pi \kappa\sigma^{-1} - \phi_\pi(\kappa\sigma^{-1})^2) \\ \kappa \end{bmatrix}.$$

Thus, the expectations-augmented Phillips Curve converges to the actual slope when supply shocks are small. However, the estimate of the discount factor is biased.

Now consider the "FIRE Phillips Curve" OLS regression:

$$\pi_t = \hat{\beta}\pi_{t+1} + \hat{\kappa} x_t + \epsilon_{t+1}.$$

The OLS estimates converge to

$$\mathbf{b}^{FIRE} = \begin{bmatrix} Var\pi_{t+1} & Cov(\pi_{t+1}, x_t) \\ Cov(\pi_{t+1}, x_t) & Varx_t \end{bmatrix}^{-1} \begin{bmatrix} Cov(\pi_{t+1}, \pi_t) \\ Cov(\pi_t, x_t) \end{bmatrix} \neq \begin{bmatrix} \beta \\ \kappa \end{bmatrix}.$$

Note that in this setting with no dynamics, all covariance terms with respect to realized inflation π_{t+1} are zero. But this regression is also biased in terms of the Phillips curve slope coefficient as well. Again assuming that supply shock volatility $\sigma_u \to 0$. In this case, we still do not recover the Phillips curve slope. Letting $\sigma_v \to 0$ as well we find in the limit,

$$\mathbf{b}^{EXP} \to \begin{bmatrix} 0 \\ \frac{\beta\chi^{-1} - \kappa\sigma^{-1}(\phi_\pi - 1)}{\sigma^{-1}(1+\phi_\pi(\kappa\sigma^{-1} - \beta - 1))} \end{bmatrix}.$$

Next, consider the "subjective Phillips Curve" OLS regression:

$$E_t\pi_{t+1} = \hat{\beta}E_t\pi_{t+k+1} + \hat{\kappa}E_t x_{t+1} + \epsilon_{t+1}.$$

In this iid setting, the term $E_t\pi_{t+k+1} = 0$ for all $k > 1$. Hence, this regression is not well-defined. However, if we assume that $VarE_t\pi_{t+k+1} > 0$ (eg, due to measurement error), then the OLS estimates converge to

$$\begin{aligned} \mathbf{b}^{SUBJ} &= \begin{bmatrix} VarE_t\pi_{t+k+1} & Cov(E_t\pi_{t+k+1}, E_t x_{t+1}) \\ Cov(E_t\pi_{t+k+1}, E_t x_{t+1}) & VarE_t x_{t+1} \end{bmatrix}^{-1} \begin{bmatrix} Cov(E_t\pi_{t+k+1}, E_t\pi_{t+1}) \\ Cov(E_t x_{t+1}, E_t\pi_{t+1}) \end{bmatrix} \\ &= \begin{bmatrix} 0 \\ -\frac{1}{\phi_\pi\sigma^{-1}} \end{bmatrix}. \end{aligned}$$

Hence, the estimated "subjective PC" slope is negative.

3 Empirical Results

As a result of information or behavioral frictions, agents in our model overweight the likelihood that news shocks are informative of supply factors rather than demand shocks. This model provides a variety of testable implications. For instance, in response to news that increase inflation expectations, output gap expectations should fall, as agents believe the news is about a negative supply. Furthermore, in response to this news, realized inflation should rise and the output gap should fall (as they would following a negative supply shock). Finally, we test our implications for the Phillips curve. In this section, we assess if the data support the testable implications of our model.

3.1 Data

We use the Michigan Survey of Consumers (MSC) to measure expectations. We obtain the exact date (day, month, year) each survey was taken from Inter-University Consortium for Political and Social Research (ICPSR) for January 1982 through December 2018.[1] The MSC collects approximately 500 responses each month. Typically, there are fewer responses in the last few days of the month; however, otherwise they are roughly uniformly distributed over the month. The MSC asks a numerical value for inflation expectations; however, the majority of their questions are categorical. For instance, respondents' are asked if over the next year unemployment will increase, stay the same, or decrease.

3.2 News Shock Construction

We calculate a news shock to inflation expectations as the difference between average inflation expectations in the five days after and the five days before a CPI release. Figure 1 plots the time series of the news shock from January 1982 to April 2023. Over this time, the median and mean are near zero (0.07 and 0.11, respectively) and the standard deviation is roughly one (1.02). The standard deviation does vary over the sample. For example, there are times of greater volatility in the news shock (the 1980s) and times of lower volatility (the 1990s to 2007).

[1] While the MSC began collecting their monthly surveys in 1978, the exact dates for the early years are unavailable. Note that the interview dates for January 2019 through April 2023 were obtained directly from the MSC.

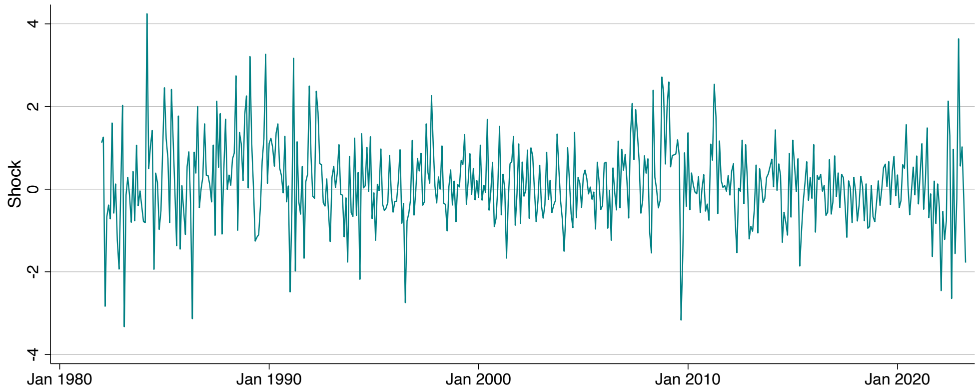

Figure 1: News Shock Time Series

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

3.3 Shock Validation

Relative to the literature that uses small, high-frequency windows to determine shocks, we need to be more concerned with endogeneity given the larger window size. To mitigate these concerns, we assess if financial or macroeconomic variables can predict the news shocks. Table 1 reports a variety of robustness checks. Our news shocks are not predictable by contemporaneous or lagged macroeconomic variables such as inflation and unemployment. News shocks are only weakly correlated with other high-frequency approaches (e.g., daily changes in Treasury yields and oil prices).

News shocks are significantly correlated with changes in other expectations of households. Specifically, a news shock which results in higher inflation expectations is significantly correlated with an increase in the fraction of households reporting they expect unemployment to rise and a decrease in households expecting unemployment to fall (Table 2 Columns 1 and 2). Furthermore, a positive inflation surprise reduces consumer sentiment, where sentiment is calculated as the first-component in a components analysis of all forward-looking expectations in the MSC (Kamdar (2019)).

Table 1: News Shock (Un)predictability

	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta y_t^{(10)}$	0.343		0.162			0.230
	(0.577)		(0.565)			(0.547)
$\Delta p_t^{(OIL)}$		0.006	0.004			0.007
		(0.020)	(0.020)			(0.020)
U_t				0.042	0.042	0.053
				(0.044)	(0.044)	(0.044)
π_t				-0.094	-0.094	-0.096
				(0.132)	(0.132)	(0.137)
U_{t-1}				-0.015	-0.015	-0.019
				(0.044)	(0.044)	(0.046)
π_{t-1}				0.112	0.112	0.118
				(0.138)	(0.138)	(0.144)
Obs.	472	437	431	479	479	431
R^2	0.001	0.000	0.001	0.004	0.004	0.006
P-val	0.553	0.777	0.940	0.632	0.632	0.630

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

Table 2: News Shocks and Expectations

	$\Delta\tilde{E}_t U^+_{t+1}$	$\Delta\tilde{E}_t U^-_{t+1}$	$\Delta\tilde{E}_t s_{t+1}$
$\Delta\tilde{E}_t \pi_{t+1}$	1.623***	-0.802**	-0.041***
	(0.378)	(0.364)	(0.008)
Obs.	490	490	490
R^2	0.039	0.013	0.062

Notes: The change in percent of households expecting unemployment to rise ($\Delta\tilde{E}_t U^+_{t+1}$) and fall ($\Delta\tilde{E}_t U^-_{t+1}$) are regressed on the estimated 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 estimated news shock ($\Delta\tilde{E}_t \pi_{t+1}$). Sentiment is calculated as the fitted first component of all forward looking variables excluding inflation.

3.4 Macroeconomic Effects

Next, we empirically test the macroeconomic effects of a news shock to inflation expectations. In particular, we assess how inflation and unemployment respond following a news shock. We conduct the following local projection:

$$y_{t+h} = \alpha_h + \beta_h \Delta \tilde{E}_t \pi_{t+1} + \boldsymbol{\gamma}_h \mathbf{X}_t + \varepsilon_{t+h}, \tag{11}$$

where y_{t+h} is the h-period-ahead realization of some macroeconomic outcome; our baseline focuses on inflation and unemployment. Our dependent variable is $\Delta \tilde{E}_t \pi_{t+1}$, the news shock constructed in Section 3.2. We also include additional controls $\mathbf{X}_t$; in our main specification, we include four lags of the inflation rate, unemployment rate, fed funds rate, oil price inflation, and the news shock are included as controls. In robustness checks, we allow for a variety of alternative choices of these controls (and find similar results).

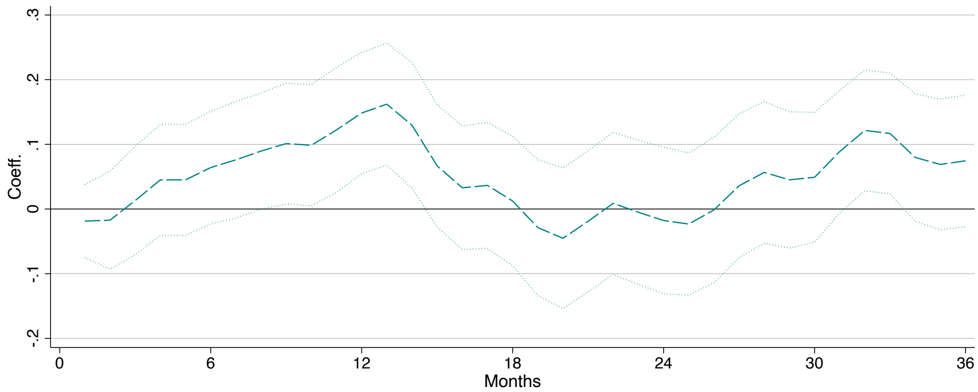

Figure 2: Local Projection: Inflation Response to a News Shock

Notes: Local projection of inflation on the estimated news shock. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

Figure 2 plots our baseline results. We find that in response to a news shock, realized inflation increases over the course of a year, reaching a peak effect of 0.15 percentage points. Then, the inflation response decreases and becomes insignificant.

In addition, following the news shock, realized unemployment steadily rises. After one year unemployment is 0.1 percentage points higher. This response continues in the following year, and peaks after two years around 0.2 percentage points higher. This increase is unemployment is large and significant. Through the lens of our model, this is precisely what we would expect. A positive news shock is perceived as coming from a negative supply shock and causes realized inflation and unemployment to both rise.

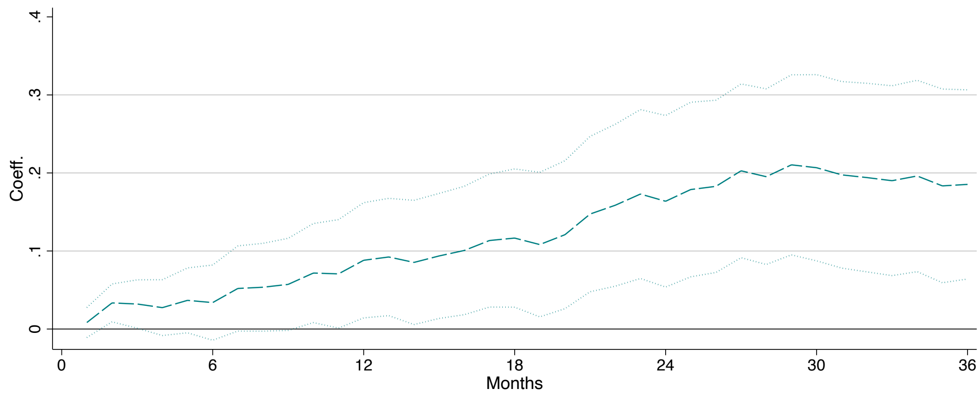

Figure 3: Local Projection: Unemployment Response to a News Shock

Notes: Local projection of unemployment on the estimated news shock. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

The qualitative results are similar in a variety of robustness checks. In the baseline, the news shock is based on 5 days before and after CPI releases; however, increasing or decreasing this window does not affect the results (Appendix A.1). Furthermore, the local projection's qualitative takeaways are similar removing all controls or limiting the controls (Appendix A.2). Finally, varying the sample also maintains the key findings.

Finally, we construct a placebo shock series: we follow the same methodology as in Section 3.2, but construct our shock 15 days after the release of the CPI (when no new macroeconomic news is systematically released). We replicate our baseline regressions, but using these placebo shocks.

The results of the placebo exercise are reported in Figures 4 and 5. As shown, our estimated local projections of inflation and unemployment using the placebo shock show that neither inflation nor unemployment respond significantly. If anything, our point estimates for both series suggest slight declines in the first few months following the placebo shock. Eventually, our point estimate for inflation becomes positive; however, point estimates for unemployment remain negative over the entire three years following the placebo shock. Our placebo estimates are

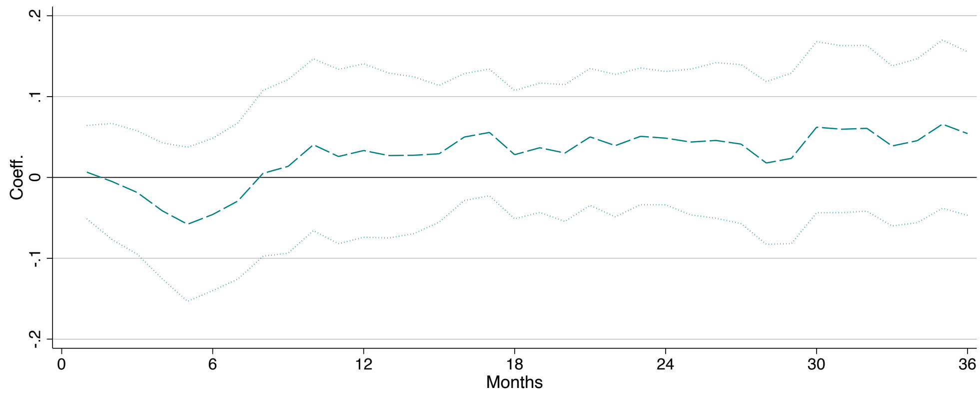

Figure 4: Local Projection: Inflation Response to a Placebo Shock

Notes: Local projection of inflation on a placebo news shock estimated around 15 days after the CPI release. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

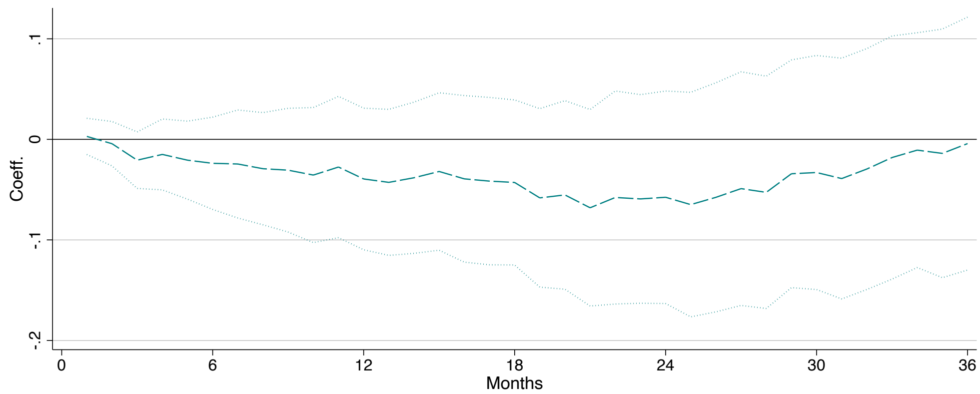

Figure 5: Local Projection: Unemployment Response to a Placebo Shock

Notes: Local projection of unemployment on a placebo news shock estimated around 15 days after the CPI release. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

economically much smaller than our baseline results, and remain insignificant over all horizons.

4 Conclusion

This paper investigates the causes and consequences of the tendency for households to overweight the importance of aggregate supply shocks relative to demand-side factors. We label this bias in beliefs "supply-side beliefs." We develop a New Keynesian model where agents receive news about future supply and demand shocks. Because supply shocks are more costly from a utility perspective, agents endogenously downweight the likelihood that news is informative about demand shocks. The model rationalizes a number of empirical puzzles, such as the observed positive correlation between survey-based inflation and unemployment expectations, and the instability of estimated Phillips curve slope coefficients.

We empirically test a key prediction of the model: following a news shock which raises inflation expectations, realized inflation and unemployment both increase. Using daily survey-based inflation expectations around CPI releases, we construct a novel measure of news-driven inflation expectation shocks. Consistent with the model, following an increase in our inflation expectation measure, both realized inflation and unemployment rise. These results are robust across a wide range of specifications and sample periods; quantitatively, we find that a one percentage point shock to inflation expectations boosts inflation by roughly 0.1% and unemployment by 0.2% over the next 2 years.

References

Andre, P., Haaland, I., Roth, C., Wohlfart, J., and Haaland, I. K. (2023). Narratives About the Macroeconomy. CESifo Working Paper Series 10535, CESifo.

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.

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

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

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.

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

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.

Kamdar, R. (2019). The Inattentive Consumer: Sentiment and Expectations. Manuscript.

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

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.

Nagel, S. and Yan, Z. (2022). Inflation Hedging on Main Street? Evidence from Retail TIPS Fund Flows. NBER Working Papers 30692, National Bureau of Economic Research, Inc.

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.

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

Woodford, M. (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annu. Rev. Econ.*, 5(1):303–346.

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

Appendix A Additional Tables and Figures

A.1 Local Projection Robustness: News Shock Days

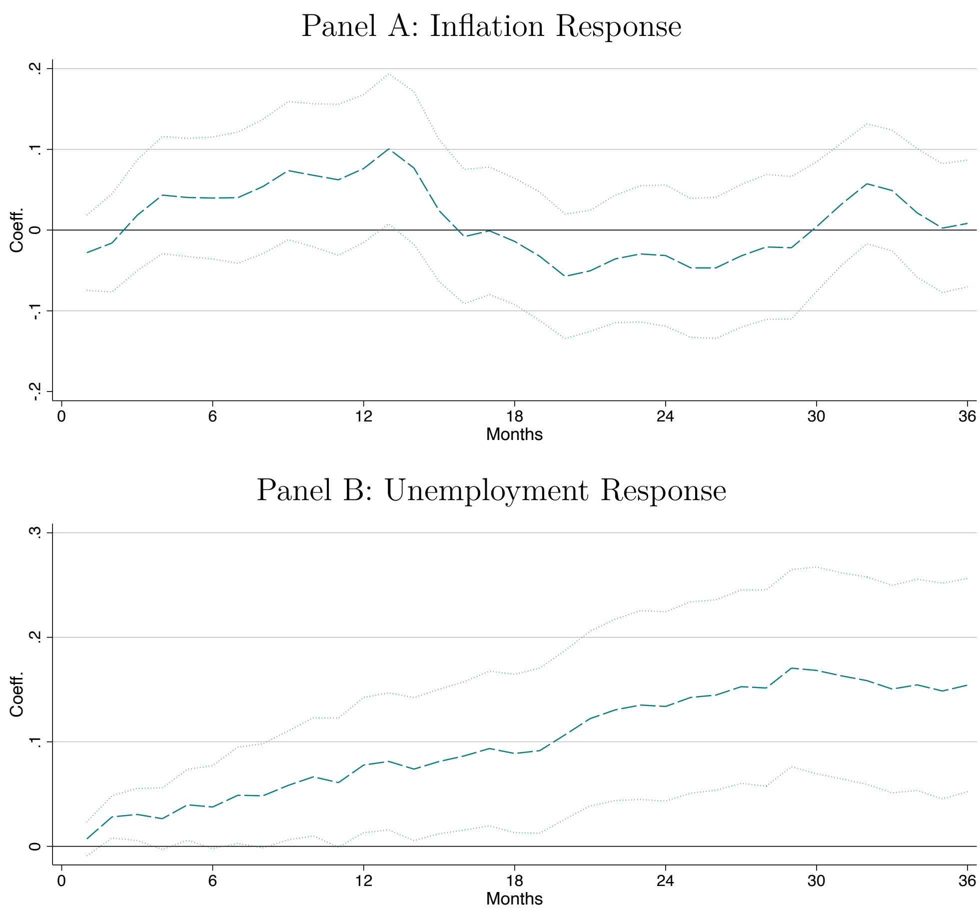

Figure A1: Local Projection: Response to a 3-Day News Shock

Notes: Local projections of inflation and unemployment on a news shock estimated based on the difference between average inflation expectations in the three days after and before CPI releases. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

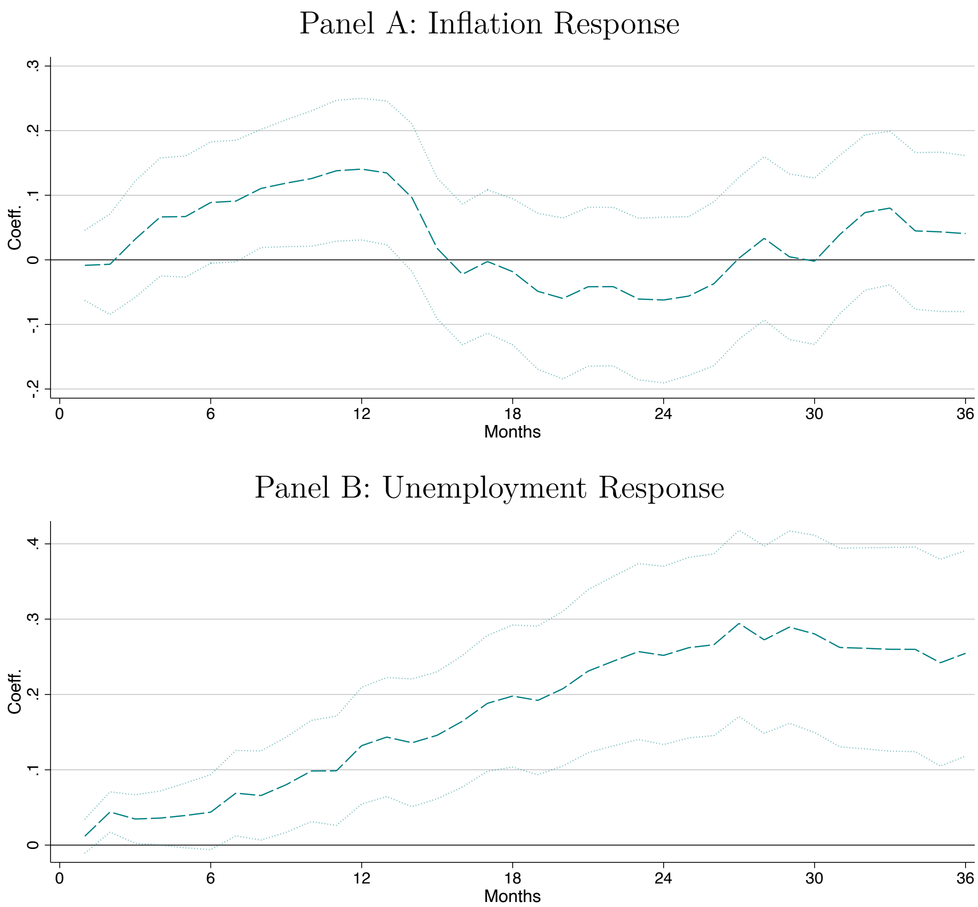

Figure A2: Local Projection: Response to a 7-Day News Shock

Notes: Local projections of inflation and unemployment on a news shock estimated based on the difference between average inflation expectations in the seven days after and before CPI releases. Four lags of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

A.2 Local Projection Robustness: Controls

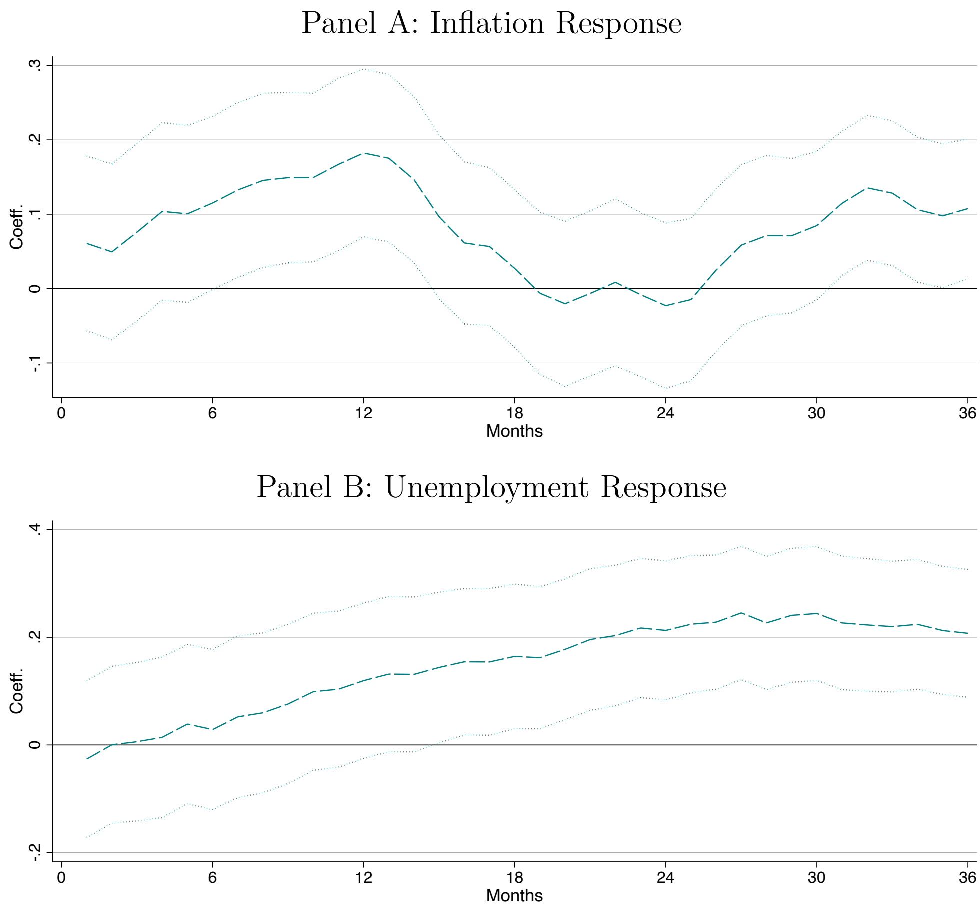

Figure A3: Local Projection: Response to a News Shock (No Controls)

Notes: Local projections of inflation and unemployment on the baseline news shock. No additional controls are included. 90% confidence intervals included.

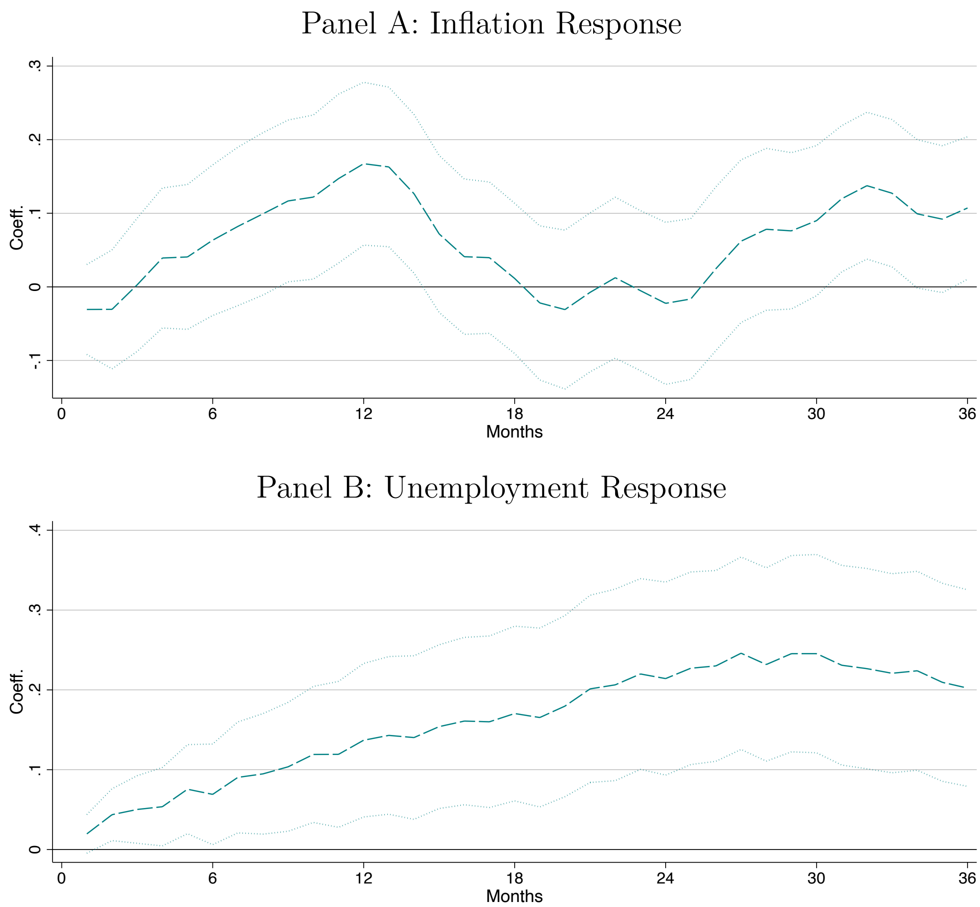

Figure A4: Local Projection: Response to a News Shock (Narrow Controls)

Notes: Local projections of inflation and unemployment on the baseline news shock. One lag of inflation, unemployment are included as controls. 90% confidence intervals included.

Panel A: Inflation Response

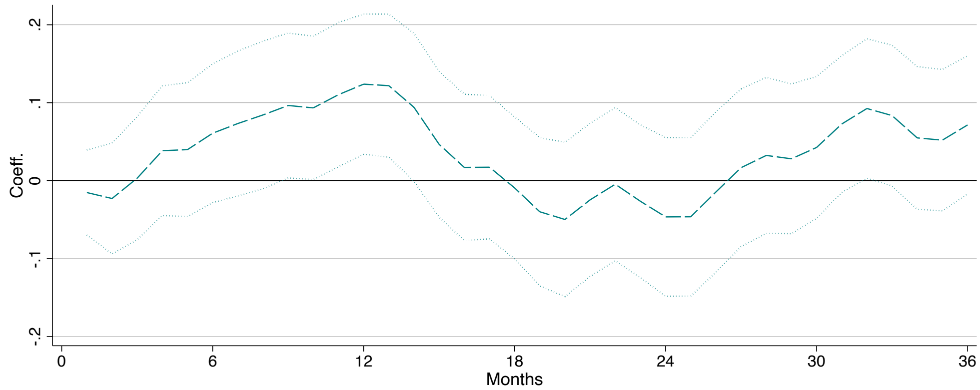

Panel B: Unemployment Response

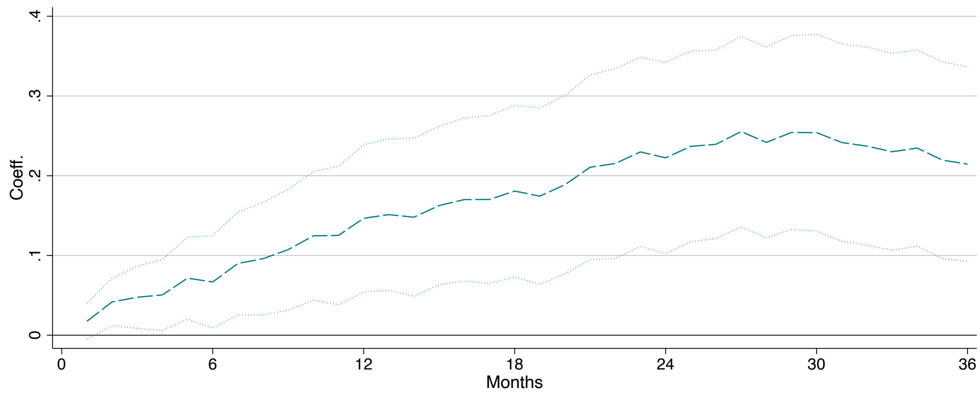

Figure A5: Local Projection: Response to a News Shock (One Lag Controls)

Notes: Local projections of inflation and unemployment on the baseline news shock. One lag of inflation, unemployment, fed funds rate, oil price inflation, and the news shock are included as controls. 90% confidence intervals included.

Appendix B Information Frictions

B.1 Model Setup

The representative household faces a standard lifetime utillity function

$$E_0 \sum_{t=0}^{\infty} \beta^t u\left(C_t, N_t\right),$$

and per-period budget constraints are given by

$$C_t + Q_t B_t = B_{t-1} + W_t N_t^S + T_t.$$

Assuming the usual functional form of per-period utility

$$u\left(C_t, N_t\right) = \frac{\left(C_t\right)^{1-\varsigma} - 1}{1-\varsigma} - \frac{\left(N_t\right)^{1+\varphi}}{1+\varphi},$$

the log-linearized aggregate dynamics are given by

$$E_t \Delta c_{t+1} = \varsigma^{-1}\left(i_t - E_t \pi_{t+1} - v_t\right),$$
$$\pi_t = \kappa c_t + \beta E_t \pi_{t+1} + u_t.$$

We assume that all information dated at time t or earlier is observable. Future shocks $\varepsilon_{v,t+1}, \varepsilon_{u,t+1}$ can be partially observed; shocks dated $t+k$ for $k > 1$ are unobservable. Thus, we derive the utility loss for the representative consumer realtive to the full-information baseline case. In this setting, full-information implies that shocks $\varepsilon_{v,t+1}, \varepsilon_{u,t+1}$ are fully observed. Losses are therefore driven by suboptimal choices c_t, n_t relative to the full-information case (which we denote by c_t^*, n_t^*. Additionally, from the intratemporal optimality conditions

$$w_t = \varsigma c_t + \varphi n_t,$$

we see that conditional on the consumption choice c_t, households make the same labor supply decision regardless of our assumption about beliefs regarding future shocks. Thus, a quadratic approximation to the loss relative to a full-information baseline is simply given by

$$-\left(c_t - c_t^*\right)^2$$

where

$$\mathbb{E}_t \Delta c_{t+1}^* = \varsigma^{-1}\left(i_t^* - \mathbb{E}\pi_{t+1}^* - v_t\right),$$
$$\pi_t^* = \kappa c_t^* + \beta \mathbb{E}\pi_{t+1}^* + u_t.$$

When demand and supply shocks are iid and the central bank chooses $i_t = \phi_\pi \pi_t$, we have that

$$c_t - c_t^* = \chi^2 \varsigma \left[\left(1 - \phi_\pi(1+\beta)\right)\left(E_t u_{t+1} - u_{t+1}\right) + \left(1 + \kappa\varsigma^{-1}(1 - \phi_\pi \beta\right)\left(E_t v_{t+1} - v_{t+1}\right)\right].$$

Further, note that under typical parameterizations, the coefficient on supply shock misperceptions $E_t u_{t+1} - u_{t+1}$ will be larger in magnitude than the coefficient on demand shock expectations $E_t v_{t+1} - v_{t+1}$. We maintain this assumption in the next sections.

B.2 Rational Inattention

Suppose that the representative household can receive information about future supply and demand shocks. The household seeks to minimize $-\left(c_t - c_t^*\right)^2$, which are driven by misperceptions regarding future shocks. When facing information costs as in Sims (2003), the optimal signal structure for the representative household is given by

$$\begin{aligned} & s_t = c_t^* + \eta_t \\ \iff & s_t = (1 - \phi_\pi(1+\beta))u_{t+1} + (1 + \kappa\varsigma^{-1}(1 - \phi_\pi \beta)v_{t+1} \end{aligned}$$

Thus, the optimal signal structure has the same form as in Section 2.

B.3 Robustness

Alternatively, suppose that agents fully understand the model structure, and receive signals as in Section 2. However, they are concerned about misspecification regarding the volatility of the structural supply and demand shocks. We assume the representative agent takes a "minmax" approach, with the following constraints:

$$\tilde{\sigma}_u^2 + \tilde{\sigma}_v^2 + \sigma_\eta^2 = \sigma_u^2 + \sigma_v^2 + \sigma_\eta^2$$

That is, agents know the total volatility of supply and demand shocks $\sigma_u^2 + \sigma_v^2$, but do not known the individual volatilities σ_u^2 and σ_v^2. They form beliefs about $\tilde{\sigma}_u^2$ and $\tilde{\sigma}_v^2$, under "minmax" approach. Thus, as shown above, because supply shocks more costly from a utility perspective, agents overweight the likelihood of supply relative to demand shocks.