Aggregate Skewness and the Business Cycle

Martin Iseringhausen
European Stability Mechanism

Ivan Petrella
University of Warwick
CEPR

Konstantinos Theodoridis
European Stability Mechanism
Cardiff Business School

July 2022

Abstract

We develop a data-rich measure of expected macroeconomic skewness in the US economy. Expected macroeconomic skewness is strongly procyclical, mainly reflects the cyclicality in the skewness of real variables, is highly correlated with the cross-sectional skewness of firm-level employment growth, and is distinct from financial market skewness. Revisions in expected skewness deliver dynamics that are nearly indistinguishable from those produced by the *main business cycle* shock of Angeletos et al. (2020). This result is robust to controlling for macroeconomic volatility and uncertainty, and alternative macroeconomic shocks. Our findings highlight the importance of higher-order dynamics for business cycle theories.

JEL classification: C22, C38, E32
Keywords: Business cycles, downside risk, skewness

The views expressed in this paper are those of the authors and do not necessarily reflect those of the European Stability Mechanism (ESM). The authors would like to thank Nicholas Bloom, Harris Dellas, Ian Dew-Becker, Marco Del Negro, Thomas Drechsel, Domenico Giannone, Michael McCracken, Silvia Miranda-Agrippino, Haroon Mumtaz, Emiliano Santoro, Laura Veldkamp, and the participants at the IAAE 2022 Annual Conference, the RCEA Conference on Recent Developments in Economics, Econometrics and Finance, and an ESM Research Seminar for helpful comments. Correspondence to: Martin Iseringhausen (m.iseringhausen@esm.europa.eu), Ivan Petrella (ivan.petrella@wbs.ac.uk), and Konstantinos Theodoridis (k.theodoridis@esm.europa.eu).

1 Introduction

> *"FOMC participants (Board members and Reserve Bank presidents) indicated that considerable uncertainty surrounded the outlook for economic growth and that they saw the risks around that outlook as skewed to the downside."*
>
> Monetary Policy Report to Congress, Federal Reserve Board, Feb. 2008 (p.2)

> *"The outlook for the UK and global economies remains unusually uncertain.* [...] *The risks are skewed to the downside."*
>
> Monetary Policy Report, Bank of England, Aug. 2020 (p.1)

Assessing macroeconomic risks and analysing their potential impact on the economy is a key focus of economic policy institutions. Such risks are often not balanced around the baseline outlook, and the concept of skewness has been a device for policy-makers to communicate their beliefs about the evolution of risks. The quotes above are examples of central bank communication about, respectively, the onset of the Great Recession and the aftermath of the Covid-19 shock. Therefore, a more precise assessment and understanding of economic asymmetries supports a better communication of potential risks and the adoption of economic policies to mitigate them. The academic literature has also used skewness to characterize the asymmetric effects of economic shocks due to, for instance, non-linearities (e.g. Petrosky-Nadeau et al., 2018; Jensen et al., 2020; Mumtaz and Theodoridis, 2020) or particular adverse events (e.g. Barro, 2009; Gourio, 2012; Fernández-Villaverde and Levintal, 2018). In this paper, we develop a new measure of expected macroeconomic skewness for the US economy, reflecting variations in the balance of risks of a large number of (nominal and real) macroeconomic and financial indicators. We contrast this measure with alternative measures of macro and micro skewness, and investigate the relationship between fluctuations in aggregate macroeconomic skewness and the business cycle.

A long-standing literature has argued that macroeconomic fluctuations are plagued by asymmetries, highlighting that recessions tend to be relatively deeper and more pronounced than expansions (Neftci, 1984; Hamilton, 1989; Sichel, 1993; Morley and Piger, 2012). More recent work has studied the asymmetry of the conditional distribution of GDP growth, documenting the presence of procyclical GDP growth skewness related to the state of macro-financial conditions (e.g. Adrian et al., 2019; Loria et al., 2020; Delle Monache et al., 2021; Forni et al., 2021).[1] These studies focus on measuring (expected) asymmetry of a single macroeconomic variable, namely GDP growth. While GDP is one of the most representative measures of the business cycle, it is unclear to what extent conditional skewness in GDP

[1]Theoretical and empirical contributions highlighting the role of time-varying skewness include, for example, Colacito et al. (2016), Dew-Becker et al. (2019), Jensen et al. (2020) and Fève et al. (2021) at the macro level, and Busch et al. (2018), Salgado et al. (2019), and Dew-Becker (2021) at the micro level.

growth summarizes fluctuations in downside risk for the broader macroeconomy. This highlights the need for an economy-wide measure covering also, for example, prices, labor market indicators and financial variables. We derive a new measure of aggregate expected skewness, which represents a common factor driving the individual conditional skewness series of the indicators included in the McCracken and Ng (2020) dataset. The latter are computed using robust asymmetry measures (Kelley, 1947), where time-varying asymmetry derives from the relative movements of the conditional quantiles of the distribution captured using quantile regression techniques (Koenker and Bassett, 1978; Engle and Manganelli, 2004). This procedure allows us i) to derive summary measures that refer to different subgroups (e.g. prices, labor market indicators and financial variables) and ii) to understand which variables contribute most to overall skewness. Finally, the simplicity of the derivation allows the timely update of the series, which can be downloaded from the authors' websites.

The common skewness factor explains only a limited part of the dynamics in expected skewness for most of the macroeconomic indicators. It explains more of the skewness variation of the real economy variables (including income, labor markets, orders and sales, and production indicators) compared to, for example, prices. Moreover, the factor accounts for a non-negligible fraction of the conditional asymmetry in some of the financial indicators, in particular non-household balance sheet indicators, whereas it is less related to the skewness in interest rates and credit measures.

The economy-wide measure is strongly procyclical and is highly, but not perfectly, correlated with the skewness of GDP growth, meaning that the latter may not always capture economy-wide risks. Our measure also comoves with the GDP growth skewness that conditions on past macro-financial data (Adrian et al., 2019). This is in spite of the fact that our measure captures common movements in conditional asymmetry across a large number of indicators, where the skewness of each variable is derived using only information contained in past observations of the variable itself. This has two advantages: i) it does not require to identify the most appropriate predictors for each of the variables and ii) it allows for the possibility that macroeconomic asymmetries are not related to – or move in tandem with – financial conditions, as it was the case during the Covid-19 pandemic crisis. Our expected skewness factor is also highly correlated with the cross-sectional skewness of employment growth computed at the firm level by Salgado et al. (2019), which is remarkable since the data and methodologies used to construct these two measures are completely different. By contrast, our measure is only very mildly correlated with indicators of financial market skewness, including stock return skewness, either computed at the market level (Dew-Becker, 2021) or the firm level (Salgado et al., 2019).

Our second contribution relates to investigating the role of our skewness factor in the US

business cycle. In recent studies, Salgado et al. (2019) and Forni et al. (2021) demonstrate that shocks to the cross-sectional skewness of firm-level stock returns and the predictive GDP growth distribution, respectively, can produce contractionary movements in macroeconomic and financial indicators. Building on these results, we show that revisions in expected skewness, which are associated with an increase in perceived downside risk, lead to a substantial contraction in output, consumption, and investment, while leaving prices and TFP broadly unaffected. Remarkably, the IRFs of such a shock are almost identical to those documented in Angeletos et al. (2020). In fact, revisions in expected skewness are strongly correlated with the *main business cycle* (MBC) shock identified in Angeletos et al. (2020). This finding is robust to various sensitivity exercises. Specifically, revisions in expected skewness are distinct from movements in aggregate volatility and uncertainty, and appear unrelated to alternative shocks capturing credit risk, productivity, fiscal policy, and monetary policy.

Our empirical results highlight that any model that has the ambition to explain the main force of macroeconomic fluctuations needs to allow for higher-order dynamics and possibly relate those to economic agents' varying perception of downside risk. In this regard, within theories that suggest that a single shock is driving the business cycle, this key driver of macroeconomic fluctuations also needs to account for the bulk of the variation in revisions of perceived macroeconomic risk. Theories allowing for i) confidence or sentiment shocks (Angeletos and La'O, 2013; Angeletos et al., 2018); ii) the possibility of rare disasters (Rietz, 1988; Barro, 2006; Barro and Ursúa, 2008; Gabaix, 2008; Barro, 2009; Gourio, 2012; Wachter, 2013; Petrosky-Nadeau et al., 2018; Jordà et al., 2020); iii) informational frictions and learning asymmetries (Veldkamp, 2005; Ordonez, 2013); or iv) left-skewed uncertainty of households or firms (Salgado et al., 2019), could provide promising avenues.

The remainder of the paper is structured as follows: Section 2 derives the aggregate expected skewness factor. Section 3 presents the VAR results while Section 4 discusses various robustness checks. Finally, Section 5 concludes.

2 A data-rich skewness measure for the US economy

This section presents a new measure of expected asymmetry based on a large dataset of macroeconomic and financial variables. To construct the skewness measure, we use the quarterly version of the McCracken and Ng (2016) dataset (FRED-QD) that contains 248 time series starting from 1959 and categorized into 14 groups.[2] All variables are transformed to make them stationary by using the transformations suggested by the authors. We remove

[2]These are *national income and product accounts (NIPA)*; *industrial production*; *employment and unemployment*; *housing*; *inventories, orders, and sales*; *prices*; *earnings and productivity*; *interest rates*; *money and credit*; *household balance sheets*; *non-household balance sheets*; *stock markets*; *exchange rates*; and *other*.

those series that have missing observations over our sample period 1960:Q1–2019:Q4, which reduces the number of variables to $N = 211$. Next, we estimate for each (de-meaned) variable y_i and each quantile level $p = \{10\%, 50\%, 90\%\}$, the following autoregressive quantile regression as developed in Engle and Manganelli (2004)

$$Q^p(y_{i,t}) = \beta_0^p + \beta_1^p Q^p(y_{i,t-1}) + \beta_2^p y_{i,t-1}\mathbb{I}(y_{i,t-1} > 0) + \beta_3^p y_{i,t-1}\mathbb{I}(y_{i,t-1} < 0), \tag{1}$$

where $i = 1, ..., N$ and $t = 2, ..., T$. This *asymmetric slope* model (Engle and Manganelli, 2004) allows for a different impact of past observations on the respective quantiles, depending on whether they lie above or below the unconditional mean of the series. This permits an asymmetric impact of contractions and expansions in each variable on the different quantiles, so that, for instance, a recession can affect downside risk without necessarily affecting upside risk. In addition, the model allows the quantiles to be persistent, which seems appropriate given the well-documented persistence of the first two moments of many macroeconomic series (see, for example Antolin-Diaz et al., 2017).[3] The model parameters are estimated by regression quantiles (Koenker and Bassett, 1978) and further details can be found in Engle and Manganelli (2004). The conditional quantile autoregressive model belongs to the class of observation-driven models, for which the trajectories of the time-varying parameters are perfectly predictable one-step-ahead given past information (Cox, 1981). Using the estimated model parameters from these quantile regressions, and assuming that agents' use Equation (1) to form their expectations, we compute for each variable the one-step-ahead expected, or predicted, Kelley skewness (Kelley, 1947)

$$\mathbb{E}_t[Skew(y_{i,t+1})] = \frac{\mathbb{E}_t[Q_{i,t+1}^{0.9}] + \mathbb{E}_t[Q_{i,t+1}^{0.1}] - 2\mathbb{E}_t[Q_{i,t+1}^{0.5}]}{\mathbb{E}_t[Q_{i,t+1}^{0.9}] - \mathbb{E}_t[Q_{i,t+1}^{0.1}]}. \tag{2}$$

Since each quantile estimate is computed as a (variable-specific) moving average of a non-linear function of the variable itself, there is no reason ex-ante to expect that the skewness of any series displays a particular cyclical behaviour or comoves across indicators. Our overall measure of expected asymmetry is then constructed as the first principal component obtained from the set of series-specific expected skewness measures, where each measure is first standardized by subtracting the series-specific mean and dividing by its standard deviation (see, e.g., Stock and Watson, 2002).[4] Since the skewness factor is based on PCA, its sign is not identified. We identify the sign by assuming a positive correlation between the

[3] Since we are interested in capturing cyclical movements in skewness rather than slow-moving trends, we restrict the degree of persistence, i.e. $0 < \beta_1^p < 0.8$.

[4] While somewhat different from our approach, Chen et al. (2021) is another example of combining quantile regression and factor analysis to measure, in their case, comovement across quantiles.

skewness factor and the skewness of GDP growth. The factor reflects common movements of skewness across a large number of macroeconomic and financial indicators and does not necessarily overlap with the skewness of any specific indicator, e.g. the skewness of GDP growth. Moreover, the common factor should be relatively immune to idiosyncrasies and noise in the measurement of expected skewness for each of the individual series arising, for instance, from the estimation of the time-varying quantiles. This relies on such noise being idiosyncratic, which is arguably a reasonable assumption since the quantile model does not include common predictors. Appendix A presents a simulation exercise showing that our two-step approach to construct the skewness factor does not yield spurious results, i.e. the factor is, on average, zero if the DGP does not feature conditional skewness.

Table 1: Descriptive statistics of skewness variation explained by first principal component (in %)

Group	**Variables**	**Mean**	**Median**	**Max.**	**Min.**
National income and product accounts	22	18.3	8.9	55.9	0.1
Employment and unemployment	44	17.9	14.1	66.7	0.0
Inventories, orders, and sales	6	15.5	17.9	26.7	0.3
Non-household balance sheets	11	14.4	12.7	28.1	0.8
Industrial production	15	12.9	9.8	43.9	0.4
Stock markets	5	12.6	8.0	34.2	0.2
Exchange rates	4	12.5	14.5	20.4	0.5
Household balance sheets	9	9.7	8.7	26.8	0.0
Housing	6	9.4	7.5	20.2	0.3
Prices	46	8.9	5.4	52.1	0.0
Interest rates	18	8.5	4.1	50.9	0.0
Earnings and productivity	10	7.2	5.4	19.9	0.0
Money and credit	14	6.2	3.0	21.5	0.2

Note: This table presents descriptive statistics for the shares of variation of the individual skewness series explained by the skewness factor (in %). The grouping follows McCracken and Ng (2020). The group *Other* has been dropped from this table as it only contains one variable with a low share of explained variation.

This skewness factor explains around 12% of the variation across the individual skewness series. While appearing low, this is not surprising given that the dataset includes many series with a small degree of asymmetry that load only weakly on the common factor.[5] Table 1 illustrates this point by showing the share of variation explained by the skewness factor for

[5]For comparison, the first principal component of the actual data accounts for around 20% of the variation, while a common volatility (GARCH) factor (see Section 4) accounts for around 26% of the variation in dispersion. Lastly, a common factor of a quantile-based dispersion measure (expected interquartile range), accounts for around 24% of the variation.

the groups of variables. The skewness factor tends to explain more of the skewness variation of the real economy variables including NIPA, labor markets, and production indicators compared to, for example, prices. Moreover, the factor accounts for a non-negligible fraction of the conditional asymmetry in some financial indicators such as non-household balance sheets, stock markets and exchange rates, whereas it accounts only for a smaller fraction of the skewness variability in interest rates and money and credit indicators.

Existing studies have largely focussed on the conditional asymmetry of a single variable, i.e. GDP growth (see, for example, Adrian et al., 2019; Loria et al., 2020; Jensen et al., 2020; Forni et al., 2021; Castelnuovo and Mori, 2022). This is different from our data-rich approach where the skewness factor reflects variation in risks of a large number of macroeconomic and financial indicators. The upper panels of Figure 1 compare this expected skewness factor with the individual (de-meaned) skewness series of GDP growth obtained from different conditional quantile models. Figure 1(a) shows that our skewness factor is strongly correlated with the individual GDP skewness series retrieved using the conditional autoregressive quantile model described above. The two series display a similar procyclical behavior. In particular, aggregate expected skewness drops strongly during recessions and moves to mildly positive values during the expansionary phases of the cycle.

Interestingly, the documented correlation with GDP growth skewness does not depend on the inclusion of GDP itself in the underlying dataset. In fact, the skewness factor appears very robust to the choice of the specific dataset. To illustrate this aspect, we have also computed an alternative skewness factor based on a subset of 101 variables that largely match those used in Stock and Watson (2012).[6] Notably, real GDP and several expenditure components are not part of this smaller dataset. Nevertheless, the correlation with GDP growth skewness remains very similar, around 0.8. In spite of their similarities, there are also important differences between our measure of aggregate expected skewness and the expected skewness of GDP growth. In particular, the latter features a distinct downward trend in the last part of the sample, which is in line with the findings of Delle Monache et al. (2021) and appears to be a feature not shared by other indicators in our dataset.

A growing literature starting from Adrian et al. (2019) has recently developed measures of downside risk and skewness for future GDP growth conditional on measures of financial conditions. Therefore, Figure 1(b) contrasts our skewness factor with the (Kelley) skewness of the fitted growth distribution based on the model specification of Adrian et al. (2019). Quantile regressions that include financial conditions imply a more asymmetric conditional growth distribution and a longer left tail during recessions.

[6]The respective variables are explicitly labelled in the McCracken and Ng (2020) dataset. Figure D-1 in Appendix D compares the two skewness factors which are highly correlated.

Figure 1: Skewness factor vs. other measures of skewness

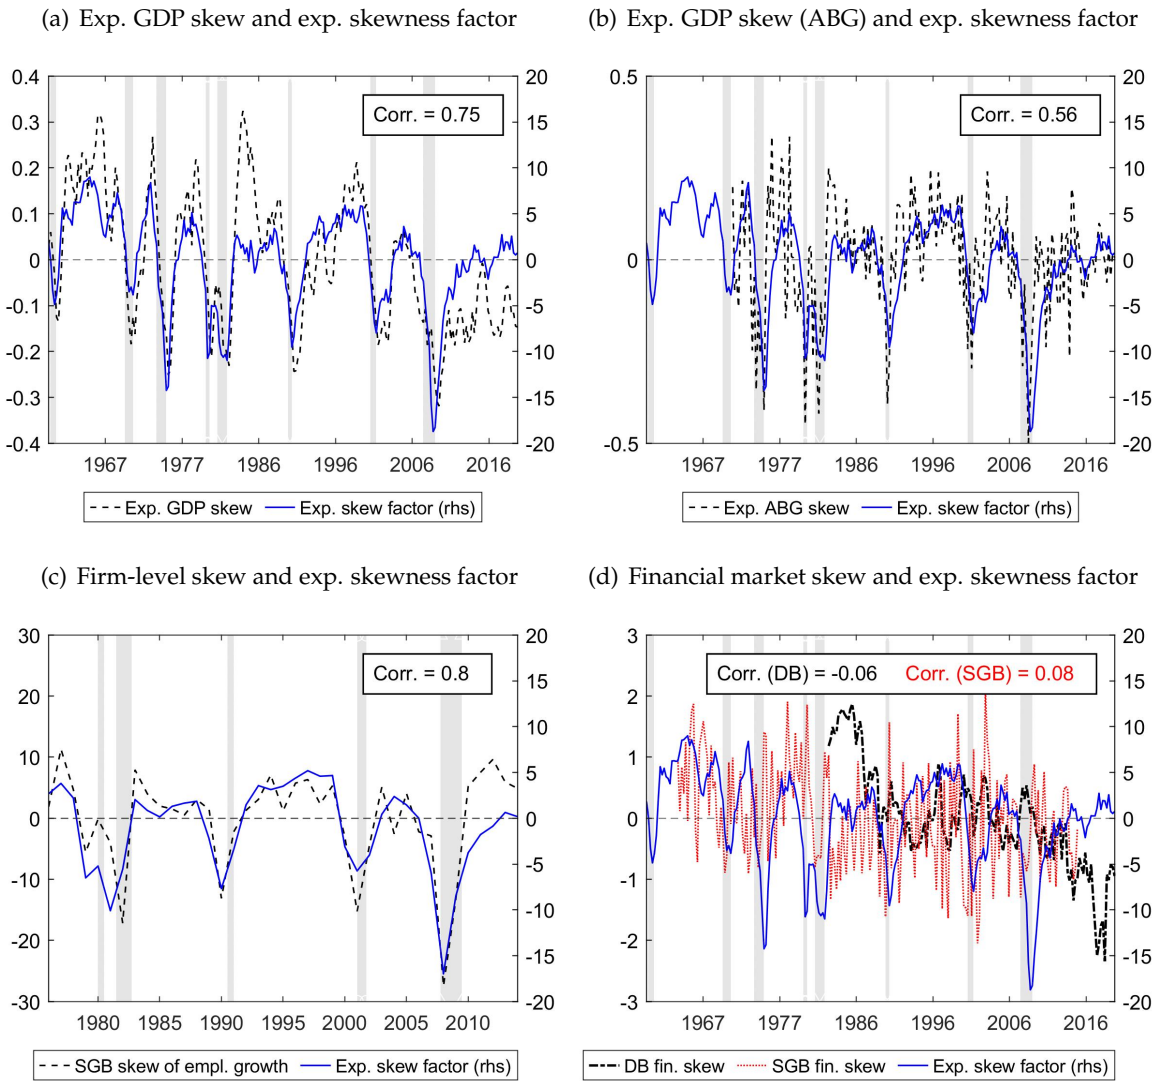

Note: Figure 1(a) shows the expected skewness factor together with the individual skewness series of quarter-over-quarter real GDP growth derived based on the quantile specification in Equations (1)-(2). Figure 1(b) shows the skewness factor together with the (Kelley) skewness of GDP growth obtained using the approach of Adrian et al. (2019) (ABG). This series is based on quantile regressions of real GDP growth on lagged growth and the lagged National Financial Conditions Index (NFCI) computed by the Chicago Fed. Figure 1(c) shows the skewness factor (annual avg.) together with the (employment-weighted) cross-sectional Kelley skewness of firms' log employment growth obtained from Salgado et al. (2019) (SGB) for the period 1976–2014. Figure 1(d) shows the skewness factor together with i) the monthly option-implied measure of market skewness for the S&P 500 developed in Dew-Becker (2021) (DB, quarterly avg., 1983:Q2–2019:Q4), and ii) the cross-sectional (Kelley) skewness of firms' daily stock returns within a month computed in Salgado et al. (2019) (SGB, quarterly avg., 1964:Q1–2015:Q1). All alternative skewness series are de-meaned and the scale of the SGB financial skewness measure is adjusted for comparability with the DB measure. Gray areas are NBER recessions.

We document a sizeable correlation between our expected skewness factor and the expected growth skewness based on the Adrian et al. (2019) approach, which is around 0.6. While this correlation is simply a stylized fact, it highlights that elevated asymmetry during downturns is a feature shared by a number of economic indicators and not necessarily related to fluctuations in financial conditions (see also Plagborg-Møller et al., 2020).

We also compare our measure of macro skewness with micro-level and financial market measures of asymmetry. First, Salgado et al. (2019) analyse the cross-sectional skewness of firm-level outcomes such as employment growth and sales, and find these series to be procyclical. Figure 1(c) compares the cross-sectional (Kelley) skewness of firms' employment growth (Salgado et al., 2019) with our expected skewness factor.[7] Both series move together closely and share a correlation of around 0.8. Given the different underlying methodologies, we interpret this result as i) potential evidence that the same shocks or mechanisms drive both firm-level and aggregate skewness and ii) an affirmation of our interpretation of the skewness factor as an economy-wide skewness measure. Second, Figure 1(d) contrasts our expected skewness factor with two measures of financial market skewness. Specifically, we show the option-implied skewness of the S&P 500 index computed at the market level by Dew-Becker (2021), and the cross-sectional firm-level series of stock return skewness of Salgado et al. (2019). The correlation between the skewness factor and, respectively, option-implied market-level skewness and cross-sectional return skewness is relatively low. This provides further support to the interpretation of the aggregate skewness factor as a measure of macroeconomic skewness which is distinct from financial market skewness.[8]

Lastly, our skewness measure correlates with – but is still quite distinct from – aggregate volatility and uncertainty.[9] Table D-1 in Appendix D shows a correlation matrix including the expected skewness factor, the first principal component of the actual data (X) and squared data (X^2) akin to Gorodnichenko and Ng (2017), a common factor of the expected interquartile ranges derived from Equation (1), an expected volatility (GARCH) factor, and two popular measures of uncertainty (Jurado et al., 2015; Ludvigson et al., 2021).[10] For all measures, except the uncertainty indices, the dataset is the same as the one used to extract skewness. Given the procyclicality of the skewness factor, it is not surprising to find negative

[7] To preserve the forward-looking character of the skewness factor, we compute the annual average for each year t over the period Q4 (t) to Q3 ($t+1$). However, this implies that for the annual series, expectations about skewness in $t+1$ are no longer formed conditional on information in year t only. The firm-level skewness series was directly taken from the replication files provided by Salgado et al. (2019). The authors compute this series based on the US Census Bureau's Longitudinal Business Database.

[8] Ludvigson et al. (2021) highlight a similar disconnect between macro and financial market uncertainty.

[9] Orlik and Veldkamp (2014) highlight how within a Bayesian learning framework, where agents attempt to learn the evolving distribution of GDP growth, uncertainty, skewness and therefore downside risk, are naturally related to one another.

[10] The fact that the quantile-based volatility measure is strongly correlated with the GARCH factor (> 0.9) and macroeconomic uncertainty (> 0.8) provides reassurance that our procedure also reliably measures skewness.

comovement with uncertainty, which moves countercyclically (see, e.g., Jurado et al., 2015).

3 Macroeconomic effects of skewness shocks

This section studies the macroeconomic effects of exogenous variation in expected skewness. We add our measure of skewness to an otherwise standard VAR model. The empirical specification, the variables included, as well as the estimation approach largely follow Angeletos et al. (2020). Within this set up, we study the relationship between revisions in expected skewness and the *main business cycle* shock of these authors.

The baseline VAR spans the period 1960:Q1–2017:Q4 and contains the following variables: the expected skewness factor (see Figure 1), real GDP per capita, real investment per capita, real consumption per capita, hours worked per person, unemployment rate, labor share, effective federal funds rate, inflation, labor productivity (non-farm business sector), and a measure of TFP.[11] Details on these and other variables used in augmented models, can be found in Appendix B. The VAR model has the following representation:

$$y_t = \sum_{p=1}^{P} \Theta_p y_{t-p} + u_t, \quad u_t \sim \mathcal{N}(\mathbf{0}, \Sigma) \tag{3}$$

where $\Theta_p \; \forall \; p = 1, ..., P$ are the matrices of VAR coefficients, and u_t is a vector of reduced-form disturbances, which are linear combinations of the underlying structural (orthogonal) shocks $u_t = A_0 \varepsilon_t$. A_0 is the matrix containing the contemporaneous responses, where $A_0 A_0' = \Sigma$. Due to the relatively large dimension of the VAR model, we adopt a Bayesian estimation approach and employ a Minnesota-type prior. The parameter that controls the tightness of this prior is set to $\lambda = 2$. This is a commonly used value in empirical studies with US data and Section 4 shows that the results hold even for looser configurations. We approximate the joint posterior distribution of the parameters using a Gibbs sampling algorithm. Appendix C contains further details on the prior specification and the estimation of the VAR model. We choose a baseline lag length of $P = 2$ and demonstrate robustness with respect to this choice in Section 4.

To identify exogenous variation in expected skewness, our baseline approach imposes zero restrictions on the matrix containing the contemporaneous responses. For this, A_0 is identified as the lower triangular matrix obtained from a Cholesky decomposition of Σ. Ordering our skewness measure first, this simple identification scheme provides us with an intuitive interpretation of the identified shock as the revision in expected skewness, i.e.

[11] We use the original dataset made available by Angeletos et al. (2020). All key results hold when extending the sample to 2019:Q4 (where the TFP series of Fernald (2014) ends) and using the latest (revised) data.

$\mathbb{E}_t[Skew_{t+1}] - \mathbb{E}_{t-1}[\mathbb{E}_t[Skew_{t+1}]]$ where the expectation $\mathbb{E}_{t-1}$ is conditional on the information set spanned by the VAR. Section 4 shows that an alternative approach (Uhlig, 2003) which identifies a shock that has the largest contribution to the variation of skewness over a one-year horizon yields very similar results. This suggests that revisions in expected skewness are also the main driver of skewness dynamics over the business cycle.

Figures 2 and 3 show the impulse response functions following a negative shock to expected skewness, i.e. a downward revision of expected skewness, and the corresponding forecast error variance contributions, together with those of the MBC shock of Angeletos et al. (2020). The latter is identified as an unemployment shock using the max-share approach of Uhlig (2003) and targeting four quarters in the time domain. Both shocks are identified within the same VAR specification. A revision in expected skewness generates business cycle dynamics that are very similar to the *business cycle anatomy* documented in Angeletos et al. (2020). These dynamics reflect a sizeable, but relatively short-lived, comovement between GDP, investment, consumption, hours worked, and unemployment, without meaningful movements in inflation and TFP. Table 2 shows that the (unconditional) correlation between the MBC shock and our skewness shock is around 0.8. Angeletos et al. (2020) use the *business cycle anatomy* to shed light on the transmission of macro shocks and, in particular, on the drivers of the business cycle. Our evidence underlines that the key source of business cycle variation in the data also accounts for short-term revisions in expected macroeconomic asymmetries as measured by our expected skewness factor.

In Section 2 we have shown that our skewness factor is correlated with alternative measures of macroeconomic skewness. It is therefore natural to ask whether revisions of these alternative measures also display a close connection with the *business cycle anatomy* and whether introducing a broader measure of skewness through our principal component approach is crucial to obtaining this result. As a first exercise, we replace the expected skewness factor with the individual expected skewness series of GDP growth shown in Figure 1(a). The results of this specification are shown in Figures D-2 and D-3 in Appendix D. In spite of the strong correlation between aggregate macro skewness and the skewness of GDP growth, revisions in the expected skewness of GDP growth do not generate meaningful comovement among the key macroeconomic variables. As a result, these revisions bear small resemblance with the *business cycle anatomy* of Angeletos et al. (2020). The correlation between revisions in expected GDP growth skewness and the MBC shock is negligible (Table 2).

Figure 2: Baseline model: Impulse response functions

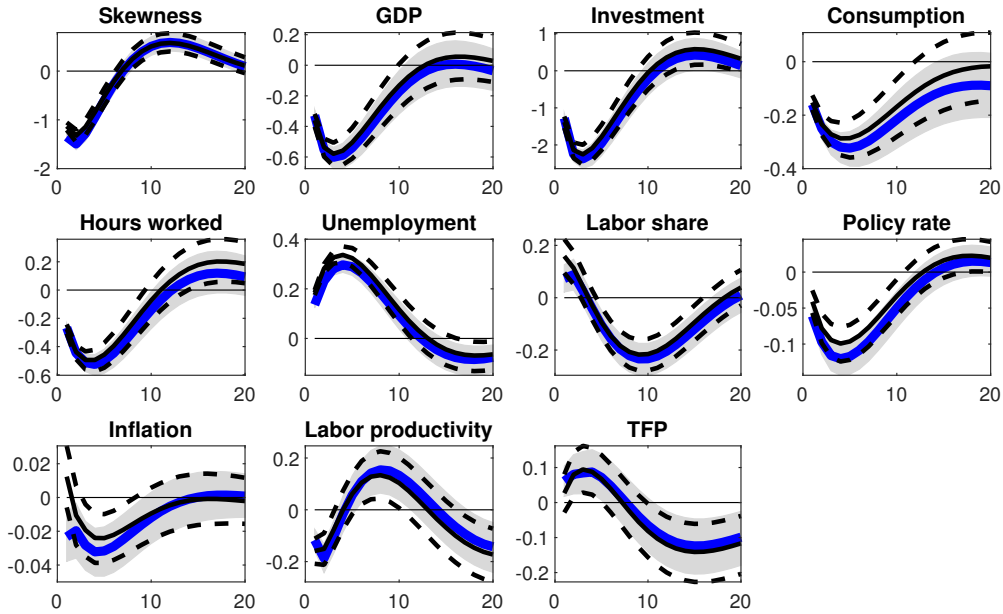

Note: The blue lines are the posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. The skewness shock is identified through a Cholesky decomposition. The black lines are the responses to a one S.D. shock to unemployment, i.e. the MBC shock of Angeletos et al. (2020). This shock is identified using the approach of Uhlig (2003). Sample period: 1960:Q1–2017:Q4.

Figure 3: Baseline model: Forecast error variance contributions

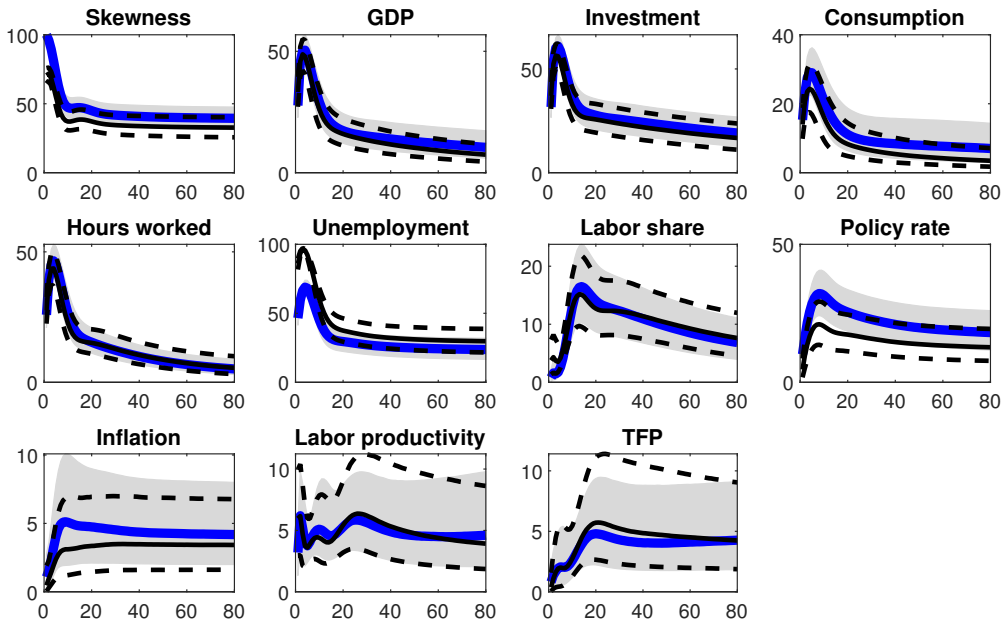

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval for a shock to expected skewness (blue) and the MBC (unemployment) shock (black).

In addition, we also compare our baseline results with the impact of revisions in expected GDP growth skewness, when computed based on the approach of Adrian et al. (2019) as shown in Figure 1(b). Figures D-4 and D-5 in Appendix D show the corresponding IRFs and variance contributions. Revisions in this measure of expected growth skewness, which can largely be interpreted as revisions associated with varying financial conditions, have a much more short-lived impact on macroeconomic asymmetry. Moreover, despite producing sizeable comovement among all the key macroeconomic quantities, the IRFs display less similarity with the *business cycle anatomy*. The correlation between revisions in this skewness measure and the MBC shock remains well-below the baseline result (Table 2). This is evidence that our expected skewness factor is a broader asymmetry measure and that this additional information is important when analysing the impact of changing risks.

We also investigate whether revisions in financial market skewness produce dynamics consistent with the ones reported above. To this end, we replace the skewness factor with the option-implied market skewness series of Dew-Becker (2021) as well as the cross-sectional stock return series of Salgado et al. (2019), both of which are shown in Figure 1(d). First, Table 2 shows that revisions to the S&P 500 skewness series are negatively correlated with the MBC shock. In fact, a downward revision in this skewness measure is associated with an expansionary response of the main business cycle indicators, and non-negligible positive inflation. The IRFs and variance contributions are shown in Figures D-6 and D-7. This result is in line with Dew-Becker (2021), who finds financial market skewness to move countercyclically. Second, when including the cross-sectional firm-level measure of stock return skewness, we only find a minor correlation between revisions in this series and the MBC shock (see Table 2, and Figures D-8 and D-9).[12]

To conclude this section, we also explore the impact of revisions in expected skewness beyond the baseline set of macroeconomic variables through augmented specifications, including selected financial variables (see Appendix E). We consider three augmented models that, in addition to the baseline variables, include either i) excess returns and the term premium (Figures E-1 and E-2); ii) real house prices and real stock prices (Figures E-3 and E-4); or iii) yields of 10-year government bonds (Figures E-5 and E-6). First, a downward revision of expected skewness decreases excess returns and increases the term premium, but the variance contributions are modest. Second, a downward revision of expected skewness decreases real stock prices, while the estimation uncertainty around the IRF of real house prices is large. Similarly, while the variance contribution for house price dynamics is relatively small, the skewness shock seems to explain a certain degree of stock price variation.

[12] We also tested VAR specifications with firm-level skewness using the option-implied skewness series of Dew-Becker (2021). Shocks to this series only share a very low correlation with the MBC shock (results not reported).

Table 2: Correlation of revisions in (exp.) skewness and MBC shock for different specifications

	Baseline model			**MBC shock**	
a)	Exp. skewness factor	**Skew. shock**	Median	**0.84**	
		(1960:Q1–2017:Q4)	95% HDI	0.78	0.89
	Other skewness measures			**MBC shock**	
b)	Exp. GDP skewness	**Skew. shock**	Median	**0.04**	
		(1960:Q1–2017:Q4)	95% HDI	-0.03	0.11
c)	Exp. GDP skewness (ABG)	**Skew. shock**	Median	**0.57**	
		(1971:Q1–2017:Q4)	95% HDI	0.48	0.66
d)	S&P 500 skewness	**Skew. shock**	Median	**-0.32**	
		(1983:Q2–2017:Q4)	95% HDI	-0.45	-0.19
e)	Firm-level stock return skewness	**Skew. shock**	Median	**0.14**	
		(1964:Q1–2015:Q1)	95% HDI	0.02	0.26
	Robustness checks			**MBC shock**	
f)	Orthog. to GARCH volatility	**Skew. shock**	Median	**0.67**	
		(1960:Q1–2017:Q4)	95% HDI	0.56	0.76
g)	Orthog. to macro and financial unc.	**Skew. shock**	Median	**0.63**	
		(1960:Q1–2017:Q4)	95% HDI	0.56	0.76
h)	Orthog. to geopolitical risk	**Skew. shock**	Median	**0.84**	
		(1960:Q1–2017:Q4)	95% HDI	0.78	0.89
i)	Orthog. to excess bond premium	**Skew. shock**	Median	**0.73**	
		(1973:Q1–2017:Q4)	95% HDI	0.64	0.81
j)	Orthog. to total factor productivity	**Skew. shock**	Median	**0.84**	
		(1960:Q1–2017:Q4)	95% HDI	0.78	0.90
k)	Orthog. to fiscal policy	**Skew. shock**	Median	**0.83**	
		(1960:Q1–2015:Q4)	95% HDI	0.77	0.89
l)	Orthog. to monetary policy	**Skew. shock**	Median	**0.84**	
		(1990:Q1–2016:Q4)	95% HDI	0.77	0.90

Note: Each row corresponds to a VAR specification and shows the correlation between downward revisions in (expected) skewness and the (contractionary) MBC shock (Angeletos et al., 2020). We report the median correlation across MCMC draws along with the 95% highest density interval (HDI). Revisions in (expected) skewness are identified through a Cholesky decomposition by ordering skewness first if no alternative shock/variable is included and second/third otherwise. Specification a) is our baseline model whereas in b), c), d) and e) the skewness factor is replaced with the exp. skewness of GDP growth, the exp. skewness of GDP growth based on the approach of Adrian et al. (2019), the option-implied skewness of the S&P 500 (quarterly avg.) computed by Dew-Becker (2021), and the cross-sectional firm-level skewness of stock returns (quarterly avg.) computed by Salgado et al. (2019), respectively. The alternative variables/shocks are: f) a data-rich measure of expected volatility based on a GARCH(1,1); g) the macroeconomic and financial uncertainty indices of Jurado et al. (2015) and Ludvigson et al. (2021); h) the (historical) geopolitical risk index (quarterly avg.) of Caldara and Iacoviello (2022); i) the excess bond premium (EBP) (Gilchrist and Zakrajšek, 2012); j) the (annualized) growth rate of the utilization-adjusted TFP measure of Fernald (2014); k) the government spending shock of Ramey and Zubairy (2018); and l) the monetary policy surprises of Jarociński and Karadi (2020).

Finally, a downward revision of expected skewness decreases the 10-year government bond yield and explains a sizeable share of its variation. Overall, a revision in expected macroeconomic skewness appears to matter somewhat more for macroeconomic than financial variables, potentially suggesting a certain disconnect between the drivers of the business cycle and those of asset prices, in line with the original evidence in Angeletos et al. (2020).

4 Robustness checks

This section discusses various checks to test the robustness of our baseline results along different dimensions. The detailed results can be found in Appendix F. First, as mentioned before, the baseline results are robust to a change in the identification scheme. In particular, to be closer to Angeletos et al. (2020), we also identify skewness shocks using the Uhlig (2003) approach which maximizes the explained share of skewness variation over four quarters in the time domain. The corresponding results are shown in Figures F-1 and F-2, and are very similar to the results based on the Cholesky identification.

Second, we augment our baseline specification with measures of macroeconomic volatility, uncertainty and geopolitical risk. Figures F-3 and F-4 present the effects of a revision in expected skewness when controlling for aggregate expected volatility, achieved by ordering this measure first in the Cholesky identification. The volatility measure is also based on a data-rich approach to match the derivation of the skewness factor. Specifically, we estimate a GARCH(1,1) model on each (de-meaned) data series of the McCracken and Ng (2020) dataset and obtain the first principal component of all standardized expected volatility (conditional standard deviation) series. We see that both the IRFs and the variance contributions in case of a skewness shock remain very similar compared to the baseline model.

In a related exercise, we control for macro and financial uncertainty (Jurado et al., 2015; Ludvigson et al., 2021). While the IRFs (Figure F-5) and variance contributions (Figure F-6) change somewhat more in this case, they still remain similar to the baseline results. The positive comovement between output and uncertainty after a downward revision in expected skewness implies that the transmission of skewness revisions is clearly distinct from the transmission of an uncertainty shock, which is generally characterized by a negative comovement between output and uncertainty. Table 2 shows that the correlation between the skewness shock and the MBC shock remains sizeable. These results are largely consistent with those in Forni et al. (2021), who show that the transmission of downside uncertainty and skewness shocks is distinct from that of a standard (symmetric) uncertainty shock. They find that asymmetry matters in the sense that uncertainty per se is not harmful, but that a

widening of the left tail is what causes economic contractions.[13] Moreover, to test whether revisions in expected skewness relate to geopolitical risk, we augment our baseline specification with the Geopolitical Risk Index of Caldara and Iacoviello (2022). Here, we find that the IRFs (Figure F-7) and variance contributions (Figure F-8), as well as the correlation with the MBC shock, remain nearly unchanged.

Third, we show that revisions in expected skewness are unrelated to other standard shocks. We show robustness when controlling for: i) (credit-)risk shocks measured as the exogenous variation in the Gilchrist and Zakrajšek (2012) excess bond premium (Figures F-9 and F-10); ii) productivity shocks measured as the exogenous variation in the growth rate of the Fernald (2014) TFP series (Figures F-11 and F-12)[14]; iii) shocks to government expenditure as identified in Ramey and Zubairy (2018) (Figures F-13 and F-14); and iv) monetary policy shocks measured by the surprise series of Jarociński and Karadi (2020), which is purged of the central bank information component (Figures F-15 and F-16). In all cases the IRFs and FVDs are similar to the baseline model and range from being nearly identical (TFP and fiscal policy) to featuring some differences (EBP and monetary policy). The skewness shock continues to be highly correlated with the MBC shock across specifications (Table 2).

Finally, we test the robustness with respect to the lag order in the VAR model and changes to the Minnesota prior. Figures F-17 and F-18 present the results using a lag order of $P = 4$. The IRFs and the error variance contributions remain very similar compared to the baseline model. Figures F-19 and F-20 show that applying an even looser configuration of the Minnesota prior ($\lambda = 10$) leaves the baseline results essentially unchanged.

5 Conclusion and direction for future research

We construct a factor that summarizes expected macroeconomic skewness. This factor is the first principal component of the time-varying expected skewness indicators of a large number of macroeconomic series. Aggregate macroeconomic skewness is strongly procyclical, comoves with the expected GDP growth skewness series based on the approach of Adrian et al. (2019), which conditions on macro-financial conditions, and is highly correlated with the cross-sectional skewness of firm-level employment growth (Salgado et al., 2019). We then document that the impulse responses of a set of macroeconomic variables associated with a revision in expected skewness, and the corresponding variance contributions, closely match the *business cycle anatomy* of Angeletos et al. (2020). Our baseline model produces an un-

[13] See also Segal et al. (2015), who distinguish between "good" and "bad" uncertainty, depending on its impact on macroeconomic growth. In related work, Castelnuovo and Mori (2022) show that skewness can respond endogenously to uncertainty shocks and amplify their impact on the business cycle.

[14] When including the growth rate of TFP, we exclude the TFP level series from the VAR specification.

conditional correlation of around 0.8 between revisions in expected skewness and the *main business cycle* shock identified in Angeletos et al. (2020). The results are robust to changes in the identification scheme, controlling for macroeconomic volatility, uncertainty, and frequently considered alternative shocks.

Our results highlight the importance of accounting for a procyclical variation in conditional skewness of macroeconomic data. Variation in conditional skewness requires the presence of non-linearities in the transmission of Gaussian shocks (see, e.g., Fernández-Villaverde and Guerrón-Quintana, 2020), or can directly derive from skewed shocks hitting the economy (as in Bekaert and Engstrom, 2017; Salgado et al., 2019). Angeletos and La'O (2013) and Angeletos et al. (2018) highlight how waves of optimism and pessimism regarding both firms' expected employment and production decisions as well as consumers' beliefs about future employment opportunities and income generate dynamics of output, employment, spending and prices akin to the business cycle patterns observed in the data. The former could potentially arise from learning asymmetries in the presence of informational frictions as in Veldkamp (2005). To the extent that fluctuations in *sentiment* or *confidence* are associated with a reassessment of upside and downside risk over the cycle, and hence shifts in expected skewness, our results provide a way of addressing the problem that "a direct, empirical counterpart to the confidence shock is hard, if possible at all, to obtain" (Angeletos et al., 2018, p. 1692). Our results are also consistent with a relevant role for expectations of rare disasters in explaining economic fluctuations (Rietz, 1988; Barro, 2006, 2009; Gabaix, 2008; Gourio, 2012; Wachter, 2013; Petrosky-Nadeau et al., 2018; Jordà et al., 2020). In particular, our results highlight the importance of allowing for time variation in the severity (Gabaix, 2008) and/or probability of such rare disasters (see, e.g., Gourio, 2012; Wachter, 2013), which could generate sizeable variation in expected skewness.[15] Most importantly, our results provide useful insights for macroeconomic theories that search for shocks and propagation mechanisms behind macroeconomic fluctuations. Any such theory will need to be able to reproduce variations in aggregate skewness whose revisions are strongly affected by the main source of business cycle fluctuations.

[15]See Giglio et al. (2021) for supporting evidence of the presence of time-varying disaster probabilities.

References

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable Growth. American Economic Review, 109(4):1263–89.

Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. Journal of Financial Economics, 110(1):110–138.

Angeletos, G.-M., Collard, F., and Dellas, H. (2018). Quantifying confidence. Econometrica, 86(5):1689–1726.

Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business-cycle anatomy. American Economic Review, 110(10):3030–70.

Angeletos, G.-M. and La'O, J. (2013). Sentiments. Econometrica, 81(2):739–779.

Antolin-Diaz, J., Drechsel, T., and Petrella, I. (2017). Tracking the Slowdown in Long-Run GDP Growth. Review of Economics and Statistics, 99(2):343–356.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector auto regressions. Journal of Applied Econometrics, 25(1):71–92.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. The Quarterly Journal of Economics, 121(3):823–866.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. American Economic Review, 99(1):243–64.

Barro, R. J. and Ursúa, J. F. (2008). Macroeconomic crises since 1870. Brookings Papers on Economic Activity, 2008(1):255–350.

Bekaert, G. and Engstrom, E. (2017). Asset Return Dynamics under Habits and Bad Environment-Good Environment Fundamentals. Journal of Political Economy, 125(3):713–760.

Busch, C., Domeij, D., Guvenen, F., and Madera, R. (2018). Asymmetric business-cycle risk and social insurance. NBER Working Papers, No. 24569, National Bureau of Economic Research.

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. American Economic Review, 112(4):1194–1225.

Castelnuovo, E. and Mori, L. (2022). Uncertainty, Skewness, and the Business Cycle Through the MIDAS Lens. Mimeo.

Chan, J. C. (2020). Large Bayesian vector autoregressions. In Macroeconomic Forecasting in the Era of Big Data, pages 95–125. Springer.

Chen, L., Dolado, J. J., and Gonzalo, J. (2021). Quantile factor models. Econometrica, 89(2):875–910.

Colacito, R., Ghysels, E., Meng, J., and Siwasarit, W. (2016). Skewness in expected macro fundamentals and the predictability of equity returns: Evidence and theory. The Review of Financial Studies, 29(8):2069–2109.

Cox, D. R. (1981). Statistical analysis of time series: Some recent developments [with discussion and reply]. Scandinavian Journal of Statistics, 8(2):93–115.

Delle Monache, D., De Polis, A., and Petrella, I. (2021). Modeling and forecasting macroeconomic downside risk. Bank of Italy Temi di Discussione (Working Paper) No, 1324.

Dew-Becker, I. (2021). Real-time forward-looking skewness over the business cycle. Mimeo.

Dew-Becker, I., Tahbaz-Salehi, A., and Vedolin, A. (2019). Macro skewness and conditional second moments: evidence and theories. Mimeo.

Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. Econometric Reviews, 3(1):1–100.

Engle, R. F. and Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. Journal of Business & Economic Statistics, 22(4):367–381.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Working Paper.

Fernández-Villaverde, J. and Guerrón-Quintana, P. A. (2020). Uncertainty shocks and business cycle research. Review of Economic Dynamics, 37:118–146.

Fernández-Villaverde, J. and Levintal, O. (2018). Solution methods for models with rare disasters. Quantitative Economics, 9(2):903–944.

Fève, P., Sanchez, P. G., Moura, A., and Pierrard, O. (2021). Costly default and skewed business cycles. European Economic Review, 132:103630.

Forni, M., Gambetti, L., and Sala, L. (2021). Downside and upside uncertainty shocks. CEPR Discussion Paper, No. DP15881, Centre for Economic Policy Research.

Gabaix, X. (2008). Variable rare disasters: A tractable theory of ten puzzles in macro-finance. American Economic Review, 98(2):64–67.

Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2021). The joint dynamics of investor beliefs and trading during the COVID-19 crash. Proceedings of the National Academy of Sciences, 118(4).

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. American Economic Review, 102(4):1692–1720.

Gorodnichenko, Y. and Ng, S. (2017). Level and volatility factors in macroeconomic data. Journal of Monetary Economics, 91:52–68.

Gourio, F. (2012). Disaster risk and business cycles. American Economic Review, 102(6):2734–66.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2):357–384.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises - the role of information shocks. American Economic Journal: Macroeconomics, 12(2):1–43.

Jensen, H., Petrella, I., Ravn, S. H., and Santoro, E. (2020). Leverage and Deepening Business-Cycle Skewness. American Economic Journal: Macroeconomics, 12(1):245–81.

Jordà, Ò., Schularick, M., and Taylor, A. M. (2020). Disasters everywhere: The costs of business cycles reconsidered. NBER Working Papers, No. 26962, National Bureau of Economic Research.

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

Kadiyala, K. R. and Karlsson, S. (1997). Numerical methods for estimation and inference in Bayesian VAR-models. Journal of Applied Econometrics, 12(2):99–132.

Kelley, T. L. (1947). Fundamentals of Statistics. Harvard University Press.

Koenker, R. and Bassett, G. (1978). Regression Quantiles. Econometrica, 46(1):33–50.

Koop, G. and Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. Foundations and Trends in Econometrics, 3(4):267–358.

Litterman, R. B. (1986). Forecasting with Bayesian Vector Autoregressions - Five Years of Experience. Journal of Business & Economic Statistics, 4(1):25–38.

Loria, F., Matthes, C., and Zhang, D. (2020). Assessing macroeconomic tail risk. Mimeo.

Ludvigson, S., Ma, S., and Ng, S. (2021). Uncertainty and business cycles: Exogenous impulse or endogenous response? American Economic Journal: Macroeconomics, 13(4):369–410.

McCracken, M. W. and Ng, S. (2016). FRED-MD: A Monthly Database for Macroeconomic Research. Journal of Business & Economic Statistics, 34(4):574–589.

McCracken, M. W. and Ng, S. (2020). FRED-QD: A Quarterly Database for Macroeconomic Research. Working Paper Series, No. 2020-005B, Federal Reserve Bank of St. Louis.

Morley, J. and Piger, J. (2012). The Asymmetric Business Cycle. Review of Economics and Statistics, 94(1):208–221.

Mumtaz, H. and Theodoridis, K. (2020). Dynamic effects of monetary policy shocks on macroeconomic volatility. Journal of Monetary Economics, 114:262–282.

Mumtaz, H. and Zanetti, F. (2012). Neutral technology shocks and the dynamics of labor input: Results from an agnostic identification. International Economic Review, 53(1):235–254.

Neftci, S. N. (1984). Are economic time series asymmetric over the business cycle? Journal of Political Economy, 92(2):307–328.

Omori, Y., Chib, S., Shephard, N., and Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. Journal of Econometrics, 140(2):425–449.

Ordonez, G. (2013). The Asymmetric Effects of Financial Frictions. Journal of Political Economy, 121(5):844–895.

Orlik, A. and Veldkamp, L. (2014). Understanding uncertainty shocks and the role of black swans. NBER Working Papers, No. 20445, National Bureau of Economic Research.

Petrosky-Nadeau, N., Zhang, L., and Kuehn, L.-A. (2018). Endogenous disasters. American Economic Review, 108(8):2212–45.

Plagborg-Møller, M., Reichlin, L., Ricco, G., and Hasenzagl, T. (2020). When is growth at risk? Brookings Papers on Economic Activity, 2020(1):167–229.

Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from US historical data. Journal of Political Economy, 126(2):850–901.

Rietz, T. A. (1988). The equity risk premium a solution. Journal of Monetary Economics, 22(1):117–131.

Salgado, S., Guvenen, F., and Bloom, N. (2019). Skewed business cycles. NBER Working Papers, No. 26565, National Bureau of Economic Research.

Segal, G., Shaliastovich, I., and Yaron, A. (2015). Good and bad uncertainty: Macroeconomic and financial market implications. Journal of Financial Economics, 117(2):369–397.

Sichel, D. E. (1993). Business cycle asymmetry: a deeper look. Economic Inquiry, 31(2):224–236.

Stock, J. H. and Watson, M. W. (2002). Forecasting Using Principal Components From a Large Number of Predictors. Journal of the American Statistical Association, 97:1167–1179.

Stock, J. H. and Watson, M. W. (2012). Generalized shrinkage methods for forecasting using many predictors. Journal of Business & Economic Statistics, 30(4):481–493.

Uhlig, H. (2003). What moves real GNP? Mimeo.

Veldkamp, L. L. (2005). Slow boom, sudden crash. Journal of Economic Theory, 124(2):230–257.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? The Journal of Finance, 68(3):987–1035.

Appendix A Monte Carlo exercise

This section addresses the concern that our two-step approach to constructing an aggregate skewness factor could yield spurious results, i.e. indicate time-varying conditional skewness in cases, where in fact there is none. For this, we conduct a Monte Carlo exercise and generate 500 datasets of size $N = 70$ and $T = 250$ from two different data generating processes (DGP), both of which do not feature conditional skewness. The first DGP has a time-varying mean and volatility, which both have a factor structure. Specifically, DGP 1 is defined as

$$y_{i,t} = \mu_{i,t} + e^{h_{i,t}/2}\varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim \mathcal{N}(0,1), \tag{A-1}$$

$$\mu_{i,t} = \lambda_i^f f_t + \omega_{i,t}, \tag{A-2}$$

$$h_{i,t} = \lambda_i^h \bar{h}_t + \nu_{i,t}, \tag{A-3}$$

$$f_t = \rho^f f_{t-1} + z_t, \qquad z_t \sim \mathcal{N}(0,\sigma_z^2), \tag{A-4}$$

$$\bar{h}_t = \rho^h \bar{h}_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0,\sigma_u^2), \tag{A-5}$$

$$\omega_{i,t} = \rho_i^\omega \omega_{i,t-1} + \epsilon_{i,t}, \qquad \epsilon_{i,t} \sim \mathcal{N}(0,\sigma_{\epsilon,i}^2), \tag{A-6}$$

$$\nu_{i,t} = \rho_i^\nu \nu_{i,t-1} + \kappa_{i,t}, \qquad \kappa_{i,t} \sim \mathcal{N}(0,\sigma_{\kappa,i}^2). \tag{A-7}$$

The parameters of the DGP are set to: $\rho^f = 0.9$, $\rho^h = 0.98$, $\rho_i^\omega = 0.9$, $\rho_i^\nu = 0.98$, $\sigma_z^2 = 1$, $\sigma_u^2 = 0.1$, $\sigma_{\epsilon,i}^2 = 1$, and $\sigma_{\kappa,i}^2 = 0.1 \; \forall \, i = 1, ..., N$. The factor loadings in the mean and log-volatility equation, λ_i^f and λ_i^h, are drawn from independent normal distributions with the moments chosen such that the average variation explained of the mean and log-volatility of the variables is 20% and 25%, respectively.

DGP 2 is similar to DGP 1 but includes the so-called leverage effect, i.e a negative contemporaneous correlation between the innovations to the mean and volatility factors, as well as the innovations to the idiosyncratic mean and volatility components. Under this assumption, it is well-known that the model remains conditionally Gaussian, but features unconditional left-skewness (e.g. Omori et al., 2007). In particular, in this case, we assume that z_t and u_t follow a multivariate normal distribution with correlation $\rho_{z_t,u_t} = -0.9$. A similar assumption is introduced for the correlation between $\epsilon_{i,t}$ and $\kappa_{i,t}$ is chosen randomly from a uniform distribution $[0, -0.9]$ for each variable $i = 1, ..., N$.

For each simulated dataset and both DGPs, we estimate the skewness factor as outlined in Section 2. Since the scale of the skewness factor is not identified, and since in this case we are interested in assessing how far the retrieved factor is from the zero line, we normalize the factor so that its standard deviation matches the mean value of the standard deviation of the individual skewness series in the dataset. Figure A-1 presents the distribution of the estimated skewness factors across the Monte Carlo samples. The results provide evidence

for the strong performance of the model and show that our two-step approach to construct the skewness factor does not capture "spurious skewness". In particular, since both DGPs do not feature conditional skewness, the distribution of the estimated factors across Monte Carlo samples is centred around the zero line with only limited dispersion.

Figure A-1: Results of Monte Carlo simulation

(a) DGP 1: (Un-)conditionally Gaussian

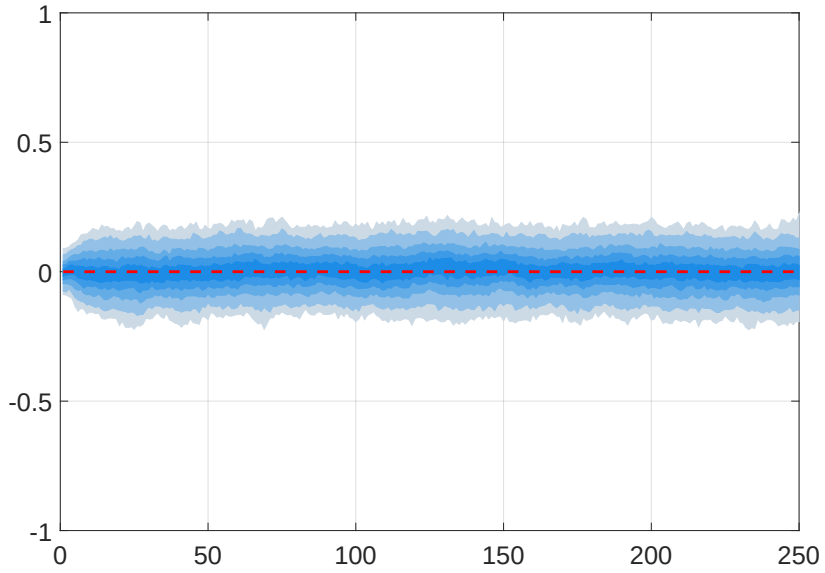

(b) DGP 2: Cond. Gaussian with uncond. skew

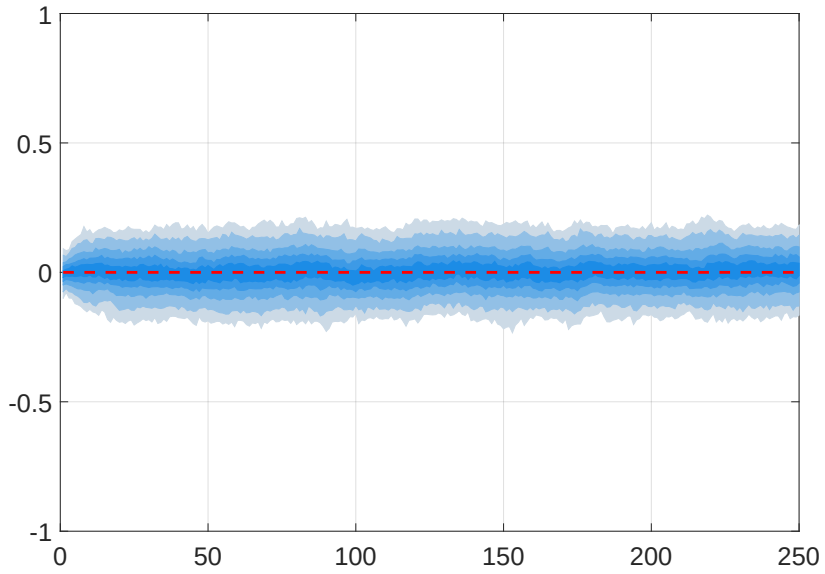

Note: The largest shaded area corresponds to the 90% confidence interval, with shades corresponding to increasing probability ranges of 10%, 20%, ..., 90%.

Appendix B Data

Table B-1: Data descriptions, transformations and sources

Dataset used to derive the skewness factor			
McCracken and Ng (2020) dataset: https://research.stlouisfed.org/econ/mccracken/fred-databases/ . Vintage: June 2021.			
Variables in baseline VAR			
Name	**Description**	**Transformation**	**Source**
GDP	Real GDP per capita	100*log(X)	FRED
Investment	Real investment per capita	100*log(X)	FRED
Consumption	Real consumption per capita	100*log(X)	FRED
Hours worked	Hours worked per person	100*log(X)	FRED
Unemployment	Civilian unemployment rate	–	FRED
Labor share	Labor share in the non-farm business sector	100*log(X)	FRED
Policy rate	Effective federal funds rate	$X/4$	FRED
Inflation	Percentage change in GDP deflator	100*log(X_t/X_{t-1})	FRED
Labor productivity	Real (non-farm) output per hours of all persons	100*log(X)	FRED
TFP	Level of total factor productivity	100*log(X)	FRED/Fernald (2014)
Variables in augmented VARs			
Name	**Description**	**Transformation**	**Source**
Excess returns	Change of S&P 500 minus 3-month treasury yield	–	FRED
Term premium	Treasury term premium for 10-year gov. bonds	–	FRBSF/Adrian et al. (2013)
House prices	Real house prices (nominal HPI divided by CPI)	100*log(X)	FRED
S&P 500	Real stock prices (S&P 500 index divided by CPI)	100*log(X)	FRED
Government bond yield	10-year government bond yield	–	FRED

Note: Data for the variables in the baseline VAR were directly taken from the replication files provided by Angeletos et al. (2020) and more details can be found in that reference.

Appendix C VAR model and prior choice

This appendix contains additional details on the VAR model used in the main part of the paper and the prior specification employed in the Bayesian estimation of this model. Since the presentation here is relatively brief and does not outline every step of the Bayesian treatment of a VAR model, we refer to standard references for further details (e.g. Koop and Korobilis, 2010; Chan, 2020). The starting point of our empirical analysis is a vector autoregressive model of order P denoted as VAR(P)

$$y_t = \sum_{p=1}^{P} \Theta_p y_{t-p} + u_t, \quad u_t \sim \mathcal{N}(\mathbf{0}, \Sigma), \tag{C-1}$$

where u_t is a $N \times 1$ vector of reduced-form errors that is normally distributed with zero mean and covariance matrix Σ. The regression-equation representation of this system is

$$Y = X\Theta + U, \tag{C-2}$$

where $Y = [y_{h+1}, ..., y_T]$ is a $N \times T$ matrix, $X = Y_{-h}$ is a $(NP) \times T$ matrix containing the h-th lag of Y, $\Theta = [\Theta_1, ..., \Theta_P]$ is a $N \times (NP)$ matrix, and $U = [u_{h+1}, ..., u_T]$ is a $N \times T$ matrix of disturbances.

The Bayesian estimation of VAR models has become standard in empirical macroeconomics. Specifically, we use a Minnesota-type prior (Doan et al., 1984; Litterman, 1986). It is assumed that the prior distribution of the VAR parameters has a Normal-Wishart conjugate form

$$\theta|\Sigma \sim \mathcal{N}(\theta_0, \Sigma \otimes \Omega_0), \quad \Sigma \sim \mathcal{IW}(v_0, S_0), \tag{C-3}$$

where θ is obtained by stacking the columns of Θ. In contrast to Litterman (1986), the covariance matrix Σ in the prior described in Equation (C-3) is not replaced by an estimated and thus known (diagonal) counterpart. Therefore, sampling from the conditional posterior distributions described below requires Gibbs sampling (see also Mumtaz and Zanetti, 2012). Our results are based on 25,000 draws and we discard the initial 5,000 draws as burn-in. The (Minnesota) prior moments of θ are given by

$$\mathbb{E}[(\Theta_p), i, j] = \begin{cases} \delta_i & i = j, p = 1 \\ 0 & \text{otherwise} \end{cases}, \quad Var[(\Theta_p), i, j] = \lambda \sigma_i^2 / \sigma_j^2, \tag{C-4}$$

and, as outlined in Bańbura et al. (2010), they can be constructed using the following T_D

dummy observations

$$Y_D = \begin{pmatrix} \frac{diag(\delta_1\sigma_1,...,\delta_N\sigma_N)}{\lambda} \\ 0_{N\times(P-1)N} \\ ..... \\ diag(\sigma_1,...,\sigma_N) \\ ..... \\ 0_{1\times N} \end{pmatrix} \text{ and } X_D = \begin{pmatrix} \frac{J_P\otimes diag(\sigma_1,...,\sigma_N)}{\lambda} \\ 0_{N\times NP} \\ ..... \\ 0_{1\times NP} \end{pmatrix}, \tag{C-5}$$

where $J_P = diag(1, 2, ..., P)$ and $diag$ denotes the diagonal matrix. The prior moments in Equation (C-3) are functions of Y_D and X_D, $\Theta_0 = Y_D X_D'(X_D X_D')^{-1}$, $\Omega_0 = (X_D X_D')^{-1}$, $S_0 = (Y_D - \Theta_0 X_D)(Y_D - \Theta_0 X_D)'$ and $v_0 = T_D - NP$. Finally, the hyper-parameter λ controls the tightness of the prior and our baseline choice is $\lambda = 2$.

Since the normal-inverse Wishart prior is conjugate, the conditional posterior distribution of this model is also normal-inverse Wishart (Kadiyala and Karlsson, 1997)

$$\theta|\Sigma, Y \sim \mathcal{N}(\bar{\theta}, \Sigma \otimes \bar{\Omega}), \quad \Sigma|Y \sim \mathcal{IW}(\bar{v}, \bar{S}), \tag{C-6}$$

where variables with a bar denote the parameters of the posterior distribution. Defining $\hat{\Theta}$ and $\hat{U}$ as the OLS estimates from Equation (C-2), the parameters of the conditional posterior distribution can be computed as $\bar{\Theta} = (\Omega_0^{-1} S_0 + YX')(\Omega_0^{-1} + X'X)^{-1}$, $\bar{\Omega} = (\Omega_0^{-1} + X'X)^{-1}$, $\bar{v} = v_0 + T$, and $\bar{S} = \hat{\Theta} X X' \hat{\Theta}' + \Theta_0 \Omega_0^{-1} \Theta_0 + S_0 + \hat{U}\hat{U}' - \bar{\Theta}\bar{\Omega}^{-1}\bar{\Theta}'$. Lastly, as in Mumtaz and Zanetti (2012), the values of the persistence parameter δ_i and the error standard deviation σ_i of the AR(1) model are obtained from its OLS estimation.

Appendix D Additional results

Skewness factor based on a smaller subset of variables

Figure D-1: Skewness factors based on different datasets

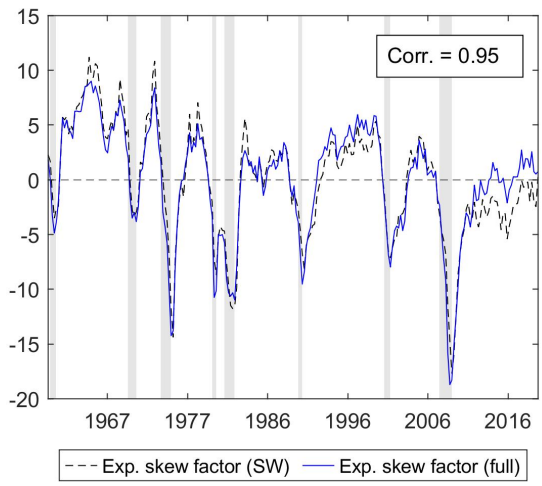

Note: This figure shows the skewness factor based on the full McCracken and Ng (2020) dataset and an alternative skewness factor based on a subset of variables similar to those used in Stock and Watson (2012). The scale of the latter is adjusted such that it is identical for both factors.

Correlation of skewness factor and volatility/uncertainty measures

Table D-1: Correlation of skewness factor and different volatility/uncertainty measures

	PC (skew)	PC (X)	PC (X^2)	PC (P_{75}-P_{25})	PC (GARCH)	Macro unc.	Fin. unc.
PC (skew)	1.00	-	-	-	-	-	-
PC (X)	0.79	1.00	-	-	-	-	-
PC (X^2)	-0.48	-0.50	1.00	-	-	-	-
PC (P_{75}-P_{25})	-0.80	-0.78	0.80	1.00	-	-	-
PC (GARCH)	-0.72	-0.53	0.80	0.92	1.00	-	-
Macro unc.	-0.72	-0.63	0.63	0.84	0.77	1.00	-
Fin. unc.	-0.46	-0.48	0.45	0.60	0.52	0.58	1.00

Note: This table contains correlations of the exp. skewness factor *PC (skew)* and different measures of volatility and uncertainty. *PC (X)*, *PC (X^2)*, *PC (P_{75}-P_{25})*, and *PC (GARCH)* are, respectively, the first principal component of the McCracken and Ng (2020) dataset, the first principal component of the squared observations (Gorodnichenko and Ng, 2017), the first principal component of the expected interquartile ranges, and the first principal component of the expected individual GARCH standard deviations. Macro unc. and Fin. unc. are the macroeconomic and financial uncertainty indices developed in Jurado et al. (2015) and Ludvigson et al. (2021).

Results of baseline model with GDP skewness

Figure D-2: Impulse response functions

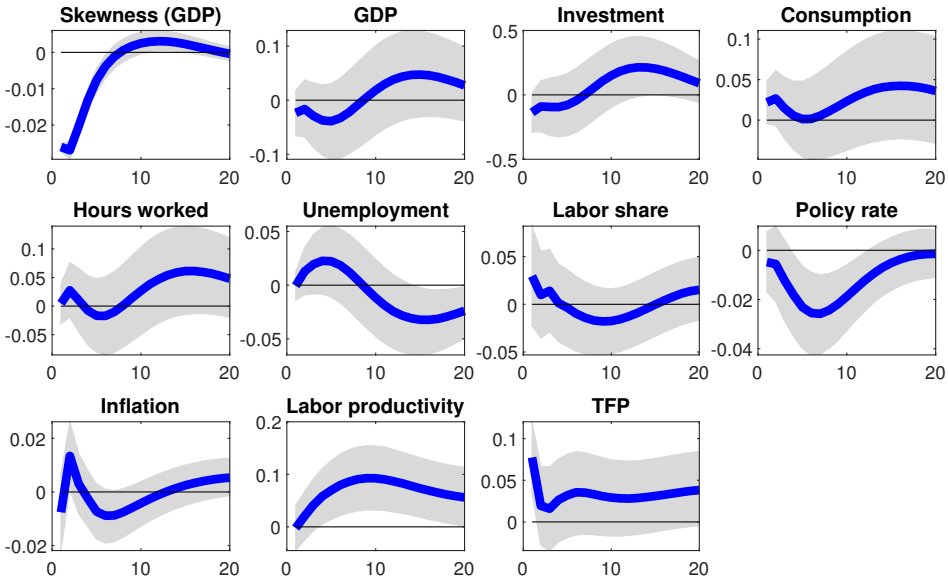

Note: Posterior mean responses to a negative one S.D. shock to expected GDP skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure D-3: Forecast error variance contributions

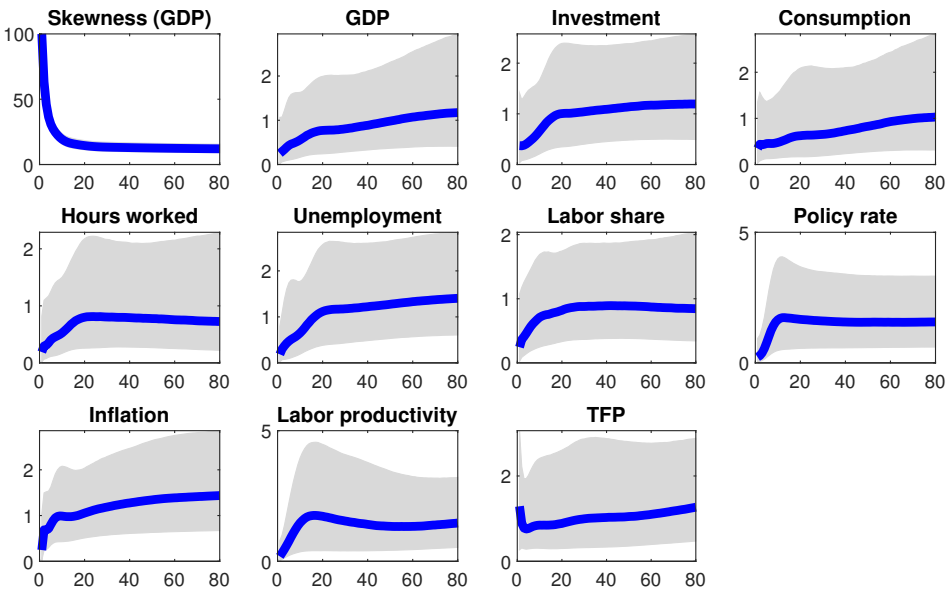

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of baseline model with GDP skewness (Adrian et al., 2019)

Figure D-4: Impulse response functions

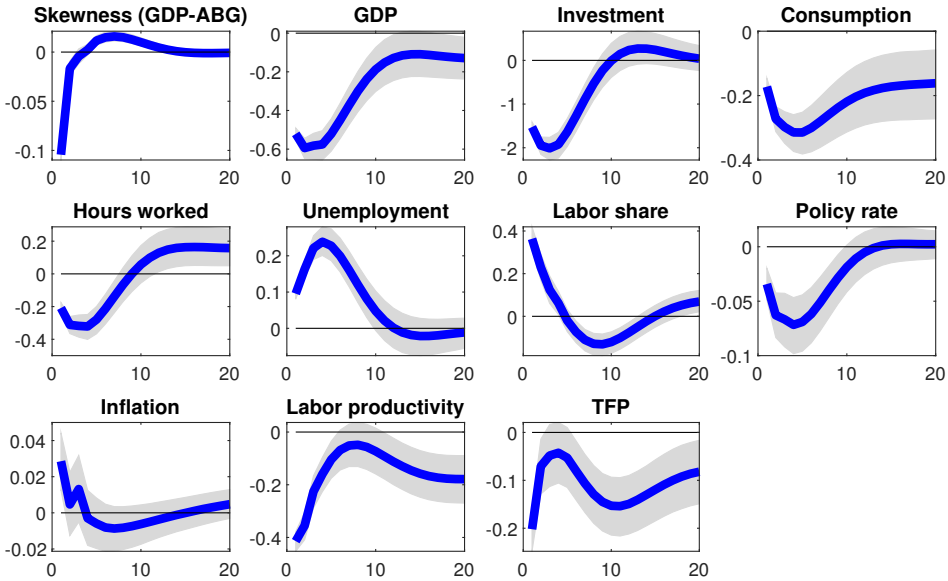

Note: Posterior mean responses to a negative one S.D. shock to expected GDP skewness (Adrian et al., 2019) along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1971:Q1–2017:Q4.

Figure D-5: Forecast error variance contributions

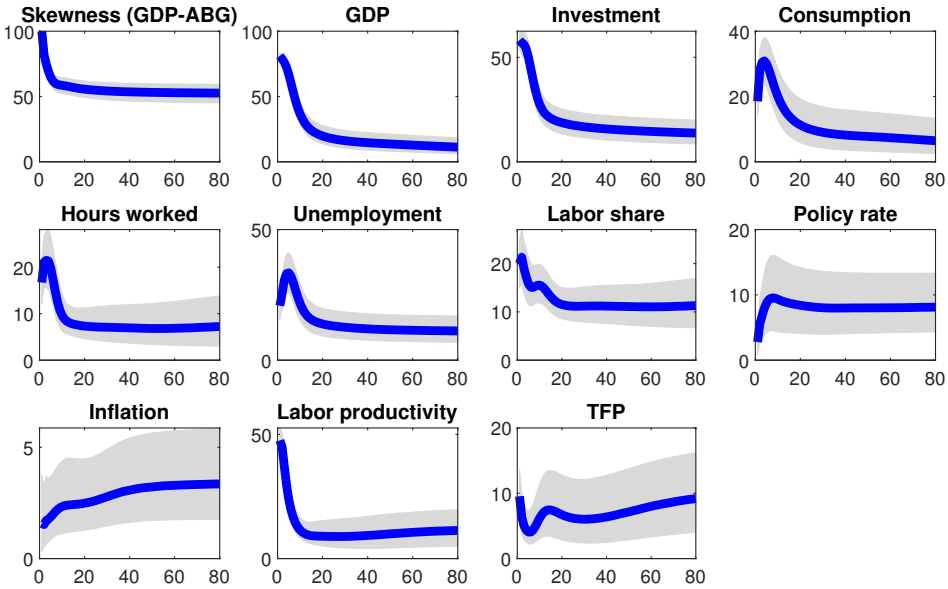

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of baseline model with S&P 500 skewness (Dew-Becker, 2021)

Figure D-6: Impulse response functions

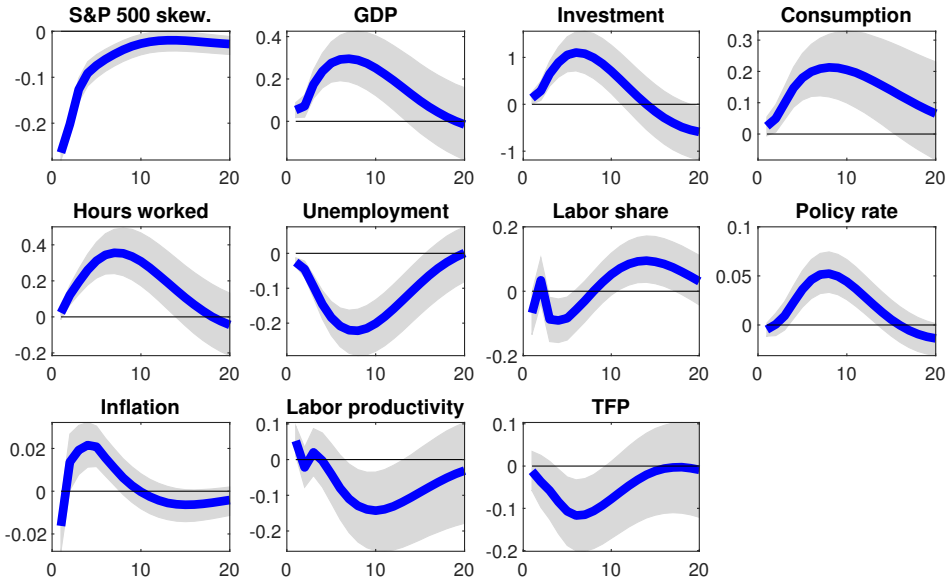

Note: Posterior mean responses to a negative one S.D. shock to option-implied S&P 500 skewness (Dew-Becker, 2021) along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1983:Q2–2017:Q4.

Figure D-7: Forecast error variance contributions

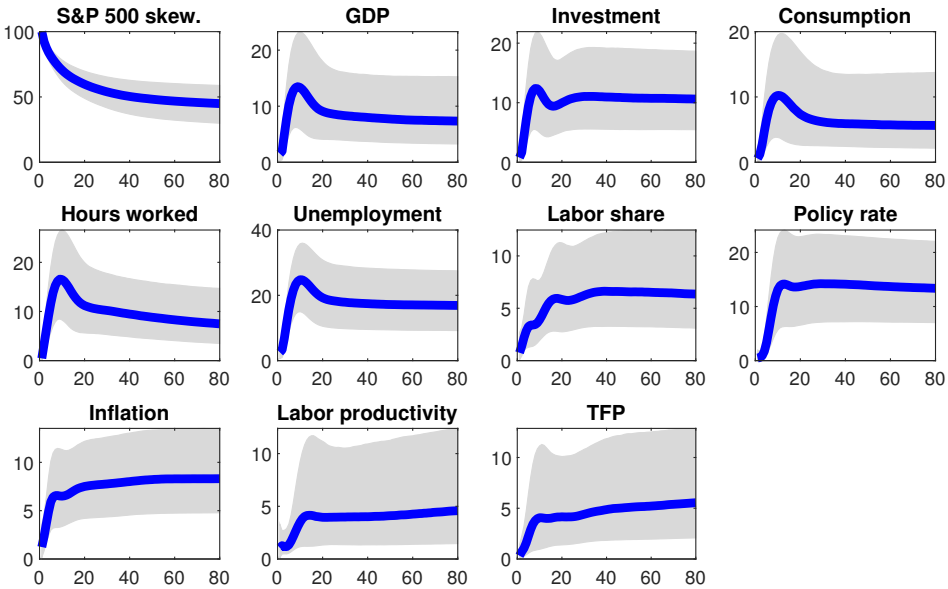

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of baseline model with firm-level return skewness (Salgado et al., 2019)

Figure D-8: Impulse response functions

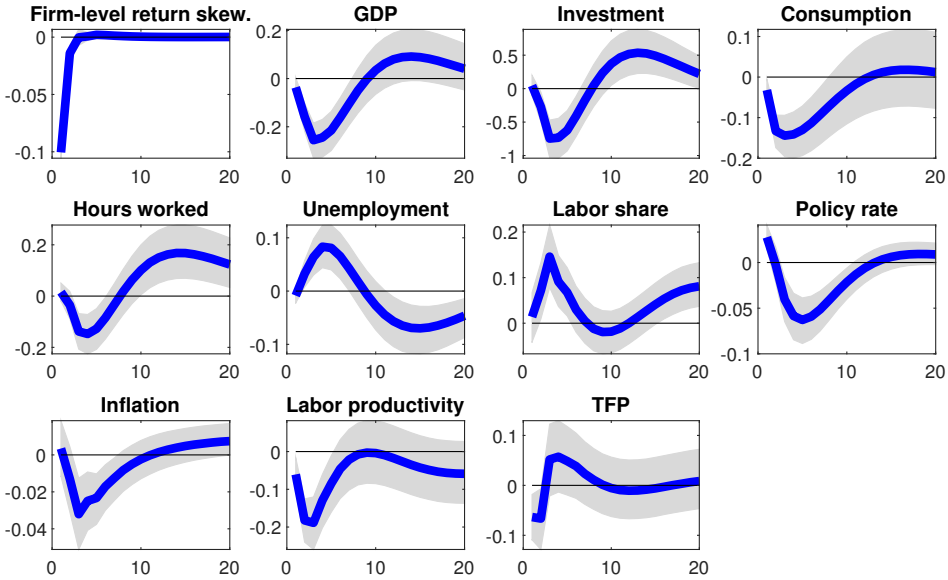

Note: Posterior mean responses to a negative one S.D. shock to the cross-sectional firm-level skewness of stock returns (Salgado et al., 2019) along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1964:Q1–2015:Q1.

Figure D-9: Forecast error variance contributions

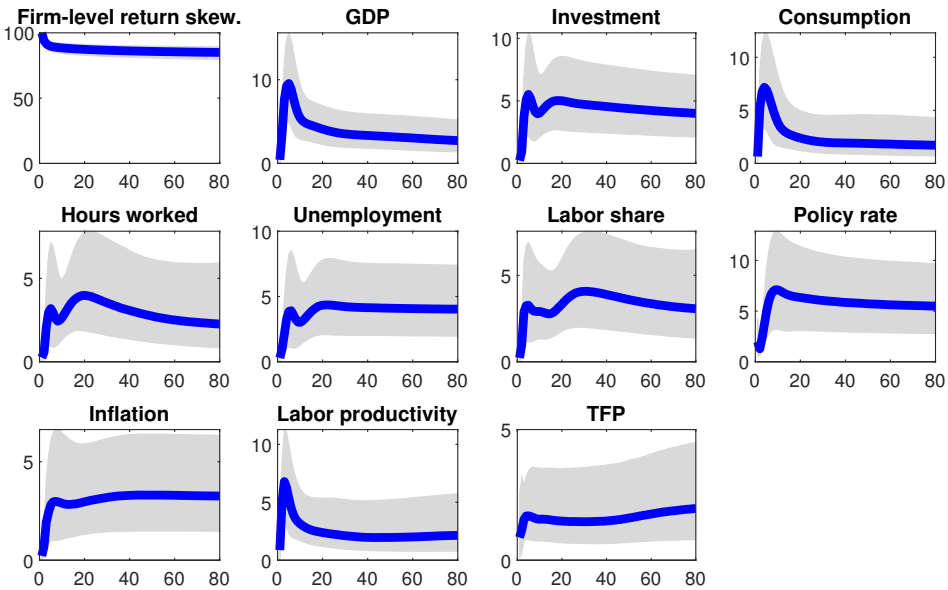

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Appendix E Augmented models including financial variables

Results of model augmented with excess returns and term premium

Figure E-1: Impulse response functions

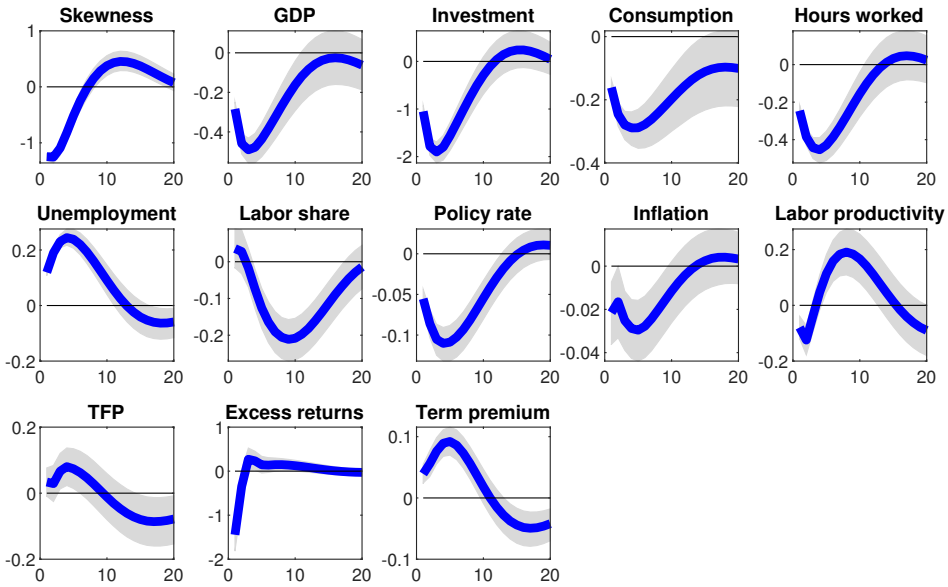

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1961:Q3–2017:Q4.

Figure E-2: Forecast error variance contributions

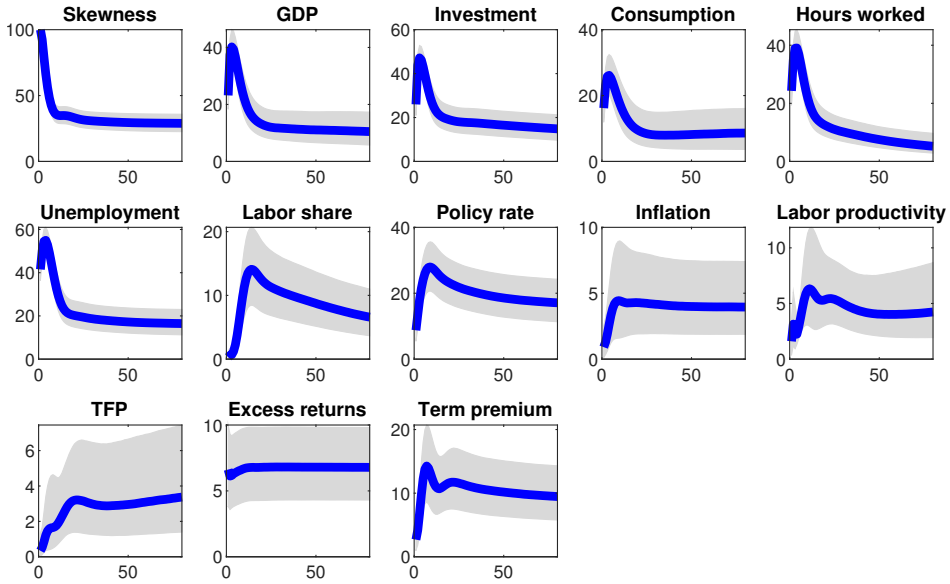

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of model augmented with house prices and stock prices

Figure E-3: Impulse response functions

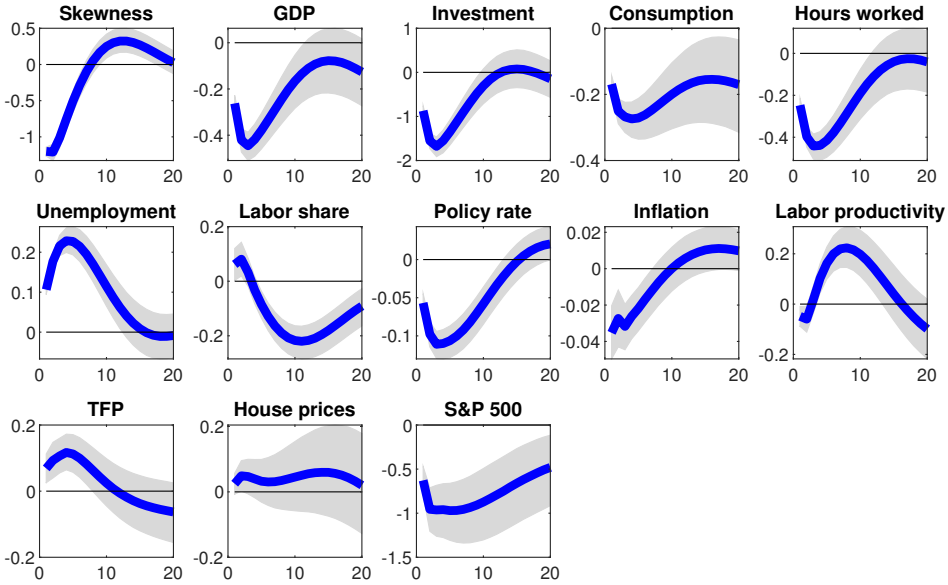

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1975:Q1–2017:Q4.

Figure E-4: Forecast error variance contributions

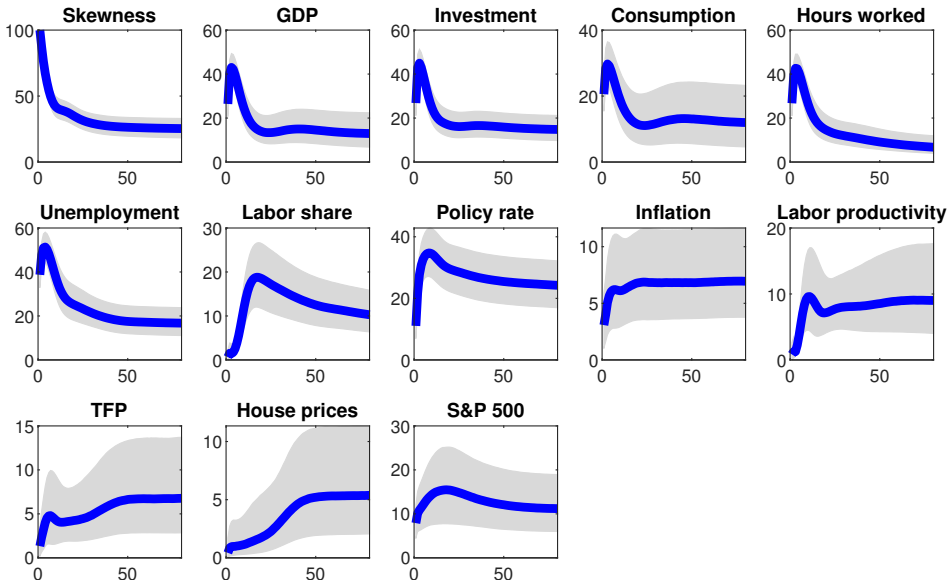

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of model augmented with government bond yields

Figure E-5: Impulse response functions

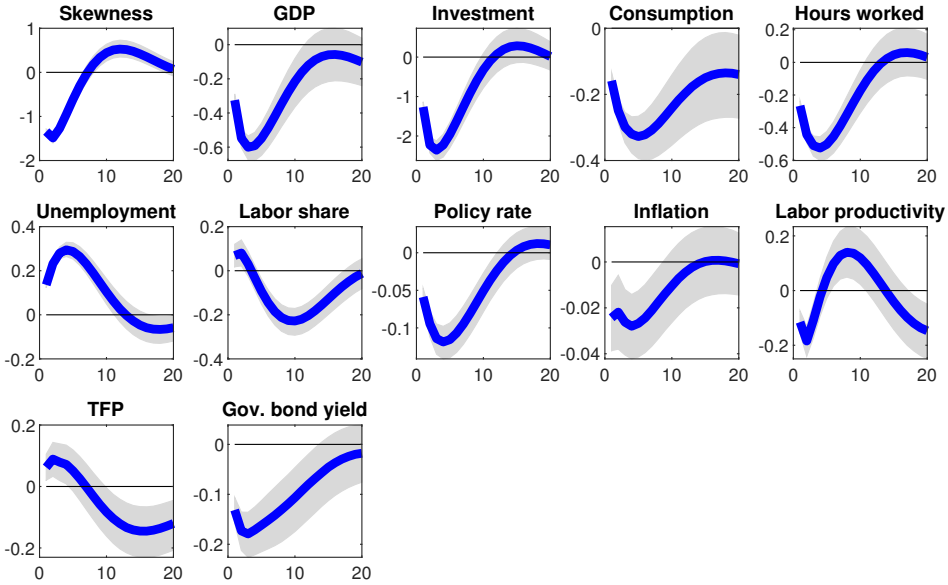

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure E-6: Forecast error variance contributions

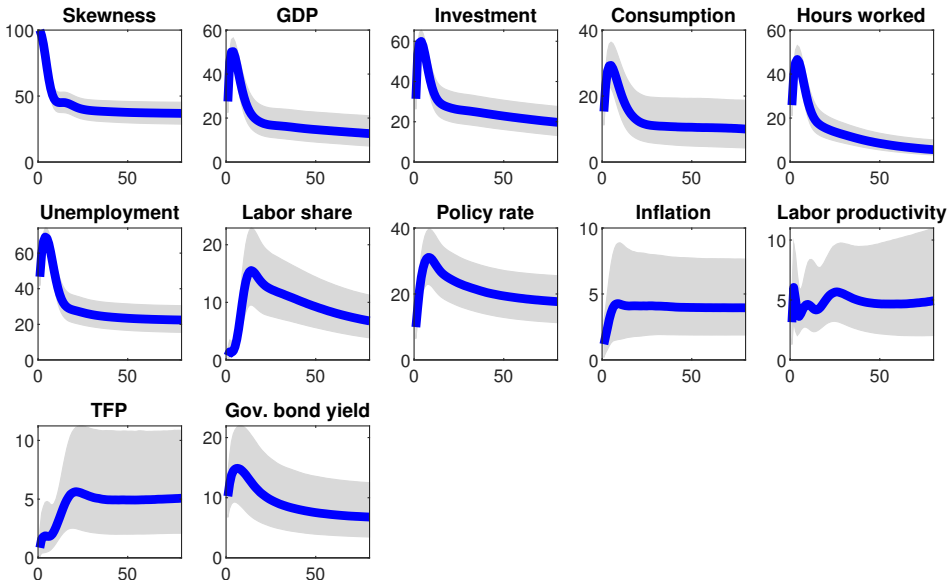

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Appendix F Robustness checks

Results of baseline model with max-share identification approach

Figure F-1: Impulse response functions

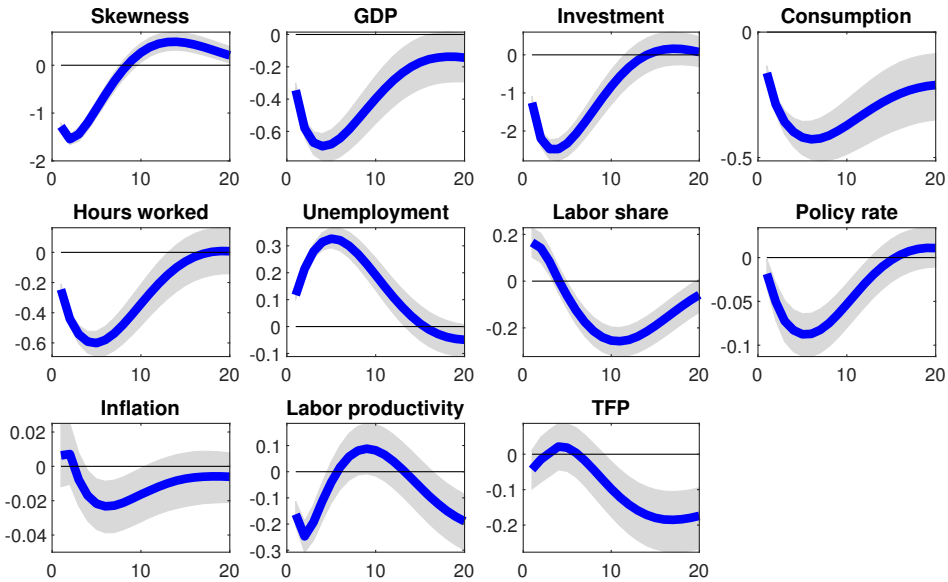

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through max-share approach (Uhlig, 2003). Sample period: 1960:Q1–2017:Q4.

Figure F-2: Forecast error variance contributions

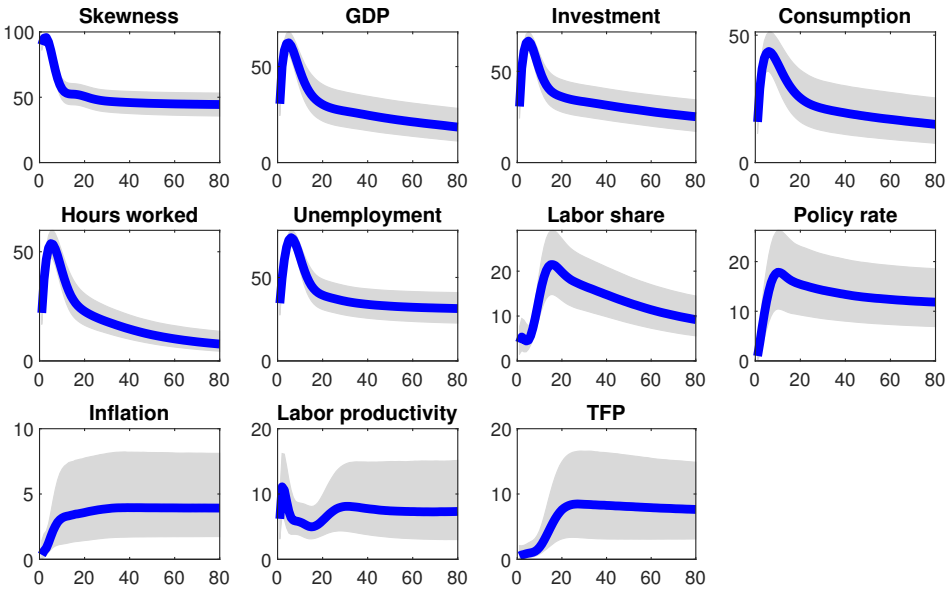

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for (GARCH) volatility

Figure F-3: Impulse response functions

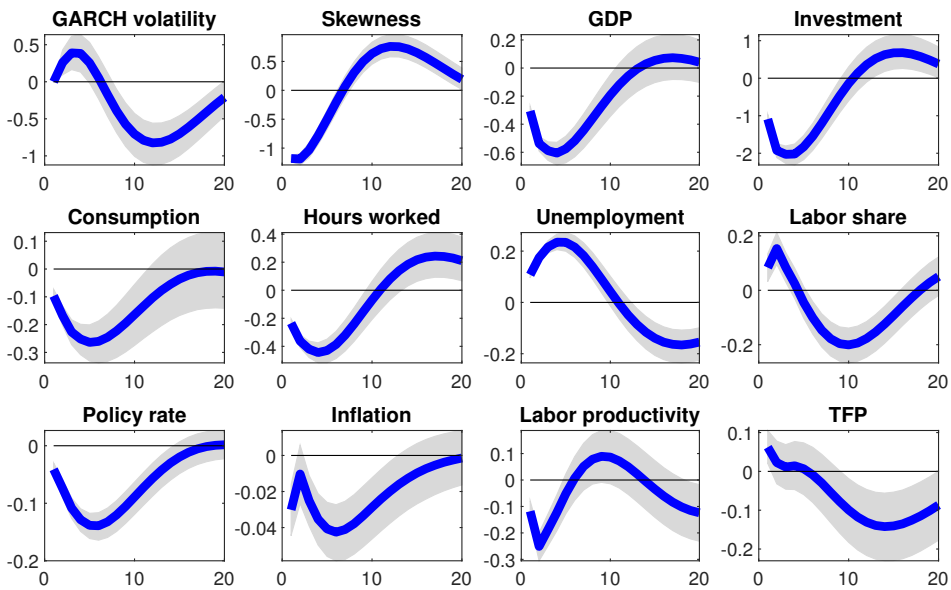

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure F-4: Forecast error variance contributions

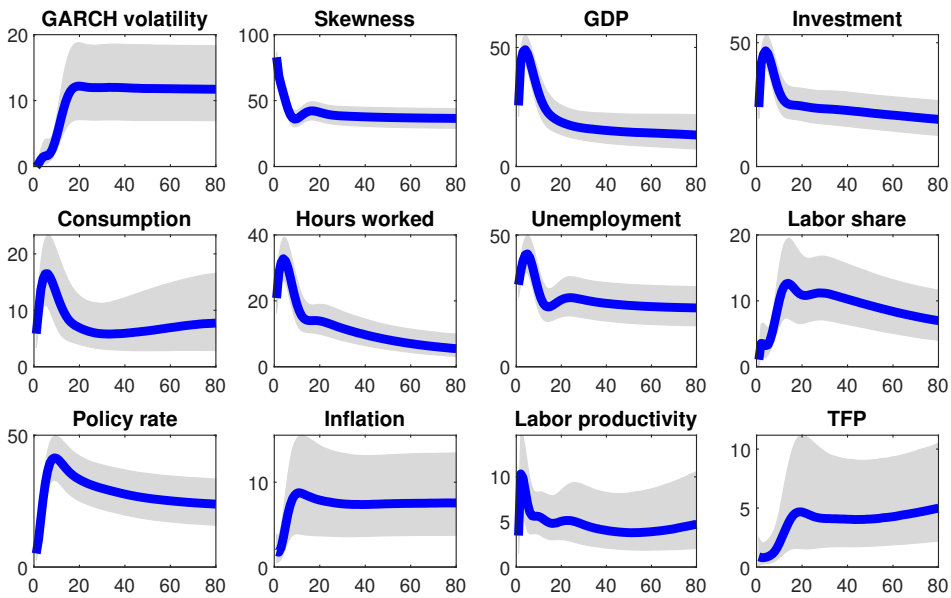

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for macroeconomic and financial uncertainty

Figure F-5: Impulse response functions

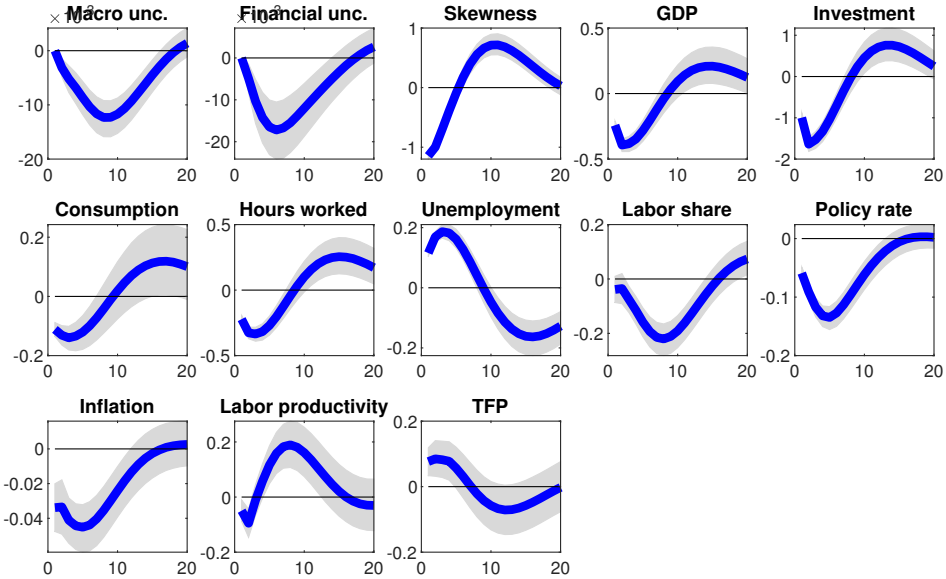

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q3–2017:Q4.

Figure F-6: Forecast error variance contributions

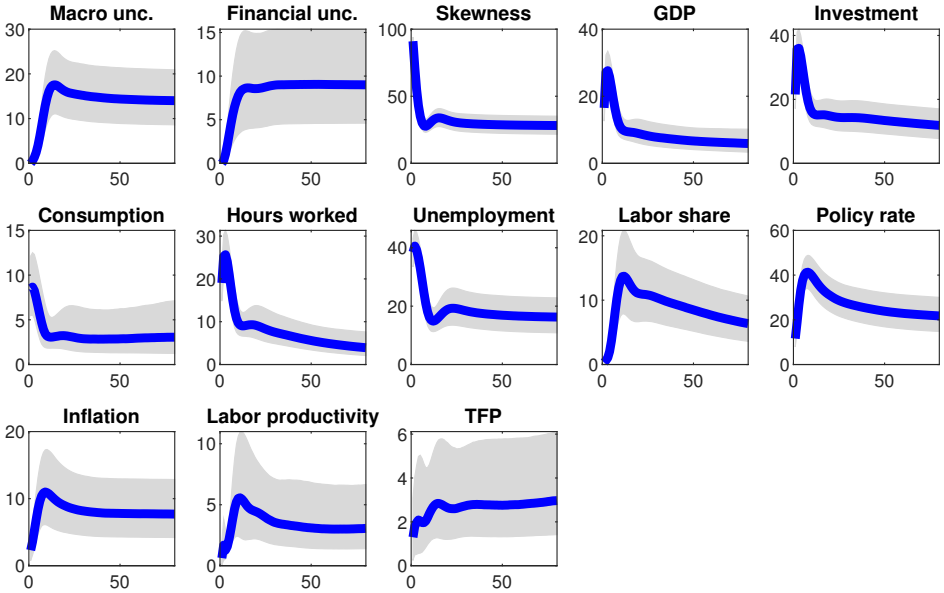

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for geopolitical risk

Figure F-7: Impulse response functions

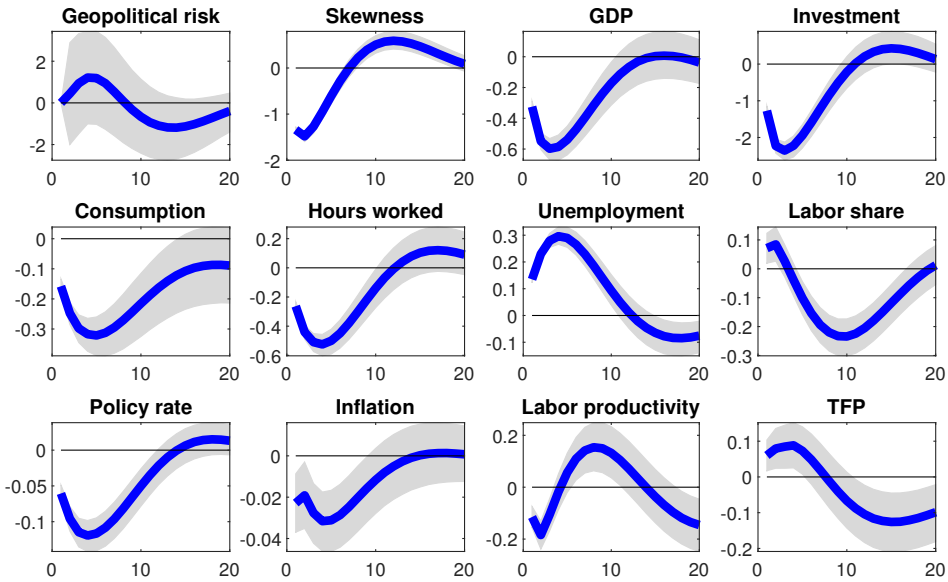

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure F-8: Forecast error variance contributions

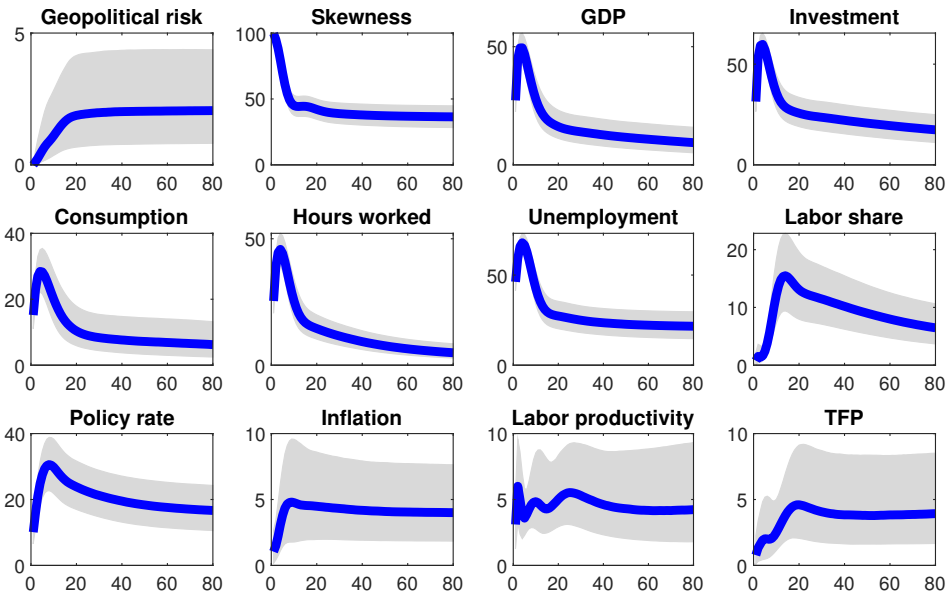

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for excess bond premium

Figure F-9: Impulse response functions

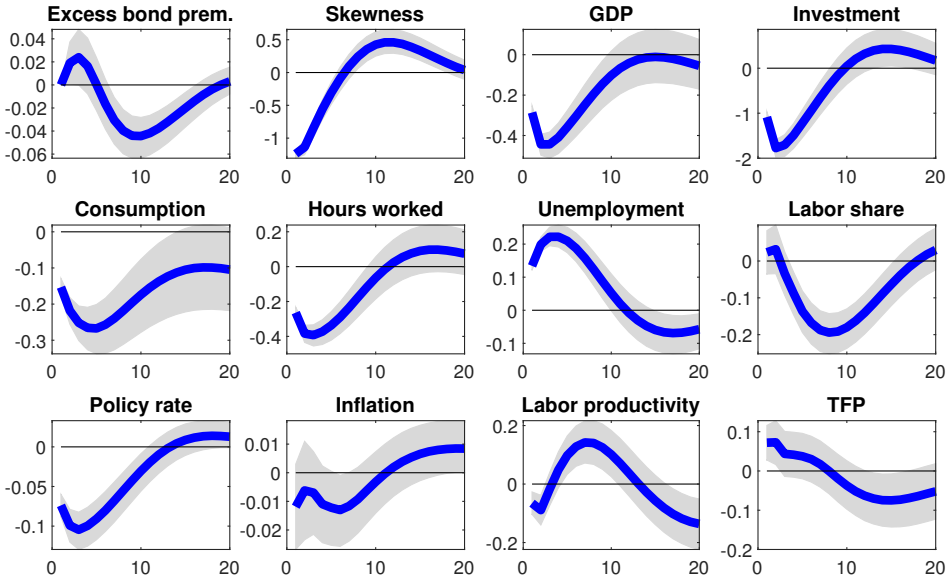

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1973:Q1–2017:Q4.

Figure F-10: Forecast error variance contributions

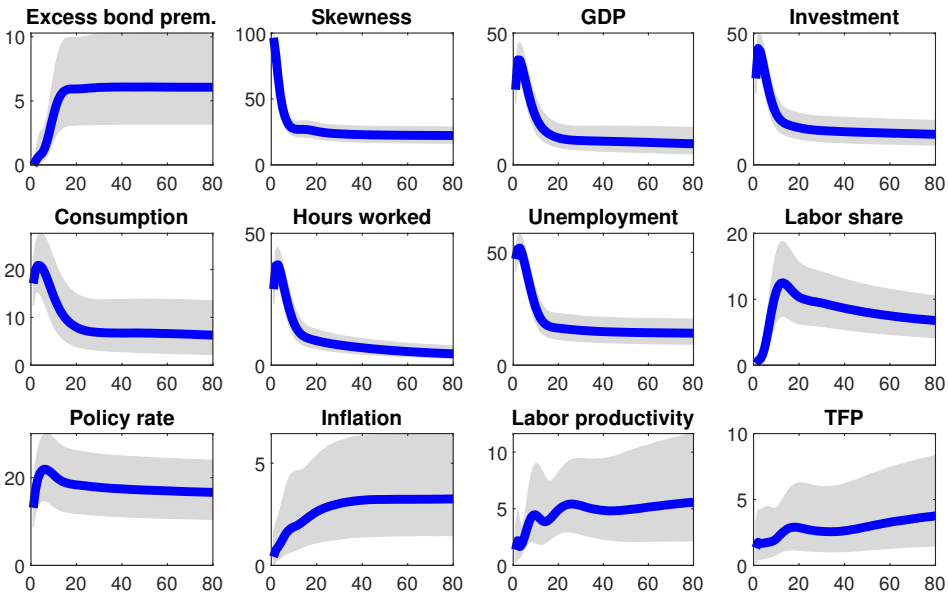

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for TFP growth

Figure F-11: Impulse response functions

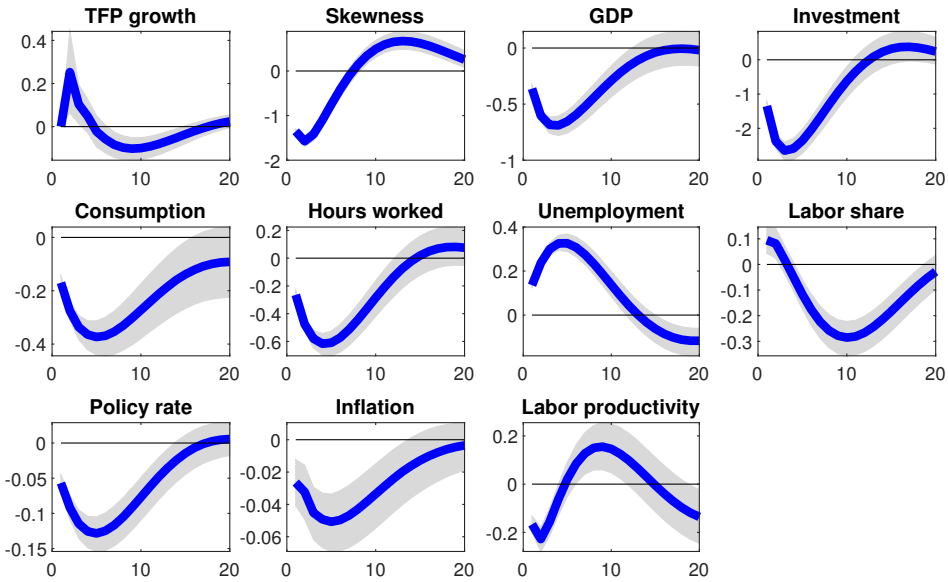

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure F-12: Forecast error variance contributions

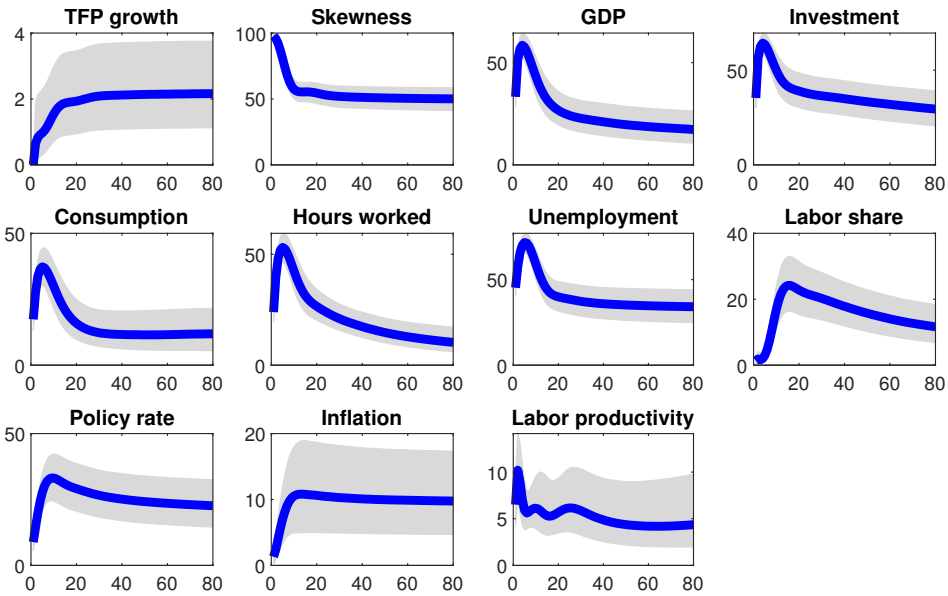

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for fiscal policy shocks

Figure F-13: Impulse response functions

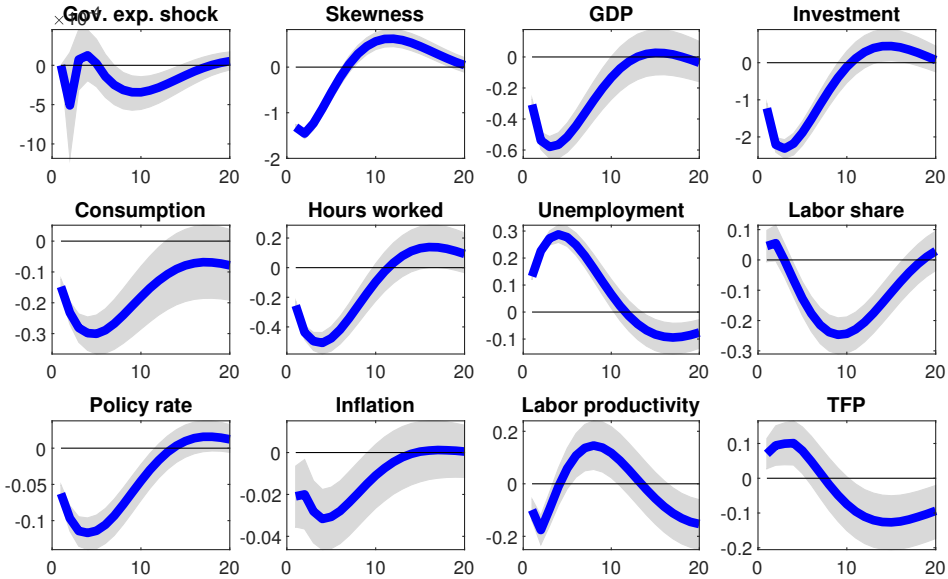

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2015:Q4.

Figure F-14: Forecast error variance contributions

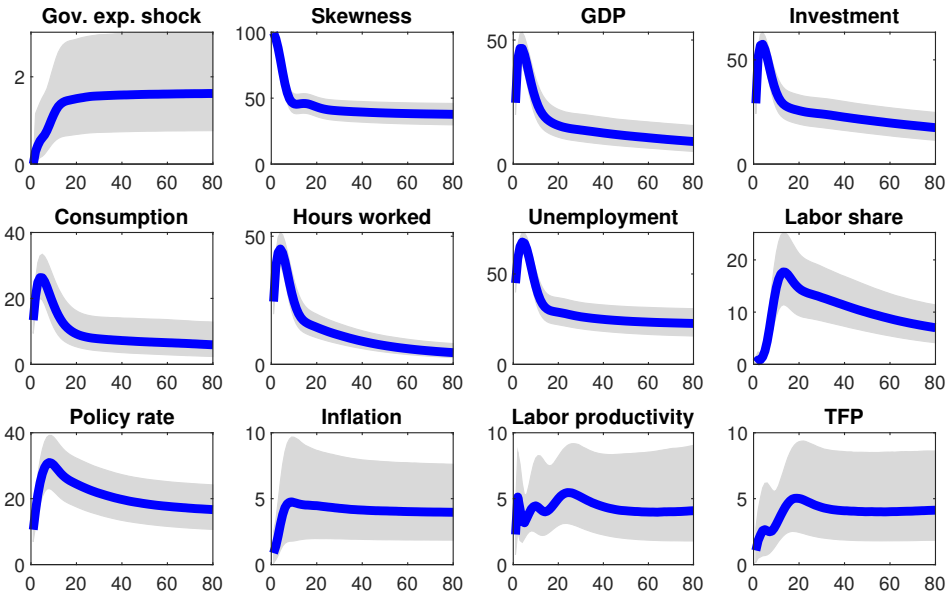

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Model results controlling for monetary policy shocks

Figure F-15: Impulse response functions

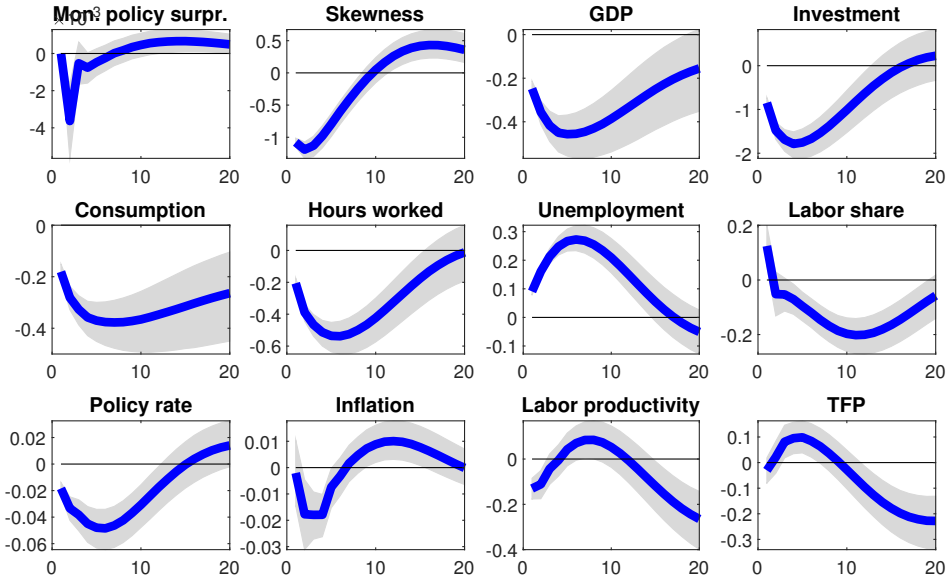

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1990:Q1–2016:Q4.

Figure F-16: Forecast error variance contributions

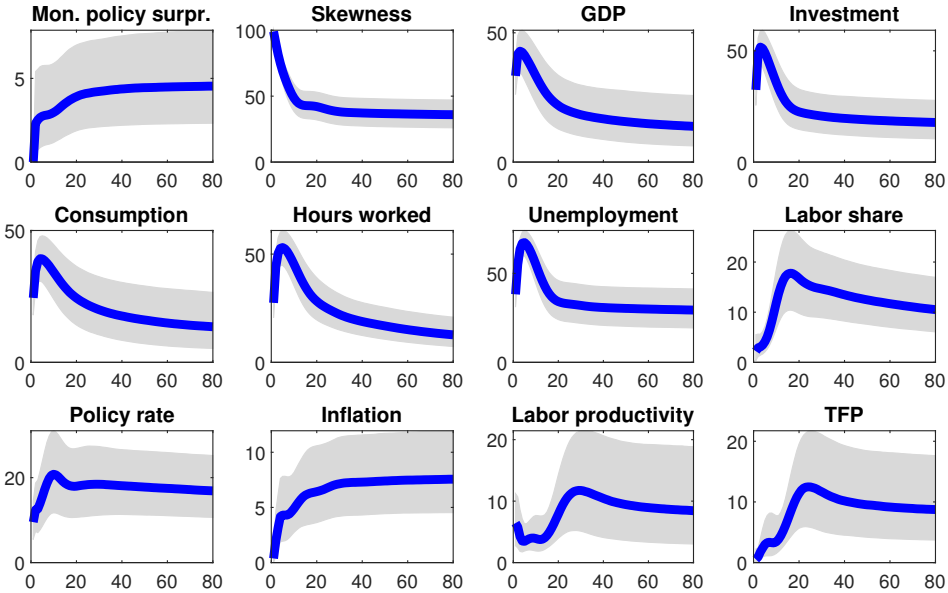

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of baseline model with lag order $P = 4$

Figure F-17: Impulse response functions

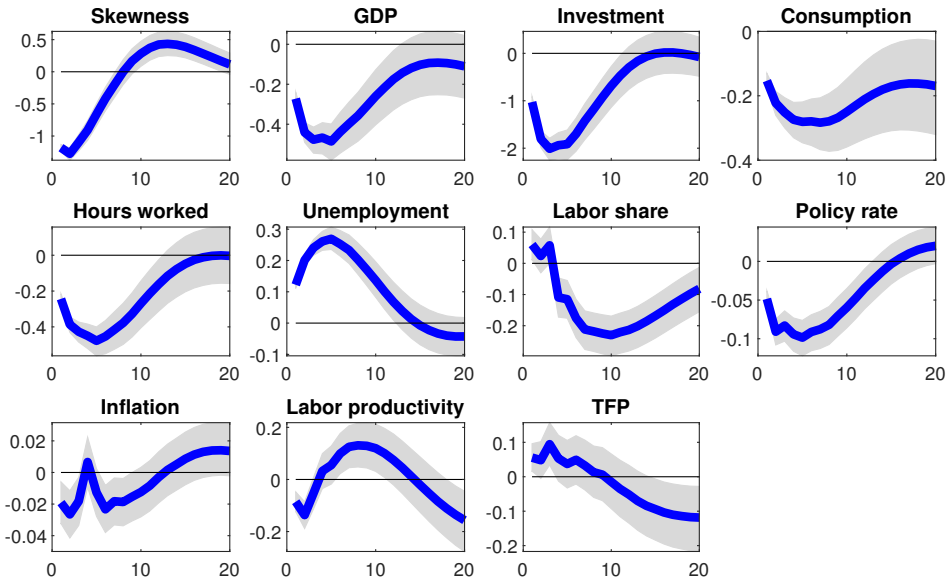

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure F-18: Forecast error variance contributions

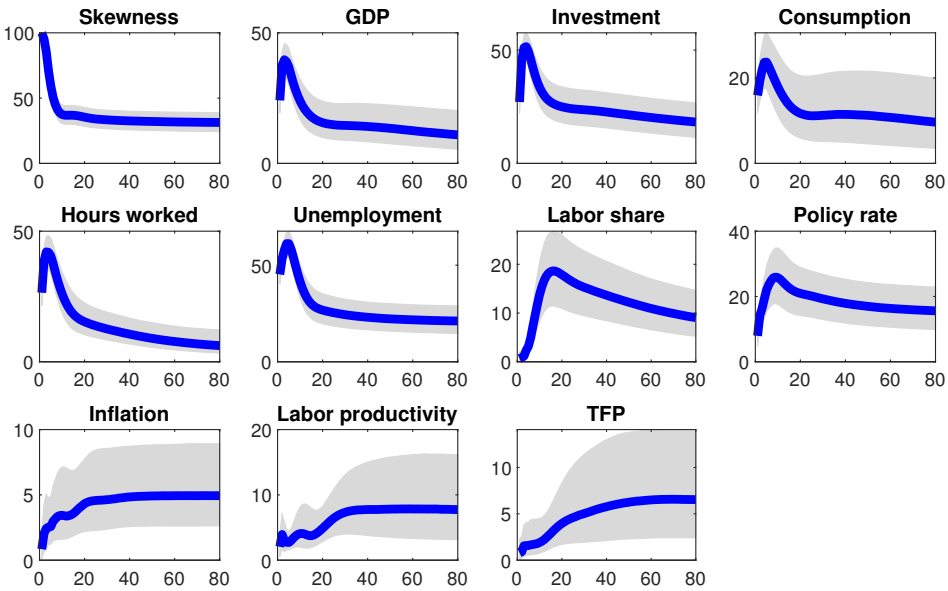

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.

Results of baseline model with looser prior configuration in VAR ($\lambda = 10$)

Figure F-19: Impulse response functions

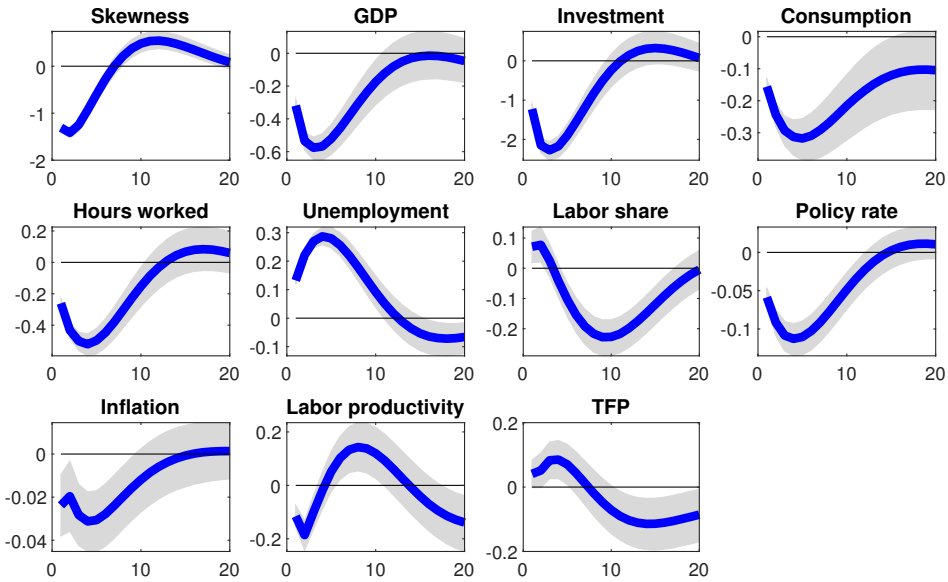

Note: Posterior mean responses to a negative one S.D. shock to expected skewness along with the 68% highest density interval. Identification through Cholesky decomposition. Sample period: 1960:Q1–2017:Q4.

Figure F-20: Forecast error variance contributions

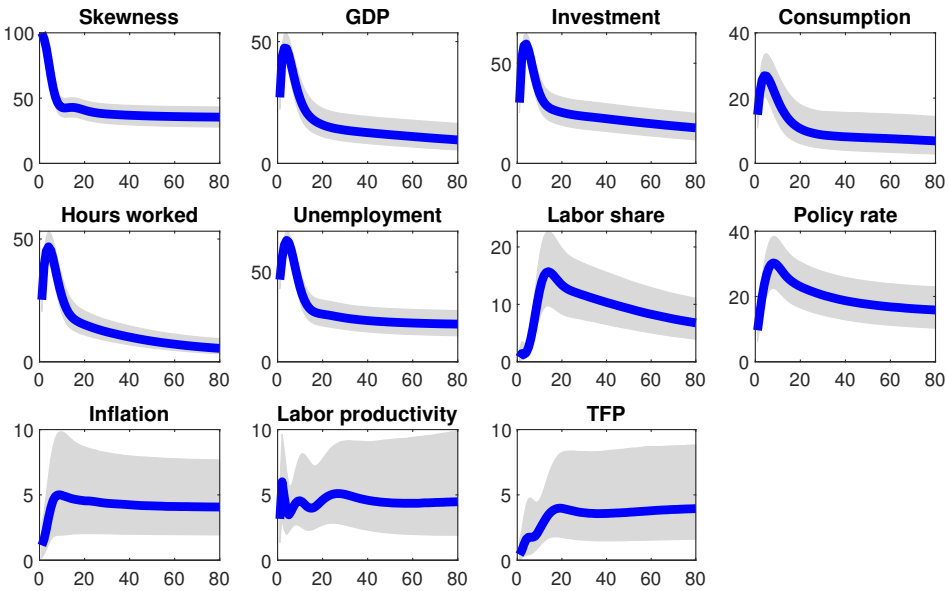

Note: Posterior mean of the forecast error variance contributions along with the 68% highest density interval.