Shocks to Inflation Expectations

Jonathan J. Adams*   Philip Barrett†

May 20, 2024

Abstract

The consensus among central bankers is that higher inflation expectations can drive up actual inflation. We assess this by devising a novel method for identifying shocks to inflation expectations, estimating a semi-structural VAR where an expectation shock is identified as that which causes measured forecasts to diverge from the rational expectation. Surprisingly, using data for the United States we find that a positive inflation expectation shock is contractionary and deflationary: output, inflation, and interest rates all fall. These results are inconsistent with the standard New Keynesian model, which predicts inflation and interest rate hikes.

JEL-Codes: D84, E31, E32, E52
Keywords: Inflation, Sentiments, Expectations, Monetary Policy

*University of Florida, email: adamsjonathan@ufl.edu, (corresponding author)
†International Monetary Fund, email: pbarrett@imf.org.
We are grateful to the editor Matteo Iacoviello and two anonymous referees for excellent feedback. We also thank Jorge Alvarez, Robert Barsky, Giovanni Dell'Ariccia, Thorsten Drautzburg, Klaus Hellwig, Oleg Itskhoki, Fabio Milani, Malhar Nabar, Diaa Noureldin, Nitya Pandalai-Nayar, Philippe Wingender, and seminar participants participants at the IMF and the Midwest Econometrics Group 2022 Conference for helpful comments and suggestions. The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

"The 1970s saw two periods in which there were large increases in energy and food prices, raising headline inflation for a time. [...] One likely contributing factor was that the public had come to generally expect higher inflation—one reason why we now monitor inflation expectations so carefully."

Jerome Powell, August 27, 2021[1]

Central bankers take great interest in inflation expectations. The concern – as illustrated by the quote above – seems to be that expectations of inflation tomorrow can in their own right drive inflation today, necessitating higher interest rates and lower output to stabilize prices.

But what *are* the macroeconomic effects of shifting inflation expectations? Despite all the concern, the effects are not well understood. At its core, this is due to a challenging identification problem. Inflation expectations may change exogenously, or they may change endogenously in response to other fundamental macroeconomic shocks. This paper disentangles the former from the latter. To this end, we address the following three questions.

First: what do we mean by shocks to inflation expectations? We define them as stochastic shocks that cause departures from rational expectations, which we refer to as "inflation sentiments".[2] We focus on this type of shock for two reasons. One is that we think that this is the closest formalization to what policymakers have in mind. There are many types of mistakes that agents might make when predicting the future – e.g. mis-forecasts about endogenous variables, policy rules, or underlying states – and so long as these mistakes are mean-zero, these will manifest in our framework as an inflation sentiment, which is any deviation of inflation forecasts from the rational expectation. The second reason is that these shocks are just a small departure from the most standard theory – full information rational expectations – and that theory says that such shocks are important for business cycles. More generally, a large literature shows that measured inflation expectations are broadly inconsistent with full information and rational expectations.[3] And standard theory

[1]Speech at Jackson Hole, https://www.federalreserve.gov/newsevents/speech/powell20210827a.htm

[2]To clarify terminology: we use "inflation expectations" as a general term used for more casual discussions, say of economic intuition or broad policy debates. We reserve the term "inflation sentiments" for narrower, more technical usage, meaning specifically the difference between economic agents' inflation expectations and the mathematical conditional expectation of inflation. This more general settings, this is sometimes referred to as a "belief distortion". When we say "sentiment shocks" we mean the exogenous unforecastable innovations to the inflation sentiment.

[3]Examples from the literature include: upward bias in firms' and households' inflation forecasts (Candia et al., 2021), variation in the bias by income level (Bruine de Bruin et al., 2010); for households and by industry for firms (Savignac et al., 2021); large disagreement in forecasts (Mankiw et al., 2003); large uncertainty about future inflation (Binder, 2017); poor understanding of recent inflation (Jonung, 1981);

predicts changes to inflation sentiments are inflationary. Yet, no work identifies and measures the macroeconomic impact of these shocks in a model-agnostic way. This paper aims to fill that gap.[4]

Second: how can we measure shocks to inflation expectations? We derive a novel strategy for identifying sentiment shocks and quantifying their macroeconomic effects. The challenge is that, because they affect fundamentals, sentiment shocks affect both the rational and non-rational components of expectations. To differentiate these effects, we estimate a vector autoregression (VAR) which includes both inflation and reported inflation forecasts. We then identify the inflation sentiment from the variation in the data where forecasted inflation departs from the conditional expectation of inflation. Because a VAR is a machine for estimating conditional expectations, we can use the reduced-form impulse responses as the cross-equation identifying restrictions that isolate the inflation sentiment shocks. This is a semi-structural VAR, identifying one shock: the inflation sentiment. We also show that this approach can be generalized to identify multiple sentiment shocks from a VAR with multiple forecasts and, under some extra assumptions, distinguish between them. We derive in closed form the estimator for the impulse response of these shocks, so that the method can be easily implemented.

This method is unique in its ability to isolate sentiment shocks from other forces that affect expectations. This includes *news shocks* – information about the about future fundamentals [5] – and *noise shocks*[6] – the errors in noisy signals about fundamentals.[7] To prove this point, we simulate a standard New Keynesian model with noise and news shocks, as well as shocks to productivity, and preferences, and sentiments about inflation, output, and productivity. Running our method on the simulated data validates our approach, as it recovers the sentiment shocks, even in small samples. Although in our empirical application we focus on using this method to identify inflation sentiments, the version with multiple

and underreaction to relevant news (Coibion and Gorodnichenko, 2015a). Weber et al. (2021) survey the literature and argue that "the precise mechanisms through which inflation expectations affect decisions... remain ambiguous."

[4]There is also a long, mostly theoretical literature on the impact of rational but non-fundamental inflation shocks, often termed "sunspots", including Kydland and Prescott (1977), Clarida et al. (2000), Benhabib et al. (2001b), and Benhabib et al. (2001a). Our focus on the non-rational component of inflation expectations is a complement to rather than a substitute for this analysis.

[5]Papers such as Cochrane (1994), Beaudry and Portier (2006), and Beaudry and Lucke (2010) estimate VARs that imply news plays a large role in the business cycle. Beaudry and Portier (2014) survey the evidence.

[6]For mixed evidence on the importance of noise shocks, see for example Barsky and Sims (2012), Blanchard et al. (2013), Forni et al. (2017a), Forni et al. (2017b), Chahrour and Jurado (2018), Gazzani (2020), or Chahrour and Jurado (2022).

[7]More generally, our sentiment shock is orthogonal to any other shock to which agents respond with rational expectations. This includes sunspots that produce rationally self-fulfilling equilibria (Clarida et al., 2000) and errors due to learning from small sample data (Farmer et al., 2021). The latter case we explicitly address in Section 5.2.

sentiments has broader applications, such as identifying as shocks to expected future GDP growth.

Third: what are the macroeconomic consequences of shocks to inflation expectations? We document that sentiment shocks are important drivers of business cycles, but have macroeconomic impacts that are inconsistent with the standard New Keynesian framework. Using data from the United States since the early 1980s, we identify shocks to household inflation forecasts, which we estimate drive only about a fifth of volatility in inflation and interest rates, and more than a third for production. Next, we show that the response of the macroeconomy to a positive realization of the structural shock to inflation sentiments is, somewhat surprisingly, deflationary: inflation falls and, despite monetary policy loosening, output declines. We also show that these results hold no matter whether we use household, market, or professional forecasters' expectations. This is a puzzle because the New Keynesian model has an expectations multiplier larger than one. In contrast, our estimated multiplier is negative.[8]

Our empirical work contributes to a growing empirical literature that attempts to identify sentiments in aggregate time series. How do our findings contrast from existing results? The most closely related work is by Leduc et al. (2007) and Leduc and Sill (2013) which identify shocks to expectations in a VAR with traditional causal ordering. In their main specification expectations are ordered first, so their identifying assumption is that the expectation shock is the only contemporaneous shock that affects forecasts. One drawback of this approach is that other primitive shocks could cause contemporaneous variation in expectations. For example, a productivity shock will affect expectations of future output gaps and thus inflation. And so a standard recursive ordering will misattribute this as an expectations shock. Our work disentangles these channels.

Our baseline results are at odds with the Leduc et al. (2007) approach, which finds that expectation shocks are inflationary and expansionary. The justification for that approach is that agents do not observe shocks when they occur, but only with a delay, because macroeconomic data are not publicly released immediately. We reproduce these findings and reconcile our results: if we assume that agents observe most series with a delay, but can immediately observe high-frequency series such as interest rates and commodity prices, then we still find that shocks to inflation expectations are deflationary and contractionary. However, this should not be interpreted as a criticism of Leduc et al. (2007). Rather, their results offer a resolution to the puzzle that we document: if forecasters do not pay

[8] We use the term "expectations multiplier" to refer to the response of 12-month inflation to a shock to year-ahead inflation expectations. This is distinct from the "expectational-passthrough" studied by Werning (2022) in two respects. First, our multiplier is the effect of expectations on future realized inflation over the same horizon, while Werning's passthrough is the effect on contemporaneous inflation. Second, the multiplier captures additional general equilibrium effects beyond the direct partial equilibrium effect measured by the passthrough.

attention to contemporaneous time series, then shocks to inflation expectations have the signs predicted by the New Keynesian model.

In recent complementary work, Ascari et al. (2023) study precisely the same type of shock that we explore: a stochastic distortion to agents' inflation expectations. They explore how these shocks affect a medium-scale DSGE model, which shares the qualitative predictions from our simple New Keynesian model. Then, they estimate the shock effects in a VAR identified with the sign restrictions implied by the DSGE model. Consistent with our results, they find that the shocks are contractionary. But in contrast to our model-agnostic VAR, they do not find evidence of a deflationary effect.

Researchers have applied alternative identifying assumptions to isolate sentiments from other other shocks that might affect forecasts. Two approaches are most closely related to our own. First, Levchenko and Pandalai-Nayar (2020) use a structural VAR to jointly identify TFP surprises, news about future TFP (following Barsky and Sims (2011)) and shocks that affect expectations, which they label the sentiment. Their identification assumption is that the sentiment is orthogonal to the TFP and news shocks, and otherwise maximizes forecast errors. They find that the sentiment shock is expansionary and drives a majority of short-run business cycle fluctuations. A drawback of their identification is that sentiments cannot necessarily be distinguished from other forces that move expectations orthogonal to TFP and news, such as discount factor shocks. Our method allows for such identification. Second, Chahrour and Jurado (2022) use a non-causal VAR to identify changes to TFP forecasts that are orthogonal to productivity at all horizons, and show that these "expectational disturbances" explain a large share of business cycle volatility. Shocks identified this way may include sentiments, but also noise shocks to which agents respond with rational expectations. Again, our identification approach can separately identify sentiments from noise or other shocks. Finally, a broader empirical literature studies "sentiments", although typically they are not cleanly identified as deviations from rational expectations. Work in this literature mainly focuses on shocks to expectations about future TFP or GDP,[9] whereas our work explicitly identifies non-rational shocks to inflation expectations.[10]

Our work also connects to a long literature on the macroeconomic effects of inflation

[9]Other strategies abound. Some papers estimate sentiments as shocks to measured consumer confidence, including Barsky and Sims (2012) and Fève and Guay (2019); these approaches typically find little role for their identified shocks to contribute to business cycles. Clements and Galvão (2021) use GDP data revisions to isolate expectation shocks. Lagerborg et al. (2020) use mass shooting fatalities to instrument for sentiments. Papers such as Milani (2011) and Milani (2017) use the structure of DSGE models to identify sentiment shocks. More generally, Angeletos et al. (2020) show that the main shock driving business cycles is nearly orthogonal to productivity.

[10]A larger literature documents that the full information rational expectations (FIRE) hypothesis fails for inflation forecasts in general. For example, Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015a) document that consensus SPF inflation forecasts underreact to news, while Bordalo et al. (2020) estimate that individual forecasters overreact.

expectations. This has typically focused on evaluating the empirical properties of expectations in the New Keynesian Phillips Curve (NKPC). In contrast, we identify *structural* shocks to expectations and characterize their general macroeconomic effects. One strand of this literature focuses on the relative importance of past versus future expected inflation in determining prices. See, for example, Galí and Gertler (1999), Rudd and Whelan (2005), and Rudd and Whelan (2006). In another strand, Roberts (1997) and Adam and Padula (2011) show that the empirical New Keynesian Phillips Curve (NKPC) more closely matches the theoretical curve when surveyed forecasts are used rather than rational expectations. Nunes (2010) and Fuhrer (2012) include both surveyed forecasts and estimates of rational expectations into the NKPC, but come to different conclusions about whether contemporaneous inflation is more sensitive to the former or the latter.[11]

2 Motivating Model

This short section aims to articulate more formally the conventional wisdom of how inflation expectations affect the macroeconomy. This allows us to be more precise about how inflation sentiments interact with standard theory, and establishes a baseline against which to compare our empirical analysis in the next section.

Specifically, we modify the canonical three-equation New Keynesian model[12] to include an explicit shock to inflation expectations. This section studies the simplest static version of the model, which allows us to clearly show that inflation sentiments generate macroeconomic responses which coincide with the central bankers' narrative in the introduction: inflation rises, monetary policy tightens, and output falls or rises depending on the policy response. Then, in Section 6, we consider the full dynamic version of the model with additional features.

The canonical three-equation New Keynesian model is given by:

$$\begin{aligned}
&\text{New Keynesian Phillips curve:} && \pi_t = \beta \pi_t^{e,1} + \kappa y_t \\
&\text{Euler equation:} && i_t = \mathbb{E}_t[\gamma(y_{t+1} - y_t)] + \pi_t^{e,1} \\
&\text{Taylor rule:} && i_t = \phi_y y_t + \phi_\pi \pi_t
\end{aligned}$$

where π_t is inflation, $\pi_t^{e,1}$ is inflation expectations, y_t is the output gap, i_t is the nominal

[11] Additionally, Brissimis and Magginas (2008) demonstrate that the NKPC fits the Great Moderation period well when using surveyed expectations, a conclusion confirmed for the Great Recession by Coibion and Gorodnichenko (2015b). Coibion et al. (2018) discuss further advantages of empirical expectations, with special attention paid to the NKPC. For a general survey of expectations in the NKPC, see Mavroeidis et al. (2014).

[12] See Galí (2008) for a textbook description.

interest rate, and $\mathbb{E}_t[\cdot]$ denotes the mathematical conditional expectation operator.

When expectations are rational, this framework is unable to model the situation we want to think about. For example, imagine economic agents expect inflation higher than what the policymaker will deliver. By definition, such beliefs cannot be rational – the outcome for inflation is lower than the expectation. The conditional expectation of inflation *is* the path that the policymaker will deliver. If agents think something different, then they must be departing from rationality.

We thus modify the canonical model to allow expected inflation to depart from the rational expectation:

$$\pi_t^{e,1} = \mathbb{E}_t[\pi_{t+1}] + \zeta_t \tag{1}$$

where ζ_t is exogenous and stochastic.[13] To distinguish it from $\pi_t^{e,1}$, we refer to ζ_t as the *inflation sentiment*. The sentiment may be autocorrelated, and we refer to the (mean-zero, white noise) innovations to the sentiment as the *sentiment shock*.

Expected inflation still has a rational component, so policymakers can still affect expectations today by communicating future policies – this will affect the $\mathbb{E}_t[\pi_{t+1}]$ part of expectations. Shocks to real fundamentals will affect the rational component too, as usual. And because ζ_t is mean-zero, expectations are still rational on average. But now fundamentals and policies do not *entirely* drive expectations; there is scope for inflation expectations to move in ways outside of the the control of policymakers. This captures a concern that policymakers have when they talk about inflation expectations.

The canonical model has clear predictions about the effects of such a sentiment shock: a positive sentiment shock should increase inflation and interest rates, and in most cases will cause real output to contract. To see why, consider the simplest case where the sentiment ζ_t is i.i.d. This implies that the rational expectation terms are zero, so the model becomes:

$$\pi_t = \beta\zeta_t + \kappa y_t \qquad \text{[AS]} \tag{2}$$

$$\phi_\pi \pi_t = -(\phi_y + \gamma) y_t + \zeta_t \qquad \text{[AD]} \tag{3}$$

These two equations are referred to as the New Keynesian "aggregate supply" and "aggregate demand" curves (Eggertsson and Krugman, 2012), due to their resemblance to the traditional Keynesian relationships.

Figure 1 plots how the AS and AD curves respond to a sentiment shock. The AS (Phillips) curve shifts upwards because sticky-price firms expect higher prices in the future,

[13] Sentiments are asusmed to distort all agents' inflation forecasts; Appendix B.2 considers an alternative case where a sentiment only affects one of households, firms, or the central bank. With this common structure for expectations, the motivating model most closely resembles Milani (2011), except in our setting the sentiment shock causes expectations to deviate from the rational expectation, versus an adaptive learning forecast.

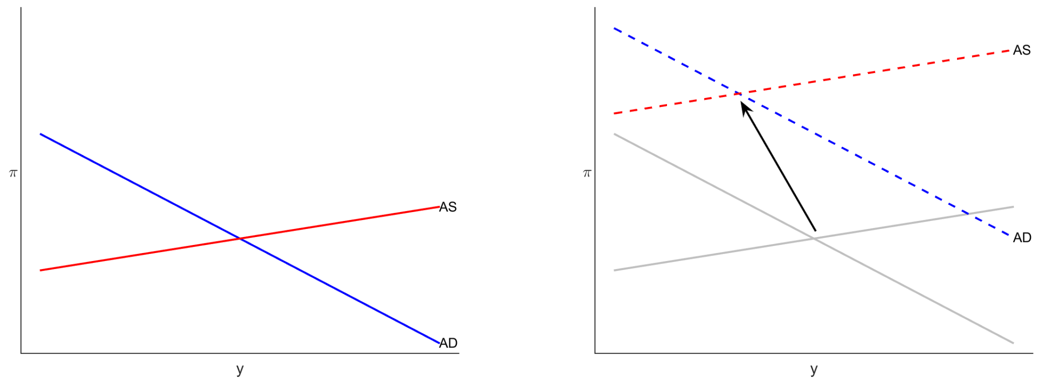

(a) Steady State

(b) Response to Sentiment Shock

Figure 1: New Keynesian Aggregate Supply and Demand after a Sentiment Shock

so they would choose higher prices today, at every level of the output gap. Absent any change in the AD curve, this shift alone would create inflation and lower output. This is because the AD curve is downward-sloping, which is due entirely to the central bank's policy response: when inflation rises, the central bank raises interest rates and pushes down on output. However, the sentiment shock also shifts the AD curve upwards because it impacts the Euler equation. By raising household expectations of future inflation, the shock lowers the effective real interest rate, increasing contemporaneous consumption and hence output.

In equilibrium, this simple framework clearly predicts higher inflation, but output is potentially ambiguous. The effect on output depends on the slope of the AD curve, which is determined by the central bank's policy response. If the central bank aggressively combats inflation, they raise interest rates enough to cause a recession. This is the case where ϕ_π is large ($\phi_\pi > 1/\beta$) so the AD curve has a small enough slope that the AS shift dominates and output declines. However, if the central bank is less aggressive (i.e. ϕ_π is small) then households' expected real interest rate declines enough to for output to expand.[14]

In this static model, the sentiment shock affects *contemporaneous* inflation. But our empirical findings are about *future* inflation. Does the intuition from a static model hold up when dynamics are introduced? Section 6 modifies the model to include persistence and dynamics. The modification loses the simple clarity of the static model, but allows for us to consider exactly the quantities that we measure in the data. And the intuition holds up: when we allow sentiments to be persistent, a positive shock increases future inflation and reduces output.[15]

In sum: the baseline New Keynesian model makes strong predictions for the effect of

[14] Appendix B.1 describes this case in greater detail.

[15] Ascari et al. (2023) confirm this finding in a richer model with firm dynamics.

a positive shock to inflation expectations. Inflation rises regardless of the policy regime. And although the central bank faces a trade-off between controlling inflation and creating a recession, nominal rates increase unambiguously. The net effect on output depends on the strength of the policy response, but is negative for standard parameter values.

3 Identifying Shocks to Inflation Expectations

Having defined what we mean by shocks to inflation expectations, we now detail a method to measure them. Our method exploits two facts. First, that inflation expectation shocks are departures from the mathematical conditional expectation of future inflation. Second, that a reduced-form VAR is a machine for estimating conditional expectations. We thus estimate a VAR which includes both inflation and a measure of inflation expectations, using the reduced-form coefficients to estimate the rational component of the response of expectations to shocks.[16] We use this to inform an identifying restriction which can recover the non-rational movement in expectations, i.e. sentiment shocks.

3.1 Basic VAR Structure

Consider the following vector autoregression (VAR). x_t denotes a $m \times 1$ vector of variables of interest (e.g. inflation) in period t, f_t^h denotes the $m \times 1$ period t forecasts of x_t over the following h periods (e.g. inflation over the next year) and y_t denotes a $n \times 1$ vector of other macroeconomic time series.[17] The VAR is given by

$$\begin{pmatrix} f_t^h \\ x_t \\ y_t \end{pmatrix} = B \begin{pmatrix} f_{t-1}^h \\ x_{t-1} \\ y_{t-1} \end{pmatrix} + A\varepsilon_t \tag{4}$$

where $\varepsilon_t \sim N(0, I)$ is an i.i.d. $(n+2m)\times 1$ vector of shocks. We consider forecasts h periods ahead as this is what appears in the data, typically reporting some agents' expected inflation over the next year.

The structural shocks are related to the VAR's reduced form "innovations" u_t by

$$u_t = A\varepsilon_t$$

[16]Our approach is related to Doh and Smith (2022), who also exploit the relationship between VAR forecasts and empirical forecasts implied by rational expectations. However, they use the rational expectations assumption to discipline the VAR coefficients and improve sign restrictions, whereas we use the assumption to identify structural sentiment shocks. In the terminology of equation (4), we study the matrix A but they focus on B.

[17]We use the term "forecast" to mean an empirical time series, distinct from the more general concept of an expectation. A forecast is a real world measure of expectations.

We subdivide $\varepsilon_t = \begin{pmatrix} \varepsilon_t^S \\ \varepsilon_t^F \end{pmatrix}$ into $n+m$ "fundamental" shocks ε_t^F that determine inflation, output, etc. consistent with rational expectations, and m "sentiment" shocks ε_t^S that independently affect expectations.

Standard VAR estimation identifies B and the autocovariance matrix of forecast errors Σ, which satisfies

$$\Sigma = AA'$$

Σ is symmetric, so it has up to $(n+2m)(n+2m+1)/2$ unique entries; we need at least $(n+2m)(n+2m-1)/2$ independent restrictions in order to identify the $(n+2m)^2$ entries in A.

3.2 Identification

In this section we describe our main innovation: an identifying assumption to isolate a single sentiment shock from fundamental shocks. Then, we describe additional assumptions that are needed when there are multiple sentiment shocks. Finally, we derive a closed form solution for the estimator.

3.2.1 Identifying Assumption with a Single Sentiment

We start with the case of a single sentiment, so $m = 1$. Then, the main identifying assumption is that *the sentiment shocks ε_t^S are the only contemporaneous shocks that causes forecasts to deviate from rational expectations.* In contrast, the fundamental shocks may only affect contemporaneous forecasts f_t^h through their endogenous effects on future variables.

The reduced-form VAR identifies the standard impulse response function (IRF) ϕ_x^k of variables in x to a reduced form innovation u_t after k periods. The rational expectation of x_{t+h} conditional on a reduced form shock u_t is given by

$$\mathbb{E}[x_{t+h}|u_t] = \phi_x^h u_t \tag{5}$$

The innovation depends on structural shocks by $u_t = A\varepsilon_t$, so the rational expectation conditional on a structural shock is

$$\mathbb{E}[x_{t+h}|\varepsilon_t] = \phi_x^h A\varepsilon_t$$

Partition the matrix A along similar dimensions as the shock vector $\varepsilon_t = \begin{pmatrix} \varepsilon_t^S \\ \varepsilon_t^F \end{pmatrix}$:

$$A = \begin{pmatrix} A_f^S & A_f^F \\ A_c^S & A_c^F \end{pmatrix}$$

where the scalar A_f^S is the contemporaneous effect of the sentiment shock on the forecast, and the row vector A_f^F is the effect of the fundamental shocks. The blocks A_c^S and A_c^F are the corresponding effects on the remaining contemporaneous variables.

The identifying assumption implies that the effects of fundamental shocks on contemporaneous forecasts is equal to their effects on the rational expectation, i.e.

$$\phi_x^h \begin{pmatrix} A_f^F \\ A_c^F \end{pmatrix} = A_f^F$$

which can be inverted to find the effect of fundamental shocks on forecasts:

$$(1 - \phi_{x,f}^h)^{-1} \phi_{x,c}^h A_c^F = A_f^F \tag{6}$$

where $\phi_{x,f}^h$ denotes the first m columns in ϕ_x^h and $\phi_{x,c}^h$ denotes the remaining columns. Equation (6) is our main identifying restriction. It says that any fundamental shock causes the forecast vector f_t^h to move by the amount that the outcome x_{t+h} will change on average. That is, the forecast responds rationally. Any variation which departs from this is loaded onto the sentiment shocks.

The implied restrictions for the matrix A are:

$$A = \begin{pmatrix} * & (1 - \phi_{x,f}^h)^{-1} \phi_{x,c}^h A_c^F \\ * & A_c^F \end{pmatrix} \tag{7}$$

where $*$ denotes unrestricted entries, of which there are $n + 2$ in the first column. This unrestricted column is the contemporaneous impact of the inflation sentiment shock.

In general, the block A_c^F is not identified. An arbitrary assumption must be made to select a A_c^F block. In our implementation, we let it be lower triangular so that the fundamental shocks have a causal ordering (as in Sims (1980) among many others). This does not affect the sentiment shock ε_t^S; every valid choice of A_c^F is just a unitary transformation of the fundamental shocks alone and yields the same sentiment shocks.

How accurate is the VAR-implied rational expectation? The VAR includes both forecasts and realizations, so any private information or news contained in the forecasts are accounted for in the VAR. By construction, the rational expectation is necessarily at least as accurate as the measured forecasts in the case where $h = 1$.[18]

3.2.2 Identifying Assumptions for Multiple Sentiments

In our baseline application (Section 4) we consider sentiment shocks about a single variable: inflation. However, it could be that there are sentiment shocks about multiple variables. In this case, our main identifying assumption (equation (6)) only identifies linear combinations of sentiment shocks. To separately identify different types of sentiment shocks, an additional identifying assumption must be made.

Define the period t *sentiment* s_t^h as the difference between the observed forecast f_t^h and the rational expectation:

$$s_t^h = f_t^h - \mathbb{E}_t[x_{t+h}]$$

The effect of structural shocks ε_t on each of these quantities is

$$\mathbb{E}[s_t^h|\varepsilon_t] = \mathbb{E}[f_t^h|\varepsilon_t] - \mathbb{E}[x_{t+h}|\varepsilon_t] = \begin{pmatrix} A_f^S & A_f^F \end{pmatrix} \varepsilon_t - \phi_x^h A \varepsilon_t \tag{8}$$

$\mathbb{E}[s_t^h|\varepsilon_t]$ represents the contribution of structural shocks to contemporaneous sentiments. Our identifying restriction in the single sentiment case assumed that only sentiment shocks can affect contemporaneous sentiments (although they also affect forecasts through the rational expectation). Substituting in this restriction (equation (6)) implies

$$\mathbb{E}[s_t^h|\varepsilon_t] = \Lambda_S \varepsilon_t^S$$

where Λ_S denotes the matrix encoding these effects, given by:

$$\Lambda_S \equiv A_f^S - \phi_x^h \begin{pmatrix} A_f^S \\ A_c^S \end{pmatrix} \tag{9}$$

Let $V_S = Var(\mathbb{E}[s_t^h|\varepsilon_t])$ denote the variance of these sentiment innovations. V_S is identifiable from the reduced-form VAR, but Λ_S is not in general. The matrices are related by

$$\Lambda_S \Lambda_S' = V_S \tag{10}$$

so if the sentiment shock ε_t^S is a scalar, then Λ_S is identified up to sign from V_S. However, if there are multiple sentiment shocks, then an additional identifying assumption is needed.

[18] The rational expectation also makes the strong assumption that the innovation u_t is observable at time t; we relax this assumption in Appendix D.

To identify the model, a practitioner must assume that these matrices are related by

$$\Lambda_S = \xi(V_S) \tag{11}$$

for some known function ξ satisfying equation (10). This says that further assumptions are needed to distinguish between the different sentiment shocks. One such is a causal ordering within the sentiments, analogous to arguments for identification of fundamental shocks via timing assumptions, which implies that Λ_S is lower triangular. In Section 5.4 we apply this sentiment shock ordering to the data, and in Section 6.3 we show that this method recovers the inflation sentiment when applied to data generated by a New Keynesian model with multiple sentiments.

3.2.3 Identification Theorem and Estimator

Theorem 1 *If the VAR satisfies the identifying assumptions* (6) *and* (11), *then the variance matrix* Σ *and impulse response function* ϕ_x^h *uniquely determine the sentiment block, whose components are given by*

$$A_c^S = \xi(V_S)^{-1}(I - \phi_{x,f}^h)\left(\Sigma_{12} - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h\Sigma_{22}\right)$$

$$A_f^S = (I - \phi_{x,f}^h)^{-1}\xi(V_S) + (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^S$$

while the fundamental block $\begin{pmatrix} A_f^F \\ A_c^F \end{pmatrix}$ *is determined up to unitary transformation of the columns.*

Proof: Appendix A

The proof is constructive. Given any assumption to map $A_c^F(A_c^F)'$ to A_c^F (e.g. Cholesky decomposition), the proof gives analytical expressions for the estimator of A implied by the estimates from the reduced form VAR.

Finally, to conclude our discussion of identification, what happens if our structural assumptions fail to hold? Then the shock identified by the VAR has a less structural but still useful interpretation: the identified sentiment shocks are precisely the dimensions of the VAR innovations that cause forecasts to depart from the VAR-implied rational expectations.

4 Baseline Results

In this section estimate the vector autoregression articulated in the preceding section using US data, outlining our baseline results.

4.1 Data

Our baseline specification includes six variables. Five of these are utterly standard, following the choices by Coibion (2012), who selects a monthly analog to the standard set by Christiano et al. (1999): inflation is the log change in the CPI, a commodity price index is included from the PPI, the industrial production index and unemployment rate measure economic activity, and the nominal interest rate is the Federal Funds Rate (FFR). Given that our results may depend on the specification of our empirical approach, it is important that we convince the reader that our statistical model is not somehow misspecified. We tackle this issue in details in Section 5.3 where we run an extensive model selection exercise. But for now we offer convention as a defence. By sticking so closely to the most commonly-used set of variables, we leave ourselves no room to manipulate our results through the choice of specification.

The one novel variable in our VAR is expectations, for which we use the median 12-month-ahead inflation forecast from the Michigan Survey of Consumers which measures expectations of households.

We also consider three alternative measures of inflation forecasts. The expectations of the median household may not correspond to those of firms setting prices, or policymakers setting interest rates, so other measurements may yield interesting insights. To measure the expectations of the market, we use the Cleveland Fed's 12-month-ahead inflation forecasts, which are published monthly. To get a sense of the views of more informed economic observers, and to comport with the standard in the literature,[19] we also run our VAR with consensus forecasts from the Survey of Professional Forecasters (SPF). This data is quarterly, so we use the direct quarterly analogues of the other series, with one exception – substituting quarterly GDP for industrial production (as too is standard). Finally, we use the Federal Reserve's Greenbook inflation forecasts as a measure of policymakers' expectations. These are released at frequencies lower than monthly but higher than quarterly. So we take the last available observation in each quarter, on the grounds that it is (as close as possible to) fully-informed by all the data available that quarter.

For use in the VAR we deseasonalize and detrend the data by removing common monthly (or, for quarterly data, quarterly) components and a linear time trend.[20] Together, our baseline sample runs from 1982:M1 - 2022:M6.[21]

[19] e.g. Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015a), and Lagerborg et al. (2020) among many others. Coibion et al. (2018) include further examples in their survey of the literature.

[20] To be consistent, we apply the same treatment to all variables in the VAR. As an alternative, we also adjust inflation forecasts using the true inflation trend and seasonal components in Appendix G.

[21] Appendix H plots both the unadjusted and transformed time series.

4.2 Impulse Responses

Figure 2 shows the impulse responses to a one-standard-deviation structural inflation sentiment shock for our baseline monthly model. We use an Akaike information criterion to choose lag length, which selects a three-lag specification. On impact, forecasted inflation rises by around 19 basis points (top middle panel of Figure 2), while actual inflation over the next year falls on average by 12 basis points (top right panel). Throughout the paper, we focus on year-ahead inflation to assess the inflationary effect of the shock, because it matches the reported forecast horizon and is robustly estimated across specifications. When we describe a shock as reducing inflation, this is the measure to which we refer.

The average effect on realized inflation is also the rational expectation, so the inflation sentiment IRF is the difference between the forecast IRF and the one-year-ahead inflation IRF.[22] Because the shock is deflationary, the sentiment rises by even more than the forecast. Characterizing the size and dynamics of this sentiment is a valuable result in its own right. We find such shocks to be moderately large – with a standard deviation of around 30 basis points – and of limited persistence – decaying almost to zero within a year. The one-year-ahead inflation response is the average of the monthly annualized inflation responses (upper left panel) over periods 1 to 12. Monthly inflation falls in the year following the shock, although the instantaneous decline is not statistically different from zero. Moreover, the sign of the instantaneous effect is not robust to the specification changes presented in the following sections; this contrasts with our main measure, the 12-month-ahead inflation, which is reliably estimated to decline.

How can the rational expectation of future inflation go down if we identify our shock as affecting the non-rational component of expectations? Here, having a conceptual framework in mind, such as the model outlined earlier, is useful. If an exogenous shock to inflation expectations has macroeconomic effects, one should expect it to affect future inflation. And if it affects future inflation, the rational component of expectations must respond.

More generally, the dynamic effects of the shock to the inflation sentiment are at odds with the predictions of the standard New Keynesian model. Future realized inflation falls on impact, and remains low for several months further. Although the decline in real activity (Figure 2, center panel) is consistent with the canonical framework, the response of monetary policy is most certainly not. Interest rates decline by around 14 basis points over the following months (bottom middle panel). While this is relatively small the effect is very persistent. This persistence is matched and even exceeded by that of real activity, which

[22]In the notation of equation (1) we are performing the decomposition $\pi_t^{e,12} = \mathbb{E}_t \sum_{j=1}^{12} \pi_{t+j} + \zeta_t^{12}$, where ζ_t^{12} is the sentiment component for one-year-ahead inflation and π_{t+j} are monthly (non-annualized) inflation rates. In Section 6, we clarify how to extend the conceptual framework of the New Keynesian model to account for the difference between a one-period and and h-period sentiment.

remains well below its starting level for several years, while unemployment remains elevated (middle right panel) for a similar duration.

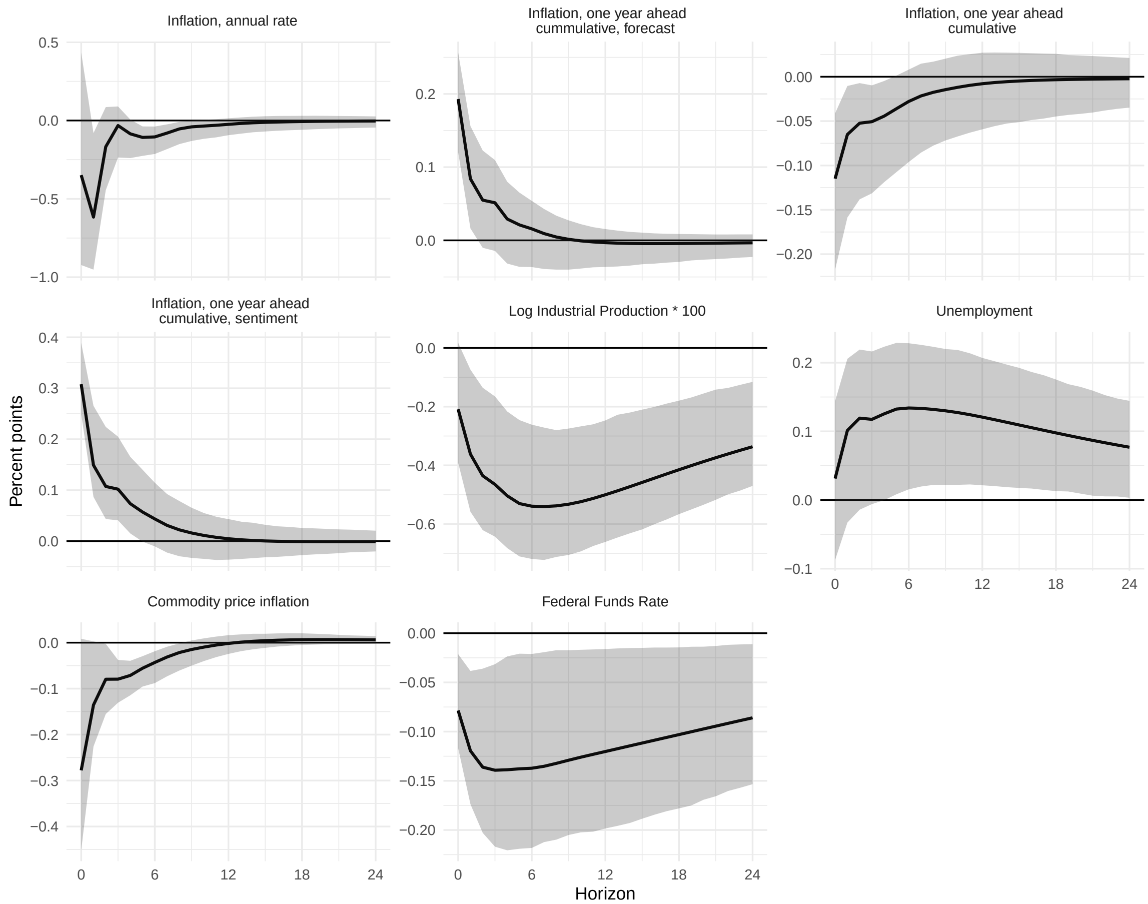

Figure 2: Impulse Responses to an Inflation Sentiment Shock

Structural impulse responses to a one standard deviation sentiment shock, baseline model. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications. The cumulative inflation over the next 12 months (i.e. the average of horizons 1 to 12 of monthly annualized inflation) is the rational expectations component of inflation expectations. The inflation sentiment is the difference between the response for the inflation forecast and the rational component. The current month's realized inflation is reported at the annualized rate. IRF units are "percent points", which is literal for time series reported as such (e.g. inflation) and 100× log values for the remaining series.

4.3 Variance Decomposition

Although the impulse responses can measure the size of the sentiment shock's effects, this does not tell us whether this shock is actually a large force in the macroeconomy. It could be that these shocks are swamped by the impact of other, fundamental shocks. To address this, we compute the variance decomposition of the structural VAR, which attributes the variation in each variable into that due to the non-fundamental and other (fundamental) shocks.[23] The first column ("Michigan") of Table 1 presents this decomposition, reporting the share of the unconditional variance[24] of several time series that are due to the non-fundamental shocks. Roughly one tenth of the variance in inflation forecasts is due to non-fundamental shocks: most of the volatility is rational, but a sizeable share is not. Non-fundamental shocks drive a similar share for interest rates, which makes sense given their determination by expectations. But sentiment shocks contribute to a greater share of the variance for realized inflation than inflation expectations. This is partially due to their additional feedback through the real economy, where sentiment shocks drive more than a quarter of volatility in real activity. Overall, the results suggest that sentiment shocks may be an important contributor to macroeconomic fluctuations.

	Michigan	Cleveland	SPF	Fed Greenbook
100 * Industrial Production	0.31	0.20	0.22	0.08
	(0.12, 0.52)	(0.04, 0.41)	(0.06, 0.45)	(0.01, 0.30)
Federal Funds Rate	0.15	0.08	0.21	0.09
	(0.02, 0.37)	(0.01, 0.26)	(0.05, 0.44)	(0.01, 0.32)
Realized inflation	0.13	0.23	0.58	0.31
	(0.03, 0.34)	(0.11, 0.39)	(0.37, 0.76)	(0.13, 0.53)
Year-ahead inf. exp.	0.09	0.09	0.24	0.30
	(0.02, 0.29)	(0.03, 0.23)	(0.07, 0.47)	(0.14, 0.53)

Table 1: Variance Shares Attributed to the Sentiment Shock, for each Forecast Measure

Bootstrapped 90 percent confidence interval are in parentheses. Log activity is industrial production for monthly series (Michigan and Cleveland forecasts) and real GDP for quarterly series (SPF and Fed forecasts).

4.4 The Estimated Sentiment

In Figure 3 we plot the time series for the estimated sentiment, and its 12-month moving average. Sentiments are largest during the 2008 financial crisis, when household inflation expectations remained persistently high. Estimated sentiments are low in the first half of 2021; before that year's rapid inflation, households expected too little.

To better understand what our estimated inflation sentiment series captures, we compare it to some other measures of economic and inflationary sentiment.

We start with the Michigan household survey consumer sentiment index. This index, which is meant to capture a broad sense of how favorably households view the economy,

[23] See Hamilton (1994) for details.

[24] The unconditional variance decomposition is the variance that is due to shocks in the long-run. We also calculate short-run variances at different horizons in Appendix C.1. We discuss the other columns of Table 1 later in Section 5.1.

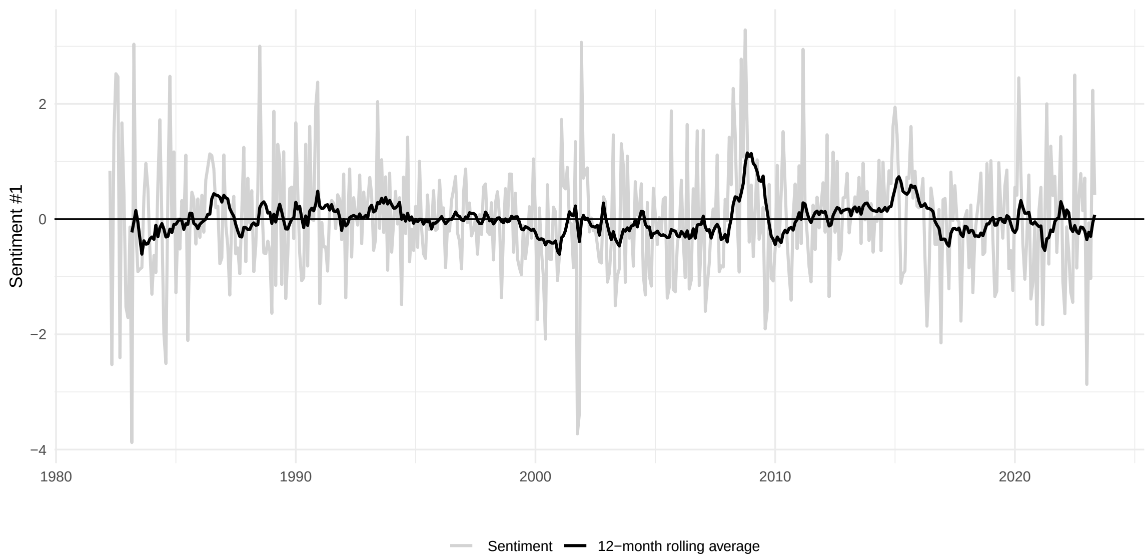

Figure 3: Inflation Sentiment Time Series

averages positive responses to five questions about both current economic conditions and the outlook.[25] In Figure 4 we plot the inflation sentiment from Figure 3 with the Michigan consumer sentiment. There is a clear and statistically significant negative relationship between the two series.[26] That is, when our inflation sentiment is generally positive, the Michigan consumer sentiment is negative.

There are many ways to interpret this relationship. In isolation, might be evidence that our inflation sentiment is mis-measured and is picking up more general sense of economic optimism or pessimism. However, we show later in Section 5.4 that including other sentiments does not change the results significantly. Taken together with that result, this correlation therefore suggests an alternative interpretation: that the extent to which households' inflation expectations are unusually positive (or or negative) informs their general economic optimism. If they expect inflation to be unusually high then this is reflected in a more negative view of economic conditions in general, and vice versa.[27]

To understand how public discussion of inflation might relate to our measured sentiment, we also compare our series to media coverage of inflation. To do this, we construct a two simple measure of press attention on inflation. Using Factiva, a press aggregation service,

[25] Paraphrasing, these questions are: are you better off than a year ago? will you be better off a year from now? will business conditions in the next year improve? will the economy perform well over the next five years? is now a good time to purchase a major household item?

[26] The correlation coefficient is -0.35 with a 95 percent confident interval of $[-0.43, -0.27]$.

[27] This interpretation is consistent with numerous studies finding that households' pessimism about inflation is connected to pessimism about the real economy. See for example Andre et al. (2022), Binder and Makridis (2022), Bernstein and Kamdar (2023), Coibion et al. (2023), or D'Acunto et al. (2023).

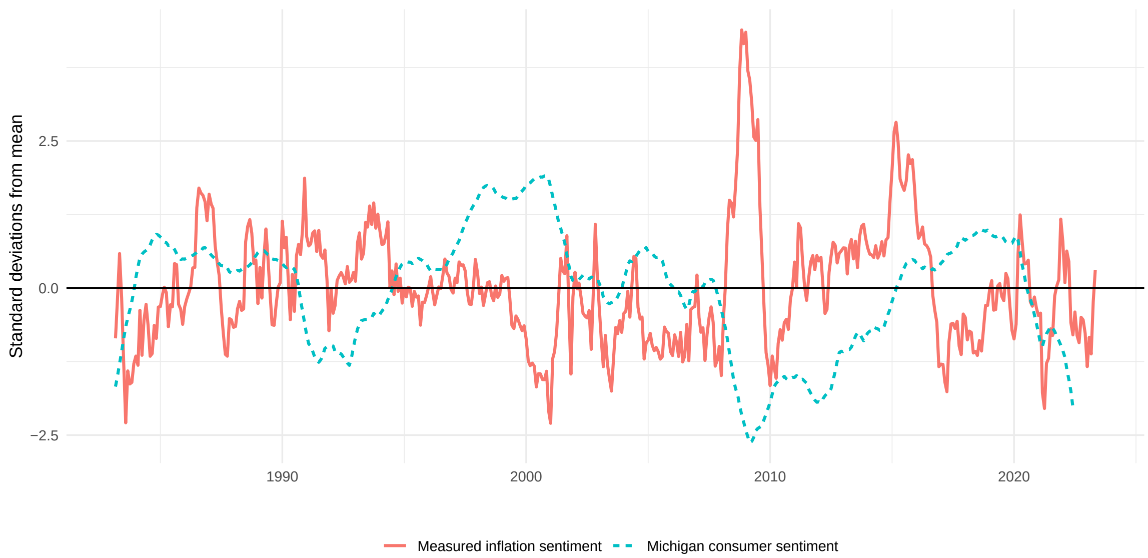

Figure 4: Estimated Inflation Sentiment and Michigan Consumer Sentiment, normalized

The "Measured inflation sentiment" line is the 12-month rolling average of our estimated inflation sentiment, show in Figure 3. The "Michigan consumer sentiment" line is the household economic sentiment measure from the Michigan Survey of Consumers, also shown as a 12-month rolling average. Both series are normalized to have zero mean and unit standard deviation.

we count the number of monthly articles containing the term "inflation" in five major US newspapers: the Wall Street Journal, the Washington Post, the New York Times, the Los Angeles Times, and USA Today. For each source we construct an index from the number of matching article counts. We focus on the results from two combinations of these, one using just the Wall Street Journal – aiming to capture business-focused press coverage – and an unweighted average across all five – aiming to reflect more general trends.

As reported in Table 2, the correlation between our inflation sentiment series and these indices is statistically indistinguishable from zero. However, it is not obvious how to interpret this. Press coverage is an unsigned magnitude, and so might plausibly reflect abnormally low expectations of inflation just as much as abnormally high ones. To test this, we compute the correlation of the absolute value of our estimated inflation sentiment with the constructed press indices. The correlation is positive and statistically significant; press coverage is higher both when inflation expectations are abnormally high and when they are abnormally low. To understand what drives this result, we also subdivide the sample into those where inflation is positive and where it is negative. Consistent with the correlation with the absolute value, the point estimates of the correlation are positive when the sentiment is positive, and negative when it is not. However, this relationship is stronger and more statistically significant for positive sentiments – suggesting that worries about abnormally high future inflation are more newsworthy than those about low inflation. Throughout the results in general are much stronger when looking at just the Wall Street Journal, consistent with the notion that the financial press provide a clearer signal of public concern about future inflation.

Series	Observations	Wall Street Journal	Five-newspaper average
Inflation Sentiment	All	0.015	0.021
		(-0.077, 0.106)	(-0.071, 0.112)
Sentiment, absolute value	All	0.149	0.021
		(0.058, 0.237)	(0.018, 0.199)
Inflation Sentiment	Sentiment≤0	-0.11	-0.055
		(-0.235, 0.019)	(-0.182, 0.074)
Inflation Sentiment	Sentiment>0	0.187	0.16
		(0.058, 0.31)	(0.03, 0.285)

Table 2: Correlation of Estimated Inflation Sentiment with Newspaper-based Indices of Inflation Coverage January 1985 - March 2023.

Table shows the correlation and 95 percent confidence intervals for the baseline estimated inflation series and two indices of inflation coverage. Indices are constructed from the normalized counts of articles including the word "inflation". Both series are 12-month rolling averages.

One concern with the preceding approaches is that the comparator series capture inflation expectations only indirectly or in combination with other factors. And so we also compare our estimated series to more formal measures of inflation sentiments estimated by Aruoba and Drechsel (2022). These series use natural language processing to measure how the Federal Reserve's views of inflation and expected inflation are associated with positive or negative language. The sign of these series are such that numbers greater than zero indicate more concern about the relevant variable. Again, there are many reasonable interpretations of how this should correspond to our series, but the most natural is that the correlation should be negative. The occurrence of a positive household sentiment seems to most likely be associated with more negative Federal Reserve communication. As Table 3 shows, the relationship is indeed negative and highly statistically significant.

FOMC Inflation	FOMC Inflation Sentiments
-0.465	-0.257
(-0.538, -0.384)	(-0.346, -0.163)

Table 3: Correlation of Estimated Inflation Sentiment with Aruoba-Dreschel Measures of Federal Reserve Inflation Concern

Table shows the correlation and 95 percent confidence intervals for our baseline estimated inflation series and the Aruoba and Drechsel (2022) Federal Reserve language-based sentiment series. Both series are 12-month rolling averages.

5 Extensions

Our baseline results are somewhat surprising: shocks to expectations were deflationary, contractionary, and reduced interest rates. Could this puzzle be due to misspecification? In this section, we estimate additional specifications to address four specific concerns.

First, in case household forecasts reported by the Michigan Survey are a poor measure of inflation expectations, we conduct our analysis with alternative forecasts. Second, our baseline estimates of the rational expectation utilize the entire time series, which contains more information than individuals have in-sample. So we repeat our estimation without this informational advantage. Third, it is crucial to our identification to consistently identify the rational expectation, and this may be impossible with the few variables in the standard VAR. Therefore we consider alternative models, using factor VARs and machine learning methods. Fourth, we check that our results hold when we apply the multi-dimensional version of our method, which admits other sentiments beyond just to inflation.

Across these tests, our findings are broadly robust. The only exception is that the immediate response of monthly inflation is not always clearly negative – something that we highlighted in our initial presentation of the results. Despite this, that the overall response is deflationary and contractionary is a robust finding; when year-ahead inflation expectations increase, prices, output, and interest rates decline over the following twelve months.

5.1 Alternative Measures of Inflation Expectations

We now repeat our main analysis with alternate measures of inflation expectations. This serves a dual purpose, both acting as a check on the broader validity of our results, but also allowing an investigation of how different types of agents' inflation expectations might have different macroeconomic consequences. In particular, we estimate three further VARs. The first uses a monthly measure of market expectations, calculated by the Federal Reserve Bank of Cleveland using bond prices, derivatives, and surveys. The second uses a quarterly measure of economists' expectations: the Survey of Professional Forecasters' (SPF) average. The third uses the central bank's own expectation, as reported in the Fed's Greenbook forecasts.

We report the results from all three exercises in Figure 5, in addition to our baseline VAR. For simplicity, we do not report here the responses for unemployment and commodity price inflation (although we do include them in the VAR) as their inclusion is principally a matter of model fit, rather than a test of economic theory.

All sentiment shocks are associated with deflation and recession, in line with our baseline findings. Although despite qualitative consistency, quantitative differences are evident. The two monthly measures also feature remarkably similar dynamics, which may be a

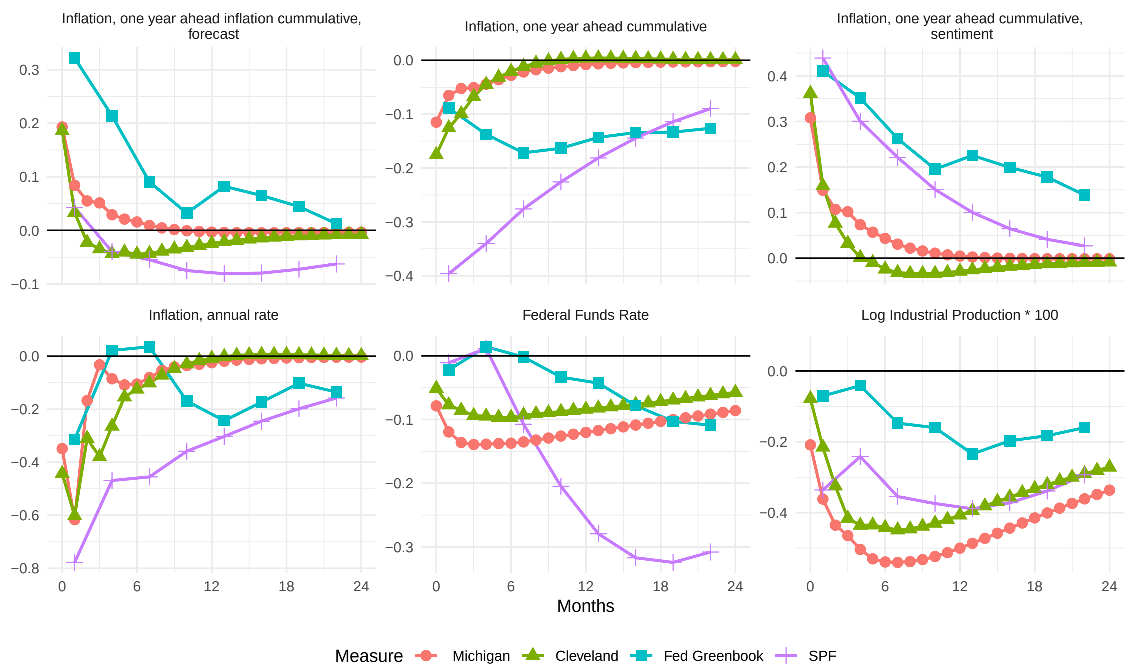

Figure 5: Impulse Responses to a Sentiment Shock for Alternative Forecast Measures

Structural impulse responses to a one standard deviation sentiment shock. For comparison across different measures, this plot combines estimates from monthly and quarterly data. For quarterly data (SPF and Fed Greenbook measures) the quarterly impulses are plotted in the middle month of each quarter, i.e. the 2nd, 5th, 8th, etc. horizon. Realized inflation is CPI inflation for monthly measures and GDP inflation for quarterly ones. Activity in industrial production for monthly estimates, and GDP for quarterly ones.

consequence of containing similar information. The results from the quarterly SPF are the most dissimilar, predicting deep deflation and persistent interest rate declines. Estimated sentiments are also especially persistent for the SPF and Fed.

The alternative expectation measures are also all responsible for sizeable shares of macroeconomic volatility, but with different effects on different time series. Table 1 reports our estimated the variance decomposition for each measure. The contribution of sentiment shocks to macroeconomic volatility vary by forecast measure. The baseline's Michigan survey-derived shocks are unique in contributing the most to real activity. The Cleveland Fed's market measure of expectations is typically responsible for similar shares as in the baseline, but drives somewhat more inflation and less real activity. The SPF sentiment shocks have a similar profile across the different series, except scaled up. In particular, the SPF shocks are the largest contributor to realized inflation, implying that nonfundamental shocks are the most important driver of inflation volatility.[28] Finally, the

[28] The SPF shock's contribution to inflation volatility is much larger than for the Michigan shock. How can this be, if professional forecasters are more accurate than households? While the SPF sentiment shocks

Fed's measure is the most dissimilar, driving only a small share of real activity, but sizeable shares of expectated and realized inflation.[29]

In summary, the puzzle we document is robust to how inflation expectations are measured. Alternate measures of inflation expectations produce very similar sentiment shocks, and macroeconomic responses which are broadly in line with our baseline results. No specification reproduces the predictions of the canonical New Keynesian model from Section 2.

5.2 Identification without the Benefit of Hindsight

One assumption of our baseline approach is that agents form rational expectations using the true statistical model: the rational expectation that we use to identify the sentiment is the conditional forecast (equation (5)) estimated from the entire sample. However, even rational agents might form different expectations if older data do not imply the same dynamics as we estimate. Our baseline approach is a fine approximation if learning is fast, but if learning is slow, a rational agent will form expectations differently in 1992 than in 2022. Some research suggests that slow learning may be responsible for many of the observed puzzles in expectations data (Farmer et al., 2021). In this section, we test if our results are robust to this concern.

To account for learning, we independently run our entire estimation at every time period τ, beginning in 1992 when 10 years of data are available. This way, we identify the sentiment in every period, using the rational expectation formed from data available at that date. The structure of the VAR (4) becomes

$$\begin{pmatrix} f_t^h \\ \pi_t \\ y_t \end{pmatrix} = B_\tau \begin{pmatrix} f_{t-1}^h \\ \pi_{t-1} \\ y_{t-1} \end{pmatrix} + A_\tau \varepsilon_t \qquad t \leq \tau \tag{12}$$

At every period τ we calculate the reduced form shock u_τ from the τ-period VAR (12), identify the coefficient matrix A_τ, and calculate the shock vector by $\varepsilon_\tau = A_\tau^{-1} u_\tau$. Then, we calculate the average responses to our identified shocks ε_τ to estimate a *learning-robust* result.

The learning-robust estimates are qualitatively similar to our baseline results. Figure

are small, they affect the real economy through different channels; Figure 5 demonstrates that the SPF shocks have more persistent propagation effects than the Michigan shocks. As a result, household confusion captured in the Michigan shocks do not matter as much for aggregate inflation as errors in professionals' forecasts.

[29] Our large VAR-based estimates of sentiments' contribution to volatility supports the DSGE model-based conclusions by Milani (2017) which finds a large share of business cycle volatility is due to sentiments in general, including modest shares for inflation sentiment shocks.

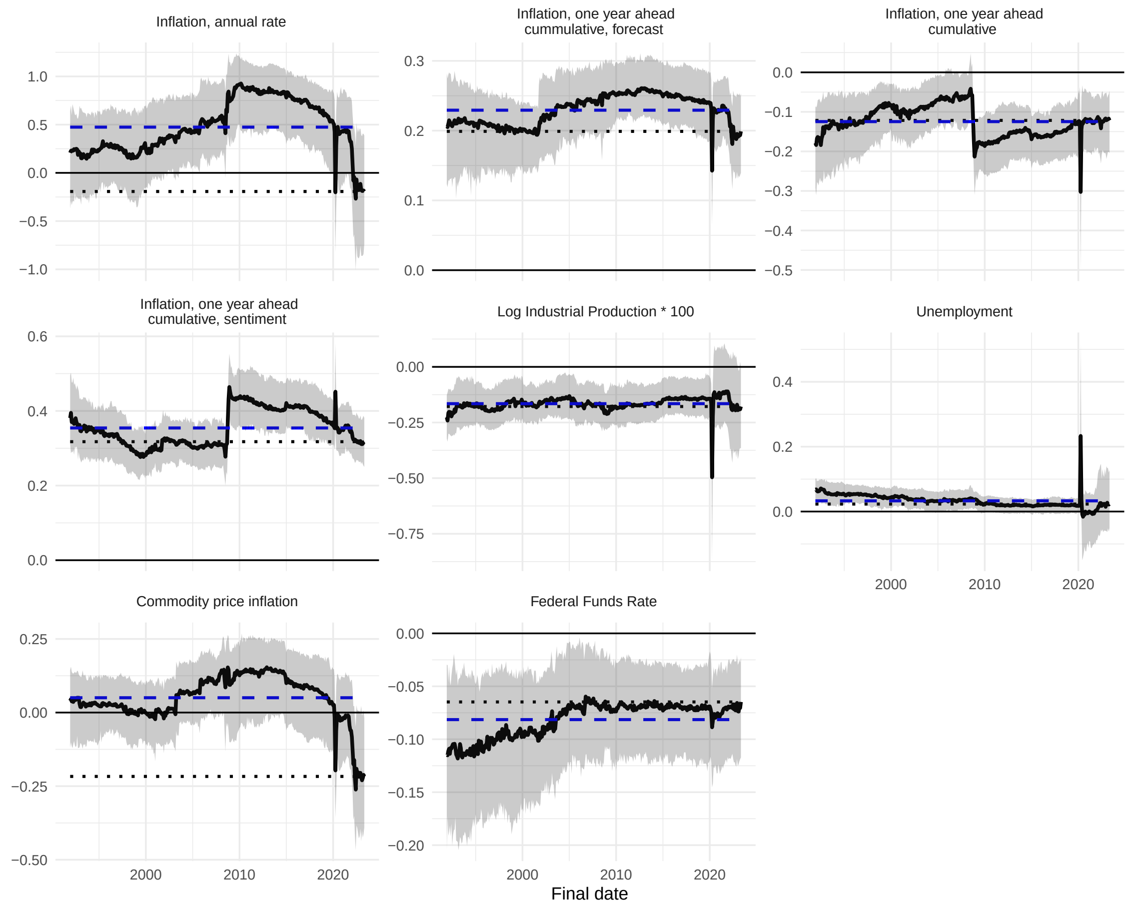

Figure 6: Time-Varying Estimation of Sentiment Shock Effects

Each subplot is an entry in the A_τ coefficient matrix, estimated using data up to the specified year. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications. The black dotted lines are the full-sample estimates and the blue dashed lines are the averages across the rolling estimates.

6 plots the first column of the time-varying A_τ. This is the contemporaneous effect of a sentiment shock on each time series, estimated using data available in the given period. For each series, the blue dashed line is the average value over the sample period: our learning-robust estimate of the instantaneous impact of a sentiment shock. The dotted line is the value at the end point: our baseline result which uses information from the entire sample to identify the shock at each point in time. These two lines are broadly similar with two main exceptions: the learning-robust result suggests the monthly contemporaneous inflation and commodity prices rise immediately; only with the latest data do the immediate impact elasticities turn negative, albeit not statistically different from zero.[30] However,

[30]We also checked whether the estimated IRFs are stable over the sample period. The only major time-

our preferred summary measure of inflation – the one-year-ahead realized inflation rate – remains robustly negative even when accounting for the benefit of hindsight.

Figure 6 also demonstrates that our results are robust to exclusion of the zero-lower-bound (ZLB) period. For example, estimates using data up to 2008 all have the same sign as our baseline results, except again for the immediate monthly inflation rate.

5.3 Models with Additional Time Series

A reasonable objection to our analysis is that we might be using a misspecified statistical model. If our model omits some variables needed to estimate accurate reduced form impulse responses, then our measurement of the conditional expectation of future inflation will be inaccurate. As a result, our claim to identify inflation sentiments will be invalid. In Section 4.1 we sidestepped this issue with an appeal to convention, including just the variables most commonly used in other VARs. Here, we return to this issue, conducting an extensive model selection exercise.

We allow for the potential inclusion of 26 other monthly macroeconomic times series in our VAR, shown in Table 4. These plausibly span a wide set of sources of macroeconomic variation missing from our baseline framework. This set includes a several forward-looking variables, both explicit – most notably, long-term interest rates – and implicit – exchange rates, stock market indices, and the like, which convert information about the future into prices. This forward-looking aspect addresses a specific case of the more general concern about misspecification. That is, that there might be information about future monetary policies not reflected in the baseline information set.

Since the number of VAR parameters grows with the square of the number of variables, one cannot simply include all the new variables in a bigger VAR. This would replace one problem with others, trading misspecification for over-fitting and imprecision. We thus adopt two methods which aim to include as much of the useful variation in the additional data while avoiding these problems: a factor-augmented VAR (FAVAR) and, in Appendix C.5, a machine learning approach.

The factor-augmented VAR, popularized by Bernanke et al. (2005), aims to extract the most important dimensions of variation in a set of possible covariates by transforming them into their principal components. A handful the most informative principal components are

varying features are the impact signs already discussed; otherwise persistence and propagation appear stable.

Group	Variables	Transformation
Prices	*Consumer Prince Index*, *Commodity Price Index*	Growth rate
Interest rates	*Federal Funds Rate*, 3-month, 1-, 2-, 5-, 10-, and 30-year US Treasury rates	None
Financial	USD vs. GBP, JPY, and CAD exchange rates, Real oil price, Willis 5000 Index	Log
Money & Credit	M2, Currency in Circulation, Bank credit, Chicago Fed financial conditions leverage index	Log (excl. leverage)
Real Activity	*Industrial Production*, New housing starts, Vehicle sales	Log
Labor Markets	*Unemployment rate*, Employment, Average hours	Log (excl. unemp.)
Fiscal	Real government surplus ratio	Unit variance

Table 4: Variables Available for FAVAR and Machine Learning Model Selection

Items in italics are those used in the baseline specification. All specifications also include the Michigan Survey of Consumers mean inflation expectation.

then included in the VAR. We extend the specification in equation (4) to:

$$\begin{pmatrix} f_t^h \\ \pi_t \\ y_t \\ F_t \end{pmatrix} = B \begin{pmatrix} f_{t-1}^h \\ \pi_{t-1} \\ y_{t-1} \\ F_t \end{pmatrix} + A\varepsilon_t \tag{13}$$

Where y_t includes only the federal funds rate and log industrial production, and F_t includes the first N principal components of the remaining series in Table 4. We use the first 4, 8, and 12 principal components, as they cover 50, 75, and 90 percent of the variance in the data respectively.

Figure 7 shows the impulse responses for the baseline and factor-augmented VARs. The main findings from our baseline model hold true. The identified sentiment is almost exactly the same. The overall effect of the shock is also deflationary: cumulative inflation in the 12 months after the shock falls in all versions, and by as much as 18 basis points in the largest factor-augmented cases. Output is also persistently lower and monetary policy loosens, although the timing and magnitude of these two is a little slower and smaller. Another difference from the baseline model is that for specifications with the most factors, the impact on contemporaneous inflation can be positive. This uncertainty over the sign of the short-run effect is something we flagged in our baseline results (and in Section 5.2). But even for specifications which give a positive contemporaneous inflation impact, the subsequent deflationary impact is deeper and more persistent than in the baseline, producing similar overall reductions in the price level.

The advantage of factor methods is that they are transparent and well-understood[31]. One downside is that it is not obvious what is the right number of factors to use. More factors improve the fit but also increases the likelihood of over-fitting. The temptation to emphasize the particular factor model which aligns with one's priors is strong. Machine learning methods can offer a better solution. In Appendix C.5 we use these techniques to check the robustness of our baseline specification. These show similar results to the FAVARs, with the overall impact deflationary and contractionary, albeit with smaller declines in output and interest rates, and positive initial responses for inflation but consistent overall price declines.

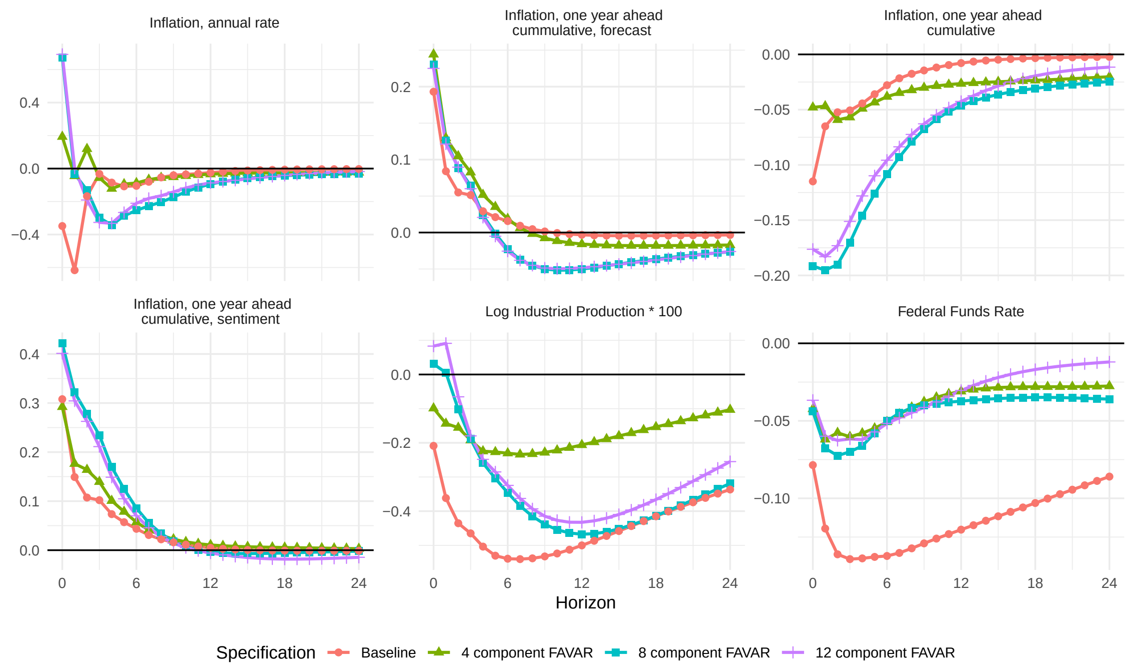

Figure 7: Factor-Augmented VAR Impulse Responses

5.4 Multiple Sentiment Dimensions

In this section, we allow for the possibility of sentiment shocks to multiple forecasts – not just inflation.

What happens if forecasters are subject to sentiment shocks about multiple quantities? Many types of sentiments might cause inflation forecasts to deviate from the rational expectation. For example, supply-side shocks can affect inflation, so sentiments about supply will look like an inflation sentiment in our baseline analysis.

[31] Stock and Watson (2005) discuss estimation and identification in FAVAR models. Stock and Watson (2011) survey the literature more broadly.

Fortunately, our identification result (Theorem 1) allows us to uniquely identify inflation sentiment shocks from other sentiments. In this case, our additional identifying assumption is that inflation sentiment shocks are the only sentiment shocks that do not cause other forecasts to deviate from their rational expectation. That is to say, inflation sentiment shocks directly distort inflation forecasts alone, while other types of sentiment shocks may distort both inflation and other forecasts.

This assumption implies that Λ_S is the Cholesky decomposition of V_S, where inflation is ordered last among both forecasts and sentiment shocks.[32] In Section 6, we validate this approach, showing that in a simulated New Keynesian model with multiple sentiments, this assumption correctly identifies the inflation sentiment shock from other sentiment shocks that also distort inflation forecasts.

We consider two versions of this approach. In the first we modify our baseline model to allow for a real activity sentiment. So that this additional sentiment can also be interpreted as a shock to household expectations – and because households tend not to make forecasts for industrial production – we replace industrial production with real personal income as the real activity measure. We can then use the Michigan household survey data to construct a year-ahead real personal income forecast.[33] This gives household expectations for both inflation and real personal income from the same source at the same horizon. The Cholesky ordering within the sentiment block produces two shocks: one which moves the inflation sentiment alone and one which moves both the inflation sentiment and the real personal income sentiment. Figure 8 presents the impulse responses for the first one, where just the inflation sentiment moves. This isolates the impact inflation sentiment alone, distinct from other expectational shocks which might cause correlated changes in believes about both inflation and real factors (note how the personal income sentiment is pinned to zero on impact). If our baseline results were substantially driven by correlated changes in beliefs about other factors, this would be mapped into the impulse responses for those other shocks, and the impulses for the inflation sentiment would change. Instead, the inflation sentiment remains contractionary. As in our baseline results, unemployment rises, and interest rates, real personal income and inflation all fall (the latter on average).

To check further that extra dimensions of sentiment shocks are not driving our results, we also extend our quarterly estimates using SPF forecasts. Unlike households, SPF participants make forecasts about a wide range of macroeconomic data. We thus include in our VAR not only forecasts for inflation but for GDP and interest rates as well. Again, we apply a Cholesky ordering to the sentiment shocks, allowing us to separate out a inflation sentiment isolated from the impact two further dimensions of shocks to beliefs. Figure 9

[32]This approach is similar to classic causal ordering (Sims, 1980), except only within the sentiment block.

[33]The survey asks about nominal personal income growth forecasts, which we can combine with the inflation forecast to back out implied real personal income one year ahead.

shows the impulse responses for this shock. If anything, the impact is even more clearly contractionary – inflation falls on impact and on average, activity declines, and unemployment rises. The only difference is that the interest rate response is not significant.

Overall, our main results appear not to be driven by other types of sentiments. When we control for sentiments to real personal income (in the baseline) and to GDP and interest rates (in the quarterly SPF version), the baseline result remains.

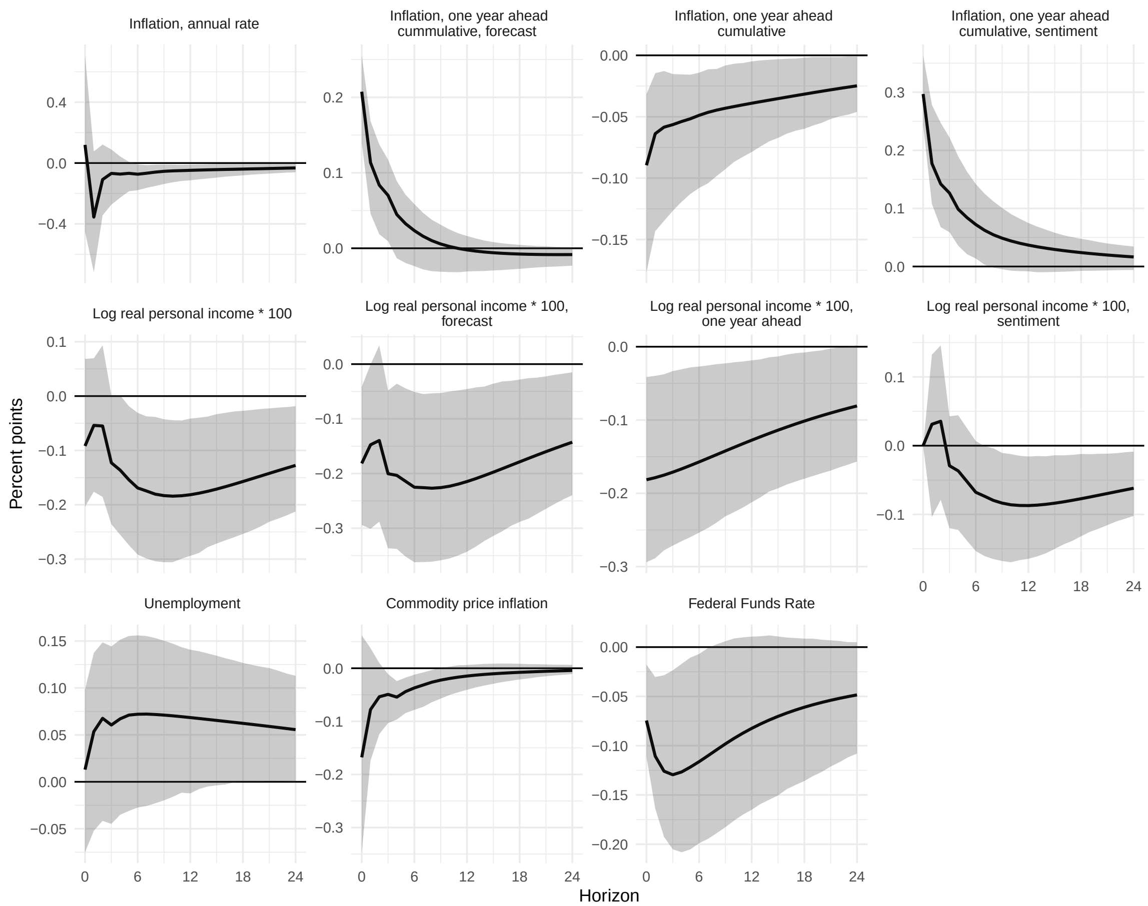

Figure 8: Impulse Responses to an Inflation Sentiment Shock: Multi-Sentiment Baseline

Structural impulse responses to a one standard deviation sentiment shock, baseline model allowing for a real personal income sentiment. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications. The cumulative inflation over the next 12 months (i.e. the average of horizons 1 to 12 of monthly annualized inflation) is the rational expectations component of inflation expectations. The inflation sentiment is the difference between the response for the inflation forecast and the rational component. The current month's realized inflation is reported at the annualized rate. IRF units are "percent points", which is literal for time series reported as such (e.g. inflation) and 100× log values for the remaining series.

5.5 Further Extensions and Robustness Checks

We perform several additional tests reported in the appendices. Across specifications, our main conclusions about inflation sentiment shocks are largely unchanged.

In Appendix C.3 we estimate impulse responses by local projections instead of relying on those implied by the VAR. In Appendix C.5 we apply several machine learning methods to augment the VAR, in order to ensure that the model-implied forecasting is robust and accurate, which is crucial to our identification. In Appendix C.6, we include oil prices in the VAR and compare the effects of inflation sentiment shocks to oil shocks.

In Appendix D we relax the assumption that absent sentiment shocks, agents have rational expectations. Instead, we identify the shocks as distortions to alternative models of expectations. Broadly, we find that forwards-looking behavioral expectations lead to results consistent with the puzzle that we identified in our baseline estimates. However, if agents have backwards-looking expectations or are missing enough data, then the puzzle disappears. Within Appendix D, we reconcile our findings with Leduc et al. (2007), who assume that agents use no contemporaneous macroeconomic data when forming forecasts, which allows them to identify inflation sentiments shocks by causal ordering. We confirm with modern data that their identified shock is inflationary. But then we show that if even if agents have no information about delayed macroeconomic series (such as inflation), but they are allowed to form forecasts using high frequency variables (such as the interest rate), then again the effects of the sentiment shock match our baseline results.

6 Dynamic Model and Monte Carlo Validation

In this section we seek to better understand our results by comparing them to a full dynamic New Keynesian model. To address two questions, we extend the standard conceptual framework beyond the simple motivating model in Section 2, adding dynamics and a bevy of additional shocks.

First: how puzzling are our findings compared to a New Keynesian DSGE model? We show that even with extra bells and whistles, the canonical New Keynesian model is still fundamentally at odds with our VAR results: a positive sentiment shock still causes inflation to rise, inducing a contractionary response by the central bank and a recession. The dynamic model reinforces the puzzle.

Second: does our estimation strategy reliably identify the effects of sentiment shocks? We show that it does. To do so, we estimate the structural VAR on simulated data, showing that it consistently recovers the true sentiment shocks even in the presence of confounding news, noise, discount factor shocks, policy shocks, and sentiments to GDP and technology, even on samples of similar length to ours. This is an important validation of our method: it

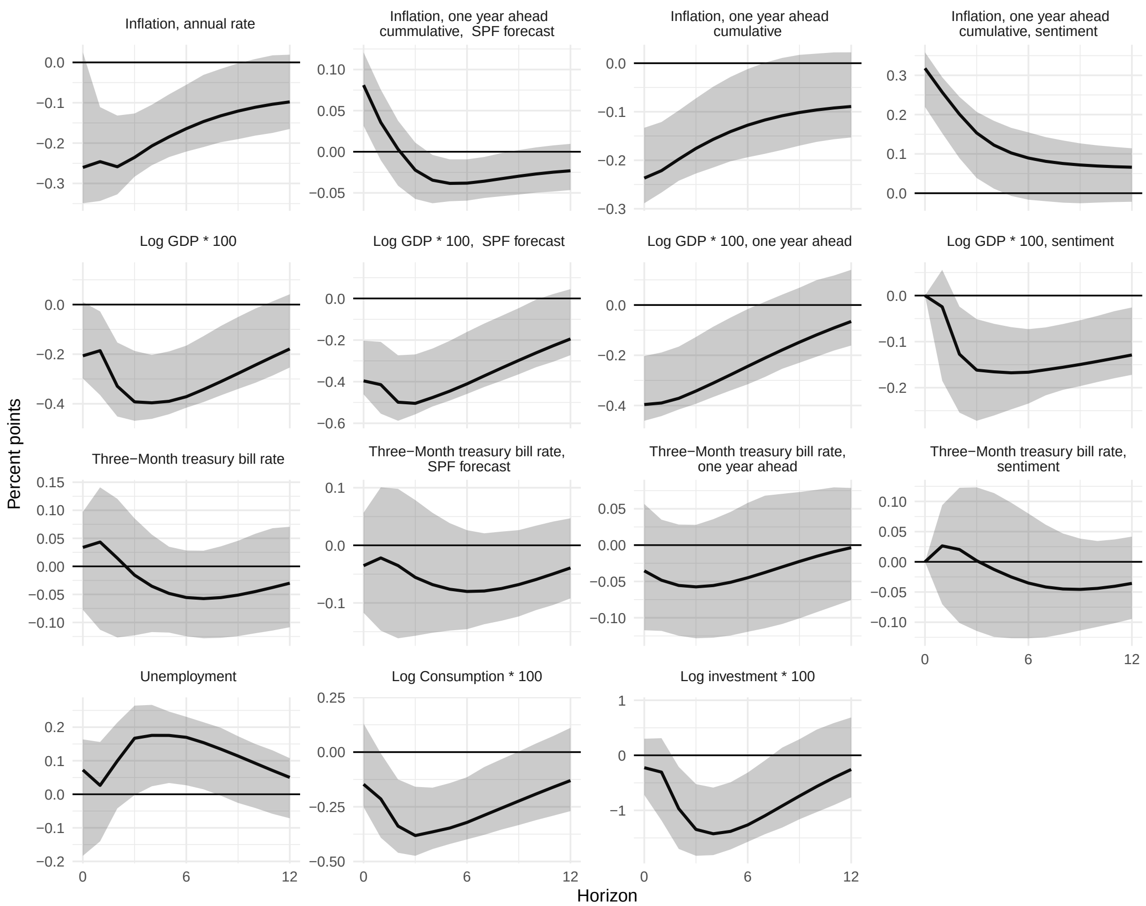

Figure 9: Impulse Responses to an Inflation Sentiment Shock: Multi-Sentiment SPF model

Structural impulse responses to a one standard deviation sentiment shock, using SPF forecasts for inflation, GDP and interest rates. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications. The cumulative inflation over the next 12 months (i.e. the average of horizons 1 to 12 of monthly annualized inflation) is the rational expectations component of inflation expectations. The inflation sentiment is the difference between the response for the inflation forecast and the rational component. The current month's realized inflation is reported at the annualized rate. IRF units are "percent points", which is literal for time series reported as such (e.g. inflation) and 100× log values for the remaining series.

shows that if inflation expectation shocks did generate macroeconomic fluctuations in line with the standard New Keynesian model, we would estimate impulse responses to sentiment shocks consistent with it.

6.1 A Dynamic New Keynesian Model

We modify the canonical New Keynesian model by introducing additional economic shocks, sentiments, and information structure. The additional structural economic shocks are standard. We include stochastic terms for productivity a_t, interest rate deviations x_t, and discount factor z_t. In Section 3.2 we show that we can use multiple forecast series to separately identify inflation sentiments from other forces, so we also include sentiment shocks to productivity a_t and output y_t. The standard three equations (the New Keynesian Phillips Curve, the Euler equation, and the Taylor rule) become

$$\pi_t = \beta \pi_t^{e,1} + \kappa(y_t - \psi a_t) \tag{14}$$

$$0 = \gamma(y_t - y_t^{e,1}) + i_t - \pi_t^{e,1} + z_t \tag{15}$$

$$i_t = \phi_y y_t + \phi_\pi \pi_t + x_t \tag{16}$$

where z_t is a stochastic deviation of the discount factor, and x_t is the central bank's stochastic deviation from the interest rate rule.

We allow expectations $\pi_t^{e,1}$ and $y_t^{e,1}$ to deviate from rationality in two ways: 1) through sentiments separately affecting inflation and output expectations directly, and 2) via a sentiment about future productivity which affects agent's views on inflation and output expectations indirectly. Agents are rational apart from these two distortions.

We denote the equilibrium policy functions for inflation and output by

$$\pi_t = b_a^\pi a_t + b_{\mathbf{m}}^\pi \mathbf{m}_t \qquad\qquad y_t = b_a^y a_t + b_{\mathbf{m}}^y \mathbf{m}_t$$

where a_t is productivity and $\mathbf{m}_t$ is the vector of the remaining state variables. Then, imposing the first of our assumptions above, agents' expectations are given by

$$\pi_t^{e,1} = b_a^\pi a_t^{e,1} + b_{\mathbf{m}}^\pi \mathbf{m}_t^{e,1} + \zeta_t^\pi$$

$$y_t^{e,1} = b_a^y a_t^{e,1} + b_{\mathbf{m}}^y \mathbf{m}_t^{e,1} + \zeta_t^y$$

That is, agents' forecasts depend on their state variable forecasts $a_t^{e,1}$ and $\mathbf{m}_t^{e,1}$, plus a distortion due to the exogenous sentiments ζ_t^π and ζ_t^y. These are the direct sentiments.

Imposing the assumption that productivity forecasts are subject to sentiments but that forecasts are otherwise rational, we can write:

$$a_t^{e,1} = \mathbb{E}[a_{t+1}] + \zeta_t^a \qquad \mathbf{m}_t^{e,1} = \mathbb{E}[\mathbf{m}_{t+1}]$$

This says that their productivity forecast $a_t^{e,1}$ is distorted by the exogenous productivity sentiment ζ_t^a, whereas forecasts of other state variables $\mathbf{m}_t^{e,1}$ are rational.

Putting these pieces together, expectations in the model satisfy

$$\pi_t^{e,1} = \mathbb{E}_t[b_a^\pi a_{t+1} + b_{\mathbf{m}}^\pi \mathbf{m}_{t+1}] + b_a^\pi \zeta_t^a + \zeta_t^\pi \tag{17}$$

$$y_t^{e,1} = \mathbb{E}_t[b_a^y a_{t+1} + b_{\mathbf{m}}^y \mathbf{m}_{t+1}] + b_a^y \zeta_t^a + \zeta_t^y \tag{18}$$

The inflation and output sentiments ζ_t^π and ζ_t^y exogenously affect forecasts directly and independently, while the productivity sentiment shock ζ_t^a indirectly and endogenously affects both forecasts.[34]

The VAR uses one-year-ahead forecasts as an input, but sentiments in the model are distortions to one-month-ahead forecasts, so one-year-ahead forecasts must be calculated as an endogenous variable in order to apply our methodology to simulated data. To do so, we assume that the law of iterated expectations holds for agents when they make forecasts. That is, they correctly assess the impacts of the sentiment on their future forecasts:

$$\pi_t^{e,12} = \mathbb{E}_t[\sum_{j=0}^{11} \pi_{t+j}^{e,1}] \tag{19}$$

We assume our exogenous stochastic terms are governed by AR(1) processes:

$$\zeta_t^\pi = \theta_{\zeta,\pi}\zeta_{t-1}^\pi + \varepsilon_t^{\zeta,\pi} \tag{20}$$

$$\zeta_t^y = \theta_{\zeta,y}\zeta_{t-1}^y + \varepsilon_t^{\zeta,y} \tag{21}$$

$$\zeta_t^a = \theta_{\zeta,a}\zeta_{t-1}^a + \varepsilon_t^{\zeta,a} \tag{22}$$

$$a_t = \theta_a a_{t-1} + \varepsilon_t^a \tag{23}$$

$$x_t = \theta_x x_{t-1} + \varepsilon_t^x \tag{24}$$

$$z_t = \theta_z z_{t-1} + \varepsilon_t^z \tag{25}$$

[34]This model structure resembles a simple version of the medium-scale New Keynesian model studied by Milani (2017), who adds additional economic features including capital, but where sentiments cause expectations to deviate from an adaptive learning forecast, rather than from the rational expectation. Even with the additional features, Milani comes to the same conclusion: inflation sentiment shocks should be inflationary.

We also impose an intertemporal information structure. At time t, agents learn imperfectly about future productivity shocks, receiving a signal v_t:

$$v_t = \varepsilon^a_{t+1} + \nu_t$$

where ν_t is an i.i.d. noise shock. This signal encodes two distinct but related concepts. The ε^a_{t+1} part is *news* about future productivity. And the ν_t is *noise*, obscuring that news. The realization of v_t will affect current forecasts of future inflation. For example, if agents learn that TFP will likely rise, they expect that future output and inflation will increase. This is a standard setup, and tests our method by introducing a shock that generates forecast errors without affecting contemporaneous fundamentals (Chahrour and Jurado, 2022).

Chahrour and Jurado (2018) show that this sort of information structure can be recast as one where agents receive incomplete (rather than noisy) news about the future. This is useful as it allows the model to be solved with standard DSGE methods. In this our case, we can model the productivity shock as a sum of i.i.d. components:

$$\varepsilon^a_t = \vartheta_t + \varsigma_{t-1} \tag{26}$$

where ϑ_{t+1} is unanticipated but ς_t is known in period t. Appendix E derives how to rewrite our noisy signal process in this form.

An equilibrium in this economy is a set of stationary time series for the 13 variables (π_t, y_t, i_t, $y^{e,1}_t$, $\pi^{e,1}_t$, $\pi^{e,12}_t$, ζ^π_t, ζ^y_t, ζ^a_t, a_t, x_t, z_t, and ε^a_t) that satisfy equations (14) - (26) conditional on time series for the 7 exogenous shocks ($\varepsilon^{\zeta,\pi}_t$, $\varepsilon^{\zeta,y}_t$, $\varepsilon^{\zeta,a}_t$, ε^x_t, ε^z_t, ϑ_t, ς_t). Appendix F describes how we solve the model.

In this section, we have modified the canonical New Keynesian model with a number of additional forcing terms – 3 types of sentiments, noise, and discount factor shocks – in addition to standard economic shocks. These additional features present a stringent test of our identification strategy, showing that it can resolve concerns raised elsewhere. The discount factor shock addresses the issue discussed in Levchenko and Pandalai-Nayar (2020): current methods have difficulties separating the effects of sentiments from other unobserved factors that move expectations, such as shocks to preferences. And we include the news signal in order to demonstrate that our sentiment is also independently identified from other shocks to agents' information sets. VAR studies such as Barsky and Sims (2011) or Chahrour and Jurado (2022) identify news or noise by disciplining how expectations respond to structural shocks. Although our method is similar in spirit, it identifies something quite different: a sentiment shock. Thus, it is important to show that our approach is also not confounded by news or noise.

6.2 Calibrated Impulse Responses

We assign a standard calibration to the economy, following Galí (2008). Table 5 reports our parameter values. To calibrate the time series for the shock processes, we adopt the monthly analogs to the quarterly MLE estimates from Smets and Wouters (2007). Additionally, we set the noise variance σ_ν^2 to be the productivity shock variance σ_a^2, so that only half of the productivity shock process is predictable. The addition of sentiment shocks introduces two parameters about which the literature is quiet. First, we set the sentiment autocorrelation $\theta_{\zeta,\pi} = 0.63$ to reproduce the monthly persistence of the inflation sentiment that we estimated in Section 4.2. Second, we conservatively set the standard deviation of the sentiment shock to $\sigma_{\zeta,\pi} = 0.15$ to explain only 11% of income volatility, the number that we estimate is due to the inflation sentiment in our multi-sentiment VAR (Section 5.4).[35] Then to avoid any special advantage for estimating inflation sentiments, we assign these same values to the output and productivity sentiment processes too.

Parameter	Interpretation	Value
β	Discount factor	0.997
γ	Risk aversion	1
κ	Phillips Curve elasticity	0.2
ψ	Output gap elasticity	0.2
ϕ_π	Policy response to inflation	1.5
ϕ_y	Policy response to output	0.1
$\theta_{\zeta,\pi}, \theta_{\zeta,y}, \theta_{\zeta,a}$	Sentiment autocorrelations	0.63
θ_a	TFP autocorrelation	0.98
θ_x	Interest rate deviation autocorrelation	0.49
θ_z	Discount factor autocorrelation	0.56
$\sigma_{\zeta,\pi}, \sigma_{\zeta,y}, \sigma_{\zeta,a}$	Sentiment shock standard deviations	0.15
σ_a	TFP shock standard deviation	0.26
σ_ν	Noise shock standard deviation	0.26
σ_x	Interest rate shock standard deviation	0.14
σ_z	Discount factor shock standard deviation	0.14

Table 5: Standard Monthly Calibration

Figure 10 plots the dynamic effects of an inflation sentiment shock in the dynamic New Keynesian model. For ease of comparison with our estimated shocks, the sentiment shock is scaled to equal one standard deviation of those estimated in the baseline results (around 0.35 percentage points, top right panel).

The intuition for the model responses (red circles) resembles that of the static model in Section 2, but the dynamics are driven by the Calvo pricing friction – that only a fixed

[35] In the multi-sentiment VAR, 50% of income variance is due to sentiments, but only 11% is due to the inflation sentiment. This contrasts with our baseline VAR, which assigns 51% of income variance to the inflation sentiment.

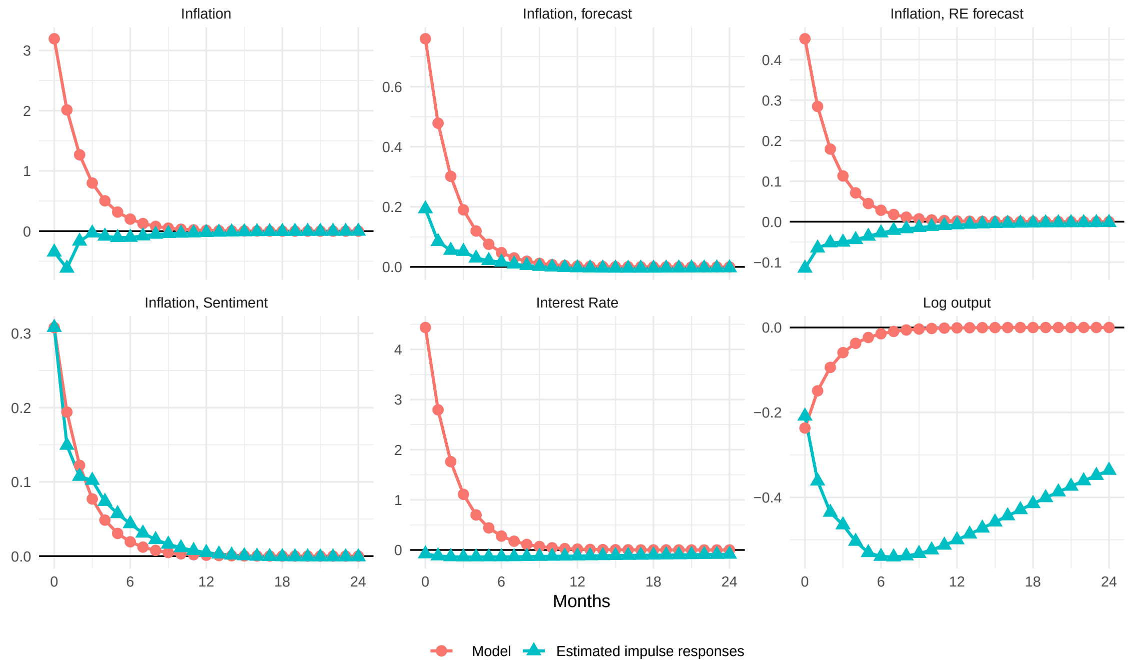

Figure 10: Inflation Sentiment Shock: Dynamic New Keynesian Model versus Baseline VAR Estimates

fraction of firms can reset their prices each period. Initially, inflation expectations rise due to the sentiment. Because prices are sticky, firms that can change their prices today will increase them, as they want to avoid their prices being too low tomorrow. Thus, inflation happens today. This is the direct impact of the sentiment shock. But there is also an indirect impact, acting through rational expectations. Inflation today means that some firms' prices will be higher tomorrow. Consumers will thus substitute away from these firms' products, increasing demand for those that cannot increase prices today. Of these firms, some will be able to reset prices tomorrow. And with higher demand, they will raise their prices, creating inflation tomorrow. But firms today anticipate this – except for the sentiment, their expectations are rational. So they respond by further raising prices today.[36] This effect on future inflation is the "expectations multiplier", and it is quantitatively large: even though the central bank raises real rates, realized inflation over the next 12 months (the top right panel) increases by around 0.8 percentage point. The multiplier is thus around 2.7 (as $0.8/0.3 \approx 2.7$) This raises the rational expectation, producing an indirect impact on forecasts around two times the size of the direct one. Indeed, this is why interest rates have to rise so much – higher inflation erodes the real interest rate. If this is the framework that

[36] Forecasted inflation in this figure is for the following year, given by equation 19, as it was in our empirical analysis. This is why the rational component does not track the inflation impulse response exactly.

central bankers have in mind when talking about inflation expectations, then they are right to be concerned about them as an source of inflationary shocks.

However, these dynamics are clearly at odds with the estimated results from Section 4.2, also shown in Figure 10 (blue triangles). For the same sentiment shock, the estimated responses are generally smaller in magnitude and of different sign for inflation and interest rates. In particular, the measured inflation expectations multiplier – which is large and positive in the model – is much smaller and negative. The ratio of realized one-year inflation to the one-year-ahead inflation sentiment shock is around -0.3 versus almost 4 in the model. In other words, the empirical puzzle we document cannot be explained by the dynamic New Keynesian model.

6.3 Validating Our Identification

We now use the extended New Keynesian model developed in the preceding section to form two tests our identification strategy.

The first test is a long-sample test, meant to uncover whether our method is at least asymptotically valid. We simulate the model for 100,000 periods and estimate shocks to inflation sentiments using the semi-structural identification procedure in Section 3, applied to a VAR featuring inflation, output, the interest rate, and year-ahead inflation expectations. We use the multiple-sentiment method explored in Section 5.4 to identify the inflation sentiment shock. The green triangles in Figure 11 plot our estimates from the large sample. The red circles plots the true impulse responses to the sentiment shock, just as in Figure 10. These coincide almost exactly. This says that given enough data, we can precisely identify the effects of inflation sentiment shocks from other sentiments, preferences, news, noise, and other structural shocks in the model.

Although our method is asymptotically valid, it could be that our empirical results are a fluke, due to statistical noise in a short sample. To assess this possibility we repeatedly simulate 39 year samples, mimicking the data in our baseline regressions. The blue squares in Figure 11 plot the median estimated impulse response function from these smaller simulations, while the gray shaded area is the bootstrapped 90 percent confidence interval. Broadly speaking, the median is very close to the true shock and its impact. Moreover, the confidence intervals at least on impact are tight enough to convincingly reject the possibility that our results – disinflation, lower interest rates, and a hump-shaped output loss – could be generated by a dynamic New Keynesian model of this sort. Further evidence for the validity of our approach comes from the estimated response of the output sentiment, which is indistinguishable from zero (bottom right panel). The rational expectation for output does respond, though, since future output falls.

In the final row of Figure 11, the estimated response of output expectations are all statistically indistinguishable from zero. This is because the effect of the shock on output decays rapidly, the 12-month-ahead rational expectation is near zero. And the inflation sentiment shock is also correctly estimated to have no effect on the output sentiment (bottom right panel).

Together these exercises represent the acid test of our identification strategy. If something was wrong with our approach, it would recover counterfactual impulse responses. That it does not, even in the presence of other potentially confounding shocks, says that our method is valid.

Nevertheless, an important question remains: *how* does our method work? The answer is that only sentiment shocks impact the non-rational part of agents' forecasts. For example, they cause inflation expectations $\pi_t^{e,12}$ to move distinctly from the mathematical conditional expectation of inflation, $\mathbb{E}_t\pi_{t+12}$. The economic shocks (news, noise, and preference shocks) also affect expected inflation $\pi_t^{e,12}$, but only because they affect the distribution of future inflation, and thus its conditional expectation, $\mathbb{E}_t\pi_{t+12}$. And so for these shocks, $\pi_t^{e,12}$ and $\mathbb{E}_t\pi_{t+12}$ move in lockstep. Because our identification strategy is a device for uncovering the covariation in the data where forecasts and rational expectations move differently, it can identify the effects of sentiment shocks.

The method then separates inflation from other sentiments via the ordering within the sentiment block. Both inflation and productivity sentiment shocks make the inflation forecast $\pi_t^{e,12}$ depart from the rational expectation $\mathbb{E}_t\pi_{t+12}$. But productivity sentiment shocks do so for the output forecast as well, while inflation sentiment shocks only affect output forecasts through the rational expectation. Therefore, within the sentiment block A_f^S, we can identify the inflation sentiment shock by ordering inflation last, and letting Λ_S (equation (9)) be lower triangular, so that equation (10) is a Cholesky decomposition.

7 Conclusion

We developed a novel method to identify sentiment shocks from a structural VAR with aggregate data and empirical forecasts. When applied to inflation forecasts, we find a puzzle: a positive shock to inflation expectations are contractionary, deflationary, and induce monetary loosening. This is inconsistent with the canonical New Keynesian model, in which such shocks are inflationary, and only cause a recession if the following policy tightening is sufficiently aggressive. Our findings suggest two avenues for additional work.

First: what explains the puzzle? Why are inflation sentiments deflationary? We explore one possible explanation in Appendix D, by considering agents who have behavioral instead of rational expectations. When we identify sentiments as distortions to alternative forward-looking beliefs such as diagnostic expectations, the puzzle persists. But, if beliefs are sufficiently backwards-looking, such as adaptive expectations, sentiment shocks can become inflationary.

Second: our method for identifying sentiments is robust, and can apply to expectations of quantities other than inflation. Do sentiments of future productivity have the effects that a long information frictions literature predicts? What of sentiments for other variables with measured expectations, such as interest and exchange rates, income, wages, and so forth?

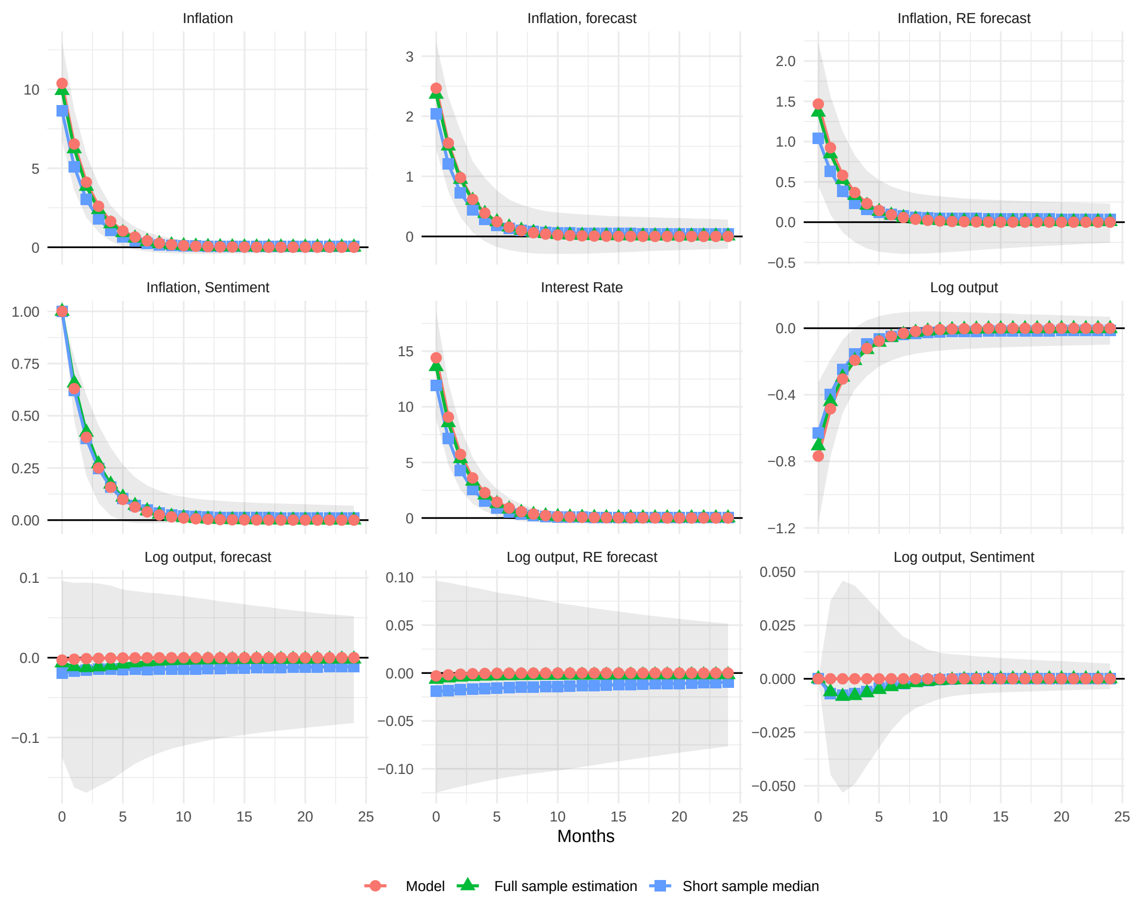

Figure 11: Validation Exercise: Structural Impulse Responses to an Inflation Sentiment Shock

The "long simulated sample" shows the point estimates of the structural VAR decomposition using a single sample of 100,000 points. The shaded range shows the 90 percent confidence interval from 1000 shorter simulations, each of 39 years, the length of our data sample. "Short sample median" is the median across these 1000 simulations.

References

Adam, Klaus and Mario Padula, "Inflation Dynamics and Subjective Expectations in the United States," *Economic Inquiry*, 2011, *49* (1), 13–25.

Adams, Jonathan J., "Behavioral Expectations Equilibrium Toolkit," Working Papers 001012, University of Florida, Department of Economics May 2024.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart, "Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples," *The Review of Economic Studies*, November 2022, *89* (6), 2958–2991.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas, "Business-Cycle Anatomy," *American Economic Review*, October 2020, *110* (10), 3030–3070.

Aruoba, S. Boragan and Thomas Drechsel, "Identifying Monetary Policy Shocks: A Natural Language Approach," March 2022.

Ascari, Guido, Stefano Fasani, Jakob Grazzini, and Lorenza Rossi, "Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations," *Journal of Monetary Economics*, April 2023.

Barsky, Robert B. and Eric R. Sims, "News shocks and business cycles," *Journal of Monetary Economics*, April 2011, *58* (3), 273–289.

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

Beaudry, Paul and Bernd Lucke, "Letting Different Views about Business Cycles Compete," *NBER Macroeconomics Annual*, January 2010, *24*, 413–456. Publisher: The University of Chicago Press.

_ and Franck Portier, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, September 2006, *96* (4), 1293–1307.

_ and _ , "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, December 2014, *52* (4), 993–1074.

Beladi, Hamid, Munit A. S. Choudhary, and Amar K. Parai, "Rational and Adaptive Expectations in the Present Value Model of Hyperinflation," *The Review of Economics and Statistics*, 1993, *75* (3), 511–514. Publisher: The MIT Press.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe, "Monetary Policy and Multiple Equilibria," *American Economic Review*, March 2001, *91* (1), 167–186.

_ , _ , and _ , "The Perils of Taylor Rules," *Journal of Economic Theory*, January 2001, *96* (1), 40–69.

Bernanke, Ben S, Jean Boivin, and Piotr Eliasz, "Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach," *The Quarterly journal of economics*, 2005, *120* (1), 387–422.

Bernstein, Joshua and Rupal Kamdar, "Rationally inattentive monetary policy," *Review of Economic Dynamics*, April 2023, *48*, 265–296.

Binder, Carola and Christos Makridis, "Stuck in the Seventies: Gas Prices and Consumer Sentiment," *The Review of Economics and Statistics*, March 2022, *104* (2), 293–305.

Binder, Carola C., "Measuring uncertainty based on rounding: New method and application to inflation expectations," *Journal of Monetary Economics*, October 2017, *90*, 1–12.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni, "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, December 2013, *103* (7), 3045–3070.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic Expectations and Credit Cycles," *The Journal of Finance*, 2018, *73* (1), 199–227. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12586.

_ , _ , Yueran Ma, and Andrei Shleifer, "Overreaction in Macroeconomic Expectations," *American Economic Review*, September 2020, *110* (9), 2748–2782.

Brissimis, Sophocles N. and Nicholas S. Magginas, "Inflation Forecasts and the New Keynesian Phillips Curve," *International Journal of Central Banking*, June 2008.

Cagan, Phillip, "The Monetary Dynamics of Hyperinflation," in Milton Freidman, ed., *Studies in the Quantity Theory of Money*, University of Chicago Press, 1956.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko, "The Macroeconomic Expectations of Firms," *Forthcoming in Handbook of Economic Expectations*, 2021.

Chahrour, Ryan and Kyle Jurado, "News or Noise? The Missing Link," *American Economic Review*, July 2018, *108* (7), 1702–1736.

_ and _ , "Recoverability and Expectations-Driven Fluctuations," *The Review of Economic Studies*, January 2022, *89* (1), 214–239.

Christiano, Lawrence J., "Cagan's Model of Hyperinflation under Rational Expectations," *International Economic Review*, 1987, *28* (1), 33–49. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

_ , Martin Eichenbaum, and Charles L. Evans, "Chapter 2 Monetary policy shocks: What have we learned and to what end?," in "Handbook of Macroeconomics," Vol. 1, Elsevier, January 1999, pp. 65–148.

Clarida, Richard, Jordi Galí, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, February 2000, *115* (1), 147–180.

Clements, Michael P. and Ana Beatriz Galvão, "Measuring the effects of expectations shocks," *Journal of Economic Dynamics and Control*, March 2021, *124*, 104075.

Cochrane, John H., "Permanent and Transitory Components of GNP and Stock Prices," *The Quarterly Journal of Economics*, February 1994, *109* (1), 241–265.

Coibion, Olivier, "Are the Effects of Monetary Policy Shocks Big or Small?," *American Economic Journal: Macroeconomics*, April 2012, *4* (2), 1–32.

_ and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, February 2012, *120* (1), 116–159.

_ and _ , "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, August 2015, *105* (8), 2644–2678.

_ and _ , "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 197–232.

_ , Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber, "Forward Guidance and Household Expectations," *Journal of the European Economic Association*, October 2023, *21* (5), 2131–2171.

_ , Yuriy Gorodnichenko, and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, December 2018, *56* (4), 1447–1491.

Curtin, Richard, "Inflation expectations and empirical tests," *Inflation Expectations*, 2010, *56*, 34–61. Publisher: Routledge New York.

D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber, "What do the data tell us about inflation expectations?," in Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, eds., *Handbook of Economic Expectations*, Academic Press, January 2023, pp. 133–161.

de Bruin, Wändi Bruine, Wilbert Vanderklaauw, Julie S. Downs, Baruch Fischhoff, Giorgio Topa, and Olivier Armantier, "Expectations of Inflation: The Role of Demographic Variables, Expectation Formation, and Financial Literacy," *Journal of Consumer Affairs*, 2010, *44* (2), 381–402.

Doh, Taeyoung and A. Lee Smith, "A New Approach to Integrating Expectations into VAR Models," *Journal of Monetary Economics*, August 2022.

Eggertsson, Gauti B. and Paul Krugman, "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach," *The Quarterly Journal of Economics*, August 2012, *127* (3), 1469–1513.

Farmer, Leland, Emi Nakamura, and Jón Steinsson, "Learning About the Long Run," Working Paper 29495, National Bureau of Economic Research November 2021.

Forni, Mario, Luca Gambetti, Marco Lippi, and Luca Sala, "Noise Bubbles," *The Economic Journal*, September 2017, *127* (604), 1940–1976.

_ , _ , _ , and _ , "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics*, October 2017, *9* (4), 122–152.

Friedman, Milton, *Theory of the consumption function*, Princeton university press, 1957.

Fuhrer, Jeff, "The Role of Expectations in Inflation Dynamics," *International Journal of Central Banking*, 2012, p. 29.

Fève, Patrick and Alain Guay, "Sentiments in SVARs," *The Economic Journal*, February 2019, *129* (618), 877–896.

Gabaix, Xavier, "A Behavioral New Keynesian Model," *American Economic Review*, August 2020, *110* (8), 2271–2327.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2008.

_ and Mark Gertler, "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, October 1999, *44* (2), 195–222.

Gazzani, Andrea, "News and noise bubbles in the housing market," *Review of Economic Dynamics*, April 2020, *36*, 46–72.

Gelain, Paolo, Nikolay Iskrev, Kevin J. Lansing, and Caterina Mendicino, "Inflation dynamics and adaptive expectations in an estimated DSGE model," *Journal of Macroeconomics*, March 2019, *59*, 258–277.

Hamilton, James Douglas, *Time Series Analysis*, Princeton University Press, 1994.

Hasenzagl, Thomas, Filippo Pellegrino, Lucrezia Reichlin, and Giovanni Ricco, "A Model of the Fed's View on Inflation," *The Review of Economics and Statistics*, July 2022, *104* (4), 686–704.

Jonung, Lars, "Perceived and Expected Rates of Inflation in Sweden," *The American Economic Review*, 1981, *71* (5), 961–968.

Jorda, Oscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

Kim, Gwangmin and Carola Binder, "Learning-through-Survey in Inflation Expectations," *American Economic Journal: Macroeconomics*, April 2023, *15* (2), 254–278.

Kydland, Finn E. and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, June 1977, *85* (3), 473–491.

Känzig, Diego R., "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, April 2021, *111* (4), 1092–1125.

Lagerborg, Andresa, Evi Pappa, and Morten O. Ravn, "Sentimental Business Cycles," SSRN Scholarly Paper ID 3661442, Social Science Research Network, Rochester, NY July 2020.

Leduc, Sylvain and Keith Sill, "Expectations and Economic Fluctuations: An Analysis Using Survey Data," *The Review of Economics and Statistics*, October 2013, *95* (4), 1352–1367.

_ , _ , and Tom Stark, "Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston Survey," *Journal of Monetary Economics*, March 2007, *54* (2), 433–459.

Levchenko, Andrei A and Nitya Pandalai-Nayar, "Tfp, News, and "Sentiments": the International Transmission of Business Cycles," *Journal of the European Economic Association*, February 2020, *18* (1), 302–341.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, November 2002, *117* (4), 1295–1328.

_ , _ , and Justin Wolfers, "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual*, January 2003, *18*, 209–248.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock, "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, March 2014, *52* (1), 124–188.

Milani, Fabio, "Expectation Shocks and Learning as Drivers of the Business Cycle," *The Economic Journal*, May 2011, *121* (552), 379–401.

_ , "Sentiment and the U.S. business cycle," *Journal of Economic Dynamics and Control*, September 2017, *82*, 289–311.

Nasir, Muhammad Ali, Daniel Balsalobre-Lorente, and Toan Luu Duc Huynh, "Anchoring inflation expectations in the face of oil shocks & in the proximity of ZLB: A tale of two targeters," *Energy Economics*, February 2020, *86*, 104662.

Nicholson, William B, David S Matteson, and Jacob Bien, "VARX-L: Structured regularization for large vector autoregressions with exogenous variables," *International Journal of Forecasting*, 2017, *33* (3), 627–651.

_ , Ines Wilms, Jacob Bien, and David S Matteson, "High Dimensional Forecasting via Interpretable Vector Autoregression.," *J. Mach. Learn. Res.*, 2020, *21* (166), 1–52.

Nunes, Ricardo, "Inflation Dynamics: The Role of Expectations," *Journal of Money, Credit and Banking*, 2010, *42* (6), 1161–1172. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4616.2010.00324.x.

Roberts, John M., "Is inflation sticky?," *Journal of Monetary Economics*, July 1997, *39* (2), 173–196.

Rudd, Jeremy and Karl Whelan, "New tests of the new-Keynesian Phillips curve," *Journal of Monetary Economics*, September 2005, *52* (6), 1167–1181.

_ **and** _ , "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?," *American Economic Review*, March 2006, *96* (1), 303–320.

Savignac, Frédérique, Erwan Gautier, Yuriy Gorodnichenko, and Olivier Coibion, "Firms' Inflation Expectations: New Evidence from France," Working Paper 29376, National Bureau of Economic Research October 2021.

Sims, Christopher A., "Macroeconomics and Reality," *Econometrica*, 1980, *48* (1), 1–48. Publisher: [Wiley, Econometric Society].

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, June 2007, *97* (3), 586–606.

Stock, James H. and Mark W. Watson, "Implications of Dynamic Factor Models for VAR Analysis," Working Paper 11467, National Bureau of Economic Research July 2005.

Stock, James H and Mark Watson, "Dynamic factor models," in Michael P. Clements and David F. Hendry, eds., *Oxford Handbook on Economic Forecasting*, Oxford: Oxford University Press, 2011.

Uhlig, Harald, "A Toolkit for Analysing Nonlinear Dynamic Stochastic Models Easily," in Ramon Marimon and Andrew Scott, eds., *Computational Methods for the Study of Dynamic Economies*, Oxford University Press, October 2001, p. 0.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "Reality Check: How People Form Inflation Expectations and Why You Should Care," SSRN Scholarly Paper ID 4058361, Social Science Research Network, Rochester, NY October 2021.

Werning, Iván, "Expectations and the Rate of Inflation," Working Paper 30260, National Bureau of Economic Research July 2022.

Wong, Benjamin, "Do Inflation Expectations Propagate the Inflationary Impact of Real Oil Price Shocks?: Evidence from the Michigan Survey," *Journal of Money, Credit and Banking*, 2015, *47* (8), 1673–1689. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jmcb.12288.

A Proof of Theorem 1

Proof. The sub-blocks of A are related to the sub-blocks of $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}' & \Sigma_{22} \end{pmatrix}$ by:

$$A_f^S(A_f^S)' + A_f^F(A_f^F)' = \Sigma_{11} \tag{27}$$
$$A_f^S(A_c^S)' + A_f^F(A_c^F)' = \Sigma_{12} \tag{28}$$
$$A_c^S(A_c^S)' + A_c^F(A_c^F)' = \Sigma_{22} \tag{29}$$

Per equation (8), the variance of sentiment innovations $V^S = Var\big((\mathbb{E}[s_t^h|\varepsilon_t]\big)$ is given by

$$V^S = Var\left(\begin{pmatrix} A_f^S & A_f^F \end{pmatrix}\varepsilon_t - \phi_x^h A\varepsilon_t\right) = Var\left(\begin{pmatrix} I & 0 \end{pmatrix} A\varepsilon_t - \phi_x^h A\varepsilon_t\right)$$

$$= Var\left(\begin{pmatrix} I - \phi_{x,f}^h & -\phi_{x,c}^h \end{pmatrix} A\varepsilon_t\right) = \begin{pmatrix} I - \phi_{x,f}^h & -\phi_{x,c}^h \end{pmatrix}\Sigma\begin{pmatrix} I - \phi_{x,f}^h & -\phi_{x,c}^h \end{pmatrix}'$$

so V^S and thus $\xi(V_S)$ are determined by Σ and ϕ_x^h.

Assumption (11) and definition (9) imply

$$\xi(V_S) = A_f^S - \phi_x^h \begin{pmatrix} A_f^S \\ A_c^S \end{pmatrix} = (I - \phi_{x,f}^h)A_f^S - \phi_{x,c}^h A_f^S$$

which implies

$$A_f^S = (I - \phi_{x,f}^h)^{-1}\xi(V_S) + (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^S \tag{30}$$

Substitute equation (6) into equation (28):

$$A_f^S(A_c^S)' + (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^F(A_c^F)' = \Sigma_{12}$$

subtract equation (29) pre-multiplied by $(I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h$:

$$A_f^S(A_c^S)' - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^S(A_c^S)' = \Sigma_{12} - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h\Sigma_{22}$$

Then substitute for A_f^S using equation (30):

$$\left((I - \phi_{x,f}^h)^{-1}\xi(V_S) + (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^S\right)(A_c^S)' - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h A_c^S(A_c^S)' = \Sigma_{12} - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h\Sigma_{22}$$

$$\implies (A_c^S)' = \xi(V_S)^{-1}(I - \phi_{x,f}^h)\left(\Sigma_{12} - (I - \phi_{x,f}^h)^{-1}\phi_{x,c}^h\Sigma_{22}\right)$$

This equation pins down A_c^S, and restriction (30) gives A_f^S.

Equation (29) only gives $A_c^F(A_c^F)' = \Sigma_{22} - A_c^S(A_c^S)'$ but not A_c^F individually; the fundamental shocks can be arbitrarily reordered or otherwise unitarily transformed. Any decomposition of $A_c^F(A_c^F)'$ to select a A_c^F then gives A_f^F by equation (6). ■

B Extensions to the Motivating Model

B.1 Passive Monetary Policy

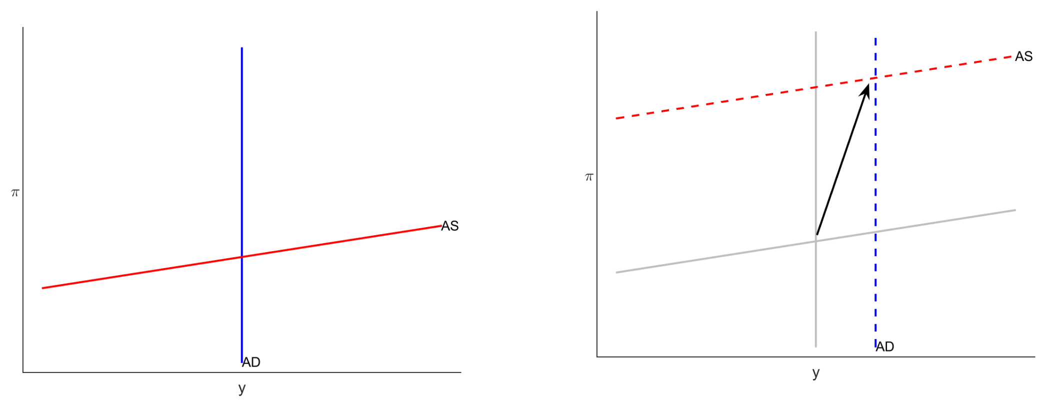

(a) Steady State without Taylor Rule

(b) Response to Sentiment Shock without Taylor Rule

Figure 12: New Keynesian Sentiment Shock without a Policy Response

In the motivating model of Section 2, an inflation sentiment shock caused inflation while lowering output and interest rates. The inflation response is unambiguous, but can the AD shift dominate so that output increases? Yes, if the central bank's interest rate policy is relatively passive.

To see why, it is helpful to think about the special case where the central bank does not respond to inflation at all: $\phi_\pi = 0$. Then the AD curve is vertical, as shown in Figure 12. Without the central bank's policy response, inflation increases by more, because the AS curve rises along a vertical AD curve. On top of this, the AS curve shifts to the right; agents expect high inflation, so real real rates are low, boosting consumption and output, even though nominal rates rise (as the Taylor rule still puts weight on the output gap). So overall, the shock has positive real and nominal effects: inflation, output, and the interest rate all rise. When $\phi_\pi > 0$, the outcome is slightly different. With a flatter AD curve, output can fall. The specific condition for an increase in ζ_t to reduce y_t is that $\frac{1-\phi_\pi\beta}{\phi_\pi\kappa+\phi_y+\gamma} > 0$. Indeed, in most standard calibrations (where $\phi_\pi \gg 1$ and $\phi_y \geq 0$), this effect dominates and the shock causes a recession.

B.2 Whose Sentiment?

What if different actors in the economy form expectations differently? In this section we vary who receives the sentiment in the static New Keynesian model, and consider how it affects the predictions of the canonical theory.

We consider shocks to three different forecasts: those of firms, households, and the central bank. We denote the inflation forecasts of each of these sectors by $\pi^{e,1}_{f,t}$, $\pi^{e,1}_{hh,t}$, and $\pi^{e,1}_{cb,t}$ respectively. In order to allow sentiments to affect the central bank directly, we modify their Taylor rule to depend on expected inflation. Together with the Euler equation and New Keynesian Phillips Curve, the 3 equation model becomes:

New Keynesian Phillips curve:	$\pi_t = \beta\pi^{e,1}_{f,t} + \kappa y_t$
Euler equation:	$i_t = \mathbb{E}_t[\gamma(y_{t+1} - y_t)] + \pi^{e,1}_{hh,t}$
Modified Taylor rule:	$i_t = \phi_y y_t + \phi_\pi \pi_t + \phi_e \pi^{e,1}_{cb,t}$

The firms' inflation forecast $\pi^{e,1}_{f,t}$ enters the New Keynesian Phillips curve, which characterizes the optimal price setting decision by sticky price firms. The households' inflation forecast $\pi^{e,1}_{hh,t}$ enters the Euler equation, which describes their optimal consumption-savings decision. Finally, the central bank's inflation forecast $\pi^{e,1}_{cb,t}$ appears in the new Taylor rule.

As before, the dynamic New Keynesian model reduces to a two-equation static model when sentiments are i.i.d. If we allow firms, households, and central banks to receive different sentiment shocks, the static New Keynesian equations become:

$$\pi_t = \beta\zeta_{f,t} + \kappa y_t \qquad \text{[AS]}$$

$$\phi_\pi \pi_t = -(\phi_y + \gamma)y_t + \zeta_{hh,t} - \phi_e \zeta_{cb,t} \qquad \text{[AD]}$$

where $\zeta_{f,t}$, $\zeta_{hh,t}$, and $\zeta_{cb,t}$ denote sentiment shocks to firms, households, and the central bank respectively.

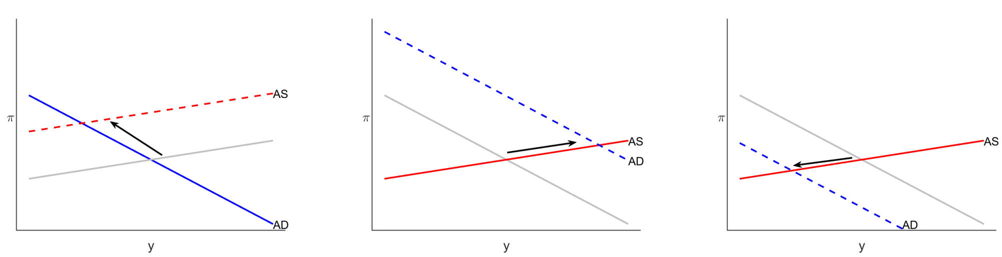

(a) Shock to Firms (b) Shock to Households (c) Shock to Central Bank

Figure 13: New Keynesian Response to Different Sentiment Shocks

Figure 13 plots the response of the macroeconomy to each type of sentiment. When firms' sentiments $\zeta_{f,t}$ increase (panel (a)), they set prices higher ceteris paribus, shifting the New Keynesian Phillips (AS) curve up. This raises inflation, which prompts policy to tighten, increasing the real interest rate and contracting output. When households' sentiments $\zeta_{hh,t}$ increase (panel (b)), they expect higher inflation, perceive real interest rates to decline, and increase consumption. This moves the economy up along the Phillips curve, increasing real output and contemporaneous inflation, despite the central bank's response of raising rates to combat the inflation. When the central bank's sentiment $\zeta_{cb,t}$ increase (panel (c)), it preemptively raises interest rates ($\phi_e > 0$), reducing current inflation and creating a recession.

Crucially, all three sentiments result in monetary policy tightening. If households or firms receive a sentiment shock, inflation increases and the central bank raises interest rates. If the central bank receives a sentiment shock, it mistakenly raises rates, creating deflation. There is no way that a sentiment shock in this static model can result in a decrease in interest rates. Furthermore, the only way that sentiments can result in deflation is if they overwhelmingly affect the central bank. How do these predictions compare to the data? Only the sentiment shocks to the Fed's forecasts are consistent with these effects.

C Additional Empirical Results

This appendix reports additional findings and robustness checks.

C.1 Dynamic Variance Decomposition

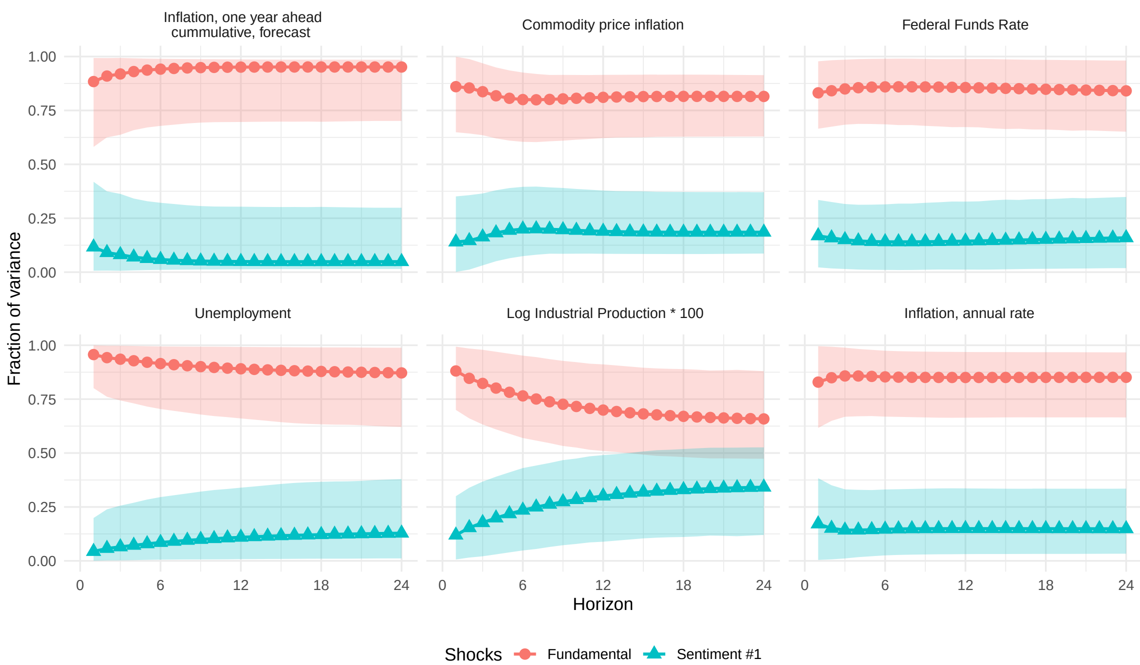

Figure 14: Variance Decomposition

The variance decomposition for each horizon in the baseline model. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications.

Figure 14 presents the variance decomposition from our baseline estimation for a variety of horizons. The variance shares are due to either non-fundamental (sentiment) shocks, or the many fundamental shocks, which are added together into one object. The figure shows that inflation sentiment shocks are responsible for a relatively large share of the variation of inflation expectations at short horizons but less at longer ones. This presumably reflects the fact that the direct (via ζ_t) and indirect (via $\mathbb{E}_t\pi_{t+1}$) effects of the sentiment shock on expected inflation offset, limiting their overall effects. In the long run, around 80 percent of the variation in inflation expectations is due to fundamental factors. Similar long-run effects hold true for most other variables, with inflation sentiments driving between around 10 and 20 percent of the variation. For industrial production, however, sentiment shocks seem to be of significant and growing importance at long horizons, consistent with the large and delayed response of industrial production to sentiment shocks. Together, the patterns in Figure 14 imply that sentiment shocks may be an important driver of macroeconomic fluctuations in both the short and long run.

C.2 Lag Length for the VAR

The lag length for our baseline VAR is rather short, at only 3 lags for monthly data. To check that this is not substantively affecting our results, Figure 15 compares the estimated responses using a variety of lag lengths. Qualitatively, the results are very similar across specifications: when the inflation sentiment increases, inflation falls, real activity contracts, and interest rates decline. Quantitatively, the results are a little different at shorter lags, with the output and interest rate responses attenuated slightly in the one- and two-lag specifications. However, more extended lag structures certainly do not overturn our results. If anything, they strengthen them, with more persistent output and interest rate responses.

Why do we end up with such a short lag structure? Our interpretation is that the inflation forecast contains quite a lot of useful information. By improving the accuracy of the VAR, including forecasts allows the model to fit just as well without so many lags of other variables. In other words, stated forecasts are a good substitute for extra lags

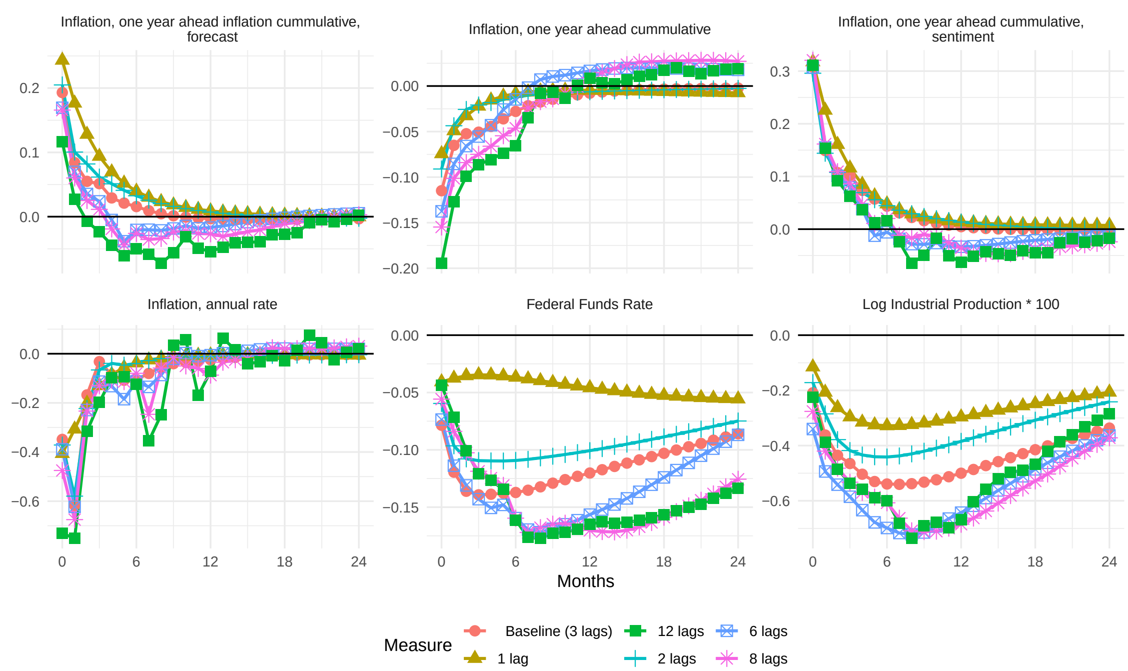

Figure 15: Impulse Responses to a Sentiment Shock for Lag Structures

C.3 Local Projections

As mentioned in the main text, as a further check on our work we also compute the impulse responses estimated by local projections in the style of Jorda (2005). For $h = 0, \ldots, H$ we estimate:

$$\begin{pmatrix} f^h_{t+h} \\ \pi_{t+h} \\ y_{t+h} \end{pmatrix} = \beta_h \hat{x}_t + \sum_{j=1}^{M} B_j \begin{pmatrix} f^h_{t-j} \\ \pi_{t-j} \\ y_{t-j} \end{pmatrix} + u_{t+h} \tag{31}$$

where $\hat{x}^M_t$ is the estimated sentiment shock from the structural VAR with lag length M. The object of interest is the vector of coefficients β_h, which traces out the local projection of the impulse response to the structural shock, conditioned on M lags of the dependent variables.

We compute this set of responses for a variety of lag lengths, just as with do with the VAR in Section C.2. For any given lag length M, we match the estimated structural shock to that derived from the VAR with the same lag length. One consequence is that the $h = 0$ response is, by construction, identical to the VAR. This can be seen in Figure 16, where the immediate responses line up identically. This is not only a useful check that we are comparing like with like across the VAR and local projections responses, but it also makes clear that this is a test of the extent to which the propagation of the identified shocks are shaped by the assumptions implicit in using a VAR. The most notable are that the true model is linear and that the model is well-specified. This seems particularly for our application given that we identify shocks from the dynamics. And so we should check that those dynamics are at least robust to these assumptions.

The local projections also offer a particularly clear opportunity to falsify our results about *average* inflation in the 12 months after the shock. We do this because inflation expectations are typically measured at 12-month horizons, and we want to match expectations of higher inflation inflation over the same horizon. But because the 12-month ahead inflation outcome is a function of the model dynamics, if our VAR gets the dynamics wrong, then this key result could be overturned.

Figure 16 compares our VAR results to the local projection in the base case, where $M = 3$. Despite some volatility in the estimates, the local projections confirm our VAR responses, with average inflation, real activity and the federal funds rate all declining in the short term. An important corroboration of our main results is in the short-term inflation response, which is not only negative on impact but broadly negative thereafter, confirming that the change in sign is not merely an artefact of the estimation method but instead something deeper.

In Figure 17 we check that the local projection specification itself is not sensitive to the lag length. It is not. The effect of lag length on the local projection is much the same as on the VAR, in that at sufficiently short lag lengths (1 or 2), the projection is sensitive to adding further lags. However, beyond 3 lags the results are broadly stable.

C.4 Confidence Intervals across FAVAR Models

Section 5.3 estimated several Factor-Augmented VARs as alternatives to our baseline specification. The impulse response functions resembled our baseline estimates, which we presented in Figure 7 without confidence intervals for readability. But how do the confidence intervals compare? 18 shows the 90 percent confidence intervals for period 0 responses for each model, which look generally similar. Sentiment shocks increase forecasts but decrease the rational component of inflation expectations. Even though 12-month inflation declines, the shock increases the immediate monthly inflation, and the Fed responds by reducing the interest rate. And while medium-run real activity contracts, the effect is delayed because industrial production does not immediately change, at least not in a way that is statistically significant.

C.5 Machine Learning

Machine learning methods for selecting VAR models allow out-of-sample forecast performance to select the appropriate statistical model. We apply four methods proposed by Nicholson et al. (2017) and Nicholson et al. (2020), each of which imposes a penalty for VAR coefficients different from zero. The first is a basic least absolute shrinkage and selection operator (LASSO) where the penalty is linear in the absolute value. The second is an elastic net regularization approach, which applies a linear combination of the LASSO and ridge (i.e. quadratic) penalties. The third, is a component-wise hierarchical VAR (HVARC), restricting candidate models to those where each row of the VAR lag matrix has non-zero coefficients up to a row-specific maximum lag. The fourth is a tapered LASSO, which downweights longer lags. In all cases the penalty functions depend on tuning parameters. These are selected by rolling cross-validation on the middle third of the sample. The final third of the the data is reserved for model evaluation, as measured by the out-of-sample mean square forecast errors.

Figure 19 shows the impulse responses from the four machine learning methods. As with the FAVAR approach, they are qualitatively very similar to our baseline estimates. Overall, the shock is well-identified as broadly deflationary in the the price level declines in the 12 months after the shock. Real activity declines, although perhaps not as soon or as far as in our baseline. The one area where the machine learning models differ from our baseline is in the policy response, which appears to be much smaller. Another benefit of the machine learning approach is that models have straightforward criteria to evaluate their efficacy. We discuss and compare model performance in Appendix C.5.

Overall, the model selection broadly validates our measurement of the inflation sentiment and our characterization of its macroeconomic impact as deflationary (one-year-ahead inflation falls) and contractionary (real activity declines). One aspect of our results that the model selection exercise does not completely confirm is the response of monetary policy, which is much more muted in all the alternative models.

	Avg.	AIC	Basic	BasicEN	HVARC	Tapered
Frac. active coefficients			0.26	0.86	0.99	0.99
Mean MSFE	87.58	8.14	6.40	7.11	6.62	6.04
MSFE st dev	5.98	3.00	2.51	2.65	2.57	2.30

Table 6: Machine Learning Forecast Evaluation

"Basic": VAR coefficients selected by LASSO; "BasicEN": LASSO, but with an elastic net loss function; "HVARC": Component-wise lag-length; "Tapered": Lag-weighted LASSO

A benefit of a machine learning approach is that one can more easily evaluate the models themselves. Unlike with the FAVAR approach, where one hopes that one has "enough" factors, cross-validation means that within a given category of model one is likely picking a near-optimal specification. And by reserving a portion of the data for out-of-sample evaluation, different categories of models can be compared. Table 6 conducts such a comparison for the four machine learning models and two benchmarks – the simple average for each variable, and the AIC baseline we use above. In all cases, the machine learning methods have superior out-of-sample performance, with the tapered-lag LASSO the best. That said, the AIC baseline still performs remarkably well, with the best machine learning model offering only a 2.2 percent improvement in forecast accuracy relative to the naïve unconditional average.[37]

C.6 Oil Prices

In this section, we study how oil prices affect inflation sentiments.

We augment our baseline VAR by including inflation in the WTI oil spot price. Figure 20 reports the IRFs to an inflation sentiment shock. The additional variable quantitatively affects the shape and scale of some of the IRFs, but the qualitative conclusions are unchanged: the sentiment shock in deflationary, contractionary, and reduces interest rates.

However, an oil shock is a particularly interesting fundamental shock because it could itself be an endogenous driver of inflation sentiments. To investigate this possibility, Figure

[37] $(8.00 - 6.06)/87.78 \approx 2.2$ percent.

?? plots the responses to an oil price shock identified by causal ordering. The assumption is that oil shocks are the only fundamental shocks that affect oil prices, and we implement this assumption by letting the matrix A_c^F be lower triangular and ordering the oil shock first among fundamental shocks. This is, of course, not the only way to identify the causal impact of an oil price shock, but is a classic method that can also be implemented easily within our existing semi-structural VAR approach.

A positive oil shock is inflationary: it increases consumer prices (top right panel) and induces a monetary policy tightening (bottom middle panel). It also is associated with higher real activity, at least over the first year (center panel); this may be because oil prices are not completely exogenous to the US economy (Känzig, 2021), which is a large producer, refiner, and consumer. Thus, like the sentiment shock, an oil price shock can move all three of inflation, activity, and the interest rate in the same direction. So are our sentiment shocks actually just oil shocks? We think this is unlikely, because of how the oil shock affects inflation expectations.

Oil price shocks have a large effect on the inflation sentiment: after oil prices rise, households raise their inflation forecasts by even more than the rational forecast (middle left panel). Thus oil shocks can be a significant driver of endogenous inflation sentiments.[38] But an oil shock differs from our identified inflation sentiment shock in one important regard: an oil shock makes the sentiment and inflation move in the same direction, while the sentiment shock makes them move in *opposite* directions.

[38] A large literature has documented that inflation expectations are especially sensitive to petroleum prices, partially due to their salience. See for example Coibion and Gorodnichenko (2015b), Wong (2015), Coibion et al. (2018), Nasir et al. (2020), Känzig (2021), Hasenzagl et al. (2022) or Kim and Binder (2023)

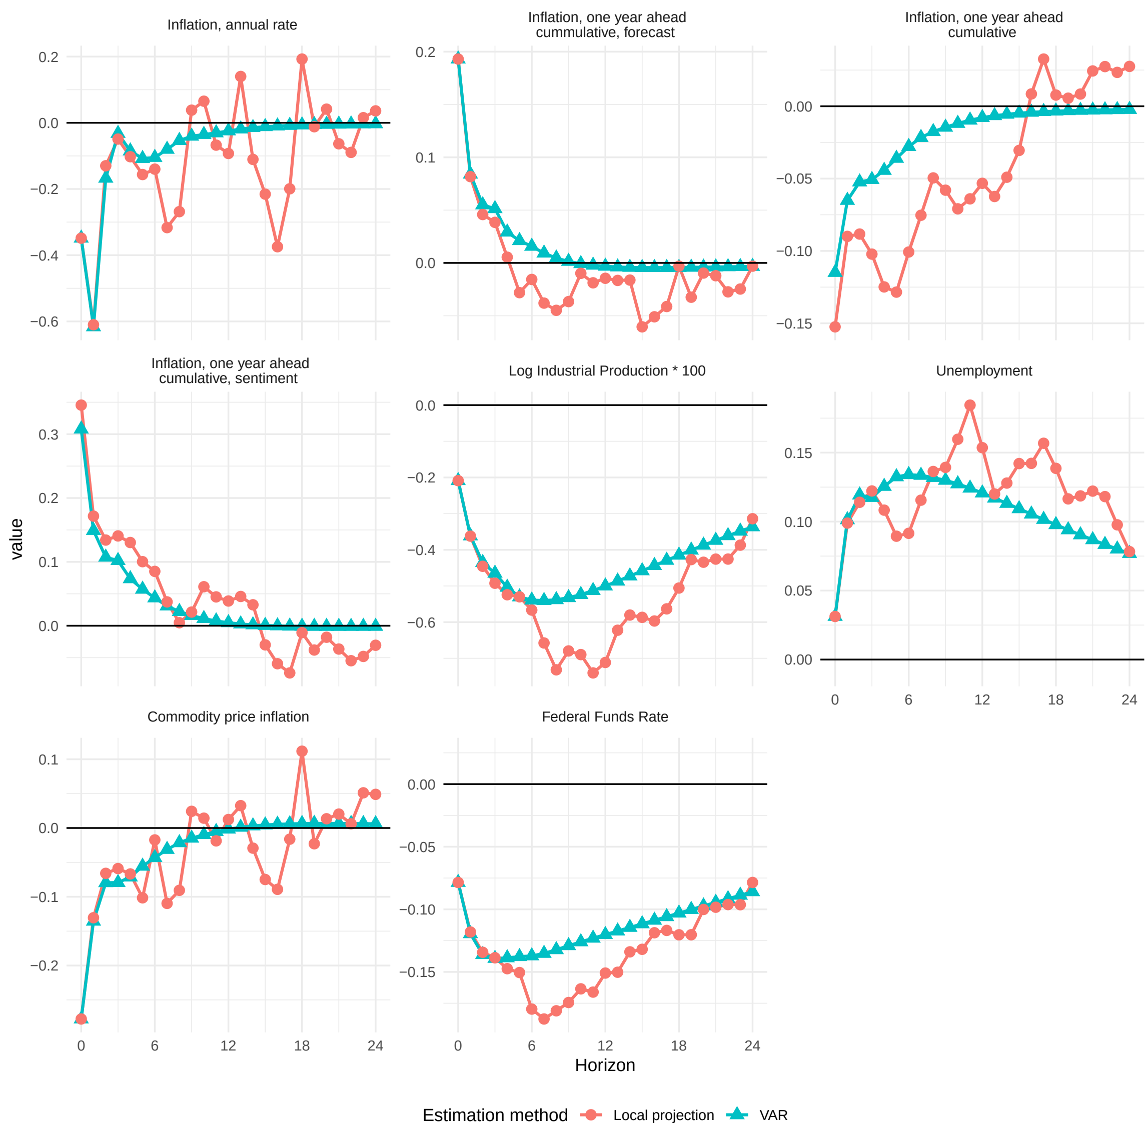

Figure 16: Local Projection versus VAR Estimates

Figure shows the baseline model point estimates along with the local projection point estimates from the equivalent (i.e. 3-lag) specification

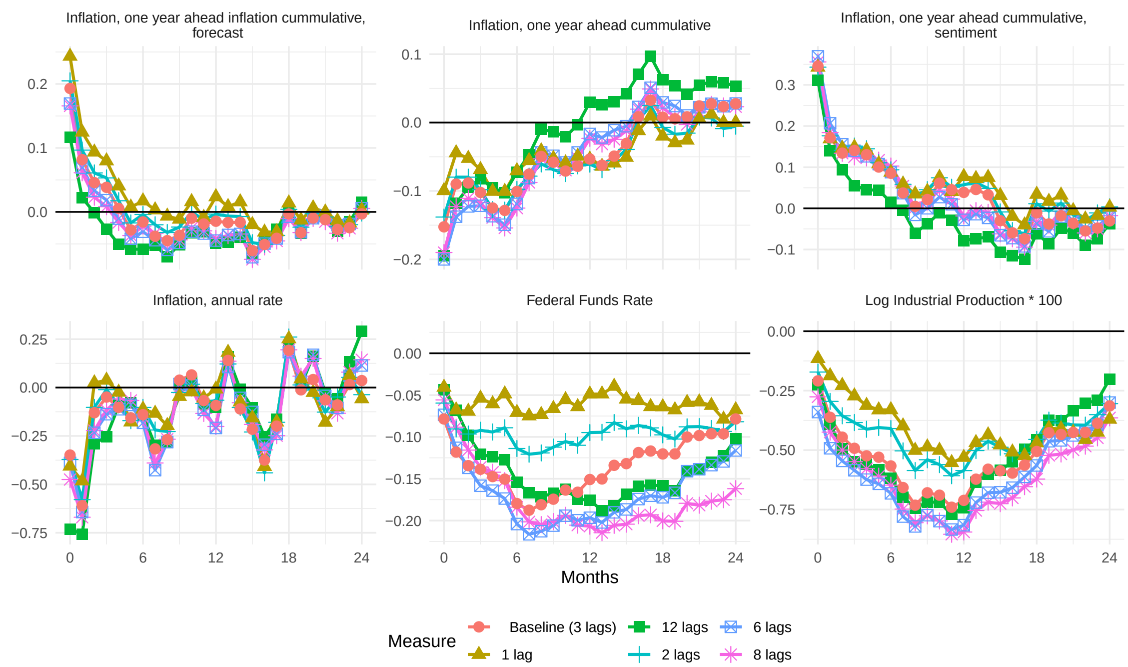

Figure 17: Local Projection Responses to a Sentiment Shock for Lag Structures

Figure the local projection point estimates from varied lag length specifications

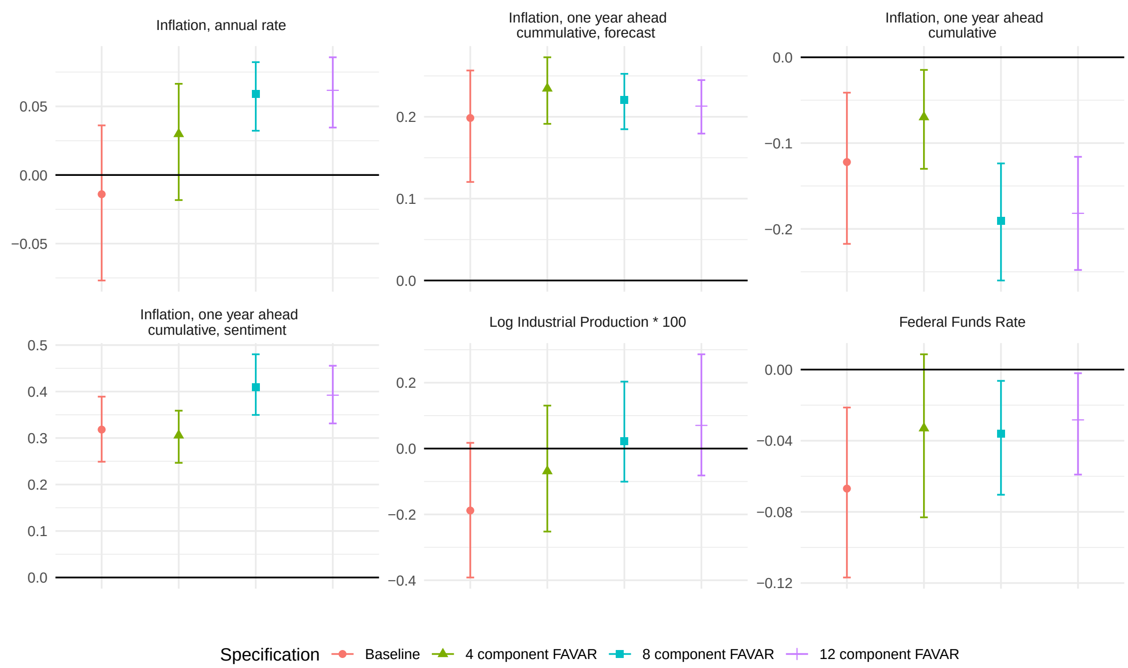

Figure 18: Sentiment Shock on Impact in Factor-Augmented VARs

Vertical bars show 90 percent confidence interval computed from 500 bootstrap replications.

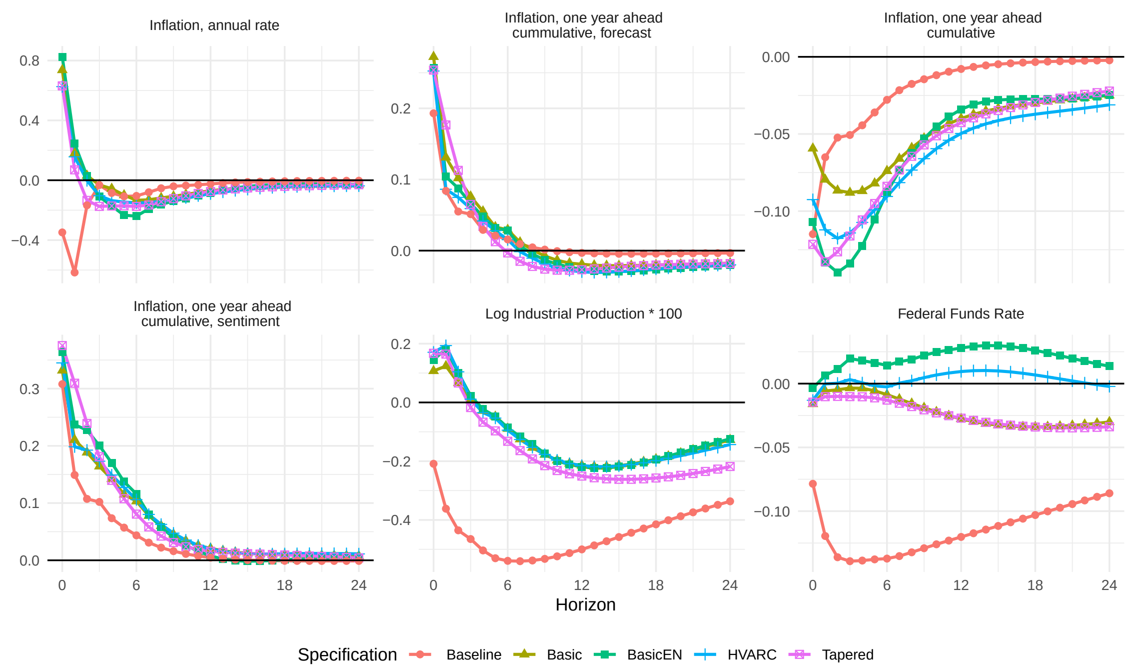

Figure 19: Model Selection by Machine learning

"Basic": VAR coefficients selected by LASSO; "BasicEN": LASSO, but with an elastic net loss function; "HVARC": Component-wise lag-length; "Tapered": Lag-weighted LASSO

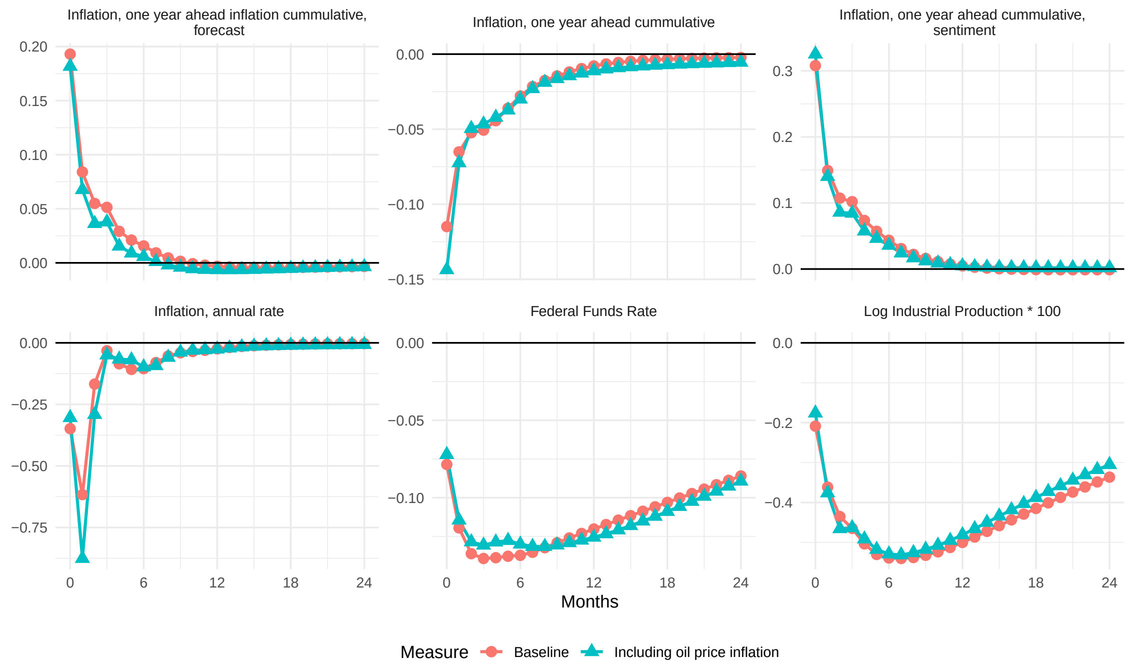

Figure 20: Impulse Responses to a Sentiment Shock with and without Oil in the VAR

Structural impulse responses to a one standard deviation sentiment shock. The compared estimation methods differ by whether they include oil price inflation as an extra variable in the VAR.

Figure 21: Response to an Oil Price Shock

Figure shows response of the VAR to a standard deviation "oil price shock", defined as the first fundamental shock when the fundamental block is ordered such that the oil price is first in the Cholesky decomposition A_c^F. Shaded ranges show a 90 percent confidence interval, computed from bootstrap with 500 replications. The cumulative inflation over the next 12 months (i.e. the average of horizons 1 to 12 of monthly annualized inflation) is the rational expectations component of inflation expectations. The inflation sentiment is the difference between the response for the inflation forecast and the rational component. The current month's realized inflation is reported at the annualized rate. IRF units are "percent points", which is literal for time series reported as such (e.g. inflation) and 100× log values for the remaining series.

D Departures from Full Information and Rational Expectations

In this section, we the investigate the extent to which strict full information rational expectations (FIRE) drive our results. In our baseline estimates, we identify sentiments as departures from FIRE. But if agents fundamentally form their expectations in a non-rational way, it may violate our identifying assumption that the sentiment shock is the only shock that can cause forecasts to deviate from the rational expectation.

To assess the sensitivity of our results to the FIRE assumption, we consider alternate identification strategies, treating sentiments as deviations from other commonly-used models of expectations. From these we conclude it is not strict FIRE that drives our results. Rather, they depend only on agents being somewhat forward-looking, even if they have very limited contemporaneous information or if they process that information irrationally. Our puzzle is overturned only in cases where agents are entirely backward-looking.

D.1 Identification with Non-Rational Expectations

We generalize equation (1), relating the sentiment ζ_t^ℓ to the measured forecast $\pi_t^{e,1}$ by

$$\pi_t^{e,1} = \tilde{\mathbb{E}}_t^\ell[\pi_{t+1}] + \zeta_t^\ell$$

where the operator $\tilde{\mathbb{E}}_t^\ell$ denotes a non-rational form of expectations, of type ℓ.

In Section 3 we calculated the rational expectation of inflation using the impulse response vector ϕ_π^h. More generally, we can construct an analogous vector $\phi_\pi^{\ell,h}$ the captures the ℓ-type expectation of the horizon h inflation response to reduced form shocks u_t such that

$$\mathbb{E}[\tilde{\mathbb{E}}_t^\ell[\pi_{t+h}^h]|u_t] = \phi_\pi^{\ell,h} u_t \tag{32}$$

The identifying assumption is now that *the type-ℓ sentiment shock $\varepsilon_t^{S,\ell}$ is the only contemporaneous shock that causes forecasts to deviate from the type-ℓ expectations.* This implies a restriction on the matrix A^ℓ analogous to equation (6).

$$A^\ell = \begin{pmatrix} * & (1-\phi_{\pi,f}^{\ell,h})^{-1}\phi_{\pi,c}^{\ell,h}A_c^{\ell,F} \\ * & A_c^{\ell,F} \end{pmatrix} \tag{33}$$

where $\phi_{\pi,f}^{\ell,h}$ and $\phi_{\pi,c}^{\ell,h}$ denote the coefficients on forecasts and contemporaneous variables respectively. As in equation (7), $*$ denotes unrestricted entries, of which there are $n+2$ in the first column. In what follows, we consider four common types of non-rational expectations, deriving the appropriate formula for $\phi_\pi^{\ell,h}$. These formulas are summarized in Table 33.

Note that equation (32) prefaces $\tilde{\mathbb{E}}_t^{\ell}[\pi_{t+1}]$ with the mathematical conditional expectation operator $\mathbb{E}$. This says that when we identify the shocks we are defining the sentiment as the shock which causes stated expectations to deviate from the assumed behavioral rule. That the impulse response is computed under the true distribution of outcomes says that we are deducing the dynamics of expected inflation under the actual law of motion, given that expectations are generated by the perceived law of motion described by $\tilde{\mathbb{E}}_t^{\ell}[\pi_{t+1}]$. Thus, there is scope for feedback between the perceived and actual law of motion when we identify the shock.[39]

Expectations	$\phi_\pi^{\ell,h}$ vector	Parameter Value
Rational Expectations	$\phi_\pi^h = \sum_{k=1}^h e_\pi B^k$	–
Delayed Observation	$\phi_{\pi,c}^{DO,h} = 0$	–
Partial Observation	$\phi_\pi^{PO,h} = \phi_\pi^h \mathbb{B}_{u,\tilde{z}} \mathbb{B}_{\tilde{z},u}$	–
Diagnostic Expectations	$\phi_\pi^{DE,h} = (1+\theta^{DE})\phi_\pi^h$	$\theta^{DE} = 0.55, -0.15$ (Bordalo et al., 2020)
Adaptive Expectations	$\phi_{\pi,f}^{AE,h} = 0 \quad \phi_{\pi,c}^{AE,h} = (1-\theta^{AE})e_\pi$	$\theta^{AE} = 0.91$ (Gelain et al., 2019)

Table 7: Restrictions for Each Type of Expectations

D.1.1 Delayed Observation

The first alternative model of expectations we consider is "delayed observation." These *DO*-type expectations assume that agents' expectations do not respond to any contemporaneous shock, relaxing the "full information" of FIRE. This may be because many macro time series are released with a delay or because agents simply do not respond to recent news.

With these expectations, the only shock that can cause forecasts to move unpredictably is the inflation sentiment shock. The is the structure assumed by Leduc et al. (2007); the response to sentiment shocks can now be identified by Cholesky decomposition of the variance matrix Σ. In our notation, *DO*-type expectations are defined by:

$$\tilde{\mathbb{E}}_t^{DO}[\pi_{t+h}^h] = \mathbb{E}_{t-1}[\pi_{t+h}^h]$$

The assumption for *DO*-type expectations is that the current forecast depends on past fundamental shocks and the current sentiment shock alone. This implies the restriction on A^ℓ that the submatrix containing the effect of current fundamental structural shocks on forecasts is $A_f^{\ell,F} = 0$. In the general expression (33) for the matrix restrictions, this is equivalent to assuming $\phi_{\pi,c}^{DO,h} = 0$.

[39] We thank a referee for pointing out the relevance of the distinction between the perceived and actual laws of motion.

D.1.2 Partial Observation

Although straightforward, the delayed observation assumption might be a little strict. Some important economic variables are observed immediately, while some that are not can be accurately inferred from from other high frequency data. A more realistic information set will contain some contemporaneous information.

We therefore also consider a modified version, "partial observation", where we continue to restrict agents from forecasting using the many time series that are only released with a delay (e.g. unemployment, CPI, etc.). But we allow them to formulate forecasts using contemporaneous data available at high frequency. The idea is that although forecasters may not know the current unemployment rate, they can observe current interest rates or other prices.

We stack the subset of contemporaneously observable high frequency variables into a vector z_t. And we continue to assume that forecasters in time t observe the prior states $\Omega_{t-1} = \{x_{t-1}, x_{t-2}, ...\}$. The PO-type expectation is given by:

$$\tilde{\mathbb{E}}_t^{PO}[\pi_{t+h}^h] = \mathbb{E}[\pi_{t+h}^h | z_t, \Omega_{t-1}]$$

Note that if z_t were empty, this expression would reduce to the DO-type expectation.

We estimate two versions of PO-type expectations. In the first, the only contemporaneous time series used in forecasts is the Federal Funds Rate. The Federal Funds Rate is one of the most widely discussed economic statistics, and rare in that is observable with no lag (unlike inflation or GDP). And even if not known precisely by all agents, it is correlated with costs which are directly observed – credit card debt, mortgages, student loan repayments, and even the general position of the business cycle. In the second version, we also let forecasters observe the commodity price index, because even though the index is released with a delay, components of the CPI correlated with it, such as gasoline or energy prices, are immediately observable.

When forecasters have partial observation, they rationally forecast using past time series and current high frequency variables z_t.

Let $\tilde{z}_t$ denote the unpredictable variation of z_t, i.e.

$$\tilde{z}_t = z_t - \mathbb{E}[z_t | \Omega_{t-1}]$$

Then the expectation becomes

$$\tilde{\mathbb{E}}_t^{PO}[\pi_{t+h}^h] = \mathbb{E}[\pi_{t+h}^h | \tilde{z}_t] + \mathbb{E}[\pi_{t+h}^h | \Omega_{t-1}]$$

$$= \mathbb{E}[\mathbb{E}[\pi_{t+h}^h | u_t] | \tilde{z}_t] + \mathbb{E}_{t-1}[\pi_{t+h}^h]$$

$$= \mathbb{E}[\phi_\pi^h u_t|\tilde{z}_t] + \mathbb{E}_{t-1}[\pi_{t+h}^h]$$

To find the relevant vector $\phi_\pi^{PO,h}$ project expectations onto the reduced form shock u_t:

$$\phi_\pi^{PO,h} u_t = \mathbb{E}[\tilde{\mathbb{E}}_t^{PO}[\pi_{t+h}^h]|u_t]$$

$$= \mathbb{E}[\mathbb{E}[\phi_\pi^h u_t|\tilde{z}_t]|u_t] + \mathbb{E}[\mathbb{E}_{t-1}[\pi_{t+h}^h]|u_t]$$

$$= \phi_\pi^h \mathbb{E}[\mathbb{E}[u_t|\tilde{z}_t]|u_t]$$

and if $\tilde{z}_t$ is spanned by u_t (e.g. if z_t only contains variables in the data vector x_t) then this vector reduces further to be $\phi_\pi^h \mathbb{E}[u_t|\tilde{z}_t]$. However, it may be useful to include other observable high frequency variables that forecasters might use to nowcast the current state of the economy (e.g. the stock market is informative about industrial production etc.)

Finally, let $\mathbb{B}_{u,\tilde{z}}$ and $\mathbb{B}_{\tilde{z},u}$ denote the matrices such that

$$E[u_t|\tilde{z}_t] = \mathbb{B}_{u,\tilde{z}}\tilde{z}_t \qquad E[\tilde{z}_t|u_t] = \mathbb{B}_{\tilde{z},u}u_t$$

Then the PO-type vector is given by

$$\phi_\pi^{PO,h} = \phi_\pi^h \mathbb{B}_{u,\tilde{z}}\mathbb{B}_{\tilde{z},u}$$

D.1.3 Diagnostic Expectations

We also relax the rational expectations part of FIRE. In this example, we assume that agents' (non-sentimental) inflation forecasts are formulated according to "diagnostic expectations." Popularized by Bordalo et al. (2018), DO-type expectations are given by

$$\tilde{\mathbb{E}}_t^{DE}[\pi_{t+h}^h] = (1+\theta^{DE})\mathbb{E}_t[\pi_{t+h}^h] - \theta^{DE}\mathbb{E}_{t-1}[\pi_{t+h}^h]$$

With this form, agents overreact (if $\theta^{DE} > 0$) or underreact (if $\theta^{DE} < 0$) to new information, but rationally respond to old information. Whether forecasters overreact or underreact depends on how the θ^{DE} is estimated. Therefore we conduct our analysis with two different parameter values estimated by Bordalo et al. (2020): they estimate $\theta^{DE} = 0.55$ using a simulated method of moments, and $\theta^{DE} = -0.15$ by running a regression following Coibion and Gorodnichenko (2015a).

To derive the DE-type vector $\phi_\pi^{DE,h}$, start with the definition, and project onto the reduced form shock u_t:

$$\phi_\pi^{DE,h} u_t = \mathbb{E}[\tilde{\mathbb{E}}_t^{DE}[\pi_{t+h}^h]|u_t]$$

$$= (1+\theta^{DE})\mathbb{E}[\mathbb{E}_t[\pi_{t+h}^h]|u_t] - \theta^{DE}\mathbb{E}[\mathbb{E}_{t-1}[\pi_{t+h}^h]|u_t]$$

$$= (1 + \theta^{DE})\phi_\pi^h u_t$$

This restriction vector is proportional to the rational expectations vector ϕ_π^h, with proportionality determined by the parameter θ^{DE}.

Furthermore, this restriction is equivalent (albeit with potentially dissimilar factors of proportionality) for any form of expectations that is a linear combination of current and past rational forecasts, including the sticky information introduced by Mankiw and Reis (2002) and the cognitive discounting studied by Gabaix (2020).

D.1.4 Adaptive Expectations

Before the rational expectations revolution, inflation forecasts were commonly modeled as the "adaptive expectations" of Cagan (1956) and Friedman (1957). Adaptive expectations are an entirely backward-looking heuristic. This contrast with diagnostic expectations which are non-rational but still forward-looking.

We define AE-type expectations as the traditional form:

$$\tilde{\mathbb{E}}_t^{AE}[\pi_{t+h}^h] = (1 - \theta^{AE})\pi_t^h + \theta^{AE}\tilde{\mathbb{E}}_{t-1}^{AE}[\pi_{t+h-1}^h]$$

Most estimated macro models assume a θ^{AE} close to 1, so that forecasts are nearly unaffected by changes to inflation. We adopt $\theta^{AE} = 0.91$ which Gelain et al. (2019) estimate in a structural DSGE New Keynesian model.[40]

AE-type expectations are entirely backwards looking. Because the rational expectation never appears, we can derive the implied restriction in the A^ℓ matrix directly.

The submatrix $A_f^{AE,F}$ denotes the effect of fundamental shocks on the forecast. Adaptive expectations assumes that this effect is proportional to the effect of fundamental shocks on current inflation:

$$A_f^{AE,F} = (1 - \theta^{AE})e_\pi A_c^{AE,F}$$

where e_π is the basis vector identifying the inflation dimension. To map this back to a vector $\phi_\pi^{AE,h}$, it implies that $\phi_{\pi,f}^{AE,h} = 0$ and $\phi_{\pi,c}^{AE,h} = (1 - \theta^{AE})e_\pi$

D.2 Results

Figure 22 reports the estimated immediate responses of the VAR to inflation sentiment shocks for each type of expectations. Results broadly fall into one of two categories.

[40] Other estimates (sometimes with different notation) include: Beladi et al. (1993) estimate $\theta^{AE} = 0.93$ in a present value model of hyperinflations, Christiano (1987) estimates 0.83 in a money demand model, Curtin (2010) estimates 0.89 in reduced form using the Michigan survey data.

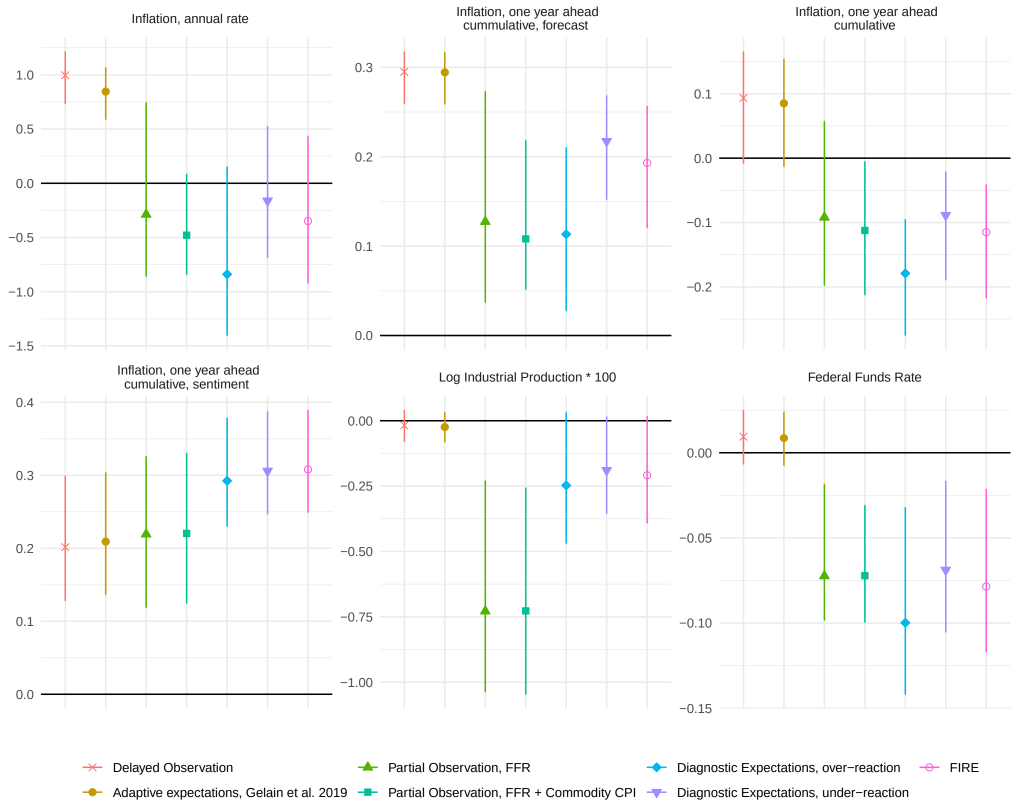

Figure 22: Sentiment Shock Effects for Different Types of Expectations

Each point is the estimate of the contemporaneous (i.e. period 0) impact of a sentiment shock for a different type of expectations. Error vars show a 90 percent confidence interval, computed from bootstrap with 500 replications. The "Predictable Inflation Error" is the difference between the measured forecast and the VAR-implied forecast; for rational expectations, this is the inflation sentiment. Diagnostic expectations "over-reaction" and "under-reaction" use $\theta^{DE} = 0.55$ and -0.15 respectively

First, all specifications with some degree of forward-looking expectations resemble the baseline. When forecasters have rational expectations (pink open circles), sentiments shocks are deflationary (top right panel), contractionary (bottom middle panel) and lower interest rates (bottom right panel). This is also the case for partial observation, where agents rationally forecast, but the only contemporaneous data they observe are high frequency variables: the Fed Funds Rate (green upward pointing triangles) and additionally with commodity prices (teal squares). When agents have diagnostic expectations (blue diamonds and purple downwards pointing triangles), sentiment shocks are also deflationary, contractionary, and reduce interest rates.

Second, the two strictly backwards-looking specifications are dissimilar from our baseline. When agents' forecasts cannot be affected by contemporaneous structural shocks (delayed observation, red X's) or can only be affected through contemporaneous inflation in a fixed functional form (adaptive expectations, orange closed circles), then the implied sentiment shocks are inflationary, and neither contractionary nor expansionary. These results reconcile our findings with those of Leduc et al. (2007), who estimate the delayed observation specification and also find that shocks to expectations are inflationary. Why is the adaptive expectations specification similar? The identifying assumption for delayed observation is that the shock-impact matrix A is lower triangular, and adaptive expectations imply that A is nearly triangular by assuming that innovations to current inflation are the only other driver of forecast innovations. Moreover, the calibration of $1 - \theta^{AE} = 0.09$ implies that the effect of inflation innovations on forecasts is small.

There are two possible interpretations of these findings. One is that our puzzle is robust. Weakening either the full information part of FIRE (the partially-observed case) or the rationality part (diagnostic expectations) preserves the puzzle. This view says that our results do not rely on agents' superhuman forecasting powers. Rather, agents need only be somewhat forward-looking: they might not use all relevant information or they could overreact or underreact to news and the puzzle still stands. The second interpretation is that there is no puzzle, just misidentification: either agents simply do not update their forecasts with new data (delayed observation) or they do not update their forecasts correctly (adaptive expectations).

Which interpretation is correct not a question we can answer here: our method cannot jointly estimate the model for expectations and the sentiment process. Instead, the correct interpretation is a matter of which assumptions one finds more plausible. That said, both ways of looking at our results agree on one thing. Identification under the assumption in the canonical model – full information rational expectations – does not deliver the expected positive multiplier on inflation.

E Implementing Noise Shocks

The dynamic New Keynesian model features "noise shocks". At time t, agents observe the noisy signal v_t of the future productivity shock ε^a_{t+1}:

$$v_t = \varepsilon^a_{t+1} + \nu_t$$

where ν_t is the i.i.d. noise shock. But this is not an equilibrium condition; rather, it is a modification of agents' information sets. How can these shocks be applied to standard solution methods that implicitly assumes expectations are projections onto current and past fundamental shocks?

When solving the model, we recast the information structure as one where agents receive incomplete news about the future, instead of receiving noisy signals. Chahrour and Jurado

(2018) prove that these noisy information structures have equivalent news structures. In our case, we can model the productivity shock as a sum of i.i.d. components:

$$\varepsilon^a_{t+1} = \vartheta_{t+1} + \varsigma_t$$

where ϑ_{t+1} is unanticipated but ς_t is known in period t.

To map from noise to news, we have to match the variance of the forecastable component of future productivity shocks. For the news structure, this is simply:

$$Var(E_t[\varepsilon^a_{t+1}]) = Var\left(\vartheta_{t+1}\right)$$

While for the noise structure, this is:

$$Var(E_t[\varepsilon^a_{t+1}]) = Var\left(\frac{Cov(v_t, \varepsilon^a_{t+1})}{Var(v_t)} v_t\right) = Var\left(\frac{Var(\varepsilon^a_{t+1})}{Var(v_t)} v_t\right)$$

$$= \frac{Var(\varepsilon^a_{t+1})^2}{Var(v_t)} = \frac{Var(\varepsilon^a_{t+1})^2}{Var(\varepsilon^a_{t+1}) + Var(\nu_t)}$$

So the news representation of our noise structure is the one where

$$Var\left(\vartheta_{t+1}\right) = \frac{Var(\varepsilon^a_{t+1})^2}{Var(\varepsilon^a_{t+1}) + Var(\nu_t)}$$

and the variance of forecast errors gives the remaining parametrization

$$Var(\varsigma_t) = Var(\varepsilon^a_{t+1}) - Var\left(\vartheta_{t+1}\right)$$

F Solving the Model

This appendix describes how we solve the dynamic New Keynesian model in Section 6 with sentiment shocks. We implement the method using the BEET toolkit (Adams, 2024), which is a wrapper for the Uhlig (2001) toolkit, so the following description uses Uhlig's same matrix notation.

The vector $\vec{x}_t$ stacks the 6 endogenous variables (π_t, y_t, i_t, $y^{e,1}_t$, $\pi^{e,1}_t$, $\pi^{e,12}_t$) while the vector $\vec{z}_t$ stacks the 7 exogenous variables (ζ^π_t, ζ^y_t, ζ^a_t, a_t, x_t, z_t, ε^a_t).

In equilibrium, $\vec{x}_t$ satisfies the matrix equation

$$0 = \tilde{\mathbb{E}}_t\left[F\vec{x}_{t+1} + G\vec{x}_t + H\vec{x}_{t-1} + L\vec{z}_{t+1} + M\vec{z}_t\right] \tag{34}$$

where the matrices (F, G, H, L, M) encode the 6 equilibrium conditions (14) - (19).

The vector $\vec{z}_t$ follows the law of motion

$$\vec{z}_t = N\vec{z}_{t-1} + \vec{\epsilon}_t \tag{35}$$

where the matrix N encodes the 7 laws of motion (20) - (26), and $\vec{\epsilon}_t$ is the vector of shocks $(\varepsilon_t^{\zeta,\pi}, \varepsilon_t^{\zeta,y}, \varepsilon_t^{\zeta,a}, \varepsilon_t^x, \varepsilon_t^z, \vartheta_t, \varsigma_t)$.

In equation (34), $\tilde{\mathbb{E}}_t[\cdot]$ is a subjective expectations operator. It differs from the rational expectations operator $\mathbb{E}_t[\cdot]$ in one way:

$$\tilde{\mathbb{E}}_t[a_{t+1}] = \mathbb{E}_t[a_{t+1}] + \zeta_t^a$$

Expectations are rational when forecasting every other series *including* y_{t+1} *and* π_{t+1}, because we directly coded distortions to those forecasts in equations (14) and (15). But, sentiments about underlying exogenous states such as a_t do not have such a simple representation.

A solution to the model is a recursive law of motion

$$\vec{x}_t = P\vec{x}_{t-1} + Q\vec{z}_t$$

Plugging into equation (34) gives

$$0 = \tilde{\mathbb{E}}_t\left[FP\vec{x}_t + FQ\vec{z}_{t+1} + G\vec{x}_t + H\vec{x}_{t-1} + L\vec{z}_{t+1} + M\vec{z}_t\right]$$

Let $\vec{\zeta}_t$ denote the vector of sentiments about exogenous variables (i.e. zeros everywhere except ζ_t^a in the a dimension) so that

$$\tilde{\mathbb{E}}_t[\vec{z}_{t+1}] = \mathbb{E}_t[\vec{z}_{t+1}] + \vec{\zeta}_t = N\vec{z}_t + \vec{\zeta}_t$$

Plugging this expectation into the equilibrium condition gives

$$0 = FP\vec{x}_t + FQN\vec{z}_t + FQ\vec{\zeta}_t + G\vec{x}_t + H\vec{x}_{t-1} + LN\vec{z}_t + L\vec{\zeta}_t + M\vec{z}_t$$

and using the recursive law of motion again gives

$$0 = FP^2\vec{x}_{t-1} + FPQ\vec{z}_t + FQN\vec{z}_t + FQ\vec{\zeta}_t + GP\vec{x}_{t-1} + GQ\vec{z}_t + H\vec{x}_{t-1} + LN\vec{z}_t + L\vec{\zeta}_t + M\vec{z}_t$$

Let Z denote the matrix isolating the sentiments about exogenous states so that $\vec{\zeta}_t = Z\vec{z}_t$. Substitute this in and collect terms:

$$0 = \left(FP^2 + GP + H\right)\vec{x}_{t-1} + \left(FPQ + FQN + FQZ + GQ + LN + LZ + M\right)\vec{z}_t$$

Finally, the matrices P and Q must satisfy the two equations

$$0 = FP^2 + GP + H \qquad\qquad 0 = FPQ + FQ(N + Z) + GQ + L(N + Z) + M$$

Thus we have translated the model into a form that can be solved as in Uhlig (2001). The main difference from a rational expectations model is that the matrix $(N + Z)$, encoding the *perceived* law of motion, takes the place of N in Uhlig's formulation.

G Alternative Treatments of Inflation and Forecasts

In this appendix, we address concerns about the relationship between the inflation and forecast series.

First, the CPI time series has regular small revisions, and forecasters might be targeting the initial inflation release rather than the revised released. To determine whether our conclusions are robust to this concern, we rerun the baseline VAR with the initial CPI inflation releases as reported by the Philadelphia Fed. Figure 23 reports the implied impulse response functions (blue squares). They are largely unchanged from the baseline specification (red circles).

Second, to be consistent we detrend and deseasonalize all of our time series individually. But if agents do not have FIRE, their detrended forecast will not necessarily be the same as their forecast of a detrended variable. To determine whether our conclusions are robust to this concern, we rerun the baseline VAR, except we transform the forecast series using the trend and seasonal factors from the true inflation series. Figure 23 reports the implied impulse response functions (green triangles). Again, they are qualitatively unchanged from the baseline VAR.

H Time Series Plots

In Figures 24-28 we plot the time series data – both unadjusted as well as deseasonalized and detrended – used in our various empirical exercises.

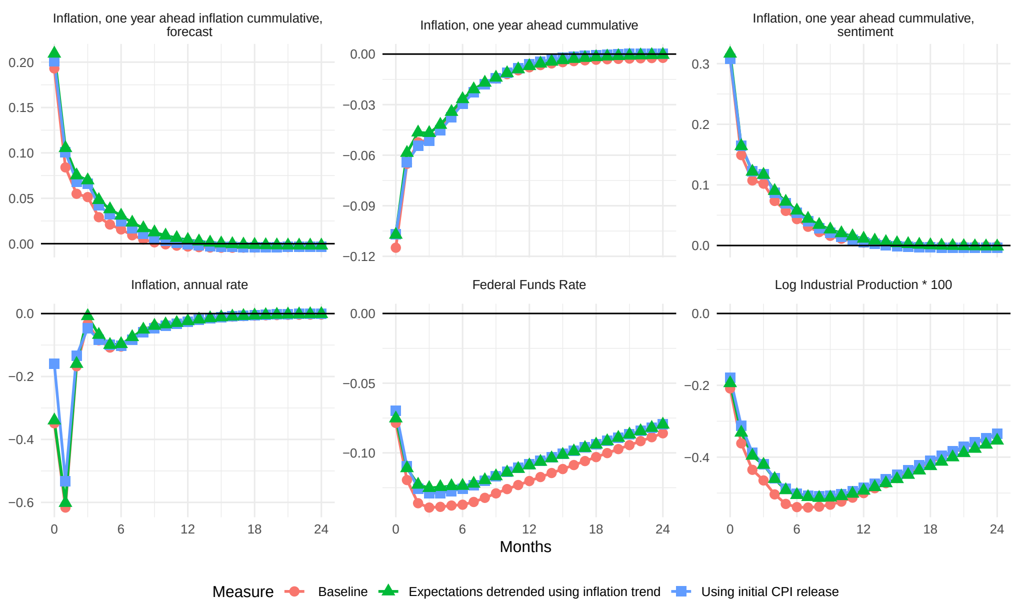

Figure 23: Impulse Responses with Alternative Treatments of Inflation and Forecasts

Figure compares the baseline results in the paper to two different versions that the referee suggests: using the initial CPI release as an approximation to real-time CPI data, and applying common seasonal factors and trend to expectations and inflation.

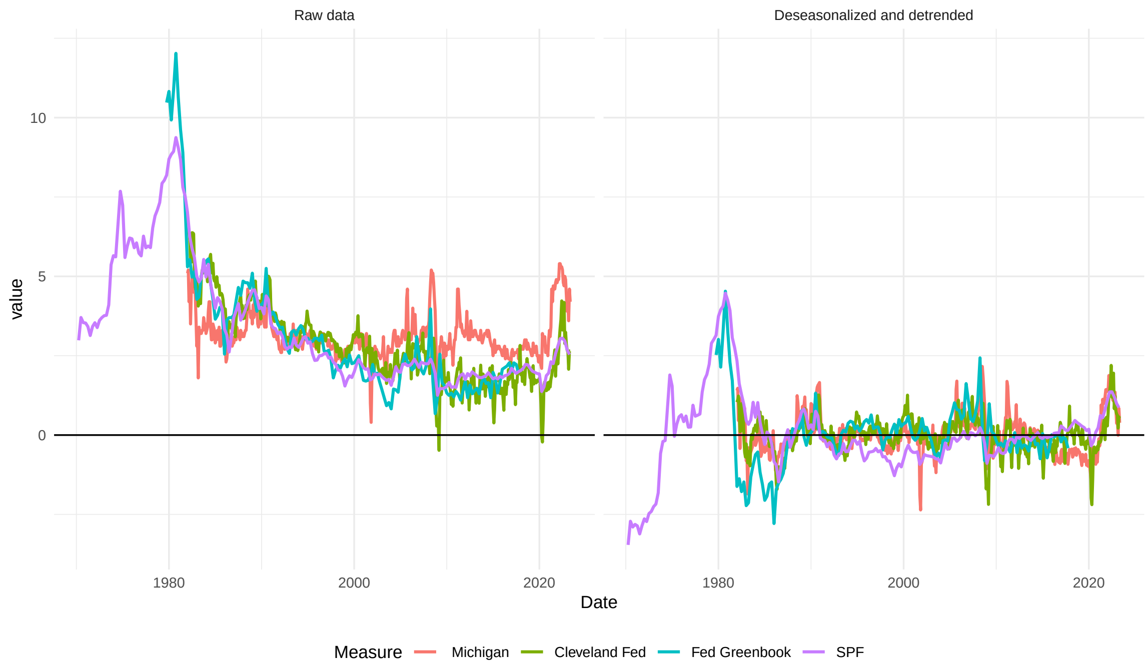

Figure 24: Time Series of Inflation Forecasts

Figure 25: Time Series of Interest Rates

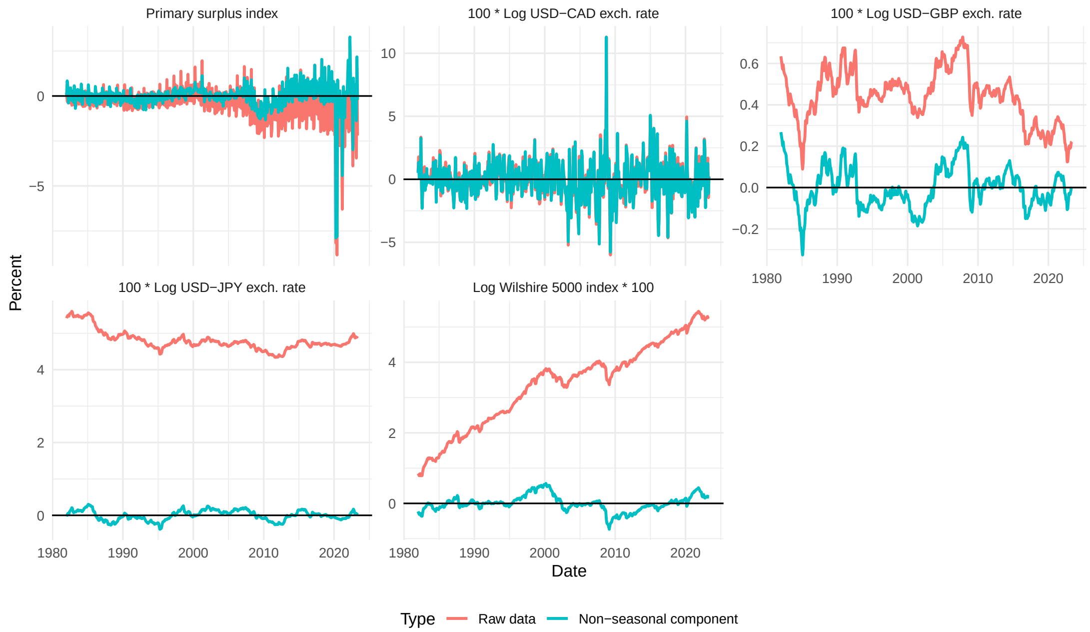

Figure 26: Time Series of Financial Variables

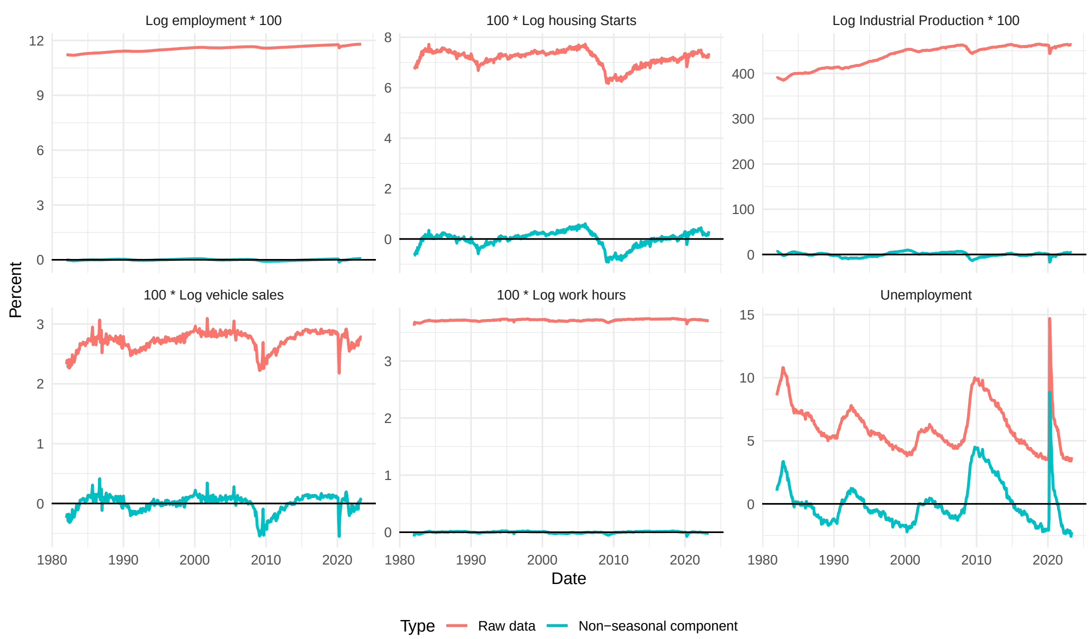

Figure 27: Time Series of Real Activity variables

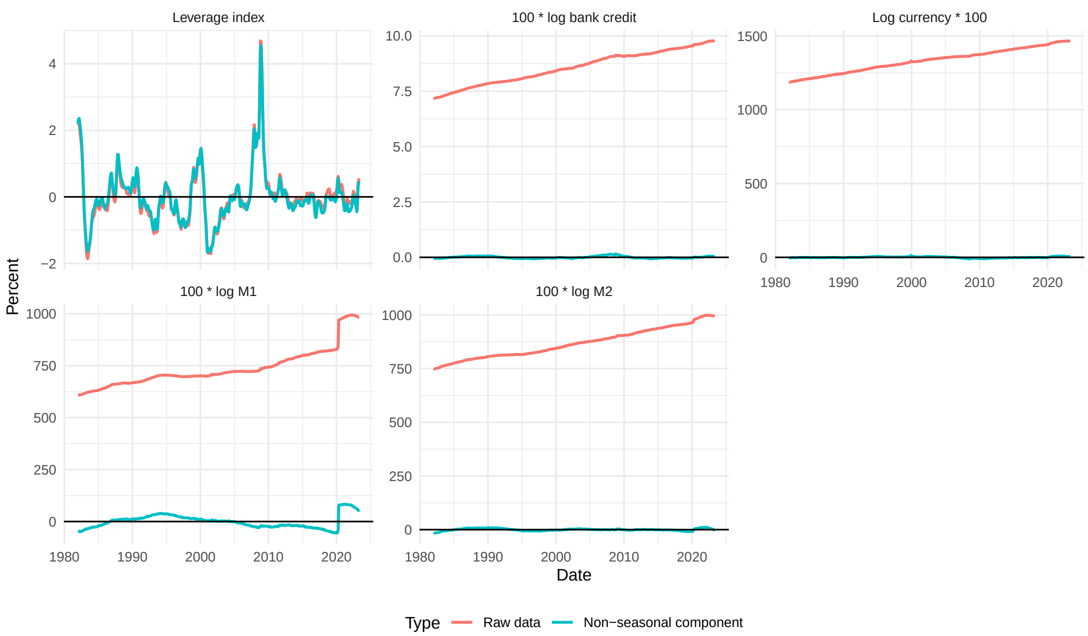

Figure 28: Time Series of Money and Credit Variables