Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1326

August 2021

Macroeconomic and Financial Risks: A Tale of Mean and Volatility

Dario Caldara, Chiara Scotti and Molin Zhong

Please cite this paper as:
Caldara, Dario, Chiara Scotti and Molin Zhong (2021). "Macroeconomic and Financial Risks: A Tale of Mean and Volatility," International Finance Discussion Papers 1326. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/IFDP.2021.1326.

Macroeconomic and Financial Risks: A Tale of Mean and Volatility*

Dario Caldara† Chiara Scotti‡ Molin Zhong§

August 18, 2021

Abstract

We study the joint conditional distribution of GDP growth and corporate credit spreads using a stochastic volatility VAR. Our estimates display significant cyclical co-movement in uncertainty (the volatility implied by the conditional distributions), and risk (the probability of tail events) between the two variables. We also find that the interaction between two shocks—a main business cycle shock as in Angeletos et al. (2020) and a main financial shock—is crucial to account for the variation in uncertainty and risk, especially around crises. Our results highlight the importance of using multivariate nonlinear models to understand the determinants of uncertainty and risk.

Keywords: Uncertainty, tail risk, joint conditional distributions, main shocks.
JEL Classifications: C53, E23, E32, E44.

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. An earlier version of this paper circulated under the title "Uncertainty and Financial Stability: a VAR Analysis." We thank for comments on this and the earlier version Todd Clark, Domenico Giannone, Matteo Iacoviello, Haroon Mumtaz, Andrea Prestipino, Giorgio Primiceri, John Rogers, Minchul Shin, and seminar participants of the XIII Conference on Real-Time Data Analysis, Methods and Applications (Madrid); Conference on Uncertainty and Economic Activity: Measurement, Facts and Fiction (Beijing); Workshop in Structural VAR models (Queen Mary University, London); Computing in Economics and Finance Conference (Milan); IAAE Conference (Montreal); the Federal Reserve Board GRUV workshop; the NBER EFSF Mid-Year Meetings (Chicago), the University of Melbourne Macroeconomic Policies and Financial Stability Workshop, the 2019 SITE-Stanford Summer Workshop, the 2019 European Economic Association Conference (Manchester), the 2020 Econometric Society World Congress, and the 2020 CFE Conference. All errors are our sole responsibility.

†International Finance Division, Federal Reserve Board, Washington DC; dario.caldara@frb.gov.
‡Division of Financial Stability, Federal Reserve Board, Washington DC; chiara.scotti@frb.gov.
§Division of Financial Stability, Federal Reserve Board, Washington DC; molin.zhong@frb.gov.

1 Introduction

In this paper, we study the joint conditional distribution of GDP growth and corporate credit spreads using a stochastic volatility vector autoregression (SV-VAR). Our measures of uncertainty—defined as the volatility implied by the conditional distributions—and risk—defined as the size and location of the tails—for these two variables display significant cyclical co-movements. Our main finding is that these dynamics are driven by the interaction of two shocks, a main business cycle driver as in Angeletos et al. (2020) and a main financial driver that captures exogenous developments in financial markets. Both shocks generate co-movement between the conditional mean and volatility of the endogenous variables, but only their interaction can account for the sizeable movements in uncertainty and risk around crises. Our results highlight the importance of using a multivariate nonlinear model that allows for the joint measurement and analysis of uncertainty and risks around the macroeconomic and financial outlook.[1]

Measurement is the foundation of our analysis. The model and estimation must produce measures of macroeconomic and financial uncertainty and risk that are accurate and plausible. To do so, we build on three features discussed in Section 2. First, in the SV-VAR, the means of the endogenous variables influence their volatilities and, similarly, the volatilities feed back into their means. The nonlinear interaction between mean and volatility gives rise to flexible conditional distributions that are state-dependent and whose mean, volatility, and skewness vary over time. For example, Figure 1 offers a visual representation, plotting the (marginal) conditional distributions for GDP growth and corporate spreads in a quarter of low volatility, 2006:Q4 in black, and in a quarter of high volatility, 2008:Q4, in blue. Second, we estimate the model using Bayesian techniques, producing conditional distributions that factor in uncertainty about the parameters of the model, the underlying volatility states, and the realization of future shocks. Finally, we estimate the model recursively using real-time data, thus utilizing only information available to economic agents and policy makers at each point in time.

We present our measurement results in Section 3. The estimation runs over the period 1947-2019. We construct real-time predictive densities starting in 1973 for GDP growth, to measure macroeconomic uncertainty and risk, and for corporate spreads to measure financial uncertainty and risk. We find

[1] Our data and additional material are available at https://caldara-data-share.github.io/CSZ/CSZ-data-update.

Figure 1: Conditional Distributions: The Global Financial Crisis

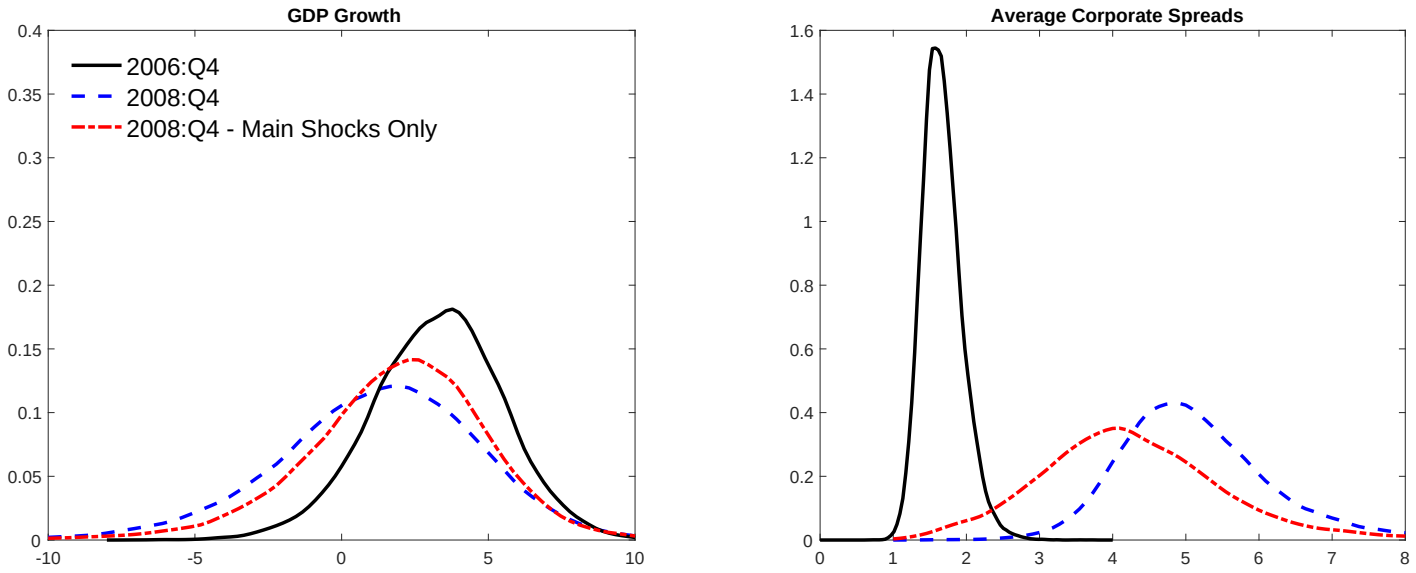

Note: The two panels plot four-quarter ahead marginal conditional distributions of average GDP growth and BAA 10-year corporate credit spreads. The black distributions are computed using the real-time data vintage in 2006:Q4, the blue distributions are computed in 2008:Q4. The red distributions are computed by running a counterfactual that keeps only main shocks active between 2007:Q1 and 2008:Q4.

that movements in volatility are an essential feature of business and financial cycles. The correlation between mean and volatility generates sizeable movements in downside and upside risks, with the former being more volatile than the latter. The emergence, importance, and correlation of macroeconomic and financial risks vary over time. State-dependence and nonlinearities amplify movements in uncertainty and downside risk around recessions, as shown by the distributions in Figure 1.

We also find that the use of real-time data influences key properties of the predictive densities. Measures of macroeconomic uncertainty and risk are larger under a real-time, recursive model estimation, compared with those of a full-sample and final-vintage estimation. Interestingly, the real-time correlation between uncertainty indicators and economic and financial conditions is lower than for final-vintage data. Those findings stress the importance of constructing real-time indicators of uncertainty and risk, especially when using those indicators for policy description and evaluation (Orphanides, 2001).

The resulting measures of uncertainty and risk compare well with two alternatives produced by larger and more flexible models such as Ludvigson et al. (2021), who use a two-step approach originally developed in Jurado et al. (2015) based on factor models to measure financial and macroeconomic

uncertainty, and Adrian et al. (2019), who also use a two-step approach based on quantile regressions and parametric distributions to quantify risks to the GDP growth outlook. The advantage of our approach, relative to these alternatives, is the ability to jointly measure uncertainty and risks around multiple variables using only one model. Through the joint conditional distribution of the endogenous variables, we can estimate the correlation structure among indicators of uncertainty and risk and their cyclical variation, we can define the probability of joint macroeconomic and financial adverse events—such as financial crises—and we can incorporate all sources of statistical and sampling uncertainty. Finally, as we discuss next, in our framework we can identify the joint determinants of uncertainty and risk for macroeconomic and financial conditions, finding that nonlinearities play a crucial role.

The strong measurement properties of the framework allow us to leverage on the key strength of the model, the ability to perform structural analysis and shock identification. Fluctuations in our model are driven by residuals in the level and volatility equations. Those residuals are correlated, as periods of unexpected negative GDP growth are often associated with high GDP volatility, and sharp rises in corporate spreads and financial volatility typically go hand in hand. Recessions and financial crises tend to combine these toxic dynamics all at once, making the identification of first and second moment shocks particularly challenging.

Accordingly, as discussed in Section 4, instead of identifying level and volatility shocks, we make use of the maximum forecast error variance identification proposed by Uhlig (2003), and more recently by Angeletos et al. (2020). As in these papers, we identify a main business cycle shock as the shock that maximizes the forecast error variance of GDP at the business cycle frequency. While details of the implementation vary, the intuition is similar: this shock, while not having necessarily a structural counterpart in the theory, generates co-movements that dominate the business cycle. Relative to these papers, we modify this identification and apply it to a nonlinear model with time-varying volatility. In addition, we sequentially identify a main financial shock, which maximizes the forecast error variance of corporate spreads. This second shock, in the spirit of Gilchrist and Zakrajšek (2012), captures co-movements attributable to financial developments that cannot be accounted for by the main business cycle shock. According to this taxonomy, main shocks capture exogenous events that can feature unexpected developments both in the level and the volatility of the variables.

Generalized impulse responses reveal that adverse main business cycle and main financial shocks generate qualitatively similar co-movements in mean and volatility, reducing activity, tightening financial conditions, and raising macroeconomic and financial volatility. These shocks, by generating adverse level and volatility conditions, can account for sizeable movements in uncertainty and risk, with their effects varying in periods of low and high volatility.

Understanding the determinants of uncertainty and risk is particularly important around recessions and crises. Accordingly, in Section 5, we study how the model parses the 2008-2009 Global Financial Crisis (GFC). Running counterfactuals, with the resulting conditional distributions displayed in Figure 1, we find that the interaction of main business cycle and main financial shocks is crucial to generate sizeable movements in uncertainty and risk—as well as of median outcomes, especially for GDP growth. As we detail in the paper, removing either shock—or considering their additive effects—leads to a substantial underestimation of uncertainty and adverse tail risks, as the importance of an individual shock in this model is amplified by its toxic interaction with other shocks through the presence of nonlinearities.

Relationship to the Literature

The key contribution of our paper is to explore the determinants of fluctuations in uncertainty and risk over the business and financial cycles, and we do so by building on the 'main shock' identification literature. In this sense, we are closely related to the work of Uhlig (2003), Angeletos et al. (2020), and Basu et al. (2021). Relative to these papers, we evaluate the ability of main shocks to produce movements in the width and tails of the conditional distributions and their state-dependent effects. We find that, in a nonlinear model, the interaction between two shocks is necessary to jointly capture fluctuations in macroeconomic and financial uncertainties and risks, and that these shocks have larger effects in recessions. In addition, the use of a nonlinear model and of the main shock identification approach differentiates us from most of the existing literature on uncertainty shocks, which mainly focuses on the identification of uncertainty shocks using linear VAR models (Bloom, 2009; Bachmann et al., 2013; Ludvigson et al., 2021).

Our paper is closely related to Adrian et al. (2021) who, using a nonparametric approach to study the joint conditional distribution of macroeconomic and financial conditions, find that a sharp tightening

of financial conditions leads to the emergence of additional modes in the distributions. Relative to their work, our results emphasize the importance of the nonlinear interaction between macro and financial shocks in generating large fluctuations in uncertainty and risk. We also aim at connecting empirical research on business cycle fluctuations, uncertainty and tail risks through a unified nonlinear parametric framework based on modelling assumptions widely used in the literature. Our work is also related to Orlik and Veldkamp (2014), who discuss the relationship between uncertainty and disaster risk. Furthermore, they provide empirical support for the use of forecasting models that feature skewness and incorporate parameter uncertainty, as these features can account for the bias in professional forecasters' GDP predictions. Gorodnichenko and Ng (2017) use factor models to explore the relationship between mean and volatility dynamics and their nonlinear interaction.

We add to the growing literature that employs SV-VARs to measure uncertainty and to study its effects on economic activity (Mumtaz and Zanetti, 2013; Jo, 2014; Creal and Wu, 2017; Carriero et al., 2018; Shin and Zhong, 2020), which builds on the work of Primiceri (2005) and Cogley and Sargent (2005). We use the econometric specification and posterior sampler of Mumtaz (2018), but focus on the joint conditional distribution as the main empirical object of interest. The conditional distributions reveal that these models have very broad appeal for the joint measurement of uncertainty and risk and to study their sources, offering insights coming from nonlinearities built in the model that cannot be directly derived from models such as linear VARs. In a recent and contemporaneous paper, Carriero et al. (2020) also discuss the performance of SV-VAR for measuring risk, although their focus is on the density and quantile forecasting properties of the model in comparison to quantile regressions.[2]

More broadly, quantifying and understanding the evolution of uncertainty and risk for macroeconomic and financial conditions is crucial for economic decisions and for policy. A vast theoretical literature discusses how consumption and investment decisions by households and firms respond to uncertainty and risk.[3] Risk management—the assessment of what could go wrong with the economy

[2]Our approach is complementary to studies that construct proxies of uncertainty and risk using textual analysis (Bloom, 2009; Baker et al., 2016; Caldara and Iacoviello, 2018; Husted et al., 2020), economic data releases (Scotti, 2016), or firm-level data (Hassan et al., 2019; Salgado et al., 2019). Bloom (2014), Fernandez-Villaverde and Guerron-Quintana (2020), and Cascaldi-Garcia et al. (2020) review the literature.

[3]See Bernanke (1983); Fernández-Villaverde et al. (2011, 2015); Christiano et al. (2014); Gilchrist et al. (2014); Basu and Bundick (2017), among many others.

and judging whether policy should be adjusted to minimize risks—is at the forefront of policy decision-making. For example, risk management considerations have been an important factor in deciding the stance of monetary policy by the Federal Open Market Committee (FOMC), as documented by Evans et al. (2015), and are also at the core of macroprudential policies.

2 Econometric Framework

In this section, we first present a class of VAR models with stochastic volatility that allows for a rich interaction between mean and volatility. We then discuss the measurement of uncertainty and risk and derive the moments of the posterior conditional distributions for macroeconomic and financial conditions generated by these models. We show that, due to the non-linear interaction between level and volatility disturbances, the densities implied by a SV-VAR feature time-varying uncertainty and risk, and contrast them to those from a standard linear VAR model. Finally, we describe the empirical setup used through the paper. Our model includes one real activity indicator—GDP growth—and one financial indicator—corporate credit spreads measured as the Baa corporate bond yield relative to the yield on 10-Year Treasuries. We employ a Bayesian environment and use real-time data.

2.1 Stochastic Volatility Vector Autoregressions

An SV-VARs can generate posterior densities that exhibit flexible and empirically relevant time-varying dynamics in uncertainty, risk, and higher-order moments such as skewness. In this section, we provide intuition using a simple univariate model.

An SV-VAR model describes the evolution of the endogenous variable z_t as follows:

$$z_t = c_z + \sum_{p=1}^{P} \beta_p z_{t-p} + \sum_{k=1}^{K} b_k h_{t-k} + H_t^{1/2} e_t, \tag{1}$$

$$h_t = c_h + \sum_{j=1}^{J} \theta_j h_{t-j} + \sum_{q=1}^{Q} d_q z_{t-q} + S^{1/2} \eta_t. \tag{2}$$

In equation (1)—the level equation—z_t depends on a constant c_z, on lags of z_t, on lags of the unobserved log volatility process h_t, and on the residual e_t. The term H_t is the defining element of SV-VAR models, a time-varying volatility process of the residual e_t. In equation (2)—the volatility equation—the volatility vector h_t depends on a constant c_h, on lags of h_t and z_t, and on a residual η_t. The literature offers various options to define the properties of the error terms in equations (1) and (2). In our paper, we follow Mumtaz (2018) and we assume that:

$$\varepsilon_t = \begin{pmatrix} e_t \\ \eta_t \end{pmatrix} \sim N(0, \Sigma), \quad \Sigma = \begin{pmatrix} \Sigma_\eta & \Sigma'_{\eta e} \\ \Sigma_{\eta e} & \Sigma_e \end{pmatrix}. \tag{3}$$

Equation (3) defines the distributions of the level and volatility residuals. The variance-covariance matrix Σ has diagonal elements normalized to 1, as we have factored out the volatility of the residuals. The term $\Sigma_{\eta e}$ in the covariance matrix denotes the correlation between the residuals in the level and volatility equations. Denoting the matrix A_0 such that $A_0' A_0 = \Sigma$, we identify orthogonal shocks by using standard identification strategies that impose restrictions on the elements of A_0.

The contemporaneous and lagged correlation between the level and volatility variables allows the model to generate endogenous volatility—as h_t can respond immediately to disturbances in the level equation and with delay to movements in z_t—and mean effects of volatility—as z_t can also respond immediately to disturbances in the volatility equation and with delay to movements in h_t. The presence of stochastic volatility in equation (1) generates a nonlinearity resulting in state-dependent effects of exogenous shocks. We illustrate these properties of the model in the next subsection.

2.2 Understanding Predictive Distributions from SV-VARs

Denote a model and the associated vector of parameters by $\mathcal{M}$ and Θ, respectively. Let $z^t \equiv \{z_\tau\}_{\tau=1}^t$ denote the data available at time t. We assume that our information set $\mathcal{I}_t$ comprises the model $\mathcal{M}$ and the data z^t. At each point in time t, we forecast the distribution of future outcomes $p(z_{t+1:t+F}|\mathcal{I}_t)$, where F represents the end of the forecast horizon and p denotes the posterior predictive distribution of the endogenous variables. We can write the posterior predictive distribution as:

$$p(z_{t+1:t+F}|\mathcal{I}_t) = \int_{\Theta}\int_{H_t}\left[\int_{H_{t+1:t+F}} p(z_{t+1:t+F}, H_{t+1:t+F}|\mathcal{I}_t, H_t, \Theta)dH_{t+1:t+F}\right]\times \\ p(H_t|\mathcal{I}_t, \Theta)p(\Theta|\mathcal{I}_t)dH_t d\Theta. \tag{4}$$

The posterior predictive distribution in equation (4) describes the entire distribution of possible future values for GDP growth and corporate spreads, conditional on the observed data and the model. Posterior distributions allow us to compute point forecasts, the typical object of interest in the forecast literature, as well as higher-order moments of the forecast, the focus of this paper. The posterior distribution incorporates three sources of uncertainty. First, there is uncertainty around the parameters of the model, reflected in the density $p(\Theta|\mathcal{I}_t)$. Second, there is uncertainty about the current unobserved volatility state, reflected in the density $p(H_t|\mathcal{I}_t, \Theta)$. Finally, there is uncertainty regarding the future realization of the random disturbances, which in turn generates uncertainty regarding the future evolution of the endogenous variables and volatility states, reflected in $p(z_{t+1:t+F}, H_{t+1:t+F}|\mathcal{I}_t, H_t, \Theta)$.[4] We discuss the algorithm used to generate the posterior predictive distributions in Appendix E.

In what follows, we describe the posterior densities generated by a simple univariate SV-VAR and compare them with those from a linear bivariate VAR model. Linear VAR models are widely used in the literature to study the effects of uncertainty shocks, and they differ from SV-VARs in two important dimensions. First, the estimation of this model requires data on both z_t and h_t. Second, the error e_t is homoscedastic and h_t is another level variable that is assumed to be observable, typically a volatility measure extracted from an auxiliary model or a proxy for uncertainty.

We compute the predictive distributions from these two models—the univariate SV and the linear VAR—by running the following experiment. We use the univariate SV model to generate 1,000 observations for z_t and h_t.[5] We assume that two orthogonal shocks drive the model dynamics. The shocks have a recursive structure, with the first shock—the main shock—explaining all the variation in e_t and some portion of η_t, and the second shock—the secondary shock—explaining the variation in η_t

[4] If either K or Q is greater than 1, there is further dependence on lags of H_t. We suppress this dependence in the main text for ease of notation.

[5] We simulate 100,000 observations and use only the last 1,000 observations. Table A.3 in Appendix B reports the values assigned to the model parameters. One choice worth emphasizing is the negative covariance between e_t and η_t in matrix Σ, which we set to $\zeta = -0.8$.

unaccounted for by the first shock. We use data on z_t and h_t to estimate the linear VAR using OLS, and identify two orthogonal shocks using the same recursive structure. For the SV model we calculate $p(z_{t+1}|z^t, \mathcal{M}^{SV}, \overline{\Theta}, \overline{H}_t)$, the one-step ahead predictive distribution for z_{t+1} conditional on setting the parameters to their population value ($\overline{\Theta}$), and treating H_t as observable data, $\overline{H}_t$.[6] For the linear VAR we calculate $p(z_{t+1}|z^t, \mathcal{M}^{LIN}, \hat{\Theta}^{OLS}, \overline{H}_t)$, the one-step ahead predictive distribution for z_{t+1} conditional on setting the parameters to the OLS estimates and treating H_t as observable data. Thus, for this simple illustrative exercise, the predictive distributions from both models include only uncertainty on the realization of future shocks.

The top row of Figure 2 shows predictive densities at different horizons from the SV-VAR calculated in a high volatility state (left panel) and in a low volatility state (the right panel). The blue distributions denote baseline forecasts conditional on the data and the estimated state, while the red distributions denote distributions conditional on adding a two-standard-deviation level shock.[7] The bottom row replicates the same experiments using the linear VAR model. Note that in all panels, we use main shocks of identical size.

The comparison of the predictive distributions between the two models reveals several interesting features. First, the predictive densities from the SV-VAR are state-dependent: there is more uncertainty and left skewness around the forecast in periods of elevated volatility relative to periods of low volatility. By contrast, the predictive densities of the linear homoscedastic VAR have constant dispersion and no skewness. Second, the main shock changes mean, variance, and skewness of the posterior distributions generated by the SV-VAR, and the effects of the shock are larger in the high volatility state. The main shock only shifts the mean of the distributions generated by the linear VAR. Third, a direct implication of the previous point is that in the SV-VAR model, the left tails of the posterior distributions are more volatile than the right tails, while the tails move in lockstep in the linear VAR. If we take z_t to be GDP growth, a stable right tail and a volatile left tail of the predictive distributions is one defining feature

[6] We obtain similar results by estimating the model on the artificial data, as the posterior distributions are centered closely around the true parameters. In Appendix F, we show that we can recover the true parameters in a simulation exercise using the full bivariate model.

[7] In a SV-VAR model, the forecast is time-varying and depends on the past realization of the shocks, as they shape the dynamics of H_t. Accordingly, we build the densities under the alternative scenario adding a two-standard deviation main shock to the simulated paths following the method discussed in Gonçalves et al. (2021).

Figure 2: Posterior Predictive Distributions
(Stochastic Volatility vs. Linear VARs)

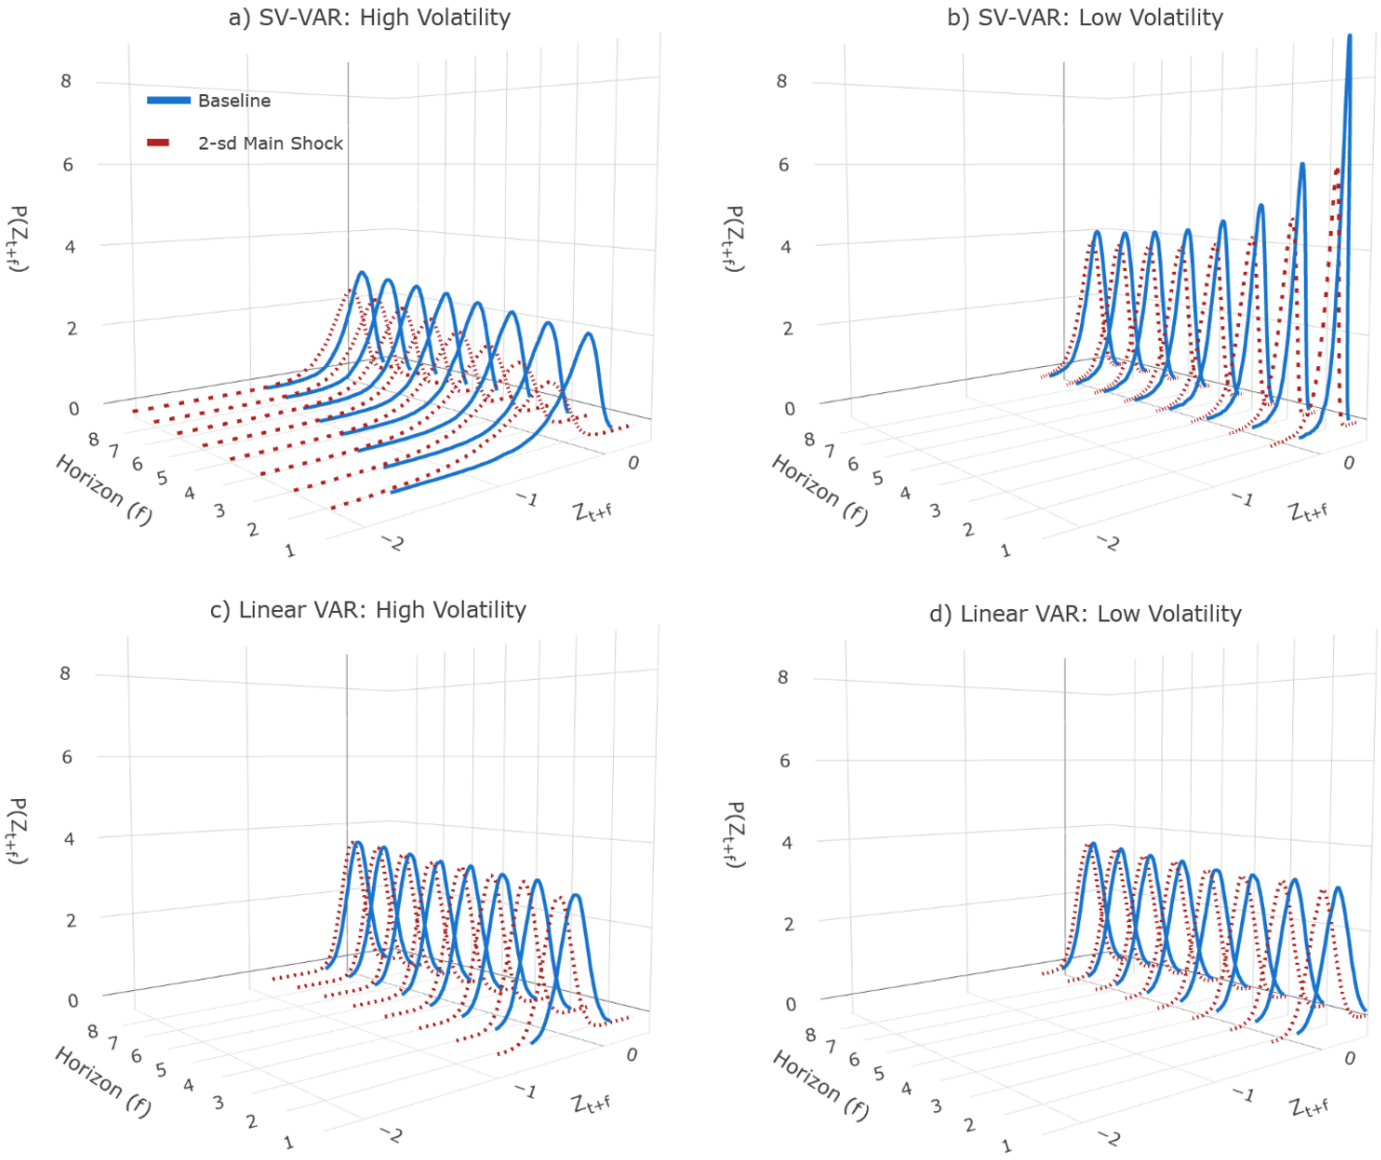

Note: The top row shows the posterior predictive distributions generated by the univariate baseline SV-VAR model in a high volatility state (left panel) and low volatility state (right panel). The bottom row shows the distributions generated by a linear bivariate VAR model in the same high and low volatility states. We compute these distributions under a baseline scenarios (blue distributions) and adding a two-standard deviation main shock to z_t and H_t (red dashed distributions).

of the growth-at-risk estimates in Adrian et al. (2019). The model generates asymmetric movements in the tails through two channels: skewness and the negative correlation between mean and volatility, as a simultaneous decline in mean and increase in volatility have a disproportionate effect on the left

tail of the distribution relative to the right tail. The intuition is that both lower mean and higher volatility contribute to increasing the likelihood of adverse events, while having opposite effects on the likelihood of positive events (a lower mean makes positive events less likely while a higher variance makes them more likely). In Section 3, when illustrating the application to U.S. data, we provide additional discussion on the role of mean, volatility, and skewness in shaping tail risk.

In Appendix C, we complement the analysis in this section and discuss how each ingredient included in the SV-VAR model affects the predictive distributions by considering restricted versions of the baseline model. We devote particular emphasis to explaining how skewness emerges through the contemporaneous feedback between mean and volatility.

2.3 Measuring Risk and Uncertainty

Our interest lies in tracking the evolution over time of the uncertainty and risk surrounding future macroeconomic and financial conditions, and in understanding their determinants. Accordingly, in our application, we use the SV-VAR models to construct density forecasts for GDP growth and corporate credit spreads. We use Bayesian techniques combined with real-time data to estimate sequentially the parameters of our model and the unobserved volatility states. The framework offers a natural way to incorporate all sources of uncertainty related to future economic and financial conditions.

To measure uncertainty and risk in either future macroeconomic or financial conditions in a multivariate version of the model, we use the *marginal* posterior predictive distribution $p(z_{i,t+f}|\mathcal{I}_t)$ for $i = 1, \dots N$, which characterizes the distribution of a particular variable z_i at forecast horizon f:

$$p(z_{i,t+f}|\mathcal{I}_t) = \int_{z_{-i,t+f}} p(z_{t+f}|\mathcal{I}_t) dz_{-i,t+f}, \text{ for } i = 1, \dots, N; \ f = 1, \dots, F. \tag{5}$$

To calculate the marginal predictive distribution, we integrate the joint density forecast $p(z_{t+f}|\mathcal{I}_t)$ over $z_{-i,t+f}$, the forecast at horizon f of all endogenous observed variables except for the variable of interest i.

From the marginal distributions, we compute the following measures of uncertainty (U), risk—expected shortfall (SF) and expected longrise (LR), and Kelley skewness (KS)—as:

$$U_{z_i,t}(f) = \sqrt{VAR_t\left[z_{i,t+f}\right]}, \tag{6}$$

$$SF_{z_i,t}(f) = E_t\left[z_{i,t+f} | z_{i,t+f} < q_\alpha(z_{i,t+f})\right], \tag{7}$$

$$LR_{z_i,t}(f) = E_t\left[z_{i,t+f} | z_{i,t+f} > q_{1-\alpha}(z_{i,t+f})\right] \tag{8}$$

$$KS_{z_{i,t}}(f) = \frac{q_{90}(z_{i,t+f}) - q_{50}(z_{i,t+f})}{q_{90}(z_{i,t+f}) - q_{10}(z_{i,t+f})} - \frac{q_{50}(z_{i,t+f}) - q_{10}(z_{i,t+f})}{q_{90}(z_{i,t+f}) - q_{10}(z_{i,t+f})}, \tag{9}$$

where $q_\alpha(z)$ is the α-percent quantile of the marginal predictive distribution of z, and E_t and VAR_t denote $E[\cdot|\mathcal{I}_t]$ and $VAR[\cdot|\mathcal{I}_t]$, respectively. Uncertainty is the square root of the conditional variance $VAR_t\left[z_{i,t+f}\right]$. This definition of uncertainty follows Orlik and Veldkamp (2014), Jurado et al. (2015) and Ludvigson et al. (2021). The risk measures $SF_{z_i,t}(f)$ and $LR_{z_i,t}(f)$ in equations (7) and (9) track the evolution of the tails of the marginal predictive distributions for variable z. The expected shortfall (longrise) at the α percent level of a variable z_i at the forecast horizon f is the expected value of z conditional on realizations in the bottom (top) α percent of the distribution. We use Kelley's measure of skewness, following Salgado et al. (2019), which calculates the right tail versus the left tail shares.[8]

In our application, shortfall and longrise for the growth rate of GDP—$SF_{GDP,t}(f)$ and $LR_{GDP,t}(f)$—represent downside and upside macroeconomic risks, respectively. These macroeconomic risk measures are closely related to the measures constructed by Adrian et al. (2019) to study GDP growth-at-risk. We complement the study of macroeconomic risk with new measures of financial risk, which we extract from corporate spreads. The shortfall and longrise measures for corporate credit spreads—$SF_{CS,t}(f)$ and $LR_{CS,t}(f)$—have the opposite interpretation. Shortfall in spreads represents upside financial risk, as the left tail of the spread distribution covers low level of spreads that typically reflect buoyant financial conditions. The longrise in spreads represents downside financial risk, as elevated levels of spreads represent tight financial conditions.

The measures of uncertainty and risk defined above—and typically reported in the literature—are based on forecasts for one variable at one specific horizon. However, business cycle dynamics vary in

[8]From equation (4), we can easily calculate probabilities associated with transformations of z, such as average GDP growth or cumulative growth.

their degree of persistence and correlation across variables. For instance, recessions are characterized by consecutive quarters of negative GDP growth, and financial crises pair recessions with very tight financial conditions. We use the joint posterior distribution in equation (4) to calculate the time-varying probabilities of adverse scenarios, defined as the joint realization of multiple events:

$$Pr(\mathbf{a} < z_{t+1:t+F} < \mathbf{b}|\mathcal{I}_t) = \int_{\mathbf{a}}^{\mathbf{b}} p(z_{t+1:t+F}|\mathcal{I}_t)dz_{t+1:t+F}, \tag{10}$$

where $\mathbf{a}$ and $\mathbf{b}$ denote vectors of thresholds that vary across variables and forecast horizons. These probabilities provide a complementary measure of risk that can be particularly relevant for policy-makers. The calculation of probabilities of events defined by the joint dynamics of multiple variables over multiple horizons requires a joint probability distribution across variables and time.[9]

2.4 Data, Estimation, and Model Specification

Our empirical exercise is based on a bivariate model estimated on U.S. GDP growth and corporate credit spreads at quarterly frequency over the 1947:Q2-2019:Q4 sample. Data on GDP are from the Federal Reserve Bank of Philadelphia Real-Time Data Set for Macroeconomists. For corporate credit spreads, from 1947:Q2 through 1953:Q2, we take the difference between Moody's Seasoned Baa Corporate Bond Yield (BAA) and the 10-Year Treasury Constant Maturity Rate (DGS10).[10] From 1953:Q3 onward, we use the Moody's Seasoned Baa Corporate Bond Yield Relative to the Yield on 10-Year Treasury Constant Maturity (BAA10Y). Quarterly values are calculated as averages of daily observations. We use corporate credit spreads because they are an important indicator of asset valuations, and spreads have emerged in empirical work as reliable and timely indicators of financial conditions (Gertler, 2012; Gilchrist and Zakrajšek, 2012).

In our baseline empirical exercise, we estimate the model recursively. In particular, we use data from 1947:Q2 through 1953:Q1 as pre-sample to calibrate the prior distributions. We then start estimating

[9] Adrian et al. (2019) do construct conditional distributions for GDP growth over time. However, by running local projections of (cumulative) GDP growth at different horizons, they cannot compute the joint probability of events involving multiple horizons and variables.

[10] Financial data were downloaded from FRED, Federal Reserve Bank of St. Louis. Mnemonics associated with the series are provided in parenthesis.

recursively the model from the data vintage for GDP growth marked as 1973:Q2 in the Real-Time Data Set for Macroeconomists. The last observation available in this vintage is for 1973:Q1. Each subsequent vintage for GDP growth differs from the previous by including one additional observation and the revisions to previous observations released during the quarter. Since credit spreads are not subject to revision, we simply add one observation per quarter to each estimation.

The baseline model is specified as follows. In the level equation (1), we include 4 lags of the level variables ($P = 4$) and 2 lags of volatility ($K = 2$). In the volatility equation (2), we include 1 lag of the volatility ($J = 1$), and 2 lags of level variables ($Q = 2$). We take $150,000$ draws from the posterior distribution, burn the first $70,000$, and sample every 20th draw to produce our posterior statistics of interest. Details about the prior distributions and the posterior sampler in Appendix D.

3 Results

This section presents our main estimation results. In doing so, we first compute predictive distributions for GDP growth and corporate spreads one-quarter and one-year ahead (calculated as average GDP growth and average corporate spreads over the year). We then extract the uncertainty and risk measures defined in Section 2.3 at these two horizons. We label GDP-related measures as *macroeconomic* (e.g., macroeconomic uncertainty) and spreads-related measures as *financial* (e.g., financial uncertainty). To understand their importance for measurement of real-time data and of parameter and state uncertainty, we present three sets of estimates: our baseline estimates, which include real-time data, and parameter and state uncertainty (the red solid lines in the figures); real-time estimates with no parameter and state uncertainty, which we obtain by fixing parameters and states to their posterior means (the dashed blue lines); and final-vintage estimates based on the data vintage 2020:Q1—with data through 2019:Q4—that include parameter and state uncertainty (the dash-dotted green lines).[11] Finally, focusing on the real-time exercise, we compute the probability of a financial crisis and we compare our real-time measures of uncertainty and risk to two widely cited papers in the literature (Adrian et al., 2019; Ludvigson et al., 2021).

[11] Final-vintage estimates of uncertainty and risk are based on only one model estimation that uses the full sample to estimate the parameters, and where the volatility processes are smoothed.

Figure 3: Macroeconomic and Financial Uncertainty

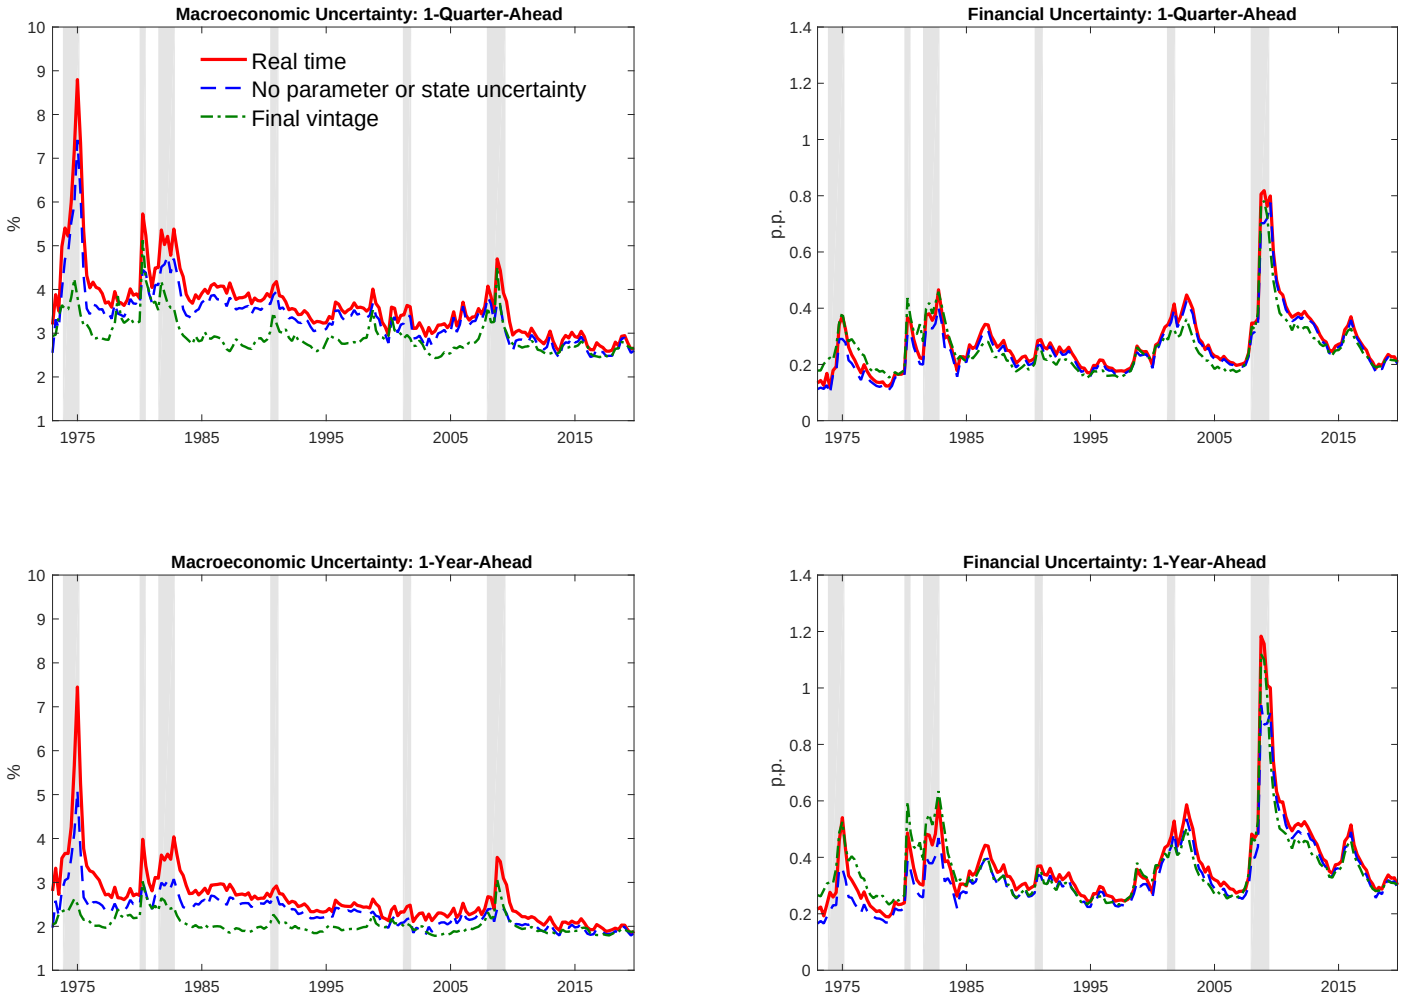

Note: The figure plots macroeconomic and financial uncertainty at the one-quarter horizon (top row) and one-year horizon (bottom row). The one-quarter ahead macro uncertainty estimates are annualized. The red line is our baseline measure, the dashed blue line is a measure that excludes parameter and state uncertainty, and the dash-dotted green line is a measure based on full-sample final-vintage data. The gray shaded areas are recession dates as defined by the National Bureau of Economic Research.

3.1 Macroeconomic and Financial Uncertainty

Figure 3 shows the estimated values of macroeconomic and financial uncertainty one-quarter ahead (top panels) and one-year ahead (bottom panels) for the three alternative measurement strategies explained at the beginning of the section. Table A.1 in Appendix A reports selected moments of these series and their correlation with the model's mean forecasts.

All measures of macroeconomic and financial uncertainty are characterized by a clear pattern of

cyclical variation, being negatively correlated with expected GDP growth and positively correlated with expected financial conditions. This correlation between the mean and the uncertainty of the forecasts, as well as the cross-correlation between macro and financial moments, plays a central role in generating time-variation in downside risk, as we discuss in the next section.

The three set of estimates have also important differences. The baseline measures of macroeconomic uncertainty (the solid red lines) have two important features. First, macro uncertainty shows a downward trend, *gradually* declining since the early-1980s from values above 5 percent to around 3.5 percent at the one-quarter horizon. Second, major peaks in macro uncertainty occurred before the beginning of the trend decline, during the recessions of the 1970s and early-1980s. During the 2008-2009 GFC, uncertainty increased to readings far higher than those immediately before and after the recession, but uncertainty was lower than its values at the beginning of the sample.

The final-vintage estimates (the dash-dotted green lines) do not share these two features. Final-vintage macro uncertainty captures the Great Moderation earlier than real-time estimates, dropping after the 1981-1982 recession and averaging at the same level thereafter. This result is a stark example of the effect of using future data when smoothing volatility, as the model can anticipate the low volatility times ahead. Moreover, the historical importance of the peak in uncertainty estimated during the GFC is much greater when using the final-vintage. In fact, at the one-year ahead horizon, the GFC is found to be the period with the largest uncertainty estimate, a fact that is plainly at odds with the real-time estimates. More broadly, the over-the-sample average of final-vintage macro uncertainty is about 20 percent lower than real-time uncertainty at the one-quarter and one-year horizons, respectively.[12]

The effect of ignoring parameter and state uncertainty (the dashed blue lines) is more modest, with macro uncertainty measures that are about 10 percent lower than our baseline measures. However, there are time periods when the differences can be quantitatively large, especially in recessions. For instance, the measure excluding parameter and state uncertainty underestimates uncertainty in the mid-1970s

[12]There are three main drivers of the differences between real-time and final-vintage estimates. First, parameter estimates differ over time when using real-time data. Parameters estimated with data ending in the 1970s are less precisely estimated and based on data with different features compared with the full sample, e.g. more volatile GDP growth as we exclude the Great Moderation. Second, volatility processes smoothed using the full sample are different than their real-time counterparts. Third, at certain time periods, such as the mid-1970s recession, large revisions to GDP growth across vintages leads to large differences in volatility and therefore uncertainty. For instance, the real-time reading of GDP growth was up to 1.5 percentage points lower than in the final-vintage for the 1973-1975 recession.

recession by over 1.5 percentage points.

Financial uncertainty behaves somewhat differently compared with macro uncertainty, and there are only minor differences between the baseline measures and those ignoring real-time and parameter/state uncertainty. For all measures, financial uncertainty does not show evidence of any long-term trend. It was elevated in the recessions of the mid-1970s and early-1980s, before reaching low levels during the Great Moderation, and it increased again around the early-2000s recession and the GFC, reaching the highest spike in the series during the latter episode.

3.2 Macroeconomic and Financial Tail Risk

Figure 4 shows the evolution of macroeconomic and financial tail risk estimates at the one-quarter (top panels) and one-year (bottom panels) horizons, and Table A.2 reports selected moments of these series.

The left panels of Figure 4 display both macroeconomic downside risk—measured by the expected shortfall in GDP growth in the lowest 5 percent of the predictive distribution (the mostly negative lines in the figure)—and macroeconomic upside risk—the expected longrise in GDP growth in the highest 5 percent of the predictive distribution (the positive lines).

Two important dynamics of downside and upside macroeconomic risk emerge when looking at the figure. First, while these risks share similar moves, their correlation is not perfect—note that the correlation between downside and upside risk is only negative 0.08 at the 1-quarter horizon and negative 0.58 at the one-year horizon. Second, the expected shortfall is more volatile than the expected longrise, in line with the evidence in Adrian et al. (2019). These dynamics are driven by two forces: simultaneous movements in the mean and variance of the distribution, and time-variation in skewness.

To understand the importance of mean and volatility shifts, it is useful to think about simpler models: the correlation between shortfall and longrise would be nearly 1 if only the mean of the predictive distributions changed over time, and it would be nearly negative 1 if only the volatility of the predictive distributions changed over time. In the first case, shortfall and longrise would move in lockstep to the left or right of the distribution and in the second case, shortfall and longrise would move by similar amounts both toward the center or both toward the tails of the distributions. Importantly, in all these

Figure 4: Macroeconomic and Financial Tail Risk

Note: The figure plots macroeconomic and financial tail risks at the one-quarter horizon (top row) and one-year horizon (bottom row), defined as the 5 percent shortfall and longrise of the forecast of real GDP growth and average corporate credit spreads. The one-quarter ahead macro risk estimates are annualized. The solid red line is our baseline measure, the dashed blue line is a measure that excludes parameter and state uncertainty, and the dash-dotted green line is a measure based on full-sample final-vintage data. The gray shaded areas are recession dates as defined by the National Bureau of Economic Research.

scenarios, shortfall and longrise would be equally volatile.[13]

To understand the importance of time-varying skewness, Figure 5 plots the evolution of Kelley skewness of the predictive distributions. As explained in Appendix C, the one-quarter ahead predictive distributions have constant skewness absent parameter and state uncertainty. The use of real-time data and the inclusion of parameter and state uncertainty leads only to small movements in skewness and

[13] An illustration of this point using simpler models is in Appendix C.

the implications for time-varying risk are quantitatively modest. However, higher-order moments play a role at the one-year horizon, with more substantial time-variation in skewness, especially for the real-time measures during recessions. For instance, during the 1973-1975 recession, 65 percent of the total distance between the 90th and 10th quantiles of the one-year ahead distribution was below the median.

Regarding the left panels of Figure 4, the dynamics of macroeconomic risk are driven by a combination of fluctuations in level, volatility, and skewness. For instance, during the first recession in the 1980s, downside risk increased while upside risk remained relatively stable, due to a simultaneous decline in the conditional mean and increase in the conditional variance amplified by a surge in left skewness. Macroeconomic tail risk shrank during the Great Moderation, and shortfall and longrise moved more tightly together, consistent with level fluctuations playing a prominent role. Tail risk rose again during the GFC when the expected shortfall reached close to negative 15 percent annualized at the one-quarter horizon.

The financial shortfall and longrise plotted in the right panels share some similarities with their macro counterparts. In particular, the expected longrise—representing an adverse risk for spreads—is more volatile than the expected shortfall. At the same time, the expected shortfall and longrise move more closely in lockstep, with a correlation at the one-quarter horizon that is close to 0.9. Time-variation in volatility and skewness plays a prominent role mostly during recessions, leading to larger movements in the longrise relative to the shortfall.

Finally, there is also a rich correlation structure between macroeconomic and financial moments. As shown in Table A.2, with tighter and more uncertain financial conditions, the macroeconomic shortfall and longrise tend to decrease, translating to larger macroeconomic downside risk (more negative) and lower macroeconomic upside risk (less positive). More adverse macroeconomic conditions (lower macro mean forecasts) are associated with larger adverse financial risk (financial longrise) and less upside risk (financial shortfall). Macro uncertainty has a small and positive correlation with both upside and downside financial risks. We parse the cross-correlation between macro and financial conditions in the counterfactual analysis presented in Section 4.

Figure 5: Macroeconomic and Financial Skewness

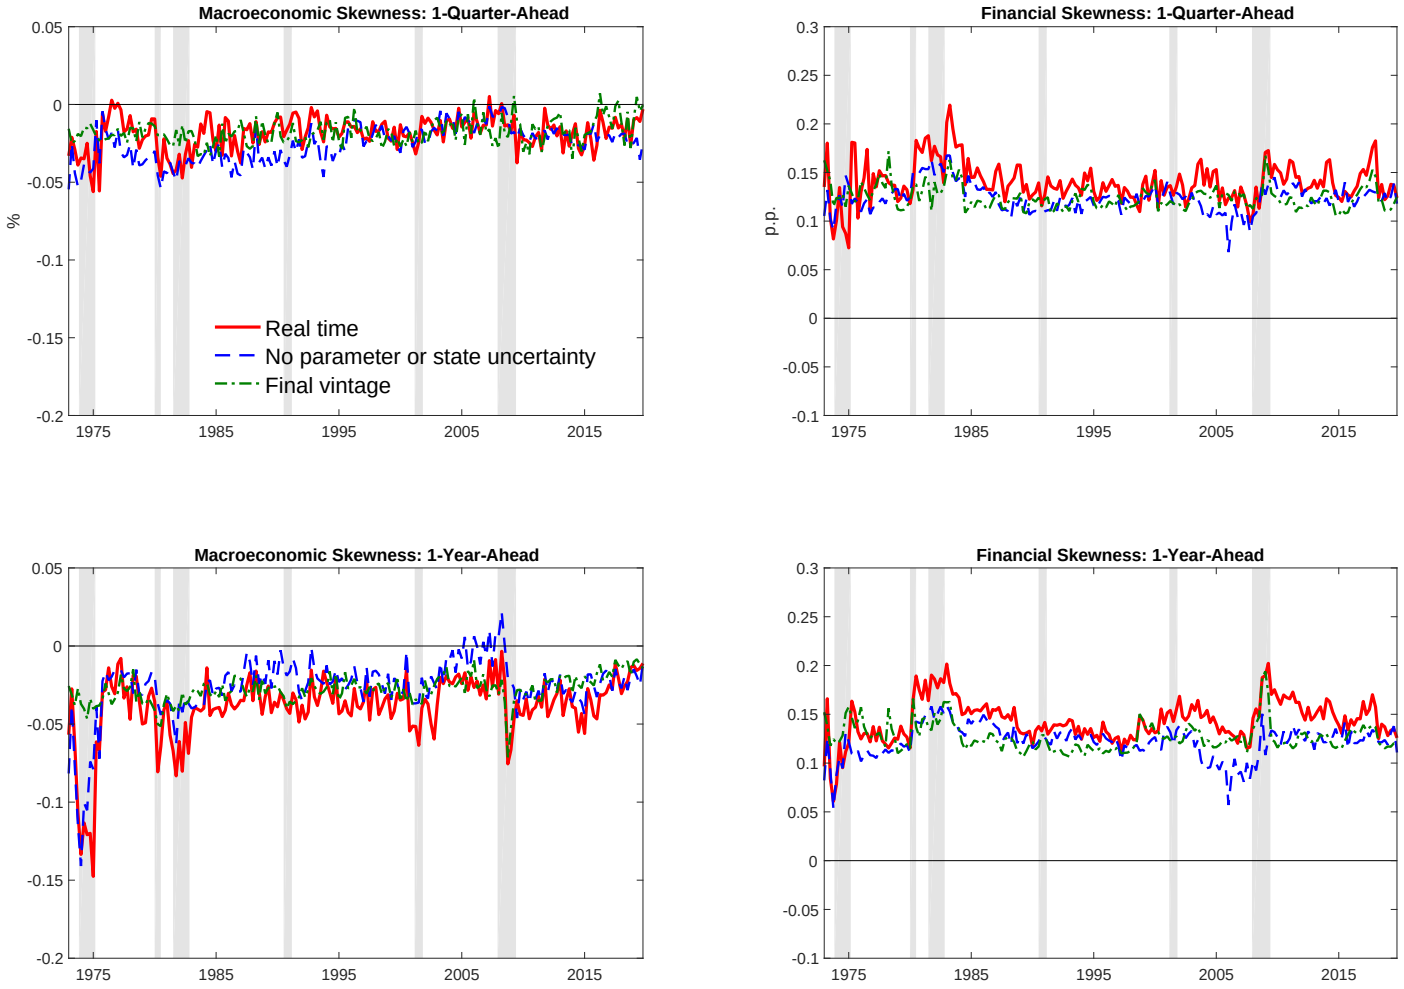

Note: The figure plots macroeconomic and financial Kelley skewness at the one-quarter horizon (top row) and one-year horizon (bottom row). The solid red line is our baseline measure, the dashed blue line is a measure that excludes parameter and state uncertainty, and the dash-dotted green line is a measure based on full-sample final-vintage data. Gray shaded areas are recession dates as defined by the National Bureau of Economic Research.

3.3 Financial Crisis Risk

This section describes an additional application that sets our model apart from existing risk measurement models: the construction of a real-time probability of a financial crisis. Financial crises are events characterized by adverse macroeconomic and financial conditions for prolonged periods, and are typically associated with elevated uncertainty. In short, they constitute a severe risk for the economy, a risk that policy-makers are interested in tracking. Our model is suitable to construct a financial crisis risk indicator for two reasons. First, the model characterizes the joint distribution across variables and

Figure 6: Financial Crisis Risk

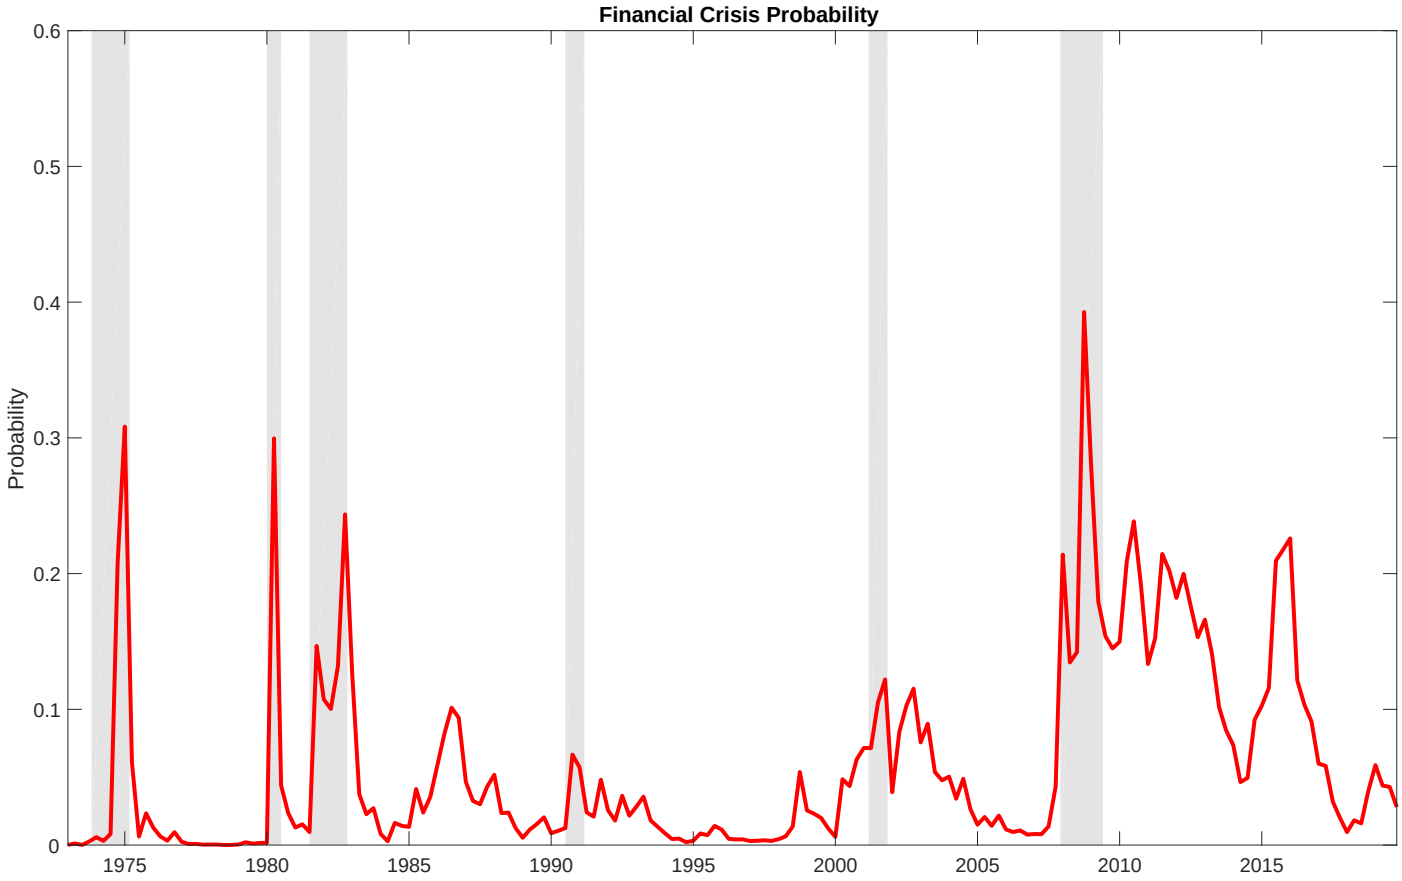

Note: The figure plots the probability of a financial crisis over the next four quarters. A financial crisis is defined as real GDP growth declining for at least two consecutive quarters and spreads rising above 300 basis points, on average, over four quarters. The gray shaded areas are recession dates as defined by the National Bureau of Economic Research.

through time, and thus can be used to define the risk of an event that satisfies multiple criteria, as described in equation (10). We view this feature as an advantage of our parametric approach relative to alternative nonparametric or local projections strategies used in the literature for risk analysis. Second, our model includes corporate spreads, a variable that previous studies have found to be informative about the likelihood of economic downturns and tight financial conditions over the medium-term horizon (Favara et al., 2016).

In such an exercise, we look at the risk of financial crises defined as situations that pair macroeconomic recessions with tight financial conditions: Real GDP growth declines for at least two consecutive

quarters, and spreads rise above 300 basis points, on average, over the scenario. Figure 6 shows the real-time estimated probability of a financial crisis. The probability increases at the onset of recessions, reaching nearly 50 percent during the GFC. The probability is not higher, even at the height of the GFC, because financial crises are periods characterized by elevated macroeconomic and financial uncertainty. Intuitively, coupling negative realizations with elevated uncertainty leads the model to assign a high probability of future negative outcomes, while allowing for the possibility of a rebound in activity and a normalization of spreads during the forecast horizon.

3.4 Comparison with Alternative Measures of Uncertainty and Risk

While the main advantage of our framework is the ability to jointly characterize uncertainty and risk over forecast horizons and across variables, our measures of uncertainty and risk are reminiscent of indicators that have been developed in the literature on uncertainty and growth-at-risk. We therefore compare our real-time versions of uncertainty and risk with measures from two widely cited papers in the literature: (1) Ludvigson et al. (2021) for macroeconomic and financial uncertainty, and (2) Adrian et al. (2019) for macroeconomic (GDP) risk—henceforth LMN and ABG, respectively.

LMN construct measures of macroeconomic and financial uncertainty using a two-step approach. The authors separately estimate factor and stochastic volatility models to quantify the volatility of the purely unforecastable components of future values of a large panel of macroeconomic and financial variables. Macroeconomic uncertainty is the average of individual uncertainty series associated with variables included in the macro dataset by McCracken and Ng (2016). Financial uncertainty is the average of individual uncertainty series associated with stock returns and other selected financial variables.

The upper panels of Figure 7 compare our macroeconomic and financial uncertainty measures with those from LMN. The correlation between macroeconomic uncertainties measures is about 0.7, and the two measures share major spikes. However, there are also clear differences in the evolution of macroeconomic uncertainty over time. Our real-time measure slowly trends down over time, while the LMN measure drops more abruptly in the mid-1980s. As with our final-vintage measure of macroeconomic uncertainty, the LMN measure captures the start of the Great Moderation earlier than a real-time mea-

Figure 7: Uncertainty and Risk Measures:
Comparison with Ludvigson et al. (2021) and Adrian et al. (2019)

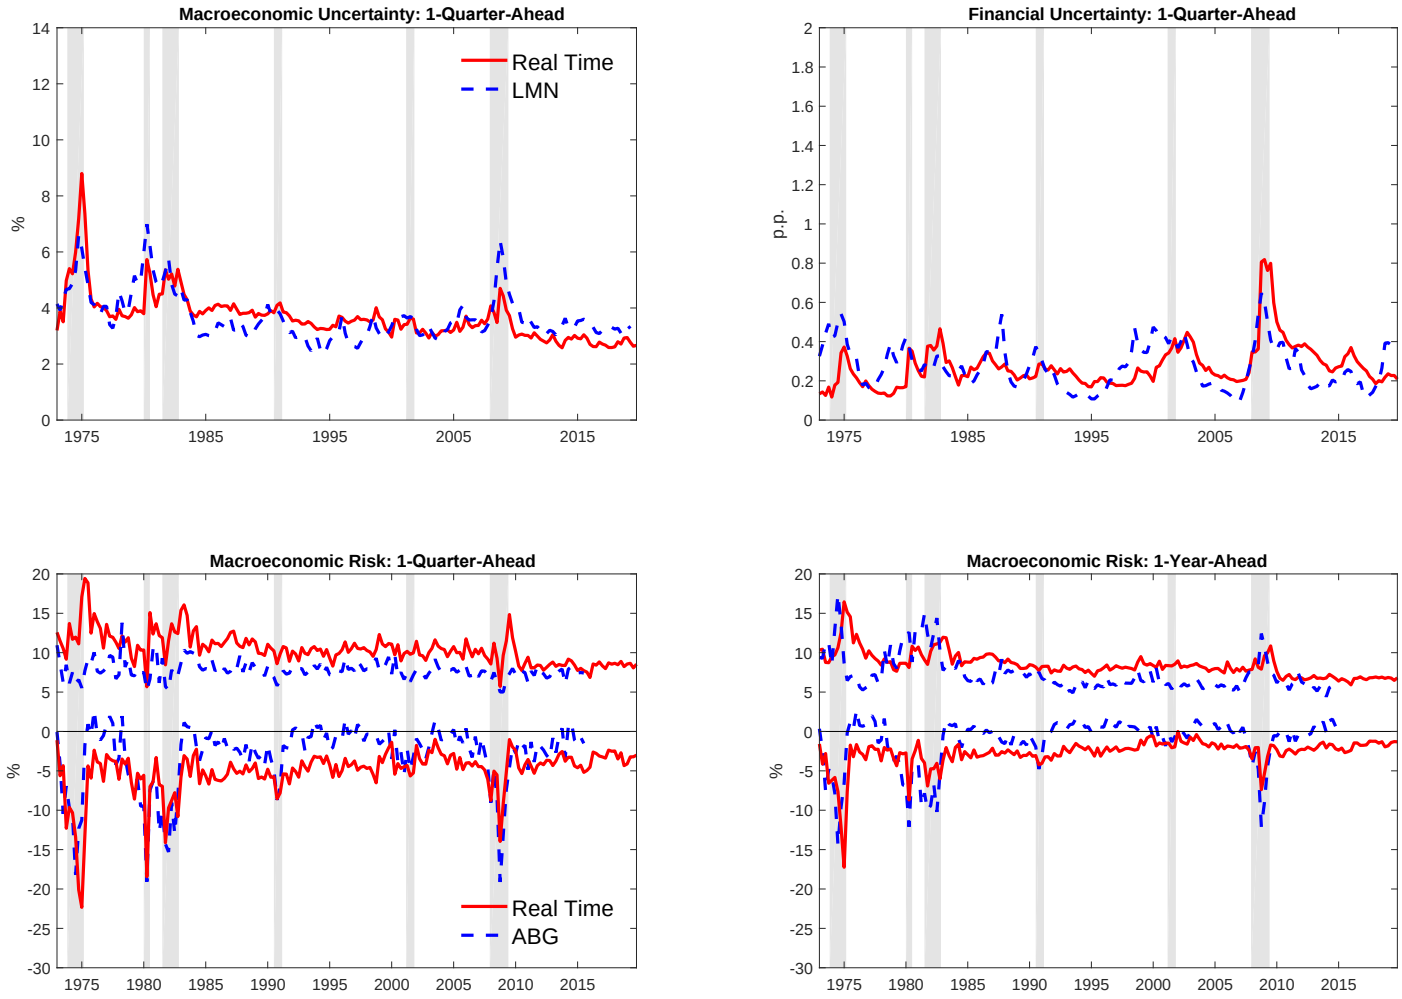

Note: The upper panels plot our estimate of one-quarter ahead macroeconomic and financial uncertainty (solid red lines) and the quarterly average of the three-month ahead real and financial uncertainty measures from Ludvigson et al. (2021) (dashed blue lines). The one-quarter ahead uncertainty estimates are annualized. We standardize the values of the Ludvigson et al. (2021) uncertainty series so they have the same mean and average standard deviation as our uncertainty series. The lower panels plot our estimates of one-quarter ahead (left panel) and one-year ahead (right panel) macroeconomic tail risk (solid red lines) against estimates of macroeconomic tail risk from Adrian et al. (2019) (dashed blue lines). The one quarter ahead tail risk estimates are annualized. The gray shaded areas are recession dates as defined by the National Bureau of Economic Research.

sure. In addition, the relative magnitude of the spikes is different between the measures. The spikes in the LMN measure are comparable in size in the 1970s and during the GFC, while, as previously discussed, the use of real-time data induces higher uncertainty in the 1970s. The correlation between

the financial uncertainties is lower, standing at 0.4. The measures share only some spikes—during the 1970s and during the GFC—while differing in the 1980s and 1990s. The difference, in part, reflects the choice of underlying financial variables. We use corporate credit spreads, while the LMN measure mostly loads on stock returns. As a result, for example, our financial volatility measure does not spike during the stock market crash of October 1986 while the LMN measure does.

ABG construct measures of macroeconomic risk (also known as growth-at-risk), also using a two-step approach. First, they run quantile regressions of average future GDP growth on current GDP and a financial conditions index. Then they build predictive distributions for future GDP growth by fitting a t-skewed distribution to the fitted values from the quantile regressions.

The lower panels of Figure 7 compare our macroeconomic shortfall and longrise measures with those from ABG. Despite the differences in models and estimation exercises—like real-time versus final-release data and out-of-sample versus in-sample—results are broadly similar. The correlation between shortfall measures is around 0.7 at both the one-quarter and one-year horizons. The correlation for the longrise is lower, at about 0.4, although the two models share some spikes.

Relative to these two papers, our approach has several benefits. Through the joint conditional distribution of the endogenous variables, we can estimate the correlation structure among indicators of uncertainty and risk and their cyclical variation, we can define the probability of joint macroeconomic and financial adverse events, and we can incorporate all sources of statistical and sampling uncertainty. Finally, as we discuss in the next two sections, in our framework we can identify the joint determinants of uncertainty and risk for macroeconomic and financial conditions, finding that nonlinearities play a critical role.

4 Understanding Fluctuations in Uncertainty and Risk

In this section, we exploit the structure of the model to study the determinants of time-variation in uncertainty and risk. Following standard practice in empirical macroeconomics and the SVAR literature, we identify orthogonal disturbances to the equations in the model. The identification is based on a statistical criterion whereby we isolate shocks that maximize their contribution to the forecast

error variance (FEV) of a target variable over a pre-specified horizon. This identification was initially developed by Uhlig (2003) to identify the main driver of real gross national product. Recent work by Angeletos et al. (2020) applies the maximum variance criterion to identify a "main business cycle shock," a shock that can drive variation in many real activity indicators at business cycle frequencies. In a similar spirit to these papers, our analysis revolves around two shocks that we label "main business cycle" and "main financial" shocks. Our goal is to quantify the degree of co-movements that these shocks generate under two dimensions: (1) co-movement between macro and financial variables, and (2) co-movement between level and volatility of GDP and corporate spreads—and consequently their ability to generate time-variation in uncertainty and risk.

4.1 Identification Strategy

So far, the maximum FEV identification has been applied to linear VARs. In what follows, we discuss how we extend the identification to a model with time-varying volatility and to sequentially identify two shocks.[14]

Variance Decomposition We begin by discussing the notion of variance decomposition that we use in our analysis. Equation (11) shows the formula that defines the contribution to the forecast error variance of variable i attributed to shock j at horizon F:

$$VD_{i,j,F}\left(z_{i,t+F}|A_0, H_t, \mathcal{I}_t, \Theta\right) = 1 - \frac{VAR\left(z_{i,t+F}|A_0, e_{-j}, H_t, \mathcal{I}_t, \Theta\right)}{VAR\left(z_{i,t+F}|H_t, \mathcal{I}_t, \Theta\right)}, \tag{11}$$

where A_0 is the matrix that imposes the identification assumptions and e_{-j} is a selector vector that denotes all structural shocks except shock j. In a linear model, the contribution of a shock to the FEV depends only on the identification assumptions imposed by matrix A_0 and on the reduced-form parameters Θ. In a SV-VAR model, because of nonlinearities and state-dependence, the contribution of a shock depends also on the state H_t and the data z^t (which is part of $\mathcal{I}_t$). Additionally, unlike in a

[14] Mumtaz (2018) uses a FEV identification in a SV-VAR abstracting from time-varying volatility in the identification. Mountford and Uhlig (2009) and Caldara et al. (2016) sequentially identify multiple shocks using the so-called penalty function approach, a similar identification that maximizes the impulse response of target variables subject to sign restrictions.

linear framework, there is no closed-form solution for the variances in our SV-VAR model. Therefore, we rely on simulation methods to compute them.

As shown in equation (11), we calculate the importance of shock j by looking at what would happen to the variance of the endogenous variables if we excluded the shock of interest. That is, we compute the FEV attributed to shock j by taking the ratio of the variance of $z_{i,t+F}$ with all shocks active except the shock of interest (the vector of shocks e_{-j}) over the FEV of $z_{i,t+F}$ when all shocks are active. We then take 1 minus this ratio. Intuitively, if the variances under the two scenarios are similar, then the structural shock in question has a small effect. Note that because our framework is nonlinear, this calculation is not equivalent to directly computing the ratio of the FEV where we shut down all of the structural shocks except shock j over the total FEV. We prefer our approach compared with this alternative because it better preserves the nonlinear interactions between the structural shocks when computing the decomposition, giving a measure of the marginal importance of the shock of interest relative to the other shocks.[15]

Max-share Identification Strategy Our max-share identification approach identifies the impact vector $A_{0,j}^{*\prime}$ associated to shock j that solves the maximization problem in equation (12).

$$\begin{aligned} &\max_{A'_{0,j}} \; VD(z_{i,t+F}|A'_{0,j}, e_j, H_t, \mathcal{I}_t, \Theta) \\ &s.t. \;\; A^*_{0,1:j-1}A'_{0,j} = 0. \end{aligned} \tag{12}$$

To identify structural shock j, we choose the column vector $A_{0,j}^{*\prime}$ that maximizes the variance share of that shock on target variable z_i at horizon F. The constraint on this maximization problem is that the column vector must be orthogonal to the columns in the structural matrix identified before it, $A^*_{0,1:j-1} = \left[A_{0,1}^{*\prime}...A_{0,j-1}^{*\prime}\right]$. We adopt the convention that for $j = 1$, this matrix is empty.

There are two points to discuss with this identification strategy—one well-known in the literature

[15]Note that, because of the nonlinear interactions, it is in general not the case that the summation of these variance shares over all structural shocks j equals 1. However, while our method preserves the nonlinear interactions when calculating the decomposition, keeping only one shock active at a time would shut down all of these interactions, leading to a summation of variance shares that severely underestimates the total variance generated by the model.

and the other specific to our nonlinear framework. First, when identifying multiple shocks, the sequence in which we run the optimization problems matters for the result. We discuss our choice at the end of this subsection. Second, since the variance decomposition in equation (11) is state-dependent, the optimization problem in equation (11) would also be state-dependent, with identifying assumptions changing over time and across volatility states. This dependence would make computation infeasible. For this reason, we fix the *initial* states by simulating the model to a steady-state value by shutting off all shocks. Note that we allow the volatility processes to evolve following the model equations throughout the forecast horizon, so that the variance decomposition reflects time-variation in the volatilities.

In our application, the first shock we identify is the main business cycle (MBC) shock, which we define as a shock that maximizes the variability of GDP growth at the four-quarter horizon.[16] The second shock we identify is the main financial (MF) shock, which we define as a shock that maximizes the variability of corporate credit spreads at the four quarter horizon that is orthogonal to the main shock. Our ordering aims at linking our identification to the existing literature. Specifically, we want to investigate whether a shock like the one identified by Angeletos et al. (2020) can generate co-movement between the level and volatility of the endogenous variables in the system. The main financial shock aims to capture co-movements left unexplained by the main business cycle shock that could potentially be attributable to changes in financial conditions, as for instance in Gilchrist and Zakrajšek (2012).

4.2 Impulse Responses

As our model is nonlinear and state-dependent, the effects of the structural shocks differ depending on the sequence in which they arrive. We therefore show generalized impulse responses (GIRFs) (Koop et al., 1996) conditioning the model on steady-state values. Figure 8 plots the median responses and 80 percent credible sets of the level variables and volatility processes to the main business cycle and main financial shocks, with the credible sets reflecting parameter uncertainty.

The first row plots the impulse responses to a one-standard-deviation negative main business cycle shock. In response to the shock, the level of GDP drops to a little below negative 0.5 percent, and

[16] Angeletos et al. (2020) maximizes the variability of the *level* of GDP between six and 32 quarters in the frequency domain. They document that maximizing the variability of the level of GDP at four quarters in the time domain approximately replicates their frequency domain results, which motivates our choice of maximization horizon.

Figure 8: Generalized Impulse Response Functions

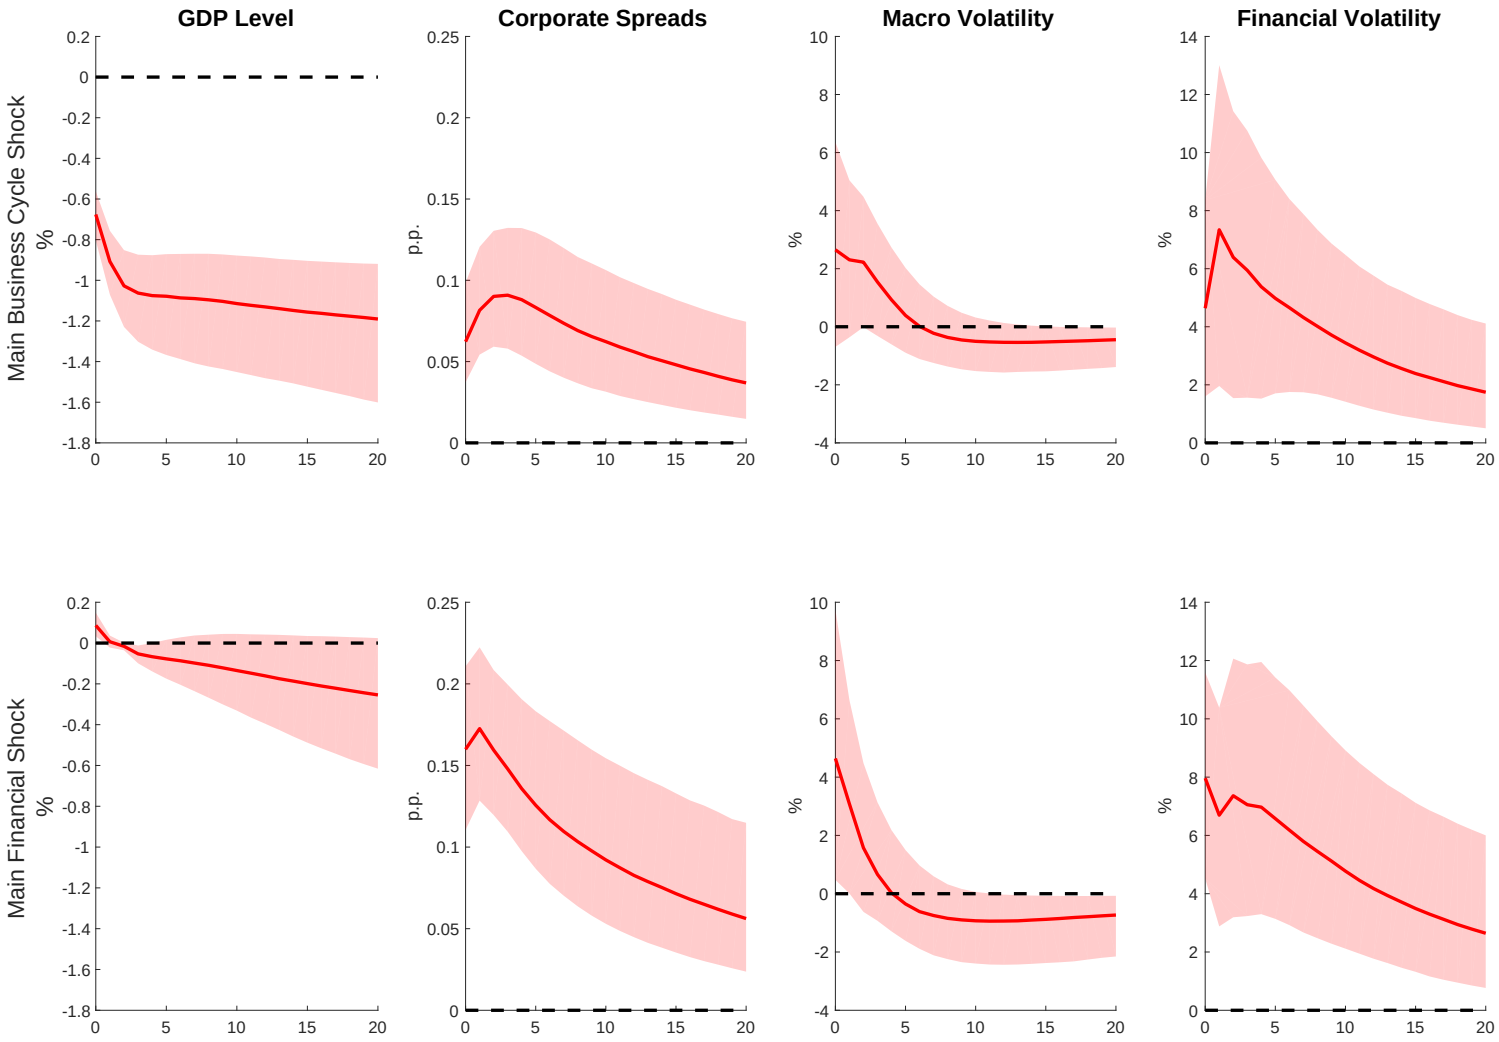

Note: The red lines depict the median generalized impulse responses of the four variables included in the baseline model to a one-standard-deviation main business cycle (top row) and main financial (bottom row) shocks. We report the level of GDP, computed by cumulating growth rates, and volatilities computed as the log standard deviations of the reduced-form innovations ($\log H_t^{1/2}$). Shaded areas denote 80 percent credible sets. All impulse responses are reported in percentages, except for corporate spreads, which are reported in percentage points.

remains at negative 1 percent thereafter. Corporate spreads increase over time, peaking at 10 basis points after four quarters. A one-standard deviation main financial shock—plotted in the second row—leads to an immediate and prolonged increase in spreads and to a prolonged decline in GDP. Both main shocks generate an increase in macroeconomic and financial volatility. Macro volatility rises in the first year after the shock, reversing to negative values in the medium term. Financial volatility remains

elevated for up to five years.

The identification of the two main shocks shows that there is sizable co-movement between level and volatility variables, and that this co-movement can be attributed to two shocks. A simple way to interpret our shocks is that they capture episodes characterized by contemporaneous adverse macro conditions and elevated macro volatility, and/or a joint tightening of financial conditions and elevated financial volatility, as for instance the recessions of the 1970s and the GFC.

4.3 Distributional Responses to Shocks

To further illustrate the nonlinearities and state-dependencies inherent in our model, we show in Figure 9 counterfactual predictive distributions at two periods in time: 2008:Q4—a time of high volatility—and 2018:Q1—a time of low volatility. The solid blue lines are the baseline predictive distributions while the dashed red lines are a counterfactual scenario in which a two standard deviation main business cycle shock hits in 2009:Q1 and 2018:Q2, respectively. This exercise is the the empirical counterpart—integrating over all sources of uncertainty—to the simulation results discussed in Section 2.[17]

In 2008:Q4, GDP growth is low, spreads are high, and macro and financial volatilities are elevated. These conditions combine to produce volatile GDP growth and corporate spreads predictive distributions that feature substantial adverse risk. Expected GDP growth is low, especially at shorter horizons. The low mean, high variance, and negative skewness of GDP growth leads to expected shortfall levels below -7% at the one year horizon. These conditions gradually fade away with a particularly large decrease in downside risk. Upside risk, on the other hand, does not move as much. A two standard deviation negative shock in 2009:Q1 generates a dramatic shift downwards in the distribution, with its tails reaching well beyond -10% at the one year horizon. The already-elevated macro volatility conditions amplify the effects of the shock, both in terms of the mean and the higher-moments of the predictive distributions.

The baseline distributions for average corporate spreads have a high mean and are wide with a positive skew. These conditions generate an expected longrise of over 8 p.p. at the one year horizon. An adverse MBC shock that hits in 2009:Q1 is amplified by the high financial volatility and produces a

[17] For GDP growth, we normalize the distributions at the various horizons so that they are annualized.

further persistent increase in expected spreads and volatility, thereby increasing upside risk. The effects of the shock are smaller for spreads than for GDP growth, as the generalized impulse response functions in Figure 8 show.

The bottom two panels show the same exercise during a time of low volatility in 2018:Q1. During that time, GDP growth is high, spreads are low, and macro and financial volatilities are low as well. In this case, a two standard deviation adverse MBC shock in 2018:Q2 still produces a decrease in expected GDP growth and increases in expected spreads, macro and financial volatilities, and therefore adverse risks, but the effects of the shock are muted.

5 Predictive Distributions during the Global Financial Crisis

In this section, we study the most extreme event in our sample—the GFC—and parse out the role of the identified shocks and their interaction.[18] To perform the decomposition, we examine the counterfactual distributions that would have been realized if only certain structural shocks were active in the buildup of and during the crisis. We use real-time vintages—rather than the final vintage—in constructing the counterfactuals, and therefore our results are best interpreted as the counterfactual decompositions of the predictive distributions made using only data up until a specific point in time. Our choice is driven by the fact that we want to use only information available to policymakers as they made policy decisions.

Figure 10 plots the predictive distributions for four-quarter ahead average GDP growth (left panels) and corporate credit spreads (right panels), characterizing the GFC by contrasting two points in time. The first is 2006:Q4 (the solid black distributions), a period right before the crisis characterized by buoyant macroeconomic and financial conditions. Accordingly, the predictive distributions anticipate strong GDP growth and low corporate spreads. The forecast is associated with low levels of uncertainty and risk. The second time period is 2008:Q4 (the dashed blue distributions), which corresponds to the height of the GFC, a period of adverse macroeconomic conditions and extremely tight financial conditions. Compared with 2006, the predictive distributions display low GDP growth and elevated credit spreads. In contrast to the previous ones, these forecasts are surrounded by much larger uncertainty and

[18] In the companion webpage to the paper, we discuss the measurement and analysis of uncertainty and risks during the COVID-19 pandemic.

Figure 9: Posterior Predictive Distributions

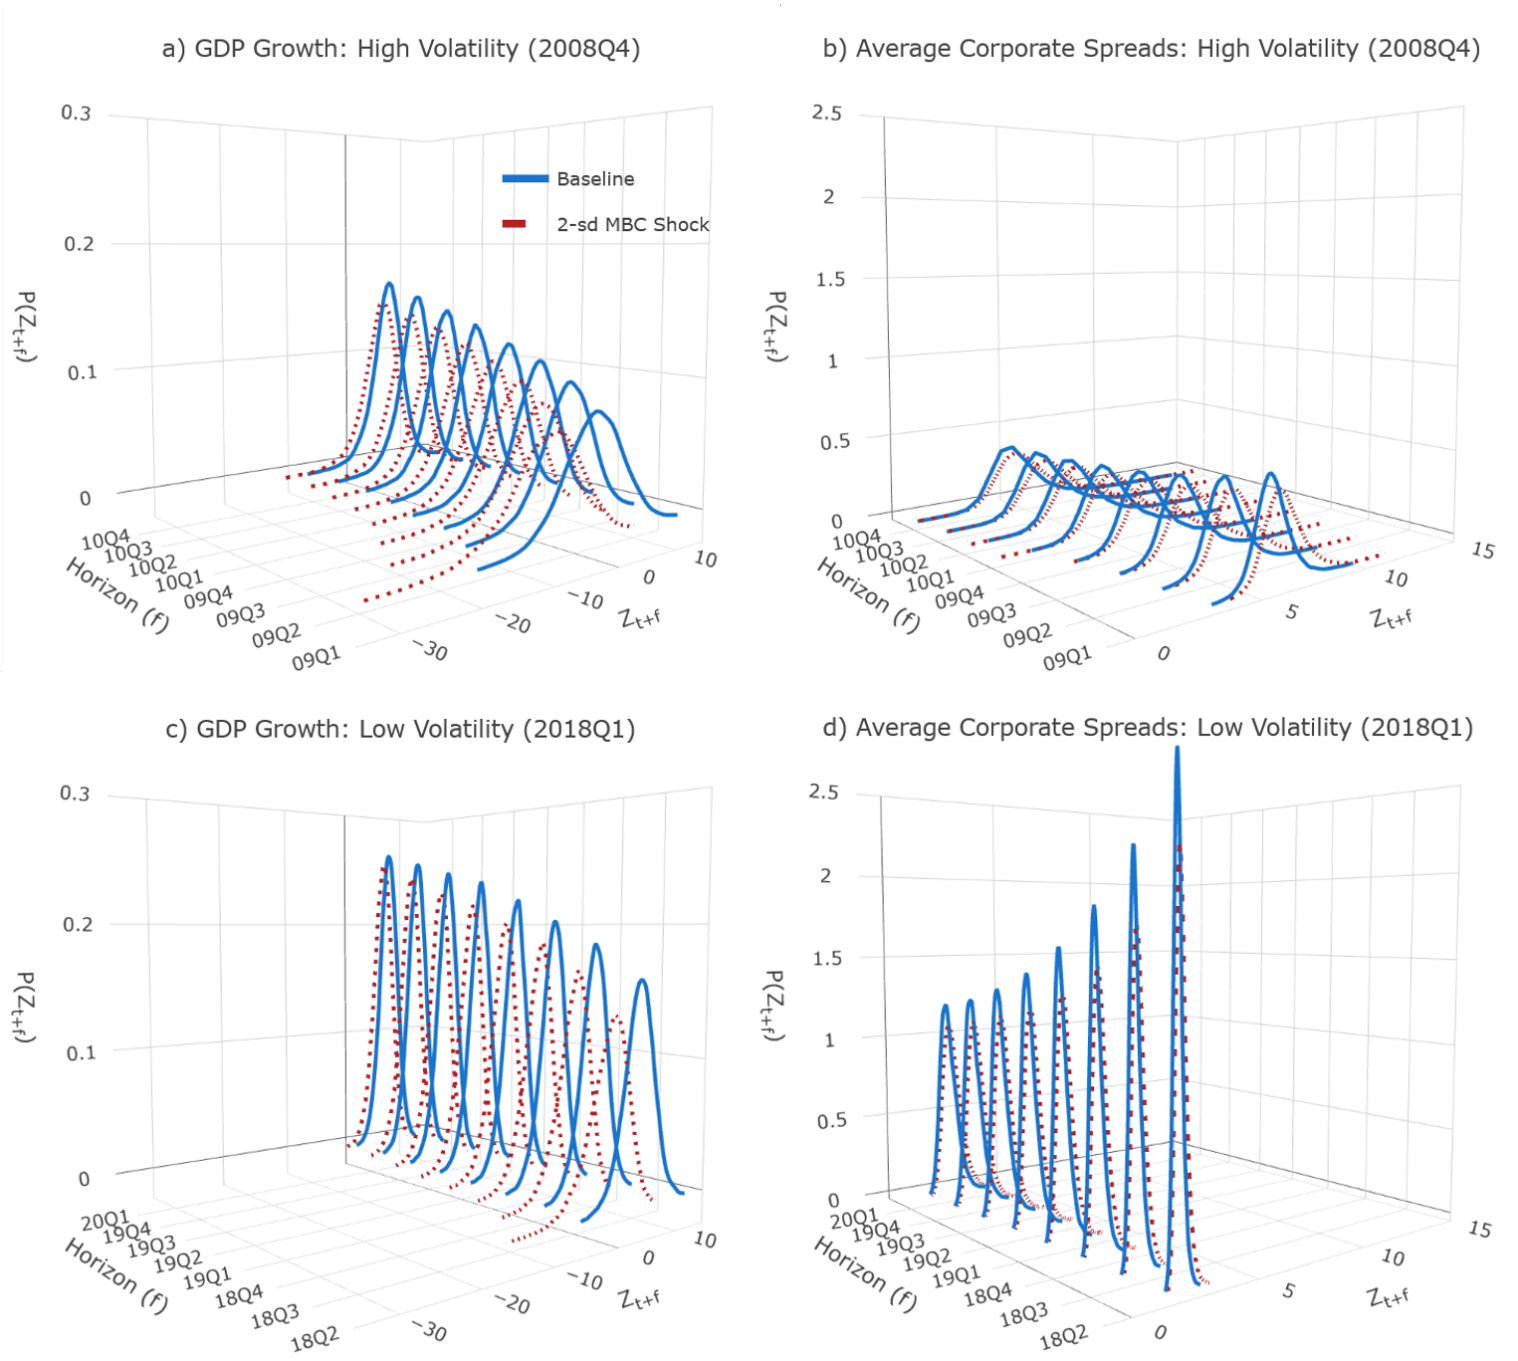

Note: The upper two panels show the predictive distributions for GDP growth and average corporate spreads generated by our model given 2008:Q4 conditions. The solid blue lines show the baseline predictive distribution while the dashed red lines show counterfactual conditions if a two standard deviation main business cycle shock hits the economy in 2009:Q1. The lower two panels show the results from repeating the exercise given 2018:Q1 conditions and the shock hitting in 2018:Q2. For GDP growth, we normalize the predictive distributions at the various horizons so that they are annualized.

Figure 10: Predictive Distributions around the Global Financial Crisis

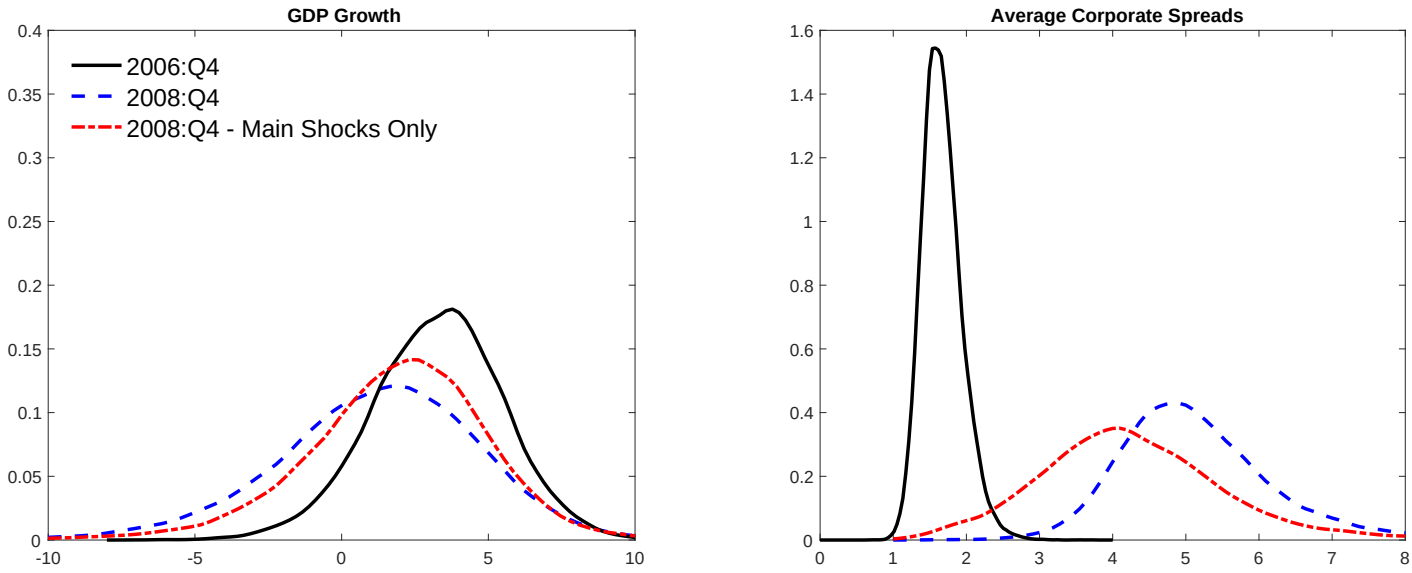

(a) Counterfactual with Both Main Shocks Active

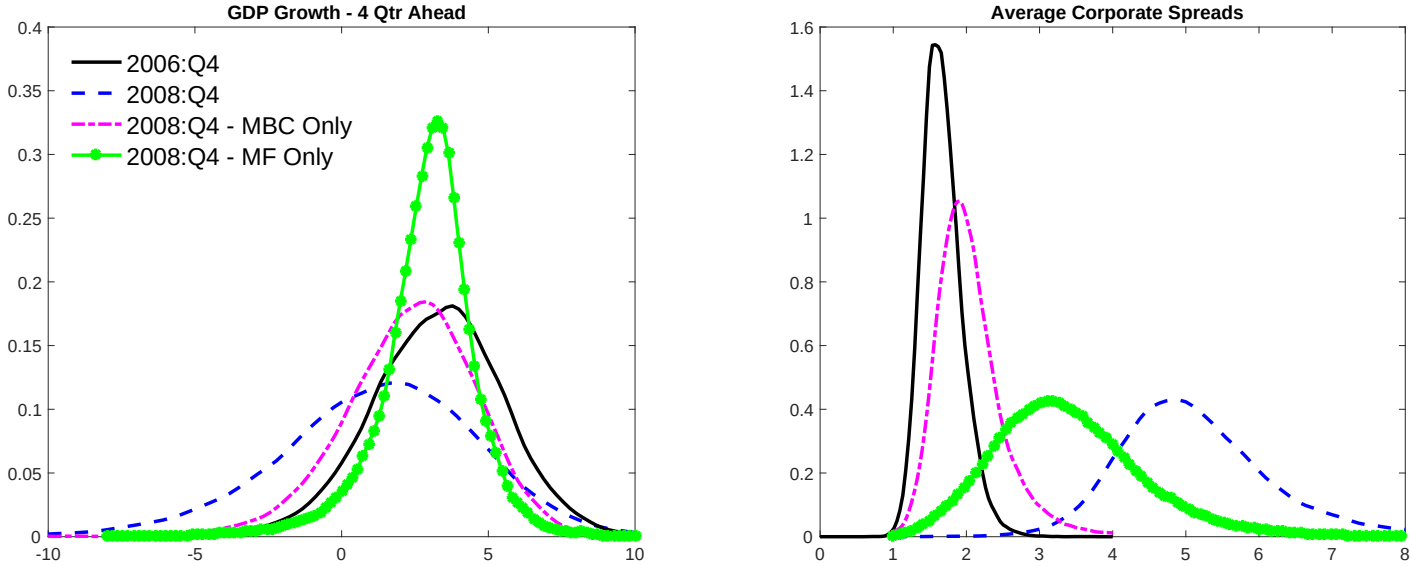

(b) Counterfactual with One Main Shock Active

Note: The two panels plot four-quarter ahead predictive distributions of average GDP growth and corporate credit spreads. The black distributions are computed using the real-time data vintage in 2006:Q4, and the blue distributions are computed at the height of the Global Financial Crisis in 2008:Q4. In panel (a), the red distributions are computed by running a counterfactual that keeps only main shocks active starting in 2007:Q1. In panel (b), the magenta (green) distributions are computed by running a counterfactual that keeps only the main business cycle (financial) shocks active starting in 2007:Q1. MBC: main business cycle; MF: main financial.

longer tails representing downside macroeconomic risks (longer left tail of the GDP growth distribution) and downside financial risks (longer right tail of the corporate spread distribution).

Our model generates the shift from the black to the blue distributions through a sequence of shocks

over the two intervening years. To understand the importance of the two main shocks, panel (a) of Figure 10 plots the predictive distributions consistent with a counterfactual simulation that keeps only the main business cycle and financial shocks active from 2007:Q1 onward (the dash-dotted red lines). Comparing the blue and the red distributions, we see that the main shocks generate a sizeable increase in macroeconomic and financial uncertainties, as well as an increase in adverse tail risks. As shown in Table 1, under this counterfactual simulation (MBC + MF case), macroeconomic uncertainty about GDP growth is 3.2 percent—only 0.3 percentage points below 2008:Q4 uncertainty. However, the absence of residual shocks leads to underestimating downside risks for macroeconomic conditions, as lower uncertainty interacts with a more muted shift in the mean of the distributions. The difference in macro shortfall between the counterfactual simulation and baseline is around 1 percentage point. The corresponding difference in macro longrise is smaller, reflecting the inherent asymmetry in the two tails of macroeconomic risks. Financial uncertainty is slightly higher in the counterfactual scenario versus the baseline estimate, although the main shocks alone do not generate enough of an increase in the mean of the distribution. Therefore, the counterfactual distribution is wider, but underestimates the longrise in corporate spreads by 0.5 percentage points and the shortfall by 1.4 percentage points.

Panel (b) of Figure 10 studies the relative importance of each of the main macroeconomic and financial shocks as well as their interaction. We obtain the dash-dotted magenta distributions by running a counterfactual simulation that keeps only the main business cycle shock active from 2007:Q1 onward, while we obtain the dotted green distributions by keeping only the main financial shock active over the same period. Both macroeconomic and financial conditions and their *interaction* are crucial to generate the large increases in uncertainties and downside risks in 2008:Q4. When only main business cycle shocks are active, macroeconomic uncertainty is 2.4 percent, one third below the full reading for 2008:Q4, and the mean prediction is too optimistic, leading to a macro shortfall that is nearly 4 percentage points smaller compared with the full reading—see the "MBC Only" case in Table 1. In addition, main business cycle shocks alone increase the volatility of the distribution of future financial conditions but do not generate much of a shift in the mean—the magenta line is centered closely to the black one in the right panel. Similarly, the model underestimates the rise in macroeconomic uncertainty and the increase in downside risk when only main financial shocks are active. While financial shocks

Table 1: Uncertainty and Tail Risks around the Global Financial Crisis

	Macro			Financial		
	UNC	SF	LR	UNC	SF	LR
Global Financial Crisis						
- *Baseline*						
2006:Q4	2.26	-1.64	7.83	0.27	1.19	2.33
2008:Q4	3.54	-7.14	7.91	1.21	3.25	8.48
- *Counterfactual*						
2008:Q4 MBC + MF	3.23	-6.08	7.85	1.43	1.87	7.94
2008:Q4 MBC Only	2.35	-3.22	6.73	0.64	1.26	3.66
2008:Q4 MF Only	1.72	-1.61	6.29	1.06	1.65	6.12

Note: This table shows uncertainty and risk associated with macroeconomic and financial conditions at the one-year horizon during the Global Financial Crisis. For 2008:Q4, we report statistics for three counterfactual distributions constructed as follows: *MBC + MF* denotes simulations where only main business cycle and main financial shocks are active from 2007:Q1; *MBC Only* denotes similar simulations where only main business cycle shocks are active; *MF Only* denotes similar simulations where only main financial shocks are active. UNC: uncertainty; SF: expected shortfall; LR: expected longrise; these three measures are defined in Section 2.

generate sizeable movements in the distribution of future financial conditions, they still underestimate the financial longrise by 2.4 percentage points.

In our model, the quantitatively relevant non-linear interaction between macroeconomic and financial conditions can be described intuitively as follows. Main business cycle shocks and main financial shocks contribute to movements in both macroeconomic and financial volatility. Thus, when we shut down one set of shocks— for instance, the main financial shocks—we reduce current and future volatility. Lower macro and financial volatility cause current and future main business cycle shocks to have less of an effect on the dispersion of the GDP growth and corporate spread distributions (see equation (1) to see how volatility states interact with innovations). As implied by the distributions shown in panel (b), the additive effects of the two main shocks (that is, in the absence of their nonlinear interaction) would lead to the underestimation of both uncertainty and tail risks.

6 Conclusion

A SV-VAR can generate time-variation and asymmetries in the distributions of future macroeconomic and financial variables to construct indicators of uncertainty and risk. We use Bayesian methods and real-time data to estimate the predictive distributions to quantify movements and correlation in measures of uncertainty and risk for both macroeconomic and financial conditions. We also find that uncertainty and risk dynamics during the Global Financial Crisis were driven by a toxic combination of main business cycle and main financial shocks.

The unified framework of the paper is best suited for the many applications where the goal is to jointly analyze uncertainty and risk around the forecasts of economic and financial indicators. While the SV-VAR estimated recursively with real-time data and Bayesian techniques produces estimates of uncertainty and risk broadly comparable to those from established papers in the literature (Ludvigson et al., 2021; Adrian et al., 2019), it presents some advantages. For example, it allows us to characterize the risk around joint events spanning multiple horizons, such as financial crises, and it avoids potential complications arising from merging statistical concepts and measurement definitions.

We look forward to a variety of modifications and extensions of this framework, including but not limited to the following ideas. One could examine the relationships between SV-VAR and DSGE models with time-varying volatility. DSGE models of stochastic volatility have proven popular in the literature (Justiniano and Primiceri, 2008; Fernández-Villaverde and Rubio-Ramírez, 2007). Moreover, models solved to higher-order approximations generate endogenous volatility through state-dependence. A particularly fruitful avenue could be to examine whether DSGE models generate predictive densities that are broadly in line with those from suitable SV-VARs. Models with occasionally-binding constraints, along the lines of Guerrieri and Iacoviello (2017), could also be a promising approach. A second area of work would be to estimate larger SV-VARs than the one considered in this paper, to measure and analyze more sources of uncertainty and risk, possibly resorting to dimension reduction techniques such as the use of a factor structure in volatility as proposed by Carriero et al. (2018).

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable Growth. *American Economic Review 109*(4), 1263–1289.

Adrian, T., N. Boyarchenko, and D. Giannone (2021). Multimodality in Macro-Financial Dynamics. *International Economic Review*.

Angeletos, G.-M., F. Collard, and H. Dellas (2020). Business Cycle Anatomy. *American Economic Review 110*(10), 3030–70.

Bachmann, R., S. Elstner, and E. R. Sims (2013). Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics 5*(2), 217–249.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics 131*(10), 1593–1636.

Basu, S. and B. Bundick (2017). Uncertainty Shocks in a Model of Effective Demand. *Econometrica 85*(3), 937–58.

Basu, S., G. Candian, R. Chahrour, and R. Valchev (2021, April). Risky Business Cycles. Boston College Working Papers in Economics 1029, Boston College Department of Economics.

Bernanke, B. S. (1983). Irreversibility, Uncertainty, and Cyclical Investment. *Quarterly Journal of Economics 98*(1), 85–106.

Bloom, N. (2009). The impact of Uncertainty Shocks. *Econometrica 77*, 623–68.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives 28*(2), 153–176.

Caldara, D., C. Fuentes-Albero, S. Gilchrist, and E. Zakrajšek (2016). The macroeconomic impact of financial and uncertainty shocks. *European Economic Review 88*, 185–207.

Caldara, D. and M. Iacoviello (2018). Measuring Geopolitical Risk. International Finance Discussion Papers 1222, Board of Governors of the Federal Reserve System.

Carriero, A., T. Clark, and M. Marcellino (2018). Measuring Uncertainty and Its Impact on the Economy. *Review of Economics and Statistics 100*(5), 799–815.

Carriero, A., T. E. Clark, and M. Marcellino (2015). Realtime Nowcasting with a Bayesian Mixed Frequency Model with Stochastic Volatility. *Journal of the Royal Statistical Society Series A 178*(4), 837–62.

Carriero, A., T. E. Clark, and M. Marcellino (2020). Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions. Working Papers 202002R, Federal Reserve Bank of Cleveland.

Cascaldi-Garcia, D., D. D. Datta, T. R. T. Ferreira, O. V. Grishchenko, M. R. Jahan-Parvar, J. M. Londono, F. Loria, S. Ma, M. del Giudice Rodriguez, and J (2020, July). What is Certain about Uncertainty? International Finance Discussion Papers 1294, Board of Governors of the Federal Reserve System (U.S.).

Chan, J. C.-C. and I. Jeliazkov (2009). MCMC Estimation of Restricted Covariance Matrices. *Journal of Computational and Graphical Statistics 18*(2), 457–480.

Christiano, L. J., R. Motto, and M. Rostagno (2014). Risk Shocks. *American Economic Review 104*(1), 27–65.

Clark, T. E. (2011). Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility. *Journal of Business & Economic Statistics 29*(3), 327–41.

Cogley, T. and T. J. Sargent (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics 8*(2), 262–302.

Creal, D. D. and J. C. Wu (2017). Monetary Policy Uncertainty and Economic Fluctuations. *International Economic Review 58*(4), 1317–54.

Del Negro, M. and F. Schorfheide (2013). DSGE Model-Based Forecasting. In G. Elliott, C. Granger, and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2, pp. 57–140. Elsevier.

Evans, C., J. Fisher, F. Gourio, and S. Krane (2015). Risk Management for Monetary Policy Near the Zero Lower Bound. *Brookings Papers on Economic Activity 46*(1 (Spring), 141–219.

Favara, G., S. Gilchrist, K. F. Lewis, and E. Zakrajsek (2016). Recession Risk and the Excess Bond Premium. Feds notes, Board of Governors of the Federal Reserve System.

Fernandez-Villaverde, J. and P. Guerron-Quintana (2020). Uncertainty Shocks and Business Cycle Research. *Review of Economic Dynamics 37*, 118–66.

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester, and J. Rubio-Ramírez (2015). Fiscal Volatility Shocks and Economic Activity. *American Economic Review 105*(11), 3352–84.

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, and M. Uribe (2011). Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review 101*(6), 2530–61.

Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007). Estimating Macroeconomic Models: A Likelihood Approach. *Review of Economic Studies 74*(4), 1059–87.

Gertler, M. (2012). Comment on 'Which Financial Frictions? Parsing the Evidence from the Financial Crisis of 2007 to 2009'. In *NBER Macroeconomics Annual*, Volume 27 of *NBER Chapters*, pp. 215–223. National Bureau of Economic Research, Inc.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2014, April). Uncertainty, Financial Frictions, and Investment Dynamics. NBER Working Papers 20038, National Bureau of Economic Research, Inc.

Gilchrist, S. and E. Zakrajšek (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review 102*, 1692–1720.

Gonçalves, S., A. M. Herrera, L. Kilian, and E. Pesavento (2021). Impulse Response Analysis for Structural Dynamic Models with Nonlinear Regressors. *Journal of Econometrics*.

Gorodnichenko, Y. and S. Ng (2017). Level and Volatility Factors in Macroeconomic Data. *Journal of Monetary Economics 91*, 52 – 68.

Guerrieri, L. and M. Iacoviello (2017). Collateral Constraints and Macroeconomic Asymmetries. *Journal of Monetary Economics 90*(C), 28–49.

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019). Firm-Level Political Risk: Measurement and Effects. *Quarterly Journal of Economics 134*(4), 2135–2202.

Husted, L., J. Rogers, and B. Sun (2020). Monetary Policy Uncertainty. *Journal of Monetary Economics 115*(C), 20–36.

Jo, S. (2014). The Effects of Oil Price Uncertainty on Global Real Economic Activity. *Journal of Money, Credit and Banking 46*(6), 1113–35.

Jurado, K., S. Ludvigson, and S. Ng (2015). Measuring Uncertainty. *American Economic Review 105*(3), 1177–1216.

Justiniano, A. and G. E. Primiceri (2008). The Time-Varying Volatility of Macroeconomic Fluctuations. *American Economic Review 98*(3), 604–41.

Koop, G., M. H. Pesaran, and S. M. Potter (1996). Impulse Response Analysis in Nonlinear Multivariate Models. *Journal of Econometrics 74*(1), 119–47.

Ludvigson, S. C., S. Ma, and S. Ng (2021). Uncertainty and business cycles: Exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics*.

McCracken, M. W. and S. Ng (2016). FRED-MD: A Monthly Database for Macroeconomic Research. *Journal of Business & Economic Statistics 34*(4), 574–89.

Mountford, A. and H. Uhlig (2009). What Are the Effects of Fiscal Policy Shocks? *Journal of applied econometrics 24*(6), 960–92.

Mumtaz, H. (2018). A Generalised Stochastic Volatility in Mean VAR. *Economics Letters 173*(C), 10–14.

Mumtaz, H. and F. Zanetti (2013). The Impact of the Volatility of Monetary Policy Shocks. *Journal of Money, Credit and Banking* (45), 535–58.

Orlik, A. and L. Veldkamp (2014). Understanding Uncertainty Shocks and the Role of Black Swans. Working Paper 20445, National Bureau of Economic Research.

Orphanides, A. (2001). Monetary Policy Rules Based on Real-Time Data. *American Economic Review 91*(4), 964–85.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies 72*(3), 821–52.

Salgado, S., F. Guvenen, and N. Bloom (2019). Skewed Business Cycles. NBER Working Papers 26565, National Bureau of Economic Research, Inc.

Scotti, C. (2016). Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro Surprises. *Journal of Monetary Economics 82*, 1–19.

Shin, M. and M. Zhong (2020). A New Approach to Identifying the Real Effects of Uncertainty Shocks. *Journal of Business & Economic Statistics 38*(2), 367–79.

Uhlig, H. (2003). What Moves Real GNP? mimeo.

Appendices – For Online Publication

A Additional Figures and Tables

Table A.1: Statistical Properties of Macro and Financial Uncertainties

	Macro Uncertainty		Financial Uncertainty	
	1 qtr	1 yr	1 qtr	1 yr
Panel A. Average Uncertainty				
Real-time	**3.68**	**2.63**	**0.28**	**0.37**
	(0.85)	**(0.66)**	**(0.11)**	**(0.15)**
No parameter/state uncertainty	3.39	2.34	0.26	0.33
	(0.66)	(0.41)	(0.11)	(0.12)
Final-vintage	2.97	2.05	0.26	0.36
	(0.42)	(0.22)	(0.10)	(0.13)
Panel B. Correlation with Mean Forecast & Cross-Correlations				
Macro mean (Real-time)	-0.30	0.13	-0.19	-0.09
Macro mean (Final Vintage)	-0.48	-0.41	-0.52	-0.64
Financial mean (Real-time)	0.06	0.05	0.90	0.92
Financial mean (Final Vintage)	0.20	0.32	0.89	0.9
Financial unc. (Real-time)	0.08	0.14		
Financial unc. (Final Vintage)	0.40	0.58		

Note: Panel A of the table shows average macroeconomic and financial uncertainty for the three measures discussed in the main text, with standard deviations reported in parentheses. Panel B reports correlation between uncertainty and mean forecasts and the cross correlation between macro and financial uncertainty for our baseline measures based on real-time data and final-vintage data.

Table A.2: Statistical Properties of Macro and Financial Tail Risks

	5% Expected Longrise				5% Expected Shortfall			
	Macro		Financial		Macro		Financial	
	1 qtr	1 yr	1 qtr	1 yr	1 qtr	1 yr	1 qtr	1 yr
Mean (std) over the sample								
Real-time	**10.45**	**8.45**	**2.88**	**3.07**	**-5.11**	**-2.69**	**1.69**	**1.48**
	(2.03)	**(1.56)**	**(0.95)**	**(1.02)**	**(2.94)**	**(1.84)**	**(0.49)**	**(0.42)**
No parameter or state uncertainty	9.58	7.69	2.81	2.93	-4.51	-2.00	1.74	1.57
	(1.71)	(1.18)	(0.91)	(0.92)	(2.49)	(1.17)	(0.52)	(0.44)
Final-vintage	8.94	7.05	2.85	3.10	-3.77	-1.71	1.76	1.57
	(1.23)	(0.63)	(0.91)	(0.95)	(2.01)	(1.01)	(0.52)	(0.44)
Correlation with macro mean forecast								
Real-time	**0.57**	**0.64**	**-0.27**	**-0.13**	**0.77**	**0.33**	**-0.34**	**-0.19**
No parameter or state uncertainty	0.70	0.77	-0.30	-0.17	0.84	0.57	-0.36	-0.15
Final-vintage	0.80	0.79	-0.59	-0.72	0.92	0.89	-0.60	-0.74
Correlation with financial mean forecast								
Real-time	**-0.20**	**-0.06**	**0.99**	**0.99**	**-0.20**	**-0.06**	**0.99**	**0.99**
No parameter or state uncertainty	-0.32	-0.18	0.99	0.99	-0.20	-0.04	0.99	0.98
Final-vintage	-0.54	-0.57	0.99	0.99	-0.51	-0.67	0.99	0.98
Correlation with macro uncertainty								
Real-time	**0.61**	**0.84**	**0.07**	**0.10**	**-0.84**	**-0.89**	**0.04**	**0.00**
No parameter or state uncertainty	0.45	0.73	-0.04	-0.12	-0.78	-0.72	0.01	-0.06
Final-vintage	0.13	0.23	0.26	0.43	-0.79	-0.78	0.11	0.18
Correlation with financial uncertainty								
Real-time	**-0.08**	**0.05**	**0.95**	**0.97**	**-0.18**	**-0.20**	**0.81**	**0.80**
No parameter or state uncertainty	-0.25	-0.23	0.94	0.96	-0.07	-0.00	0.81	0.83
Final-vintage	-0.29	-0.29	0.94	0.96	-0.54	-0.74	0.80	0.79

Note: The table shows statistics related to macro and financial tail risks. Bold numbers refer to our baseline specification. Numbers in parentheses are standard deviations.

B Details on the Simulation Exercise

In Section 2.2, we discuss the results of a Monte Carlo exercise that compares posterior predictive distributions from an univariate version of the SV-VAR to those from a bivariate linear VAR. Here, we give details about the simulation exercise. Our data generating process is:

$$\begin{aligned} z_t &= c_z + \beta z_{t-1} + b_1 h_{t-1} + \exp\left(h_t\right)^{1/2} e_t \\ h_t &= c_h + \theta h_{t-1} + d_1 z_{t-1} + S^{1/2}\eta_t \\ \varepsilon_t &\sim N\left(0, \Sigma\right); \;\; \varepsilon_t = \left[e_t, \eta_t\right]. \end{aligned} \tag{13}$$

Table A.3 reports the parameter values used to simulate the data. Increases in volatility today lead to future declines in the level variable ($b_1 < 0$) – the volatility-in-mean effect, increases in the level variable today lead to declines in the volatility of its future innovations ($d_1 < 0$) – the level-in-volatility effect, and level and volatility innovations are contemporaneously negatively correlated ($\Sigma_{12} < 0$).

We adopt the identifying assumption that the main shock can contemporaneously affect both the level and volatility whereas the secondary shock can only contemporaneously affect volatility. The implications of this identification assumption for the parameters in the simple example are shown in equation (14), with $u_{z,t}$ as the orthogonal level shock and $u_{h,t}$ as the orthogonal volatility shock.

$$\begin{pmatrix} e_t \\ \eta_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \Sigma_{12} & \sqrt{1-\Sigma_{12}^2} \end{pmatrix} \begin{pmatrix} u_{z,t} \\ u_{h,t} \end{pmatrix}, \; \begin{pmatrix} u_{z,t} \\ u_{h,t} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right) \tag{14}$$

We generate a sample of $1,000$ observations from the model in equations (13) and (14), after burning $99,000$ observations. We then treat the volatility series h_t as an observable and use OLS to estimate a bivariate linear VAR(1) model:

$$\begin{pmatrix} z_t \\ h_t \end{pmatrix} = c_{lin} + \begin{pmatrix} \beta_{lin} & b_{1,lin} \\ d_{1,lin} & \theta_{lin} \end{pmatrix} \begin{pmatrix} z_{t-1} \\ h_{t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix}, \; \epsilon_t \sim N(0, \Sigma_{lin}) \tag{15}$$

The linear VAR does not allow for time-variation in distributions of ϵ_t. Therefore, in contrast to the data generating process, where changes in volatility affect both the conditional mean of z_t and the higher moments of the predictive distribution of z_t, changes in volatility only affect the conditional mean of z_t in the linear VAR by assumption. We adopt the same identifying assumption as in the data generating process to decompose ϵ_t into main and secondary shocks.

The predictive distributions found in Figure 2 are conditional on the values of z and h at high and low volatility states in the simulation. These states are in the top and bottom 5th percentiles of the volatility distribution, respectively. We are interested in the following predictive distributions:

$$p^i\left(z_{t+1} | z^t, \mathcal{M}^m, \Theta^m, H_t\right), \; m \in \{SV, LIN\} \tag{16}$$

where Θ^m is the vector that collects the parameters of the model. For the stochastic volatility model, we use the true parameter values as specified in Table A.3. For the linear VAR model, we use the OLS

point estimates for the parameter values. We take $50,000$ draws from the $t+1$ shocks to generate the predictive distributions. When constructing the scenarios, we select two points in the simulated path corresponding to a high and low volatility state.

Table A.3: Calibration of the Univariate Stochastic Volatility Model

Parameter	c_z	β	b_1	c_h	θ	d_1	S	Σ_{12}
Calibration	-0.15	0.3	-0.02	-2.21	0.7	-0.65	0.5	-0.8

Table A.4: Moments of One-Step Ahead Predictive Distributions – High Volatility State
(Stochastic Volatility vs. Linear VARs)

Simulation	Mean	U_{z_1}	Skewness	SF_{z_1}	LR_{z_1}	$LR_{z_1} - SF_{z_1}$
(a) *Stochastic Volatility VAR*						
Baseline	-0.35	0.26	-0.34	-1.09	-0.02	1.07
Main Shock	-0.81	0.42	-0.34	-2.03	-0.29	1.75
(b) *Linear VAR*						
Baseline	-0.32	0.13	0.00	-0.58	-0.05	0.53
Main Shock	-0.47	0.13	0.00	-0.74	-0.21	0.53

NOTE: Panel (a) in the table reports summary statistics for the one-step ahead posterior predictive distributions, $p(z_{t+f}|z_{1:t};\mathcal{M}_{SV})$, and panel (b) reports summary statistics for the distributions generated by the linear VAR model described in equation (15), $p(z_{t+f}|z_{1:t};\mathcal{M}_{LIN})$. U, SF, and LR denote uncertainty, shortfall and longrise, and are defined in equations (6)-(9), while skewness is Kelley's measure of skewness. See text and the note to Figure 2 for additional details.

C Properties of the Predictive Distribution

To explain how each ingredient included in the baseline model shapes the posterior predictive distributions, and consequently the measures of uncertainty and risk, we inspect restricted versions of the model. Figure A.1 shows the associated predictive distributions. We focus on three cases, and close the section discussing the properties of skewness in the full model.

In the paper, we show Kelley skewness, which is dependent on the 10th, 50th, and 90th quantiles of the distribution. This measure is less sensitive to extreme outliers, which we find makes the simulated results more stable when compared with the moment-based definition of skewness. In this Appendix section, we use the moment-based measure to derive analytical results. We verify by simulation that these results also hold for the Kelley skewness measure.

Conditional Mean Shifts A homoscedastic version of the model, shown in equation (17), can generate movements in the predictive distribution through shifts in the conditional mean:

$$\begin{aligned} z_t &= c + \beta z_{t-1} + \exp\left(\frac{1}{2}\alpha\right) e_t \\ e_t &\sim N(0,1) \end{aligned} \tag{17}$$

The important parameter that governs the strength of this effect is β. If we set $\beta > 0$, the effects on z_t of shocks to e_t linger, which shifts around the conditional mean of the predictive distribution. Therefore, time periods of higher-than-average downside risks follow strings of negative shocks to e_t, when z_t is below its long-run average. The top panel of Figure A.1 shows how the predictive distributions and downside risks change following a negative level shock. The blue line shows the one-step ahead predictive distribution of z_{t+1} when z_t is at its unconditional mean, which is 0. The orange line shows a predictive distribution assuming a lower value for the mean of z_{t+1}. The lower conditional mean of the predictive distribution naturally leads to a higher downside risk.

Assuming a constant volatility, however, implies that the shape of the distribution beyond the conditional mean stays fixed. All one-step ahead density predictions from the model in equation (17) are normally distributed with variance of $\exp(\alpha)$. Therefore, the predictive distributions are always symmetric around the conditional mean, with shortfall and longrise (downside and upside risks) moving in lockstep.

Conditional Volatility Shifts A heteroscedastic version of the model with no feedback between mean and volatility, shown in equation (18), can generate movements in the predictive distribution through shifts in the conditional volatility:

$$
\begin{aligned}
z_t &= c + H_t^{1/2} e_t \\
\tilde{h}_t &= \alpha + \theta \tilde{h}_{t-1} + S^{1/2} \eta_t \\
H_t &= \exp\left(\tilde{h}_t\right) \\
\begin{pmatrix} e_t \\ \eta_t \end{pmatrix} &\sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right)
\end{aligned}
\tag{18}
$$

The volatility of the level shock $\epsilon_t = H_t^{1/2} e_t$ is now stochastic and persistent. To isolate the effects of stochastic volatility alone on the predictive distributions, we shut down any dependence, contemporaneous or lagged, between the level and volatility equations. We also remove lags of z_t from the level equation.

Volatility is the key driver of changes to the predictive distribution. The second panel of Figure A.1 shows an example. The blue distribution has a conditional mean of zero and assumes that h_t is at its unconditional mean. The orange line shows the predictive distribution conditional on a large positive volatility shock that realizes at time t. The added volatility raises the left tail risk. This increase in risk is symmetric, however, so the upper tails of the distribution get pushed out simultaneously.[19]

Conditional Mean and Volatility Shifts through Lags Feedback. We next consider a version of the model that allows for feedback between the level and volatility equation through the lag structure:

$$
\begin{aligned}
z_t &= c + \beta z_{t-1} + H_t^{1/2} e_t \\
\tilde{h}_t &= \alpha + d_1 z_{t-1} \\
H_t &= \exp\left(\tilde{h}_t\right) \\
e_t &\sim N(0, 1)
\end{aligned}
\tag{19}
$$

The restricted model in equation (19) allows for GARCH-type effects, and is closely related to the parametric model estimated in Adrian et al. (2019) and Carriero et al. (2020). Level shocks are allowed to affect volatility with a lag. In this model, a level shock can generate shifts in the conditional mean and volatility of the one-step ahead predictive distribution if $\beta, d_1 \neq 0$. This is because a movement in z_t generates changes in the $t+1$ volatility of e_t and the conditional mean of z_{t+1}.

The bottom panel of Figure A.1 shows in orange the one-step ahead predictions generated by this model with $d_1 < 0$ and following a negative shock to e_t. For comparison purposes, the blue distribution is repeated from the top panel. Following the negative level shock, the mean of the predictive distribution shifts down and the variance increases simultaneously. This shift is such that the downside risk in the variable increases, as evidenced by the leftward move in the 5th percentile relative to the blue distribution. At the same time, however, the upper tail does not change, as the 95th percentiles from

[19] The predictive distributions are not normally distributed because of the independent randomness from $t+1$ volatility (η_{t+1}) that multiplies the normally distributed component e_{t+1}.

the two distributions lie on top of each other. This is because, while the mean of the orange distribution is lower than the blue distribution, its volatility is greater as well. The two effects cancel out at the upper tail of the distribution, leading to no movement in the 95th percentile.

In this model, the one-step ahead predictive distributions are normally distributed, while the multi-step ahead forecasts can be asymmetric. We characterize the asymmetry of the predictive distribution by examining its skewness. There are two points that we emphasize in discussing the skewness implied by the model: its direction and its time-varying nature.

We begin by discussing the direction of the skewness, which is determined by the parameter d_1.[20] We can analytically show this asymmetry by examining the two-step ahead predictive distribution, shown in equation (20):[21]

$$z_{t+2} = \beta^2 z_t + \beta H_{t+1}^{1/2} e_{t+1} + H_{t+2}^{1/2} e_{t+2}, \tag{20}$$

where z_t is observed. We are interested in the properties of the conditional distribution of z_{t+2}. To investigate the asymmetries in this distribution, we focus on the conditional skewness. In the context of this model, the skewness formula for the two-step ahead prediction is:

$$Skew_t\left(z_{t+2}\right) = \frac{E_t\left[\left(\beta H_{t+1}^{1/2} e_{t+1} + H_{t+2}^{1/2} e_{t+2}\right)^3\right]}{\left(E_t\left[\left(\beta H_{t+1}^{1/2} e_{t+1} + H_{t+2}^{1/2} e_{t+2}\right)^2\right]\right)^{3/2}}. \tag{21}$$

The numerator of the skewness expression determines its sign because the denominator is always positive, By expanding out the cubic function in the numerator, passing expectations through, and substituting in for H_{t+1} and H_{t+2}, we end up with the following expression:

$$3\beta \exp\left(\frac{3}{2}\alpha + d_1 c + \left(\frac{1}{2} + \beta\right) d_1 z_t\right) E_t\left[\exp\left(d_1 \exp\left(\frac{1}{2}\left(\alpha + d_1 z_t\right)\right) e_{t+1}\right) e_{t+1}\right]. \tag{22}$$

If $\beta > 0$, indicating positive autocorrelation of z_t, the term outside of the expectation is always positive. The term inside of the expectation is positive if $d_1 > 0$ and negative if $d_1 < 0$.[22] Therefore, the predictive distributions for z_{t+2} are negatively skewed if negative realizations of z_t tend to produce a future increase in volatility.

The intuition for this result is as follows. Take the example of $d_1 < 0$. Consider two components in the two-step ahead predictive distribution: the component conditional on a positive e_{t+1} realization and the component conditional on a negative one. A positive e_{t+1} realization tends to be relevant for the upper tail of the predictive distribution because of the persistence implied by $\beta > 0$, and vice versa. At the same time, a positive value of e_{t+1} *decreases* the volatility of the $t+2$ innovation through its effects on H_{t+2}, thereby shrinking the upper tails of the overall predictive distribution when compared

[20] This result is assuming that $\beta > 0$.

[21] For ease of notation, we assume that $c = 0$.

[22] The expectation term is of the form $E_t\left[\exp\left(X e_{t+1}\right) e_{t+1}\right]$. We know that for the case of $X = 0$, where the weight put on different values of e_{t+1} is always equal to 1, the expectation becomes $E_t\left[e_{t+1}\right] = 0$. If $X > 0$, positive values of e_{t+1} are weighted more than 1 and negative values of e_{t+1} are weighted less than 1. In turn, this means that the expectation will be greater than 0. By following the same logic, $X < 0$ leads to the expectation being less than 0.

with a case where volatility does not depend on the level. By the same argument, a negative value of e_{t+1} *increases* the volatility of the $t+2$ innovation and lengthens the lower tails of the predictive distribution. The combination of these effects leads to a negatively skewed distribution.

We move on to discussing the time-varying nature of skewness. Equation (23) shows the formula for the conditional skewness in terms of model parameters and z_t.

$$Skew_t\left(z_{t+2}\right) = \frac{3\beta \exp\left(\frac{3}{2}\alpha + d_1 c + \left(\frac{1}{2}+\beta\right) d_1 z_t\right) E_t\left[\exp\left(d_1 \exp\left(\frac{1}{2}\left(\alpha + d_1 z_t\right)\right) e_{t+1}\right) e_{t+1}\right]}{\left(\beta^2 \exp\left(\alpha + d_1 z_t\right) + \beta \exp\left(\alpha + d_1 c + d_1 \beta z_t + \frac{1}{2} d_1^2 \exp\left(\alpha + d_1 z_t\right)\right)\right)^{3/2}}. \tag{23}$$

Clearly, the conditional skewness is dependent on the value of z_t. Our conjecture is that, through trying a variety of parameter value combinations, skewness is always monotonically increasing in z_t. Therefore, if $d_1 < 0$, indicating that the distribution is always negatively skewed, a lower value of z_t leads to an even more negatively skewed predictive distribution.

In discussing this channel, we emphasized the effects of level shocks that can affect the volatility with a lag to be as close as possible to the Adrian et al. (2019) framework. Our SV-VAR model, however, can also generate this effect by specifying exogenous volatility shocks that affect the conditional mean with a lag, as in equation (24).

$$\begin{aligned} z_t &= c + b_1 \tilde{h}_{t-1} + H_t^{1/2} e_t \\ \tilde{h}_t &= \alpha + \theta \tilde{h}_{t-1} + S^{1/2} \eta_t \\ H_t &= \exp\left(\tilde{h}_t\right) \\ \begin{pmatrix} e_t \\ \eta_t \end{pmatrix} &\sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right) \end{aligned} \tag{24}$$

Much of the same intuition holds. The main difference is that exogenous volatility shocks are now the driving force of the asymmetry. If $b_1 < 0$, a positive shock to η_t will lower $E_t\left[z_{t+1}\right]$ and raise $\tilde{h}_{t+1}$ if volatility is persistent ($\theta > 0$). Therefore, there is a simultaneous shift in the conditional mean and volatility for the one-step ahead prediction. There is no skewness in the one-step ahead predictions, however, because of the independence between η_{t+1} and e_{t+1}.[23] For multi-step ahead forecasts, the sign of b_1 determines the direction of the asymmetry. If $b_1 < 0$, the predictive distributions is negatively skewed.

Contemporaneously Correlated Mean and Volatility Shocks. Finally, we look at the implications of allowing contemporaneous correlation between the mean and volatility innovations in the full model. We focus the discussion on showing that the skewness of the one-step ahead predictive distribution is time invariant. For convenience, we shut down all conditional mean and volatility dynamics in the model:

[23] Similar to the case with exogenous stochastic volatility, the one-step ahead predictive distribution is no longer normally distributed.

$$
\begin{aligned}
z_t &= c + H_t^{1/2} e_t \\
\tilde{h}_t &= \alpha + S^{1/2} \eta_t \\
H_t &= \exp\left(\tilde{h}_t\right) \\
\begin{pmatrix} e_t \\ \eta_t \end{pmatrix} &\sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \zeta \\ \zeta & 1 \end{pmatrix}\right)
\end{aligned}
\tag{25}
$$

The contemporaneous correlation between level and volatility innovations leads to asymmetries even in the one-step ahead forecasts for z_{t+1}. The sign of the skewness depends on the direction of the correlation. If $\zeta < 0$, a negative realization of e_{t+1} is associated with higher volatility η_{t+1}. A positive realization of e_{t+1} is associated with lower volatility. The combination of these two effects generates negative skewness. The opposite holds for $\zeta > 0$. This effect can be shown formally. The formula for the conditional skewness in this model is given by:

$$
Skew_t\left(z_{t+1}\right) = \frac{E_t\left[H_{t+1}^{3/2} e_{t+1}^3\right]}{\left(E_t\left[H_{t+1} e_{t+1}^2\right]\right)^{3/2}}. \tag{26}
$$

As before, the key determinant of the sign of the skewness is the sign of the numerator. Without loss of generality, we can decompose the correlated shocks e_{t+1} and η_{t+1} into uncorrelated shocks e_{t+1} and v_{t+1} using a Cholesky decomposition. For the purposes of this exercise, which is to understand the relationship between the sign of ζ and the asymmetry of the predictive distribution, the Cholesky decomposition is not an identification scheme. It is simply an orthogonalization meant to produce independent innovations.

$$
\begin{pmatrix} e_{t+1} \\ \eta_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \zeta & \sqrt{1-\zeta^2} \end{pmatrix} \begin{pmatrix} e_{t+1} \\ v_{t+1} \end{pmatrix}, \quad \begin{pmatrix} e_{t+1} \\ v_{t+1} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right) \tag{27}
$$

The orthogonalization allows us to rewrite the numerator of the skewness equation as:

$$
\exp\left(\frac{3}{2}\alpha\right) E_t\left[\exp\left(\frac{3}{2} S^{1/2} \zeta e_{t+1}\right) e_{t+1}^3\right] E_t\left[\exp\left(\frac{3}{2} S^{1/2} \sqrt{1-\zeta^2} v_{t+1}\right)\right]. \tag{28}
$$

The first and third terms in equation (28) are positive. Therefore, the sign of the skewness is dependent on the sign of the middle expectation term. To establish the relationship between ζ and the sign of this expectation, note that $E_t\left[e_{t+1}^3\right] = 0$. Therefore, in the case of $\zeta = 0$, the middle expectation is 0. If $\zeta > 0$, then realizations of $e_{t+1} > 0$ are weighted more heavily than realizations of $e_{t+1} < 0$, and the expectation is positive. The opposite case applies for $\zeta < 0$, leading to a negative skewness.

The skewness generated for the one-step ahead forecasts is not time-varying, even in the case of autoregressive mean and volatility dynamics. Equation 26 shows that any conditioned values affecting H_{t+1} in the model cancel out, leaving the skewness of the predictive distribution unaffected.

Figure A.1: Predictive Distributions: Restricted Versions of the Baseline Model

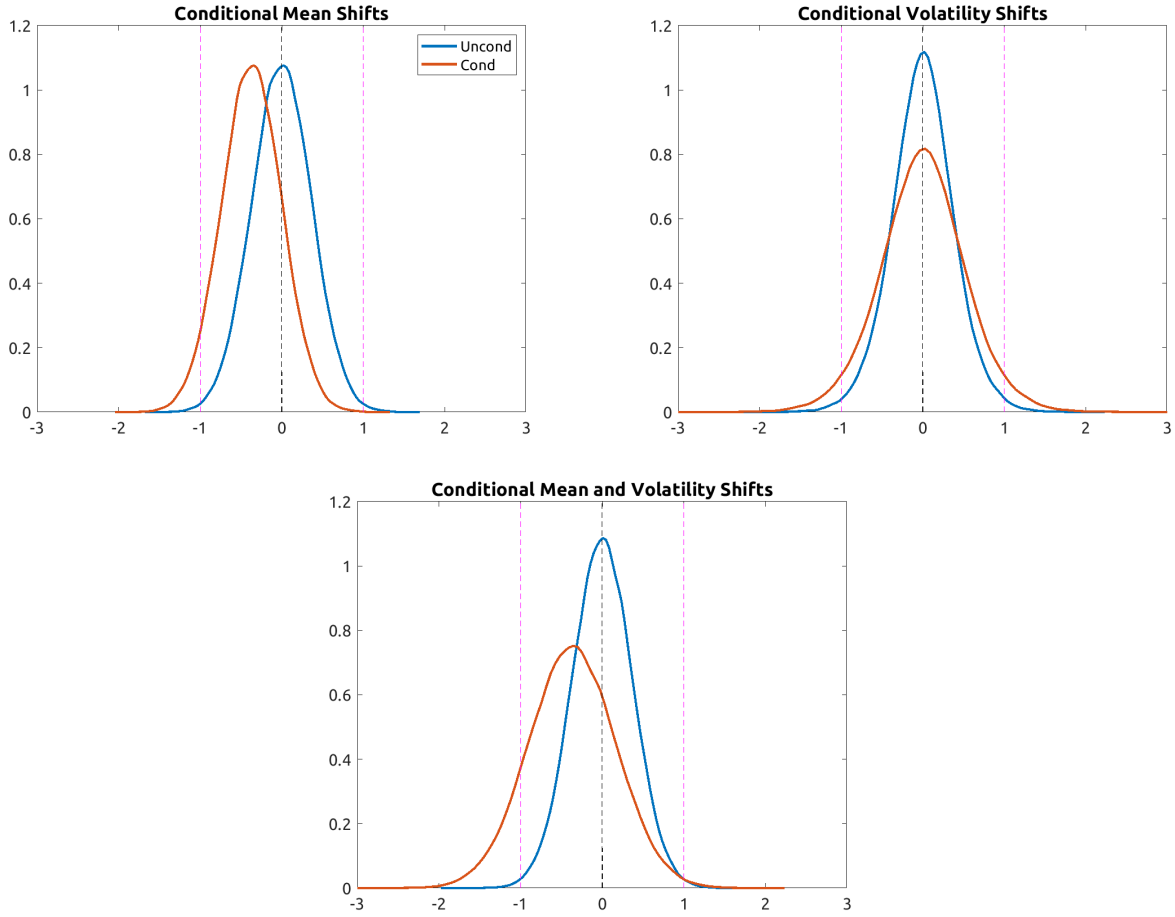

Note: This figure plots one-step-ahead predictive densities for three restricted versions of the baseline model. The top left panel shows predictive distributions generated by the model in equation (17). The blue line is the unconditional distribution while the orange line is the conditional distribution after a negative 2 standard deviation level shock. The top right panel shows predictive distributions generated by the model in equation (18). The blue line is the unconditional distribution with volatility fixed at its mean value while the orange line is the conditional distribution after a positive 2 standard deviation volatility shock. The bottom panel shows predictive distributions generated by the model in equation (19). The blue line is a repeat of the blue line in the first panel for comparison purposes, while the orange line shows the conditional distribution after a negative 2 standard deviation level shock. In all cases, the dashed black line shows the mean of the blue distribution while the dashed pink lines show the 5th and 95th percentiles.

D Details on the Estimation Procedure

Our baseline model estimation procedure largely follows Mumtaz (2018). Here, we provide the basic steps of the estimation procedure, and refer to Mumtaz (2018) for details of the posterior sampler.

Prior distributions Denote $B = \{c_z, c_h, \beta_{1:P}, b_{1:K}, \theta_{1:J}, d_{1:Q}\}$. We use Minnesota-type independent and normally distributed priors for the B parameters. For the parameters in the level equation, we center the constants and first lags at the $AR(1)$ OLS estimates on a presample of data from 1947:Q2 through 1953:Q1. We center the higher-order lags and all cross-lags around 0. For c_z, we set the variance to be $1000\sigma^2_{i,pre}$, where $\sigma^2_{i,pre}$ is the OLS estimates of the innovation variances in the presample. For the lags of the level variables, the variances are of the form $\frac{0.1^2}{l^2}$ for own-lags and $\frac{0.1^2}{l^2}\frac{\sigma^2_{i,pre}}{\sigma^2_{j,pre}}$ for the lag of variable j in equation i. For $b_{1:K}$, we set the prior means to be 0 and the prior variances to be $\sigma^2_{i,pre}$ in equation i.

For the parameters in the volatility equation, we first run an estimation using a VAR model with exogenous $AR(1)$ stochastic volatility on our entire data sample. We do not assume that there is any relationship between the level and volatility equations in this specification. We then take the posterior mean of the volatility estimates and fit $AR(1)$ models with OLS to them. We center the constants and first lags of the volatility estimates around 0.5. We center all higher-order lags and cross-lags around 0. We set the variances of the constant to be $\sigma^2_{i,h}$, where $\sigma^2_{i,h}$ is the OLS estimates of the stochastic volatility innovation variances. For the lags of the volatility variables, the variances are of the form $\frac{0.1^2}{l^2}$ for own-lags and $\frac{0.1^2}{l^2}\frac{\sigma^2_{i,h}}{\sigma^2_{j,h}}$ for the lag of variable j in equation i. For $d_{1:Q}$, we set the prior means to be 0 and the prior variances to be $0.1\sigma^2_{i,h}$ in equation i.

We set the priors for the diagonal elements of S as inverse Gamma with 5 degrees of freedom and a mean of 0.04. This is consistent with the literature (Clark, 2011; Carriero et al., 2015).

We put priors on transformed elements of $\Sigma = L^{-1}DL^{-1}$ where L^{-1} is lower triangular and D is diagonal. The priors on elements of L^{-1} are independent $N(0,1)$ and the prior on elements of D is implicit from the restriction that the diagona1 elements of Σ equals 1.

We follow Mumtaz (2018) and fix the initial condition h_0 at the corresponding estimate from the initial stochastic volatility estimation run.

Posterior Sampler Steps

1. Draw $B|S, \Sigma, h_t$

   Conditional upon knowing the sequence of stochastic volatilities, drawing the regression coefficients of the model boils down to a standard problem of linear regression estimation with stochastic volatility. We use the Kalman filter to determine the mean and volatility of the regression parameter draws.

2. Draw $S|B, \Sigma, h_t$

   Drawing S is a nonstandard problem because of the correlation amongst η_t. We use a Metropolis step assuming a mixture of an inverse Gamma distribution using the moments derived from an

assumption of independent elements in η_t and a random walk proposal using an inverse Gamma distribution centered on the previous draw. We assume a 20% probability of sampling from the random walk proposal and 80% probability of sampling from the inverse Gamma distribution. For the Metropolis step, we scale the variance of the proposal distribution by 0.01 times the posterior mean of an initial volatility estimate assuming independent AR(1) stochastic volatility processes and no level and volatility interaction.

3. Draw $\Sigma|B, S, h_t$

   We draw Σ using the algorithm of Chan and Jeliazkov (2009). We use a scaling parameter in the Metropolis step of 1.

4. Draw $h_t|B, S, \Sigma$

   We use the particle smoother with ancestor sampling using 20 particles to draw h_t. The further details of the state space form and implementation algorithm can be found in Mumtaz (2018) and references therein.

E Details on the Construction of the Posterior Predictive Distributions

We discuss in this section how we construct the posterior predictive distributions. These are nonstandard objects and we use simulation methods to generate them, following Del Negro and Schorfheide (2013).

The predictive distributions of our model are state-dependent, depending on the volatility h_t. A higher volatility amplifies the effects of shocks on the observables. As discussed in Section C of the Appendix, this state-dependent behavior carries over to the higher moments as well, including skewness. Acknowledging this fact, we are careful in differentiating the conditioning set when we form the predictive distributions.

We have two types of objects that we show in the paper: those that are conditional on a particular time period, such as the uncertainty, risk, financial crisis risk, and counterfactual predictive distributions and those that are not, such as impulse response functions and variance decompositions.

We take N draws from the posterior distribution of the parameters. For a given draw, for the objects that are conditional on a particular time period, we also construct $h_{1:T}$ for each draw. Then, we simulate out M paths for z and h. We repeat this exercise for all of the draws in the posterior distribution. Then, we compute our statistics of interest.

Algorithm to Generate Conditional Time T Distributions

1. For $n = 1, ..., N$:

   Take a parameter draw from the posterior distribution $\{B^n, S^n, \Sigma^n, \{h_t^n\}_{t=1}^T\}$.

2. For $m = 1, ..., M$:

   Simulate a path of length F $\{z_{T+1:T+F}^{sim,n,m}, h_{T+1:T+F}^{sim,n,m}\}$.

3. Collect all of the simulations $\{\{z_{T+1:T+F}^{sim,n,m}, h_{T+1:T+F}^{sim,n,m}\}_{m=1}^M\}_{n=1}^N$

4. Compute posterior statistics of interest, such as uncertainty, tail risk, and financial crisis probability.

When generating the counterfactual distributions in Section 4, we compute counterfactual paths and predictive distributions after shutting off certain structural shocks. In this case, we assume that some structural shocks are shut off beginning at some time $S < T$.

Algorithm to Generate Counterfactual Time T Distributions

1. For $n = 1, ..., N$:

   Take a parameter draw from the posterior distribution $\{B^n, S^n, \Sigma^n, \{h_t^n\}_{t=1}^T\}$.

2. Given an identification assumption, back out the implied sequence of structural shocks $\{u_t^n\}_{t=1}^T$.

3. Generate counterfactual paths $\{z_{S:T}^{counter,n}, h_{S:T}^{counter,n}\}$ by setting the structural shocks that are turned off to 0 from time S to T.

4. For $m = 1, ..., M$:

   Simulate a counterfactual path of length F $\{z_{T+1:T+F}^{sim,counter,n,m}, h_{T+1:T+F}^{sim,counter,n,m}\}$. This simulation takes as initial conditions $\{z_{S:T}^{counter,n}, h_{S:T}^{counter,n}\}$ and sets all of the structural shocks that are turned off to 0.

5. Compute posterior statistics of interest, such as predictive distributions.

For the objects that are not conditional on a certain time period, such as the impulse response functions and variance decompositions, we condition on steady state h_t levels. After taking a particular draw of the parameters, we first propagate 1000 periods of $z_t^{sim,burn}$ and $h_t^{sim,burn}$ with all shocks turned off before conditioning on the final period to compute the impulse response functions and variance decompositions.

Algorithm to Generate Impulse Response Functions

1. For $n = 1, ..., N$:

   Take a parameter draw from the posterior distribution $\{B^n, S^n, \Sigma^n\}$.

2. Simulate a burn-in path of length 1000 with all shocks turned off. Condition on $z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}$ where L is the max lag length of the VAR, which is defined as the maximum over P, K, J, and Q, which are the various lags in our VAR as defined in Equations 1 and 2.

3. Compute the conditional expectation for a path of length U

$$\{E_1\left[z_{1:U}^{sim,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n\right], E_1\left[h_{1:U}^{sim,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n\right]\}.$$

4. At time 1, consider a 1 unit structural shock of interest $u_1^{j,*,n} = 1$.

5. Compute a counterfactual conditional expectation of length F

$$\{E_1\left[z_{1:F}^{sim,counter,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n, u_1^{j,*,n}=1\right], E_1\left[h_{1:F}^{sim,counter,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n, u_1^{j,*,n}$$

6. To form the IRF, subtract the original path from the counterfactual path:

$$IRF_{F,z}^{j,n} = E_1\left[z_{1:F}^{sim,counter,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n, u_1^{i,*,n}=1\right] - E_1\left[z_{1:F}^{sim,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n\right]$$
$$IRF_{F,h}^{j,n} = E_1\left[h_{1:F}^{sim,counter,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n, u_1^{i,*,n}=1\right] - E_1\left[h_{1:F}^{sim,n}|z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}, \Theta^n\right]$$

7. Collect all of the IRF simulations $\{IRF_{F,z}^{j,n}, IRF_{F,h}^{j,n}\}_{n=1}^N$ and compute posterior statistics of interest.

Note that given the conditional values of $z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}$, it is possible to compute the conditional expectations of the system in closed form at any horizon, so we do not have to rely on simulation methods.

For the variance decompositions, we again rely on simulation methods.

Algorithm to Generate Variance Decompositions

1. For $n = 1, ..., N$:

   Take a parameter draw from the posterior distribution $\{B^n, S^n, \Sigma^n\}$.

2. Simulate a burn-in path of length 1000 with all shocks turned off. Condition on $z_{1000-L:1000}^{burn,n}, h_{1000-L:1000}^{burn,n}$ where L is the max lag length of the VAR, which is defined as the maximum over P, K, J, and Q, which are the various lags in our VAR as defined in Equations 1 and 2.

3. For $m = 1, ..., M$:

   Simulate a path of length F with all structural shocks turned on $\{z_{T+1:T+F}^{sim,n,m}, h_{T+1:T+F}^{sim,n,m}\}$.

4. Compute the variance at horizon F: $VAR\left(z_{T+F}^{sim,n,m}\right)$, where the variance is taken with respect to the m simulations.

5. For $m = 1, ..., M$:

   Simulate a path of length F shutting off the structural shock of interest j: $\{z_{T+1:T+F}^{counter,sim,n,m}, h_{T+1:T+F}^{counter,sim,n,m}\}$.

6. Compute the variance at horizon F: $VAR\left(z_{T+F}^{counter,sim,n,m}\right)$, where the variance is taken with respect to the m simulations.

7. Compute variance decomposition with respect to structural shock j: $VD^{j,n}\left(z_{t+F}\right) = 1 - \frac{VAR\left(z_{T+F}^{counter,sim,n,m}\right)}{VAR\left(z_{T+F}^{sim,n,m}\right)}$.

8. Collect all of the variance decomposition simulations $\{VD_{F,z}^{j,n}\}_{n=1}^N$

Table A.5: Results from Estimation on Simulated Data

Parameter	c_z	β	b_1	c_h	θ	d_1	S	Σ_{12}
90% CS Sample size 275	$[0.61, 1.13]$	$[0.23, 0.50]$	$[-0.52, -0.12]$	$[-0.75, -0.14]$	$[0.43, 0.72]$	$[-0.31, -0.03]$	$[0.07, 0.24]$	$[-0.76, -0.39]$
90% CS Sample size 975	$[0.81, 1.05]$	$[0.37, 0.53]$	$[-0.26, -0.10]$	$[-0.49, -0.07]$	$[0.52, 0.75]$	$[-0.34, -0.09]$	$[0.08, 0.20]$	$[-0.68, -0.43]$
Calibration	1	0.4	-0.2	-0.15	0.7	-0.2	0.2	-0.5

F Estimation of Baseline Model on Simulated Data

We test our estimation strategy by generating simulated data from a univariate version of baseline model and then using that data to estimate the parameters of the model.

We simulate data from the model in Equation 13 with parameters in the bottom row of Table A.5. Our calibration implies that the volatility-in-mean effect, level-in-volatility effect, and correlation between level and volatility innovations are all active. We generate a sample length of $50,000$ from this model and use the both the final 299 and 999 data points to estimate our model.

We use the same prior distributions as in our main estimation strategy. Therefore, we reserve the first 24 data points as a presample. We take $250,000$ draws from the posterior distribution with a burn-in of $50,000$ draws. The posterior sets are computed saving every 25th draw.

The middle two rows in Table A.5 show the results from this exercise. Regardless of whether the sample size is 275 or 975, the 90% posterior credible sets contain the true parameter in all cases. Moreover, as expected, the credible sets shrink as we increase the sample size.

The most difficult parameter to estimate in the system is S, which is known in the literature to be difficult to identify and influenced in part by the prior specification. In this case, our prior mean of 0.04 may be downwardly affecting the estimates. It is reassuring, however, that the credible sets still contains the true parameter value and that it is shrinking as we increase the sample size. Moreover, we tried a second estimation on a new set of data with 999 total observations, where the 90% credible set of $[0.08, 0.23]$ safely cover 0.2.