Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Soft Landing or Stagflation? A Framework for Estimating the Probabilities of Macro Scenarios

Eric Engstrom

2025-047

Please cite this paper as:
Engstrom, Eric (2025). "Soft Landing or Stagflation? A Framework for Estimating the Probabilities of Macro Scenarios," Finance and Economics Discussion Series 2025-047. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.047.

Soft Landing or Stagflation?

A Framework for Estimating the Probabilities of Macro Scenarios

2025

Federal Reserve Board[1]

Introduction

Recent changes in trade policy and ongoing global economic developments have introduced new sources of uncertainty into the U.S. macroeconomic outlook. These developments—ranging from shifts in tariff schedules to geopolitical conflict—may affect both inflation dynamics and the trajectory of real economic activity. Against this backdrop, forecasters and policymakers may find it useful to assess the probabilities of various macroeconomic scenarios. Two salient scenarios at the time of this writing are a soft landing, in which inflation declines to target levels of around 2 percent while growth remains positive, and a stagflationary episode, characterized by elevated inflation alongside subdued or negative output growth.

This paper uses a quasi-structural framework to quantify the relative likelihood of these two scenarios over a four-quarter horizon. The analysis builds on the approach developed by Bekaert, Engstrom, and Ermolov (2025, BEE henceforth)[2], which generates joint density forecasts for key macroeconomic variables by modeling the underlying structural supply and demand shocks. This methodology is particularly well suited to the present environment for several reasons. First, it enables a decomposition of observed macroeconomic variation into supply-driven and demand-driven components, allowing a clearer interpretation of inflation and output risks. Second, the model incorporates time-varying, asymmetric risks, capturing important features such as fat tails and skewness that are often present during periods of heightened uncertainty. Third, the model delivers joint predictive distributions—rather than just point forecasts—allowing for probabilistic assessments of scenarios such as stagflation or soft landings.

[1] The views expressed in this document do not represent those of the Federal Reserve System, its Board of Governors, or staff.

[2] Bekaert, Geert, Eric Engstrom, and Andrey Ermolov. "Uncertainty and the Economy: The Evolving Distributions of Aggregate Supply and Aggregate Demand Shocks." AEJ: Macroeconomics forthcoming.

These distributional forecasts may be informative for both financial market participants and policymakers. Market participants can use them to better understand the balance of risks around inflation and growth, with implications for pricing interest rate derivatives, inflation-linked assets, and macro-sensitive equities. Policymakers, particularly those operating under dual mandates, may benefit from tools that clarify whether risks to inflation and growth mandates are aligned, or in conflict. In that sense, this framework helps quantify potential policy tradeoffs under evolving macroeconomic conditions.

The main results of this paper are estimates of the relative probabilities of a stagflationary scenarios versus a soft landing in the U.S., and how those probabilities have evolved over time. To preview, the model estimates that the probability of a stagflation scenario was substantially elevated coming out of the pandemic. For instance, at the end of 2022 after inflation had reached peak levels and the Federal Reserve had begun an aggressive path of raising interest rates, the probability of at least mild stagflation –defined as inflation exceeding 3 percent while real GDP growth registers below 1 percent on a four-quarter basis—is estimated to have been about 35 percent. At the time, the probability of a soft landing over the coming four quarters–defined as inflation returning to near 2 percent while growth remained solid— was estimated to be low, below 5 percent. Over the next couple of years as growth remained solid and inflation slowly declined towards the Federal Reserve's 2 percent goal, the model-implied probability of stagflation is estimated to have fallen steadily by the end of 2024 to about 5 percent, while the probability of a soft landing rose to about 30 percent. By mid-2025, however, that trend had reversed, with the probability of at least mild stagflation rising precipitously and the probability of a soft landing declining notably. The shift in 2025 can be attributed mostly to the uncertain effects of tariffs on the outlooks for inflation and growth. However, the probability of severely stagflationary scenarios did not increase materially.

Summary of Methodology

A detailed description of the methodology, data, estimation, and validation techniques used in this analysis is provided in the appendix. A brief summary is presented here for readers who may not require the full technical detail.

The methodology aims to produce predictive distributions for key U.S. macroeconomic variables—specifically, four quarter headline and core inflation, four quarter real GDP growth, and the unemployment rate—at a four-quarter horizon. The approach combines time series forecasting, structural economic identification, and non-Gaussian modeling of uncertainty. The method follows and extends the framework of BEE, proceeding in three primary steps.

Step 1: Forecasting Means and Identifying Reduced-Form Shocks

The first step estimates the conditional mean forecasts for the four macro variables using ordinary least squares (OLS) regressions. These models incorporate two predictors for each macro series: (i) the four-quarter lag of the dependent variable, and (ii) the corresponding Survey of Professional Forecasters (SPF) forecast. This design thus leverages both historical data and forward-looking expectations.[3] In addition to providing conditional mean forecasts, the OLS residuals from these projections represent reduced-form forecast errors, which capture unexpected deviations from predicted paths.

Step 2: Identifying Underlying Structural Shocks

The second step maps the reduced-form shocks into four structural shocks using Keynesian sign and exclusion restrictions. The identified shocks include aggregate supply shocks, demand shocks, and two idiosyncratic shocks: one for headline inflation and one for the unemployment rate. Consistent with Keynesian intuition, supply shocks are assumed to move inflation and growth in opposite directions (e.g. oil price shock), while demand shocks affect both in the same direction (e.g. financial crisis). The model assumes a linear structure where reduced-form shocks load onto the underlying structural shocks. The loading parameters are estimated using a classical minimum distance (CMD) estimator, matching second-moment statistics (variances and correlations) of the reduced-form shocks.

[3] As documented in the appendix, SPF forecasts tend to be very significant predictors for macro variables. Moreover, the mean predictions in our framework rarely deviate substantially from the corresponding raw SPF forecast. To the extent that the model's mean forecasts coincide with the SPF, one could regard the value-added of our framework as its ability to estimate time-varying and non-Gaussian distributions around the SPF forecasts.

Step 3: Modeling the Conditional Distributions of Structural Shocks

The third step models the conditional distribution of each structural shock using the Bad Environment–Good Environment (BEGE) distribution, which allows for time-varying volatility, skewness, and kurtosis. The BEGE framework assumes that each structural shock is composed of two asymmetric components—a "good" and a "bad" shock—each drawn from a centered gamma distribution. The variance and skewness of the shocks depend on time-varying shape parameters for the gamma distribution, which evolve according to GARCH-type dynamics.[4] The model's structure allows for dynamic asymmetries in the risks to all the structural shocks, with persistent variation in both upside and downside risk.

To evaluate model fit and select between specifications, the framework relies on a combination of in-sample and out-of-sample criteria. The final model generates full multivariate predictive distributions for inflation, growth, and unemployment by integrating the forecast means from step 1 with the time-varying distribution of structural shocks. These distributions are used to compute not only central forecasts but also the probability of outcomes falling into specific macroeconomic scenarios—such as stagflation or a soft landing.

Results

The results from the estimation of the model, including parameter estimates, inference and model validation are relegated to the appendix. Here we present summary results from the model regarding the evolution of the outlooks for real GDP growth and inflation over the past two and a half years.

The first set of results that we present is the evolution of the multivariate predictive distribution of inflation and real GDP growth. In Figure 1 we plot bivariate distributions for real GDP growth and headline inflation. Each panel depicts predictions from the model at different points in time. In each panel, forecasts for four-quarter inflation are plotted on the horizontal axis, and forecasts for four-quarter real GDP growth rate are on the vertical axis. The asterisk denotes the model's mean

[4] The two time-varying shape parameters of the BEGE distribution govern time-variation in the variance, skewness, kurtosis and all higher-order moments of the shocks. This restriction implies some dependence in time-variation across the moments.

point forecast and the square denotes the mode of the distribution. The yellow area indicates the most likely outcomes, with 50 percent of all outcomes expected to fall within it. The green area depicts the next most likely 40 percent of outcomes, which we interpret as a region of material risk.[5] For reference, axes are drawn at 2 percent for inflation and 2 percent for real GDP growth. The model-implied probabilities of outcomes in each of the four resulting quadrants are listed in the corner of each quadrant.

We begin in the upper left panel, at the end of 2022 when the economy was still suffering from the aftereffects of the pandemic. Much higher than expected readings on headline inflation had prevailed for several quarters, with a peak rate of over 7 percent earlier in the year. Economic growth had remained solid but the Federal Reserve had embarked on an aggressive campaign of increases in the target range for the federal funds rate in order to bring demand and supply into better balance and to attempt to move inflation toward the Fed's 2 percent goal. As indicated by the hollow square, the model's modal expectation at the time was for inflation to moderate over the next four quarters, but only to around 4 percent, and for real GDP growth to remain around 2 percent. To a large extent, the model's modal predictions reflect survey expectations at the time. However, uncertainty was extremely elevated, with the model seeing substantial risk of stagflation scenarios, as indicated by the long green tail in the southeast quadrant, where inflation outcomes are greater than 2 percent and real GDP growth is below 2 percent. Although not all outcomes in that quadrant may be considered severe enough to merit the label of stagflation, the probability of landing in the southeast quadrant was estimated to be 54 percent. What is going on here is that, based on recent expereince of adverse supply shocks and the historical tendency for large shocks to be followed by additional large shocks, the model saw elevated risk of further adverse supply shocks that could push up on inflation and down on real activity. At the same time, the model had also observed recent large positive demand shocks, which it interpreted as indicating increased risk of demand shocks going forward. To see this, note the more modest but still notable green tail of the distribution in the southwest quadrant. The model does not offer economic narratives to go along with its worrisome predictions but it is easy to imagine that a worsening of supply chain

[5] There are an infinite number of ways to partition a distribution to isolate a fixed probability mass. We use the partition that encloses a fixed probability mass (e.g. 50 percent) in the smallest possible region, sometimes called the highest density region. This partition is unique under certain regularity conditions (e.g., a unimodal and continuous pdf).

disruptions or a resurgence of Covid may have been top of mind at the time, which could be compounded by large demand shocks.

We now fast forward one year to the end of 2023, with the model implied results for that period illustrated in the upper right panel of Figure 1. For background, inflation had indeed moderated over the year but remained nearly 3 percent. Real GDP growth had accelerated a bit, registering nearly 3 percent. As indicated by how the yellow portion of the distribution is centered near the "cross hairs" of the axes, both inflation and real GDP growth were expected to register around 2 percent looking forward under the model. However, the uncertainty bands were still wide. A still-notable tail extended deep into the southeast quadrant characterized by stagflation, with the predicted probability of landing in that quadrant in four quarters' time around 33 percent. What is going on here is that both recent economic data and survey data pointed to inflation and real GDP growth registering around 2 percent, and the large supply and demand shocks of 2021 and 2022 were further in the rear view mirror, so the model estimated that the probabiltiy of large shocks going forward had declined somewhat but remained elevated.

We now skip forward once again to the end of 2024, with the model-implied distribution for that period illustrated in the lower left panel. Over the year, inflation had continued to moderate, on net, registering 2.5 percent, still somewhat above the Fed's target. Real GDP growth had remained solid, registering around two percent. The model expected those benign conditions to maintain, with both real GDP growth and inflation expected to register around 2 percent over the subsequent four quarters. It had been quite some time since the period of large supply and demand shocks, and the model reacted by seeing reduced uncertainty in the outlook; the confidence bands by the end of 2024 had shrunk to relatively subdued levels by historical standards. The probabilities of landing in each of the four quadrants were roughly balanced, around 20-30 percent.

The outlook shifted once again by mid 2025. The model's modal expectation was for inflation to register around 3 percent, a reacceleration relative to recent data, and for growth to fall to around 1 percent. This shift reflects survey respondents lowering their forecasts for growth and increasing their forecasts for inflation, with the prevailing narrative from private sector forecasters pointing the to the expected effects of increases in tariff rates. Market commentary also focused on the large amount of uncertainty accompanying the potential effects of tariffs. However, the width of the

model's uncertainty bands remained modest; Since large shocks to "hard" data regarding inflation or GDP had not occurred recently, the model saw no reason to anticipate higher uncertainty going forward. Still, the probability of landing in the southeast "stagflation" quadrant moved up to 59 percent, higher even than the level at the end of 2022. That said, much of the mass of the distribution is near the crosshairs of the axes, suggesting that extreme stagflation scenarios with very high inflation and/or very low growth may not be particularly likely under the model. We refine our estimates of the probabilities of mild versus severe stagflation scenarios using Figure 2.

Figure 1: Conditional Bivariate Distribution for Real GDP Growth and Inflation

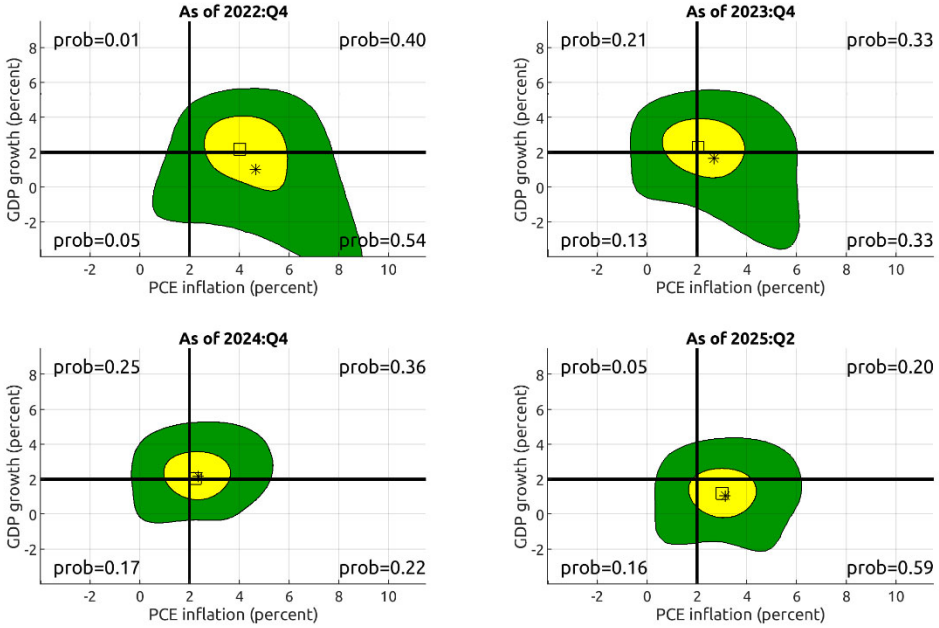

Notes: Estimated predictive distribution of headline inflation and real GDP outcomes using the model developed in Step1-3 in the appendix. Each panel depicts predictions from the model at different points in time. In each panel, forecasts for four-quarter inflation are plotted on the horizontal axis, and forecasts for four-quarter real GDP growth rate are on the vertical axis. The asterisk denotes the model's mean point forecast, and the square denotes the mode. The yellow area indicates the most likely outcomes, with 50 percent of all outcomes expected to fall within it. The green area depicts the next most likely 40 percent of outcomes, which we interpret as a region of material risk. For reference, axes are drawn at 2 percent for inflation and 2 percent for real GDP growth. The model-implied probabilities of outcomes in each of the resulting four quadrants are listed in the corner of each quadrant.

Figure 2 depicts the probabilities of some specific scenarios using the same four quarter ahead probability distributions that are shown in Figure 1. We formally define a "soft landing" scenario as four-quarter inflation registering between 1½ percent and 2½ percent while real GDP growth registers greater than 1 percent. We define the scenario "severe stagflation" as four-quarter

inflation registering over 4 percent and four quarter real GDP growth registering less than 0 percent. We define "mild stagflation" as four-quarter inflation registering over 3 percent and four quarter real GDP growth registering less than 1 percent but excluding severe stagflation scenarios. As shown by the green bars, the probability of a soft landing within four quarters was very small in late 2022 –less than 5 percent. However, amid a continuing recovery of supply conditions and better balance between supply and demand, that probability of a soft landing grew steadily through the end of 2024 to about 30 percent. However, the probability of a soft landing fell notably through the first half of 2025 amid concerns about the effects of tariffs on inflation and real activity. The probabilities of severe and mild stagflation are shown by the brown and red bars, respectively. These probabilities moved opposite to those of a soft landing, falling steadily from high levels at the end of 2022 to very low levels at the end of 2024. Interestingly, though the probability of a mild stagflation rose sharply in 2025 to levels even higher than were seen at the end of 2022, the probability of a severe stagflation rose by comparatively less. This pattern is consistent with a notable but modest shift in modal expectations for inflation and real activity as a result of the tariffs, and the fact that "hard" data have thus far shown no signs of elevated volatility that would indicate, in the context of the model, a sharp increase in downside supply shock risk.

These shifts in scenario probabilities may have useful implications for both policymakers and financial market participants. For policymakers—particularly those operating under dual mandates—an elevated probability of stagflation could underscore the potential for conflicting objectives, where efforts to support growth could exacerbate inflation pressures, and vice versa. In such an environment, a flexible approach to policy setting and communication may be preferred. For market participants, the reemergence of stagflation risk alters the perceived distribution of future macro-outcomes, which may warrant repricing of assets sensitive to inflation and growth conditions, such as nominal and inflation-linked bonds, equity sectors with exposure to inflation, and interest rate derivatives. The lower probability of a soft landing reinforces the asymmetry in current macro risks, suggesting that hedging strategies and forward-looking risk assessments should account for a tail risk environment in which inflation remains elevated even as growth softens.

Figure 2: Probabilities of Macroeconomc Scenarios

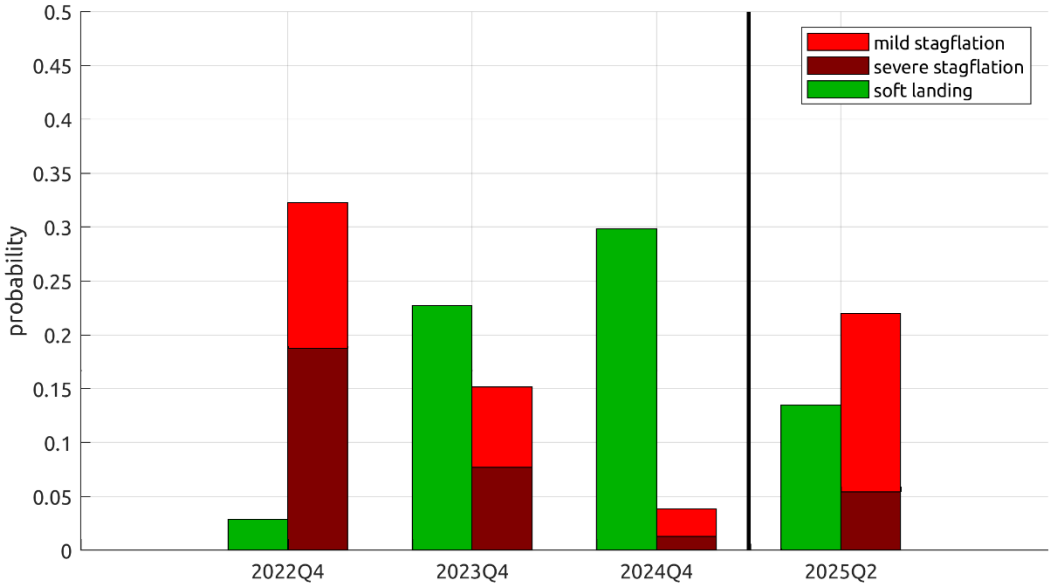

Notes: Estimated probabilities of macroeconomic outcomes at the four-quarter horizon using the model developed in in the appendix. The scenario "soft landing" is defined as four-quarter, four-quarter ahead headline inflation registering between 1½ percent and 2½ percent while four quarter four quarter ahead real GDP growth registers greater than 1 percent. The scenario labeled "severe stagflation" is defined as four-quarter inflation registering over 4 percent and four quarter real GDP growth registering less than 0 percent. The scenario labeled "mild stagflation" is defines as four-quarter inflation registering over 3 percent and four quarter real GDP growth registering less than 1 percent, excluding severe stagflation scenarios. The dates on the horizonal axis correspond to when the forecast is made, predicting outcomes four quarters hence.

There are some important caveats to these results in particular and the modeling approach more generally. First, given the small sample sizes of macroeconomic time-series data, it is challenging to be precise about the likelihood of low-probablity events, and the results are certainly sensitive to various modeling assumptions governing the estimation of uncertainty. In particular, because the assessment of uncertainty is grounded in the historical behavior of reduced-form shocks, the model cannot detect an increase in forward-looking risk unless it is reflected in recent large forecast errors. Like all GARCH-based approaches, it assumes a stable relationship between past shocks and the conditional uncertainty about future outcomes. This assumption may be violated during periods of structural change, leading to potentially misleading projections. In addition, the identification of structural shocks in the model relies on fixed loading parameters of reduced-form shocks onto structural shocks; if the true economic structure evolves—such that, for example, the effects of a given policy shift are misattributed to a supply shock—the model's decomposition and forecasted risk profile could be distorted. As with any empirical model, these limitations underscore the importance of complementing statistical forecasts with judgment and real-time contextual analysis.

Appendix: Methodology and Data

Our primary goal is to estimate the predictive multivariate distribution for several macro variables. This task is accomplished in three steps. The first step uses OLS to identify the time series of predictive means of the four variables, and to identify a series of reduced-form shocks relative each of those mean predictions –measured as the OLS regression residuals. The predictive means in the OLS regressions are informed by macro variables and survey data. The second and third steps, together with the predictive means from the first step, assemble estimates for the full multivariate predictive density of the four variables. In particular, the second step identifies structural shocks that underlie the reduced-form OLS shocks, using simple Keynesian intuition. In addition to providing a realistic treatment of the interdependence of the reduced-form shocks, the structural model with Keynesian intuition adds a bit of economic context to the exercise. This step yields time series estimates for four structural shocks. The third step models the conditional distribution of the structural shocks. Armed with (a) time series forecasts of the mean of the economic variables, (b) time series forecasts of the conditional distribution of the structural shocks, and (c) a mapping between the structural shocks and the reduced-form shocks, we can construct density forecasts for the macroeconomic variables of interest. Much of the methodology follows Bekaert, Engstrom and Ermolov (2025, BEE henceforth), to which readers are referred for more details.[6] First, we describe the data used in this exercise

Data

The data we employ for this study include realized outcomes for four macroeconomic time series as well as survey forecasts of those four variables. The four macroeconomic series are cumulative four-quarter headline PCE inflation π_t^4, cumulative four-quarter core PCE inflation, πc_t^4, cumulative four-quarter real GDP growth, g_t^4 and the average quarterly level of the unemployment rate, u_t. Historical data for these series are gathered from Fred, a data provision service of the Federal Reserve Bank of St. Louis.[7] The survey data are from the Survey of Profession Forecasters

[6] Bekaert, Geert, Eric Engstrom, and Andrey Ermolov. "Uncertainty and the Economy: The Evolving Distributions of Aggregate Supply and Aggregate Demand Shocks." AEJ: Macroeconomics forthcoming.

[7] See https://fred.stlouisfed.org/

(SPF), with data taken from the website of the Federal Reserve Bank of Philadelphia.[8] We use SPF forecasts for four-quarter cumulative headline inflation, core inflation, real GDP growth and the four-quarter ahead level of the unemployment rate. For example, standing in 2024Q4, the measured four quarter ahead SPF forecast for inflation covers the period 2025Q1-2025Q4. Our data sample is quarterly. The SPF data is the limiting factor for the start of our sample. Although some series in the SPF are available from 1968Q4, our sample period (after accounting for the need to calculate multiperiod lags for each variable) extends from 1971Q3-2025Q2.[9] For all estimation steps, we exclude data from 2020Q1-2021Q4 so that the extreme variation in the data from the Covid-19 pandemic era and its immediate aftermath, which may not be representative of the rest of the sample, does not dominate parameter estimates and model selection criteria. (We do still report results for that period.)

Step 1

We begin by using regressions to estimate time series for predictive means of the macroeconomic variables, as well as reduced-form historical shocks for four macroeconomic time series, all at the four-quarter prediction horizon. To do so, we use simple forecasting regressions with OLS estimation of the equations:

$$\begin{aligned} \pi c^4_{t+4} &= \pi c_0 + x_{\pi c,t}\beta_{\pi c} + u^{\pi c}_{t+4} \\ g^4_{t+4} &= g_0 + x_{g,t}\beta_g + u^g_{t+4} \\ u_{t+4} &= u_0 + x_{u,t}\beta_u + u^u_{t+4} \\ \pi^4_{t+4} &= \pi_0 + x_{\pi,t}\beta_\pi + u^\pi_{t+4} \end{aligned} \qquad (1)$$

Notice that the explanatory variables (denoted by "*x's*") are lagged four quarters so that we identify shocks to the variables at the four-quarter horizon. As explanatory variables, each regression includes the lagged dependent variable as well as the ex-ante forecast of the corresponding variable from the SPF. By including survey-based expectations as explanatory variables, we leverage the wide set of information that is available to survey respondents. As can be seen from Equation (1), the four-quarter ahead reduced-form shocks for four-quarter inflation,

[8] See https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters

[9] SPF forecasts for headline and core PCE inflation are available starting only in 2007. Prior to that date, we use SPF forecasts for headline GDP inflation to stand in for SPF forecasts of headline and core PCE inflation. Forecasts for GDP inflation are available starting from 1968Q4.

core inflation, the unemployment rate and real GDP growth are labeled u_{t+4}^{π}, $u_{t+4}^{\pi c}$, u_{t+4}^{u}, and u_{t+4}^{g}, respectively. In addition, each regression model in Equation (1) provides a time series of the conditional mean for each variable. For example, $E_t\left[\pi c_{t+4}^4\right] = \hat{\pi}_0^c + x_{\pi c,t}\hat{\beta}_{\pi c}$, where hats denote OLS estimates for parameters. Because our data are overlapping, comprising for instance quarterly observations for four-quarter growth, OLS standard errors are likely to be biased. We instead use bootstrapping techniques to calculate standard errors.

Table A1: OLS forecasts of macro variables at the four-quarter horizon

	πc_{t+4}^4	g_{t+4}^4	u_{t+4}	π_{t+4}^4
const	-0.03	0.33	0.05	0.20
	(0.27)	(0.72)	(0.48)	(0.44)
SPF forecast (t)	0.74	0.85	1.23	0.49
	(0.31)	(0.26)	(0.22)	(0.28)
lagged dependent (t)	0.25	0.02	-0.21	0.43
	(0.25)	(0.10)	(0.21)	(0.19)
$\overline{R}^2$	0.80	0.20	0.72	0.64

Note: OLS parameter estimates for models described by Equation (1). Standard errors, in parentheses, are from a block XY bootstrap routine with block length of 20 quarters. The sample period for estimation corresponds to realizations in 1971Q3-2025Q2 excluding the period 2020Q1-2020Q4.

As can be seen from Table A1, SPF forecasts are quite informative for all four variables with positive coefficients that are statistically distinct from zero and in most cases near unity, while the lagged dependent variables generally contribute with smaller coefficients and less statistical significance. The R-squared statistics across the four variables range from 0.20-0.80 but should be interpreted with caution due to the overlapping nature of the data.

Figure A1 depicts the regression model-implied mean forecasts for the four macroeconomic series. These will be used as the means for the predictive distributions. Unsurprisingly, expected GDP growth is relatively low during recessions, and the expected unemployment rate is higher than normal in recessions. Expected headline and core inflation have generally declined over the sample period.

Figure A1: Estimated conditional means from predictive regressions

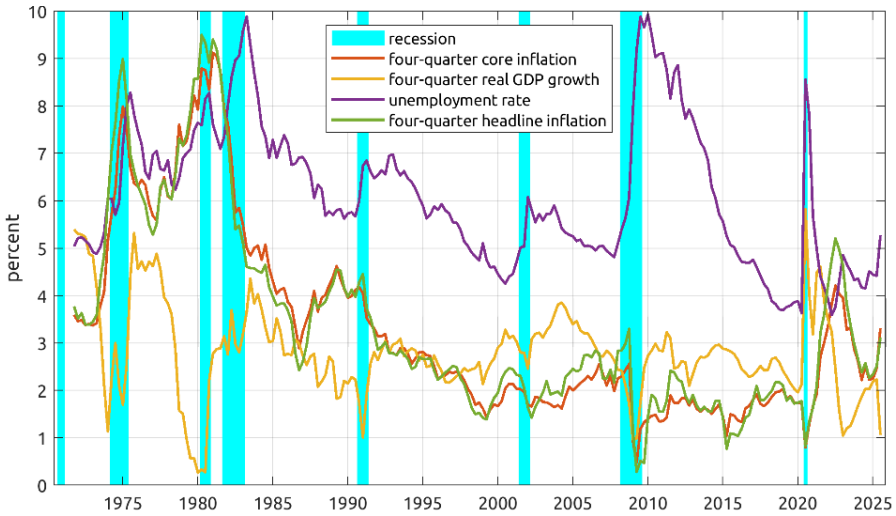

Notes: OLS estimates for the conditional mean of macro variables as described by the model in Equation (1). The sample period for estimation corresponds to realizations from 1971Q3-2025Q2 excluding the period 2020Q1-2021Q4. NBER recession periods are shaded in blue.

Figure A2 depicts the regression model-implied reduced-form shocks to the four series. Deeply negative shocks to GDP are evident for all recessions, while the inflation shocks are sometimes negative and sometimes positive during recessions. These distinct patterns highlight the need to distinguish between supply driven and demand driven risk to the outlook, which are modeled explicitly in the next step.

Figure A2: Reduced-form shocks from predictive regressions

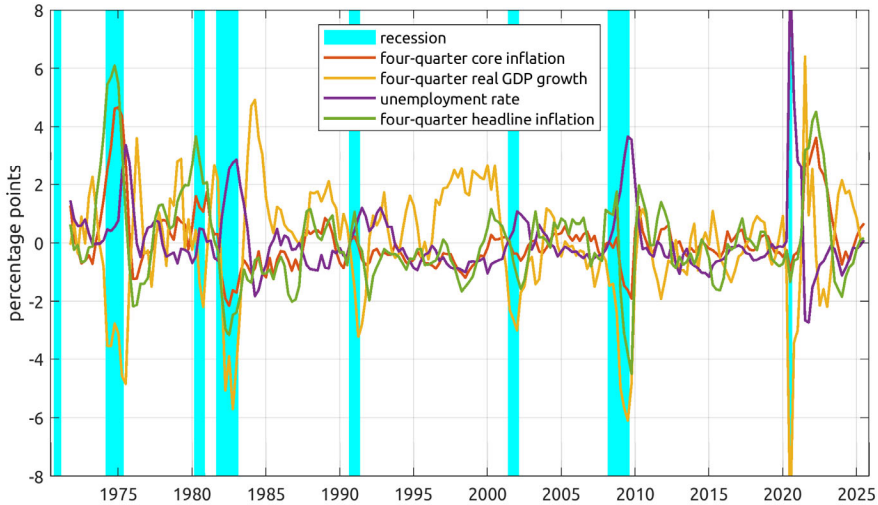

Notes: OLS estimates of reduced-form shocks for the model described by Equation (1). The sample period for estimation corresponds to realizations from 1971Q3-2025Q2 excluding the period 2020Q1-2021Q4. NBER recession periods are shaded in blue.

Step 2

Step 1 provides time series estimates of the conditional means of our four variables as well as time series estimates for the reduced-form shocks to the variables. Our ultimate goal is to generate multivariate predictive densities for the four macro variables. If the reduced-form shocks were drawn from simple i.i.d. distributions, we could produce simple static error bands around our forecasted means, by bootstrapping or a similar methodology. However, as we will see, a more elaborate model is able to fit the data substantially better. In particular, there are rich patterns of time-varying volatility, skewness, and correlation among the macro variables that can be matched to produce more accurate predictive densities.

In this step we use Keynesian intuition to model the interdependence of the reduced-form shocks. In particular, we assume that the reduced-form shocks are simple functions of deeper structural supply and demand shocks, among other structural shocks. Concretely for this step, we identify a mapping between the reduced form shocks and "structural" supply and demand shocks. The point of establishing this mapping is that once we have estimated its parameters, we can recover times series for the structural shocks. To begin, we assume a linear mapping between reduced-form and structural shocks as follows:

$$\begin{bmatrix} u_t^{\pi c} \\ u_t^{g} \\ u_t^{u} \\ u_t^{\pi} \end{bmatrix} = \begin{bmatrix} -\sigma_{\pi c,s} & \sigma_{\pi c,d} & 0 & 0 \\ \sigma_{g,s} & \sigma_{g,d} & 0 & 0 \\ -\sigma_{u,s} & -\sigma_{u,d} & \sigma_{u} & 0 \\ -\sigma_{\pi,s} & \sigma_{\pi,d} & 0 & \sigma_{\pi} \end{bmatrix} \begin{bmatrix} \varepsilon_t^{s} \\ \varepsilon_t^{d} \\ \varepsilon_t^{u} \\ \varepsilon_t^{\pi} \end{bmatrix} \quad (2)$$

Above ε_t^s are ε_t^d defined as "supply" and "demand" structural shocks, respectively. As part of our minimalist strategy to identify structural shocks, we use only sign restrictions and zero restrictions to define supply and demand shocks in a manner consistent with Keynesian intuition. Specifically, all the "σ" parameters are assumed to be positive. For example, supply shocks are defined so that they push inflation and real activity in the opposite direction. A positive supply shock reduces core inflation (as governed by the parameter $\sigma_{\pi c,s}$) but increases real GDP growth (via $\sigma_{g,s}$). A quintessential example of a (negative) supply shocks is the onset of stagflation caused by an oil price shock. Such a shocks would push up on inflation and down on real activity. Conversely, demand shocks push core inflation and real GDP growth in the same direction: A positive demand

shock increases core inflation (via $\sigma_{\pi c,d}$) and increases real GDP growth (via $\sigma_{g,d}$). A quintessential example of a (negative) demand shock is a financial crisis, which may push down both inflation and real activity.

As can be seen in the bottom two rows of Equation (2), similar sign restrictions are imposed for the relationships between supply and demand shocks and shocks to headline inflation and the unemployment rate. We also allow for idiosyncratic shocks to affect headline inflation and the unemployment rate, ε_t^{π} are ε_t^{u}, respectively. These are intended to accommodate variation in headline inflation that may not affect core inflation, and shocks to the unemployment rate which may be specific to the labor market and orthogonal to aggregate supply and demand shocks, respectively. As is common in the literature identifying structural shocks, the four structural shocks are assumed to be independent, and without loss of generality, the structural shocks as assumed to all have unit unconditional variance.

Naturally, we must estimate the σ parameters in Equation (2), of which there are 10. To do so we use a classical minimum distance (CMD) approach. Similar to the perhaps more familiar generalized method of moments methodology, CMD works by choosing parameters to best match a set of sample statistics. For our case, we utilize information in the estimated reduced form shocks from Equation (1), estimates of which are plotted in Figure A2. For each series, we calculate the unconditional second-order moment statistics: the four unconditional standard deviations the reduced-form shocks and the six pairwise correlations among them. Notice that under Equation (2), the 10 σ parameters exactly determine these 10 unconditional standard deviations and correlations for the reduced-form shocks. This can be seen as follows. Let the 4x4 matrix in Equation (2) be called *M*. Then the covariance matrix, Ω, of the reduced-form shocks is equal to *MM'* (recalling that the structural shocks all have unit variance and are uncorrelated). The 10 second-order statistics to be fit are all nonlinear functions of the 10 unique elements of Ω. It follows that that estimated Ω can be used identify the σ parameters. Furthermore, once we have estimated the σ parameters in *M*, we can use Equation (2) to invert the structural shocks from the estimated reduced-form shocks. We simply calculate the inverse of *M* and multiply that inverse by the vector reduced-form shocks. To be concrete, the statistics to be fit in this estimation procedure are shown in Table (2). As shown in the table, the unconditional standard deviations of the shocks are rather precisely estimated, with standard errors that are small relative to the point estimates. Of

the six correlations reported in Table A2, two are strongly distinct from zero: the correlation of shocks to real GDP growth and the unemployment rate are strongly negatively correlated, as expected. And the shocks to headline inflation and core inflation are strongly positively correlated.

Table A2: CMD estimation: statistics to be fit

	$u^{\pi c}_{t+4}$	u^{g}_{t+4}	u^{u}_{t+4}	u^{π}_{t+4}
std dev	0.99	1.85	0.91	1.52
	(0.21)	(0.22)	(0.14)	(0.24)
correlation with…				
u^{g}_{t+4}	-0.19			
	(0.23)			
u^{u}_{t+4}	-0.06	-0.69		
	(0.22)	(0.08)		
u^{π}_{t+4}	0.88	-0.06	-0.20	
	(0.02)	(0.21)	(0.17)	

Note: Sample statistics for OLS residuals from the model described in Equation (1) and parameter estimates in Table A1. Standard errors, in parentheses, are from a block bootstrap routine with block length of 20 quarters. The sample period for estimation corresponds to realizations in 1971Q3-2025Q2 excluding the period 2020Q1-2021Q4.

Table A3 shows the results for the CMD estimation step with estimated parameters corresponding to those in Equation (2).

Table A3: CMD estimated loadings

$$\begin{bmatrix} u^{\pi c}_{t+4} \\ u^{g}_{t+4} \\ u^{u}_{t+4} \\ u^{\pi}_{t+4} \end{bmatrix} = \begin{bmatrix} -0.76 & 0.63 & 0 & 0 \\ (0.22) & (0.11) & & \\ 1.45 & 1.15 & 0 & 0 \\ (0.24) & (0.26) & & \\ -0.36 & -0.56 & 0.62 & 0 \\ (0.11) & (0.16) & (0.07) & \\ -0.89 & 1.03 & 0 & 0.69 \\ (0.22) & (0.18) & & (0.08) \end{bmatrix} \begin{bmatrix} \varepsilon^{s}_{t+4} \\ \varepsilon^{d}_{t+4} \\ \varepsilon^{u}_{t+4} \\ \varepsilon^{\pi}_{t+4} \end{bmatrix}$$

Note: Loadings estimated from the CMD estimation of the model described by Equation (2). Standard errors, in parentheses, are from a block bootstrap routine with block length of 20 quarters. The sample period for estimation corresponds to realizations in 1971Q3-2025Q2 excluding the period 2020Q1-2020Q4.

Consistent with our identification assumptions, all the signs in Table A3 match those in Equation (2). The top row shows that shocks to core inflation $u^{\pi c}_{t+4}$ load materially onto both shocks to supply and demand, ε^{s}_{t+4} are ε^{d}_{t+4}, respectively, with slightly greater loading onto supply shocks.

Shocks to real GDP growth, u_{t+4}^{g}, load positively onto demand shocks and supply shocks, also with slightly more loading on supply shocks. Shocks to the unemployment rate load negatively onto both supply and demand shocks, which is intuitive since "good" news of either type lowers the unemployment rate, but with a relatively larger loading on demand shocks, consistent with the intuition that the unemployment rate reflects the cyclical position of the economy with respect to demand conditions. The unemployment rate also has a large exposure the idiosyncratic component, loading onto ε_{t+4}^{u} more strongly than it does to either supply or demand shocks. Finally, headline inflation loads with the expected signs onto supply and demand shocks, more strongly so than core inflation does. This is to be expected because headline inflation is more volatile than core inflation. In addition, headline inflation loads meaningfully onto an idiosyncratic component, ε_{t+4}^{π}. This is likely capturing idiosyncratic variation in food and energy inflation, which are stripped out of the core inflation measure. In results not shown, the CMD model-implied statistics corresponding to those in Table A2 suggest a near perfect fit of those statistics, consistent with the exactly identified nature of the CMD exercise.

Figure A3: Estimated supply and demand structural shocks

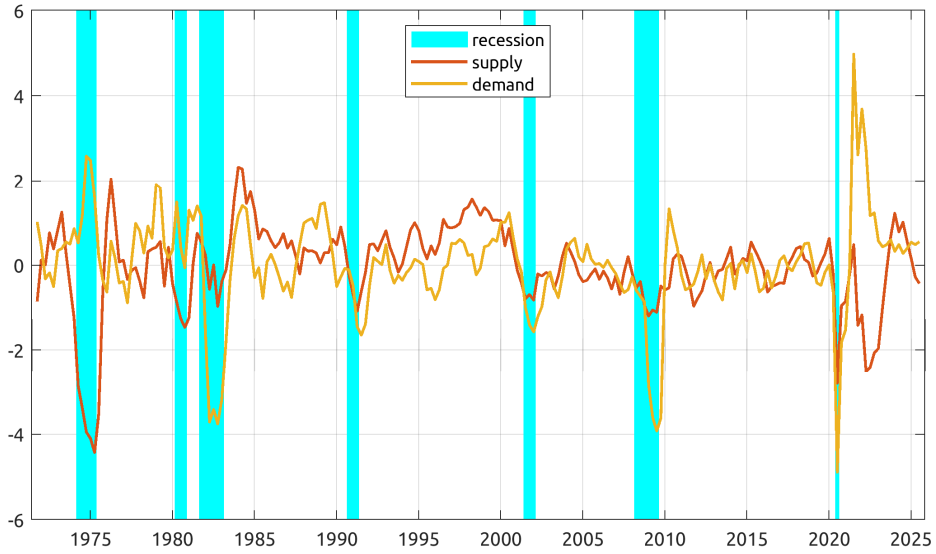

Note: Supply and demand shocks inverted using Equation (2), the reduced-form shocks from Figure A2 and the parameter estimates in Table (2). NBER-defined recessions are shaded in blue.

Figure A3 shows the estimated structural supply and demand shocks inverted from the reduced-form shocks using the parameter estimates from Table A3. Both supply and demand shocks tend to register negative values during recessions, with particularly pronounced negative supply shocks

during the recession in the early 1970s and the Covid period and its aftermath. Most other recession feature predominant negative demand shocks.

Figure A4: Estimated idiosyncratic inflation and unemployment rate structural shocks

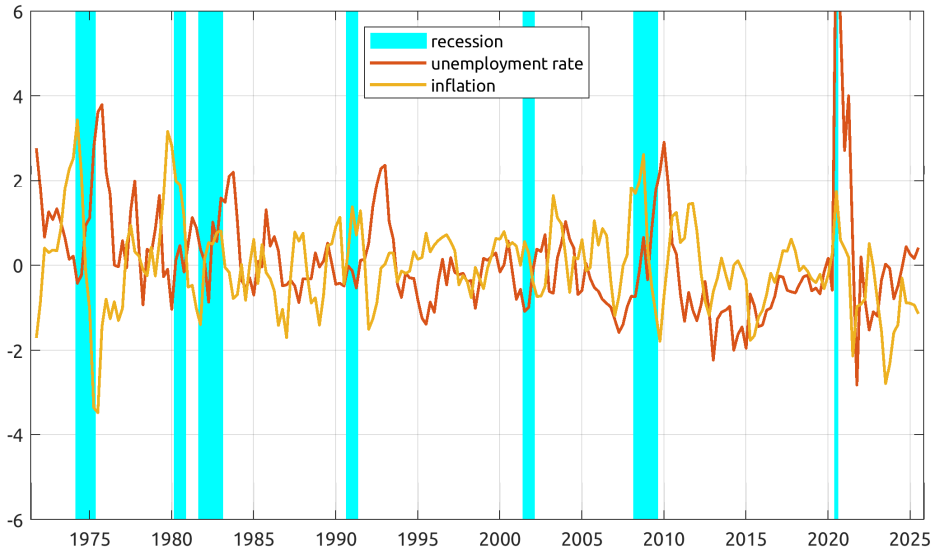

Note: Idiosyncratic shocks to headline inflation and the unemployment rate inverted using Equation (2), the reduced-form shocks depicted in Figure 2 and the parameter estimates in Table A2. NBER-defined recessions are shaded in blue.

Figure A4 shows idiosyncratic shocks to the unemployment rate and inflation. The idiosyncratic shocks to the unemployment rate were particularly large during the peak of the Covid period while the largest peaks for the idiosyncratic inflation component were registered during the 1970s and 1980s and the Covid period.

Step 3

The final and most complex step involves estimating the time-varying conditional distributions of the structural shocks identified in the previous step. Exploiting the assumption that the structural shocks are independent, we estimate univariate processes for each of them. To flexibly accommodate potentially time-varying volatility, skewness and kurtosis of the structural shocks, we use the BEGE distribution developed by Bekaert and Engstrom (2017)[10] and the BEGE-

[10] Bekaert, G. and E. Engstrom, 2017, "Asset Return Dynamics under Habits and Bad Environment–Good Environment Fundamentals," Journal of Political Economy, vol 125.3.

GARCH framework of Bekaert, Ermolov and Engstrom (2015).[11] Concretely, consider a generic structural shock, ε_{t+4} (e.g., a supply or demand shock) to be realized at time (t + 4). We model the shock as $\varepsilon_{t+4} \sim BEGE\left(p_t, n_t;\ \sigma_p, \sigma_n\right)$ where $BEGE$ denotes the distribution. To unpack the BEGE distribution a bit, the BEGE model assumes that ε_{t+4} has two components:

$$\varepsilon_{t+4} = \sigma_p \omega_{t+4}^p - \sigma_n \omega_{t+4}^n \qquad (3)$$

where ω_{t+4}^p and ω_{t+4}^n are individual component shocks. The volatility parameters σ_p and σ_n are restricted to be positive. The component shocks are independent and distributed as centered gamma:

$$\omega_{t+4}^p \sim \tilde{\Gamma}\left(p_t, 1\right); \qquad \omega_{t+4}^n \sim \tilde{\Gamma}\left(n_t, 1\right) \qquad (4)$$

Figure A5 provides some examples to illustrate features of the BEGE distribution. The upper right panel shows that the probability density function of ω_{t+4}^p, which we label the "good" component. It is bounded from the left and has an unbounded right tail. The volatility, skewness and kurtosis of the good component are governed by the shape parameter, p_t. Similarly, as shown in the upper left panel, the probability density function of $-\omega_{t+4}^n$ (the "bad" component) is bounded from the right and has an unbounded left tail. The lower panel illustrates examples of possible shapes of the overall BEGE distribution for ε_{t+4} in Equation (3), which could arise as a result of different configurations of the shape parameters p_t and n_t. In particular, the BEGE probability density function may be positively or negatively skewed, or symmetric and may have excess kurtosis depending on the shape parameters. The Gaussian distribution is a special case of the BEGE distribution (as the shape parameters tend to infinity).

[11] Bekaert, G., Engstrom, E, and A. Ermolov, 2015, "Bad environments, good environments: A non-Gaussian asymmetric volatility model," Journal of Econometrics, vol. 186.1, ppg. 258-275.

Figure A5: Illustrations of BEGE densities

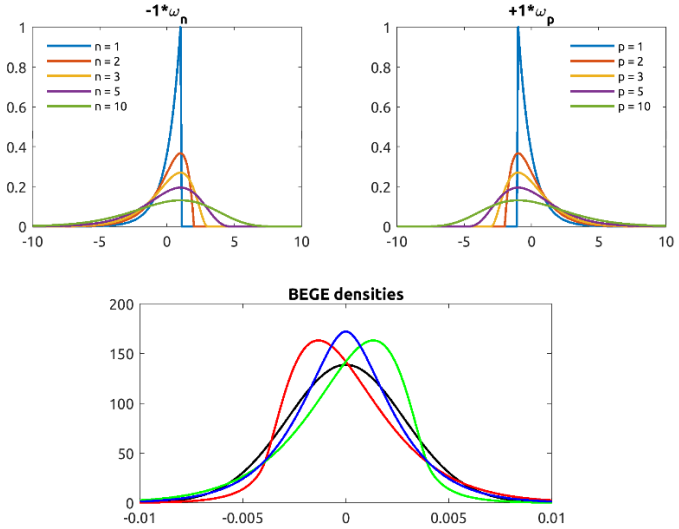

Note: Examples of gamma distribution and the BEGE distribution under various parameter configurations.

We assume that each of our structural shocks follow a potentially time-varying BEGE distribution in that the shape parameters may evolve over time. Taking the example of the supply structural shock, we assume that

$$\varepsilon_{t+4}^{s} \sim BEGE(p_t^s, n_t^s; \sigma_{sp}, \sigma_{sn}) \tag{5}$$

where p_t^s and n_t^s are time-varying shape parameters for the BEGE distribution for supply shocks, and σ_{sp} and σ_{sn} are static parameters. As above, p_t^s governs the level of "good variance" for supply shocks and n_t^sgoverns the level of "bad" variance. Specifically, under this formulation of the BEGE distribution, the conditional variance of ε_{t+4}^{s} follows

$$VAR_t\left[\varepsilon_{t+4}^{s}\right] = \sigma_{sp}^2 p_t^s + \sigma_{sn}^2 n_t^s \tag{6}$$

So that both p_t^s and n_t^s increase the variance of the supply shock. However, under the BEGE distribution p_t^s and n_t^s have opposite effects on skewness

$$SKW_t\left[\varepsilon_{t+4}^{s}\right] = 2\sigma_{sp}^3 p_t^s - 2\sigma_{sn}^3 n_t^s \tag{7}$$

Clearly, p_t^s increases skewness while n_t^s decreases skewness. For this reason, $p_t^s$is referred to as "good variance" while n_t^s is referred to as "bad variance."

To complete the model, we must specify the dynamics of the time-varying good and bad volatility variables, p_t^s and n_t^s. To do so, we use a simple GARCH specification:

$$p_t^s = p_0^s + \rho^{ps} p_{t-1}^s + \phi^{ps} \left|\varepsilon_t^s\right|^j \qquad (8)$$
$$n_t^s = n_0^s + \rho^{ns} n_{t-1}^s + \phi^{ns} \left|\varepsilon_t^s\right|^j$$

In a nutshell, p_t^s has autoregressive dynamics, with shocks proportional to the magnitude of realized shocks to supply, through the term $\phi^{ps} \left|\varepsilon_t^s\right|^j$. We explore j=1,2 to allow for either the absolute value of shocks or squared shocks to drive the future distribution. The dynamic BEGE model is estimated for all four of the structural shocks. Estimation is carried out by MLE, and we test a number of restricted models including distributions with p_t^s and/or n_t^s, being constant. Inference and specification tests for our data is complicated by the fact that we are using overlapping quarterly observations. To deal with this complication, we use tests and criteria that account for our use of overlapping data, focusing mostly on out-of-sample techniques. In contrast to in-sample model selection criteria, out-of-sample techniques for model selection are generally robust to overlapping data. We use an "expanding window" out-of-sample procedure in which the sample period is sequentially extended in blocks of five years. For each iteration, we re-estimate the model and then calculate an out-of-sample log likelihood for the five years of data subsequent to the end of each window.[12] This technique provides blocks of true out-of-sample log likelihoods, which we sum for all the out-of-sample periods and then compare across various specifications. For model selection, we also consider the in-sample Akaike criterion modified to account for small samples, AICc. Alas, the AICc is not designed for overlapping observations because it assumes conditionally independent likelihood observations, whereas sequential observations are strongly dependent in our case due to overlapping data. To be conservative, albeit informal, we introduce a modification in which we reduce the number of effective observations in the AICc calculation by a factor of four to account for the four-quarter overlap.

[12] In the expanding window estimations, the initial, minimum sample uses the first half of the full sample, so that we examine out-of-sample results for only the second half of the sample.

Table A4: BEGE model selection results

structural shock		*p const n const*	*p vary n const*	*p const n vary*	*p vary n vary*
supply	*insmp AICc*	522.0	520.1	498.2	503.6
	oos loglike	-114.5	-115.9	-94.0	-102.2
demand	*insmp AICc*	525.3	529.6	517.7	524.1
	oos loglike	-108.0	-108.1	-101.1	-100.9
idiosyncratic unemp	*insmp AICc*	556.5	555.8	557.0	559.8
	oos loglike	-136.6	-139.3	-140.2	-139.7
idiosyncratic infl	*insmp AICc*	583.7	579.1	581.9	584.0
	oos loglike	-126.4	-125.1	-126.7	-126.4

Note: In-sample and out-of-sample model selection criteria for the general model described by Equation (8) and estimated by MLE, as described in the test. The four specifications are estimated for each of the structural shocks that were identified in previous steps. For the AICc statistic, we use an effective number of observations of 54, which is the total number of quarterly observations, 216, divided by 4.

Table A4 summarizes the results of these specification tests for each of the structural shocks. For each shock, we compare the in-sample AICc criteria (the minimum being optimal) and the out-of-sample loglikelihoods (the maximum being optimal). The in-sample and out-of-sample criteria agree for the supply shock, the optimal model has a constant p_t but a time-varying n_t. This suggests that the level of "bad variance" for supply shocks (largely governing the lower tail of the distribution) varies significantly over time, "good variance" (largely governing the upper tail) is adequately modeled as being constant. For demand shocks, the in-sample criterion prefers the n-only specification, while the out-of-sample criterion mildly prefers a model in which both good and bad volatility. We use the more parsimonious and conservative p-only specification. For idiosyncratic shocks to the unemployment rate, a simple static model with constant shape parameters is preferred on an out of sample basis, whereas a p-only specification is preferred by the in-sample criterion. We use the former, again for parsimony. Finally, for the idiosyncratic shock headline inflation, the optimal specification has p_t varying, but n_t constant by both the in-sample and out-of-sample criteria. Of course, for the case of inflation, positively skewed variance governed by p_t may not be regarded as "good". Overall, these results suggest that accommodating

time-varying upside and downside risks to the structural shock generally helps improve the performance of the density forecasts.[13]

Table A5: BEGE parameter estimates

	ε_t^s	ε_t^d	ε_t^u	ε_t^π
p_0	13.13	3.12	0.93	0.01
	[1.03,20.0]	[1.77, 10.52]	[0.30,4.61]	[0.01, 5.70]
σ_p	0.15	0.34	1.01	1.02
	[0.05, 0.41]	[0.19,0.45]	[0.45,1.90]	[0.14,1.51]
ϕ_p	--	--	--	0.25
	--	--	--	[0.01,6.38]
ρ_p	--	--	--	0.64
	--	--	--	[0.01,0.82]
n_0	0.01	0.02	19.99	2.94
	[0.01,4.84]	[0.01, 0.15]	[0.49,20.0]	[0.71,15.2]
σ_n	1.99	1.84	0.11	0.42
	[0.05, 3.78]	[0.68, 3.60]	[0.02,0.41]	[0.09,0.72]
ϕ_n	0.29	0.11	--	--
	[0.08,7,98]	[0.01,0.29]	--	--
ρ_n	0.01	0.50	--	--
	[0.01,0.77]	[0.07,0.73]	--	--

Notes: Parameter estimates for BEGE model described by Equations (5) and (8) with restrictions as selected by specification criteria shown in Table A4. Bootstrapped 90 percent confidence intervals are shown in square brackets. All parameters are bounded below by 0.01 and p_0 and n_0 are bounded from above at 20.0.

Table A5 report the parameter estimates for each model selected as optimal in Table A4. As shown in the first two columns, both supply and demand shocks are modeled as having constant p_t dynamics, but time-varying n_t dynamics. In contrast, as shown by the third column, the idiosyncratic shock to the unemployment rate has constant volatility for both p_t and n_t. As shown by the right column, idiosyncratic shocks to headline inflation have time-varying p_t dynamics by

[13] For all specifications, we use j=1 in Equation (8), which is preferred using both in-sample and out-of-sample selection criteria.

constant n_t dynamics. These parameters are reported for completeness, but few of them have a straightforward interpretation, so we report more intuitive statistics for each model in Table A6.

Table A6: BEGE parameter-implied statistics

	ε_t^s	ε_t^d	ε_t^u	ε_t^π
$mean(VAR_t)$	1.10	0.92	1.22	1.19
	[0.36,2.10]	[0.49,1.94]	[0.71,1.78]	[0.59,1.54]
$std(VAR_t)$	0.89	0.35	--	0.44
	[0.01,1.94]	[0.00,0.87]	--	[0.01,066]
$acorr(VAR_t)$	0.52	0.80	--	0.90
	[0.43,0.93]	[0.55,0.92]	--	[0.53,0.93]
$mean(SKW_t)$	-3.05	-1.75	1.90	0.93
	[-11.7,0.17]	[-8.82,4.22]	[0.57,5.46]	[-0.26,1.86]
$std(SKW_t)$	3.54	1.29	--	0.90
	[0.01,13.68]	[0.01,4.22]	--	[0.01,1.87]
$acorr(SKW_t)$	0.52	0.80	--	0.90
	[0.43,0.93]	[0.55,0.92]	--	[0.53,0.93]

Notes: Statistics implied by BEGE model estimates reported in Table A5 and described by Equations (5) and (8). The top panel presents results for conditional variance, VAR_t. The top three rows report the unconditional mean, standard deviation, and first-order autocorrelation for VAR_t , respectively, for each shock. The bottom three rows report the same statistics for conditional skewness, SKW_t. Conditional variance and (unscaled) skewness for the BEGE model, VAR_t and SKW_t, respectively, are derived in Equations (6) and (7). Derived statistics are calculated by simulation with bootstrapped 90 percent confidence intervals are shown in square brackets.

Table A6 report statistics for the properties of the conditional variance and conditional skewness of each shock as implied by the parameters in Table A5. Recall that Equations (6) and (7) derive the formulas for conditional variance and skewness for the BEGE model, labeled VAR_t and SKW_t, respectively. The statistics in Table A6 are calculated by simulation, in part, to account for the overlapping nature of the data. As shown in the first column, for supply the mean level of conditional variance is around 1. This is by assumption – all reduced-form shocks were assumed to have unit variance in population. More interesting is that the unconditional standard deviation of conditional variance is 0.89, so that VAR_t has a substantial degree of variation relative to its mean. Moreover, the autocorrelation of VAR_t for skewness is 0.52 indicating a fair degree of persistence. The bottom half of the table in the first column shows the statistics for SKW_t for supply. The mean level of SKW_t is deeply negative at -3.05, indicating that large supply shocks tend to be negative.

However, the confidence range for this statistic is fairly wide, spanning from -11.7 to 0.2. SKW_t is also quite volatile for supply, with an unconditional standard deviation of 3.54. SKW_t has the same persistence as VAR_t for supply because there is a single factor, n_t, driving variation in all the conditional moments. The second column show results for the demand shock BEGE specification. The results are similar to those for supply shocks: an appreciable level of variation in both VAR_t and SKW_t, deeply negative average levels for SKW_t, and even more persistence of these conditional moments compared to supply shocks, with an autocorrelation for both VAR_t and SKW_t of 0.80. The results in the third column are for the idiosyncratic shocks to the unemployment rate. Because the optimal model for these shocks had no dynamics for either good or bad volatility, both VAR_t and SKW_t are constant. It is notable however that the mean level for SKW_t is substantially positive at 1.90, indicating once again that large shocks tend to be of the unfavorable sort. The fourth column shock results for the idiosyncratic shock to headline inflation. These tend to be positively skewed with a large degree of variation and positive autocorrelation in both VAR_t and SKW_t.

We now turn to the conditional moments that the model implies over our sample period. Figure A6 plots the fitted good and bad variance for each shock variable under the optimal specification for that variable. The top panel plots the results for supply shocks. Bad variance, the red line, varies substantially over the business cycle, with bad variance peaking prominently in recessions in the 1970s and again after the Covid-induced recession. Peak in the 1970s are consistent with the large supply-like shocks and instances of stagflation observed during that period, for example, during the famous oil price shocks. The high levels of bad variance for supply during the Covid period and its aftermath are consistent with supply chain disruptions from that period, which pushed up on inflation and down on real activity. The results for demand shocks are shown in the second panel. Bad demand variance frequently peaks during recessions, with the Great Financial Crisis of 2008-2009 standing out in particular. Large adverse shocks persistently lowered both inflation and real activity during that period. The results for the conditional moments of the idiosyncratic unemployment rate shock are rather uninteresting, because the optimal specification suggests a constant distribution for these shocks. Consistent with the optimal specification indicated in Table A4, both good and bad variance for the unemployment rate are constant. Finally, the bottom panel shows results for the idiosyncratic shock to headline inflation. Here, the p_t driven variance, the

black line, is more reasonably referred to as "bad", and it shows large peaks during the recessions of the 1970s and 1980s as well as in the aftermath of the Covid recession.

Figure A6: Estimated predictive variance for structural shocks

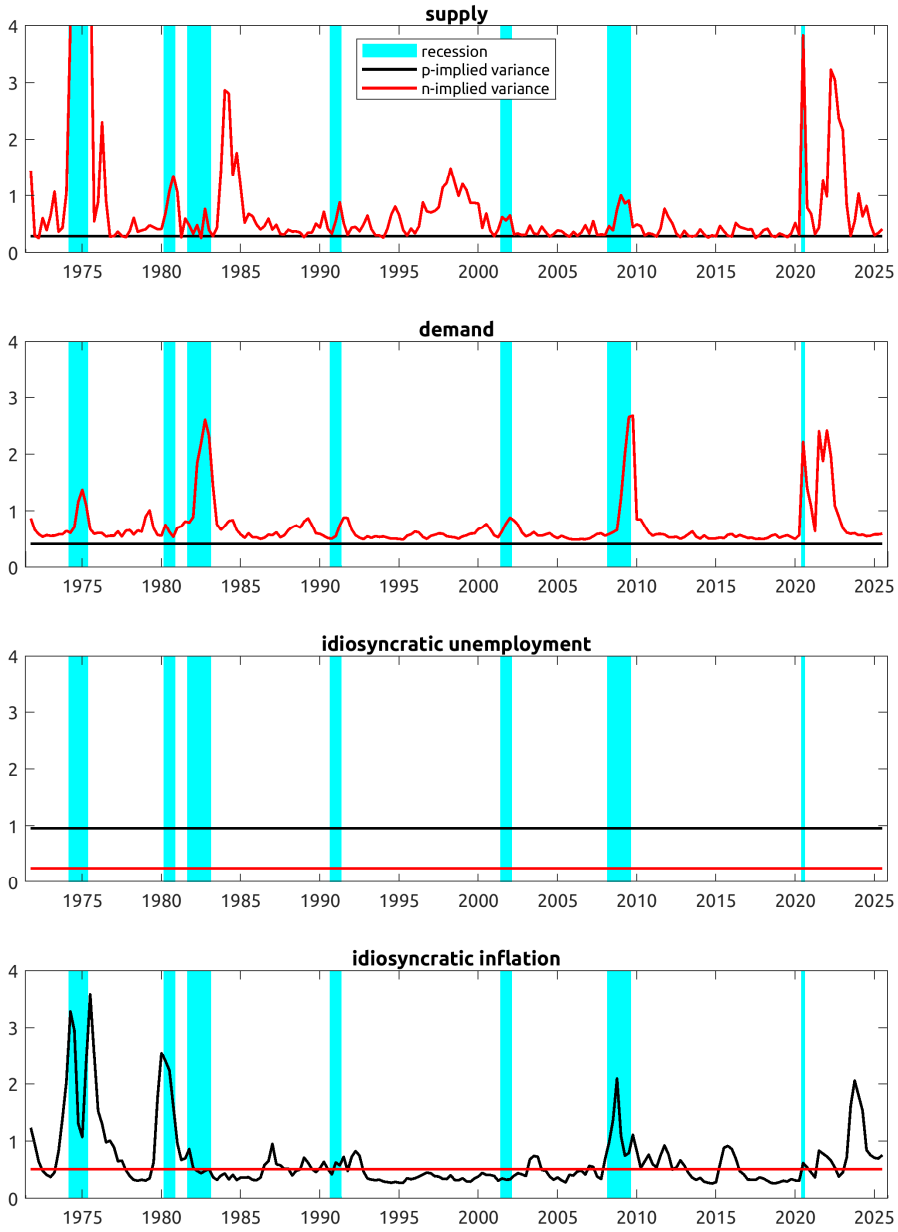

Notes: Estimated "good" and "bad" variance for each structural shock under the optimal specification for that shock as identified in Table (4) and parameters in Table A5. In each panel, p_t and n_t are plotted scaled by their respective estimated parameters. For example, for supply the plotted series are $\sigma_{sp}^2 p_t^s$ and $\sigma_{sn}^2 n_t^s$, the two components of conditional variance for supply shocks. NBER recession periods are shaded in blue. The model employed is described by Equations (5) and (8).

To examine the degree of non Gaussianity that is implied by the parameters estimates in Table A5, Figure A7 plots the model-implied conditional skewness for each of the structural shocks.

Figure A7: Estimated conditional skewness for structural shocks

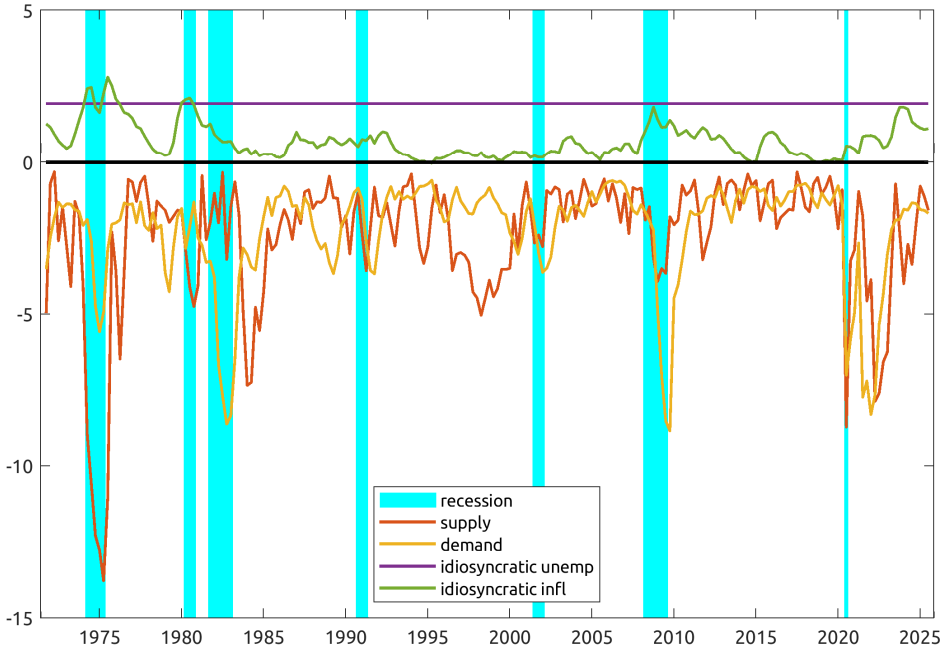

Notes: Estimated conditional unscaled skewness for each shock is calculated as, for the example of supply shocks, as $SKW_t\left[\varepsilon^s_{t+4}\right] = 2\sigma^3_{sp}p^s_t - 2\sigma^3_{sn}n^s_t$. NBER recession periods are shaded in blue. The model employed is described by Equations (5) and (8), with the relative parameters estimates in Table A5 for each shock.

Evidently, both supply and demand shocks are negatively skewed on average. Skewness typically reaches negative peaks during recessionary periods. Negative skewness for supply shocks was particularly pronounced during recessions in the 1970s and 1980s, whereas demand shocks reach negative peaks in recessions throughout the sample. Idiosyncratic shocks to inflation are positively skewed, albeit more mildly than for supply and demand shocks. Idiosyncratic shocks to the unemployment rate are positively skewed, and that degree of skewness is constant, consistent with the optimal model for that shock.

Putting it all together

Armed with the estimated conditional distribution for the structural shocks, it is straightforward in light of Equation (2), although computationally cumbersome, to calculate the predictive

distribution for the four reduced-form shocks, which capture the outlooks for the macroeconomic variables of interest.

Under the assumed linear structure in Equation (2), the conditional univariate moments for the macroeconomic shocks map linearly onto the conditional moments for the structural shocks. For example, the conditional second and third conditional moments for shocks to core inflation are:

$$\begin{aligned} VAR_t\left[u_{t+4}^{\pi c}\right] &= \sigma_{\pi c,s}^2 VAR_t\left[\varepsilon_{t+4}^{s}\right] + \sigma_{\pi c,d}^2 VAR_t\left[\varepsilon_{t+4}^{d}\right] \\ SKW_t\left[u_{t+4}^{\pi c}\right] &= -\sigma_{\pi c,s}^3 SKW_t\left[\varepsilon_{t+4}^{s}\right] + \sigma_{\pi c,d}^3 SKW_t\left[\varepsilon_{t+4}^{d}\right] \end{aligned} \tag{9}$$

Intuitively, higher supply and demand variance both drive up the variance of core inflation. In contrast, positive skewness for supply shocks drives down the conditional skewness for core inflation, whereas higher conditional skewness of demand shocks drives up the conditional skewness of core inflation shocks.

As depicted in the left-hand panels in Figure A8, all of the endogenous shocks feature conditional variance estimates that vary strongly over the business cycle, increasing sharply during recessions. The sharpest peaks in variance for the inflation series occur at the beginning and at the end of the sample period. In contrast, the conditional variance of real GDP growth and the unemployment rate rise in a fairly consistent manner during recessions across the full sample. Estimates of the conditional skewness of endogenous variables are shown in the panels on the right. Once again, strong variation over the business cycle is evident. Of note, core inflation and especially headline inflation exhibit sign-switching in their conditional skewness estimates. When supply variance dominates, such as in the 1970s and 1980s, skewness for inflation tends to be positive – the balance of risks is to the upside for inflation. However, when demand uncertainty dominates, such as during the Great Financial Crisis of 2008-2009, the balance of risks to inflation moves sharply to the downside.

Figure A8: Univariate conditional moments of endogenous shocks

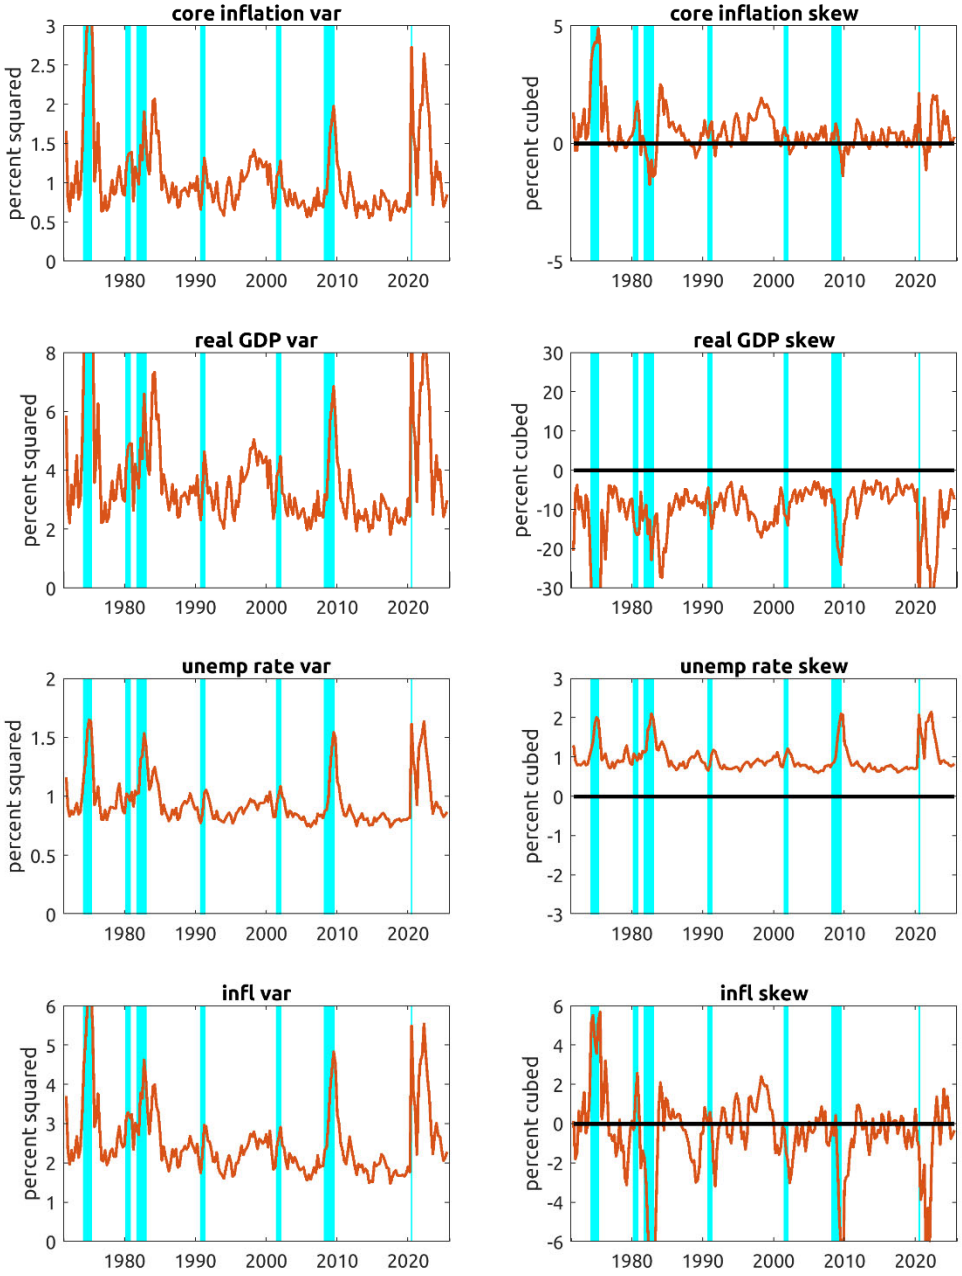

Notes: Estimated conditional variance and unscaled skewness for each endogenous shock is calculated as in Equation (9) using the estimated series of good and bad variance for the structural shocks as reported in Figure A6 and the parameter estimates in Table A2. NBER recession periods are shaded in blue. The BEGE model employed is described by Equations (5) and (8), with the relative parameters estimates in Table A5 for each shock.

The conditional moments of the structural shocks also map linearly into cross-moments for the reduced-form shocks. For example,

$$COV_t\left[u^{\pi}_{t+4}, u^{g}_{t+4}\right] = -\sigma_{\pi,s}\sigma_{g,s}VAR_t\left[\varepsilon^{s}_{t+4}\right] + \sigma_{\pi,d}\sigma_{g,d}VAR_t\left[\varepsilon^{d}_{t+4}\right] \quad (10)$$

Intuitively, higher supply variance pushes the covariance between inflation and real GDP growth towards negative territory because inflation and real GDP growth load with opposite signs onto supply shocks. In contrast, higher demand variance increases the covariance of inflation and real GDP shocks because both of them load positively onto demand shocks. Figure A9 depicts the time series for the estimated covariance between real GDP growth and headline inflation. Early in the sample period when supply shocks dominate, the covariance tends to be negative, especially during supply-driven recessions such as in the early 1970s. In contrast, the covariance moves sharply into positive territory during periods of elevated demand volatility such as during the GFC and early in the pandemic period. This sign-switching behavior should be of keen interest to monetary policy makers: When the covariance is positive, risks to the Fed's dual mandate are not in conflict because in such circumstances the predominant risk is that both inflation and real activity will fall below the Fed's goals. In contrast, when the covariance is negative, policy makers face a situation in which, for example, inflation may be too high while real activity is depressed – a stagflation scenario.

Figure A9: Conditional covariance between real GDP and inflation shocks

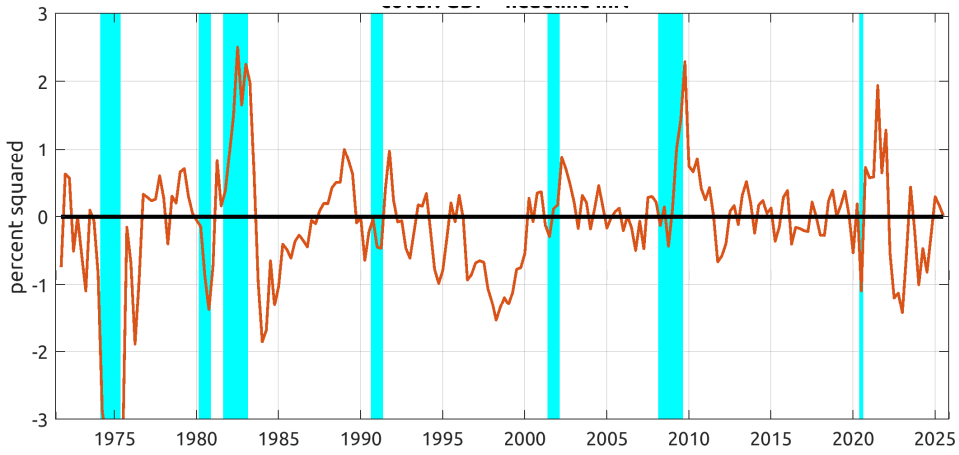

Notes: Estimated conditional covariance calculated as in Equation (10) using the estimated series of good and bad variance for the structural shocks as reported in Figure A6 and the parameter estimates in Table A2. NBER recession periods are shaded in blue.

The final application of this methodology that we present is the multivariate predictive distribution of the endogenous variables, which uses the conditional mean from Step 1 and the conditional distribution of shocks from steps 2 and 3. The conditional joint distributions can be generated by

numerical integration or simulation techniques, as described in BEE. In Figure A10 we plot bivariate distributions for real GDP growth and headline inflation (although bivariate plots with any two of the endogenous variables are of course possible and more elaborate graphical techniques may be able to illustrate three- or four-dimensional distributions). Each panel depicts predictions from the model at different points in time. In each panel, forecasts for four-quarter inflation are plotted on the horizontal axis, and forecasts for four-quarter real GDP growth rate are on the vertical axis. The asterisk denotes the model's mean point forecast from step 1, and the square denotes the mode of the distribution. The yellow area indicates the most likely outcomes, with 50 percent of all outcomes expected to fall within it. The green area depicts the next most likely 40 percent of outcomes, which we interpret as a region of material risk. For reference, axes are drawn at 2 percent for inflation and 2 percent for real GDP growth. The model-implied probabilities of outcomes in each of the four quadrants are listed in the corner of each quadrant.

Figure A10 shows that a wide variety of central locations and shapes for the predicted density have occurred over the sample period. The top left panel shows a relatively quiescent period in the early 1970s when the distribution of had a relatively small footprint, with the yellow region centered roughly around 5 percent real GDP growth and 4 percent inflation. In contrast, the distribution had widened by large amount in both dimensions by late 1973, as shown in the top right panel, amid a series of adverse supply shocks. The pronounced tail to the southeast indicates an elevated risk of stagflation – periods of elevated inflation and lackluster or negative GDP growth. The panels on the second row, in contrast, indicate periods of elevated demand risk, with the predominant tail of the distribution pointed to the southwest region characterized be low growth and low inflation indicating increased risk of a deflationary recession. This was true during the Fed-induced recession in 1982 and the financial crisis in 2009. The third row contrasts the benign distribution in late 2019 right before the onset of the pandemic, with the very wide distribution and dual prominent supply and demand risks that persisted even through 2022 in the wake of the pandemic. The final row shows two recent distributions, at the end of 2024 and the last available distribution in 2025Q2. The spread of the distributions is relatively modest, indicating a moderate amount of risk, but the distribution shifted towards the region of stagflation in 2025Q2 amid expectations that tariff policy may lower growth and increase inflation over the next year.

Figure A10: Predictive joint distributions for headline inflation and real GDP growth

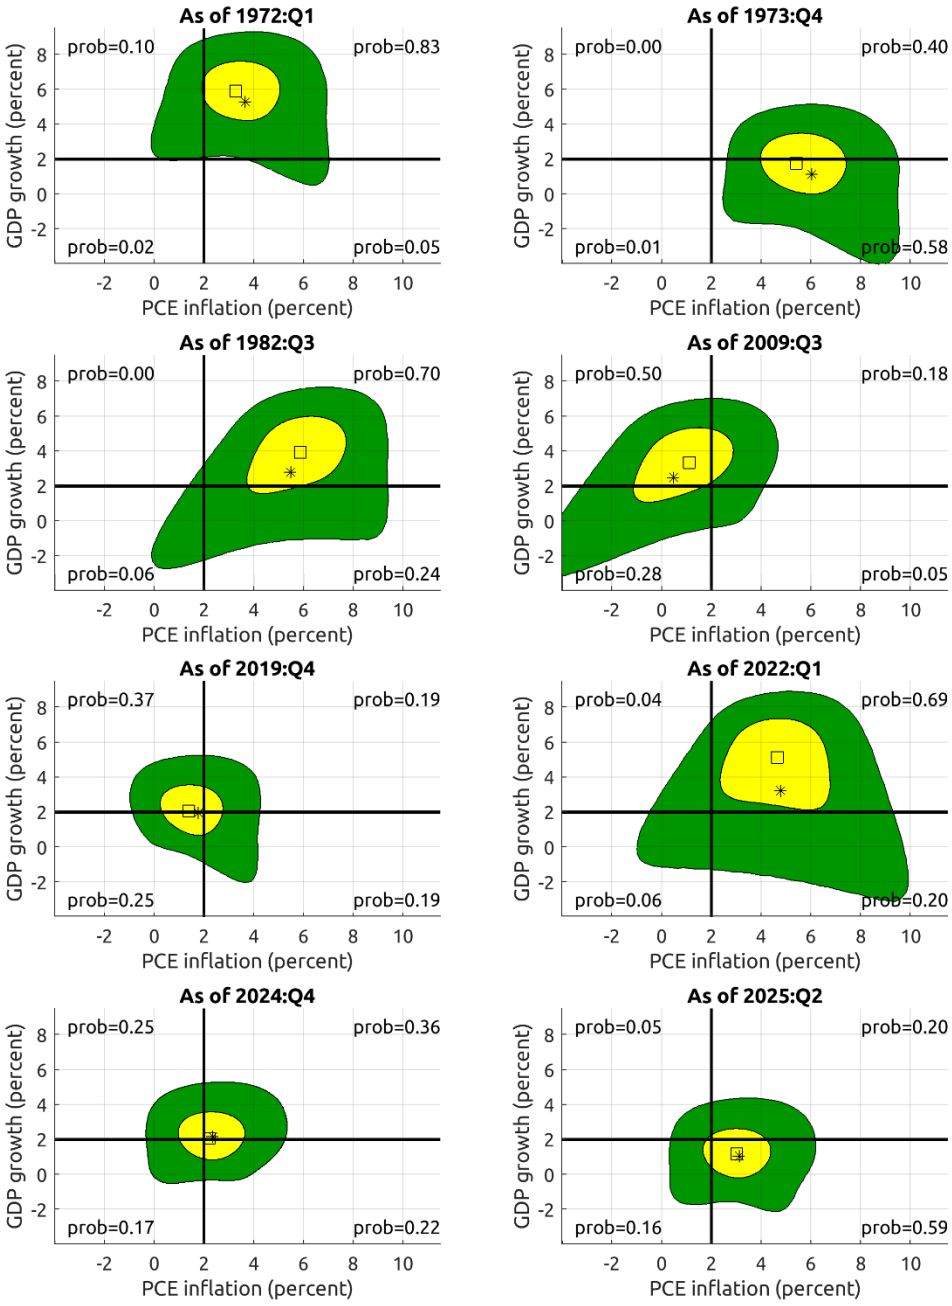

Notes: Estimated predictive distribution of headline inflation and real GDP outcomes using the model developed in Step1-3 in the appendix. Each panel depicts predictions from the model at different points in time. In each panel, forecasts for four-quarter inflation are plotted on the horizontal axis, and forecasts for four-quarter real GDP growth rate are on the vertical axis. The asterisk denotes the model's mean point forecast, and the square denotes the mode. The yellow area indicates the most likely outcomes, with 50 percent of all outcomes expected to fall within it. The green area depicts the next most likely 40 percent of outcomes, which we interpret as a region of material risk. For reference, axes are drawn at 2 percent for inflation and 2 percent for real GDP growth. The model-implied probabilities of outcomes in each of the resulting four quadrants are listed in the corner of each quadrant.