Uncertainty and the Economy: The Evolving Distributions of Aggregate Supply and Demand Shocks*

Geert Bekaert, Columbia University and the Centre for Economic Policy Research,

Eric Engstrom, Board of Governors of the Federal Reserve System

Andrey Ermolov, Gabelli School of Business, Fordham University

February 28, 2023

Abstract

We estimate the time-varying distribution of aggregate supply (AS) and aggregate demand (AD) shocks defined in the Keynesian tradition. In modeling the time variation in higher order moments, we distinguish between traditional Gaussian uncertainty and "bad" uncertainty, associated with negative skewness. The Great Moderation is driven by decreases in Gaussian (not bad) uncertainty. The increased role of bad uncertainty implies that the conditional skewness of GDP growth and inflation and the correlation between level and uncertainty shocks in macro data decreased over time.

Keywords: uncertainty shocks, business cycles, Great Moderation, AS/AD shocks, skewness, deflation risk

JEL codes: E31, E32, E44

*Contact information: Geert Bekaert - gb241@gsb.columbia.edu, Eric Engstrom - eric.c.engstrom@frb.gov, Andrey Ermolov - aermolov1@fordham.edu. The authors thank seminar participants at Bank of Finland, Baruch College, Bilkent University, University of British Columbia, Bundesbank, City University of London, Duke, Fordham, University of Illinois at Urbana-Champaign, Imperial College, Lord Abbett, Oxford, Riksbank, Sabanci Business School, Tulane, and University of North Carolina at Chapel-Hill and conference participants at the 2021 Society for Economic Dynamics Annual Meeting for useful comments. All errors are the sole responsibility of the authors. The views expressed in this document do not necessarily reflect those of the Board of Governors of the Federal Reserve System, or its staff.

1 Introduction

A growing literature presents equilibrium models where uncertainty shocks, that is, increases in the standard deviation of economic shocks, are important drivers of the business cycle (e.g., Justiniano and Primiceri, 2008, Bloom, 2009, Fernández-Villaverde and Rubio-Ramírez, 2013, or Fernández-Villaverde et al., 2015). Researchers generally use either econometric models of time-varying volatility for macro variables (with stochastic volatility and GARCH models perhaps most popular), or proxies for uncertainty, such as the VIX index (see Bloom, 2009).[1]

In a recent survey of the literature on uncertainty shocks and business cycles, Fernández-Villaverde and Guerron-Quintana (2020) highlight the lack of research on skewness shocks, citing the prevalence of negative one sided shocks, which can help create deep recessions. Building on new research in finance,[2] in this paper we decompose macro-uncertainty into "bad" uncertainty, which is accompanied by negative skewness, and standard Gaussian uncertainty.

We start by decomposing macroeconomic shocks into aggregate demand and aggregate supply shocks using only minimal economic structure. We define aggregate supply (AS) shocks in the Keynesian tradition as shocks that move inflation and real activity in the opposite direction, and aggregate demand (AD) shocks as those pushing inflation and real activity in the same direction (see also Blanchard, 1989). This distinction is important, for example, because the appropriate monetary and fiscal policy responses may be quite different for adverse demand versus supply shocks. Our identification of AS and AD shocks builds on Bekaert, Engstrom, and Ermolov (2022), who exploit higher-order moments in the macro data to resolve the identification problem for the structural AS/AD shocks. Despite the weak identification assumptions, our structural shocks exhibit some

[1] Kozeniauskas, Orlik, and Veldkamp (2018) distinguish between uncertainty shocks measured from micro dispersion, belief heterogeneity or macro uncertainty, but show that volatile macro outcomes can create all three types of uncertainty consistent with the data correlations.

[2] Patton and Sheppard (2015) advocate the use of semi variances (which separately uses positive and negative returns) to create "bad" and "good" volatility, and their methodology has been widely applied (e.g. Kilic and Shaliastovich, 2019). Bekaert and Engstrom (2017) introduce a component model with positively and negatively skewed shocks, which is the inspiration for our model.

salient properties that are postulated in the classic Blanchard and Quah (1989) paper in which demand shocks affect output temporarily, whereas supply shocks have a permanent effect on output.

We next formally estimate macro risk factors as the state variables that govern the time-varying volatility, skewness and higher-order moments of supply and demand shocks. To model these dynamic factors, we use the Bad Environment-Good Environment model of Bekaert and Engstrom (2017, "BEGE" henceforth) where each shock consists of a "good environment" and a "bad environment" component shock. In the model, four separate factors drive "good" (positively skewed) and "bad" (negatively skewed) uncertainties of AS and AD shocks. As good uncertainty increases, the distribution for the shock becomes more positively skewed.[3] Increases in the bad-type of uncertainty may pull skewness into negative territory. Thus, the model can easily accommodate asymmetric business cycles (Sichel, 1993; Morley and Piger, 2012). The BEGE model accommodates a wide set of distributions, such as a simple Gaussian or extreme rare disaster distributions. In addition, our model allows for a flexible time-varying correlation structure between shocks that only drive the level of macroeconomic variables versus shocks that affect uncertainty alone.

We use the estimated model to derive three sets of results regarding: 1) the conditional distribution of macro variables, 2) the correlation of level and volatility shocks, and 3) the Great Moderation. First, the data suggest that the good component for both demand and supply shocks is Gaussian, featuring standard persistent volatility dynamics that have been used extensively in other studies. However, the data also strongly support a "bad environment" demand component that is highly negatively skewed, which spikes in recessions and features a more transient volatility process. The supply "bad environment" component is similar but slightly less skewed and its volatility process is slightly more persistent. Overall, both AS and AD shocks are strongly non-Gaussian, at times. Our macro risk measures generate different higher-order (>2) moments for real activity

[3]This distinction opens the possibility of expansionary uncertainty shocks, such as observed during the adoption of the internet in the late 90s.

and inflation depending on whether Gaussian or “bad” risks dominate. In recent years, the conditional distributions for GDP growth and inflation show substantial negative skewness, suggesting increased macro vulnerabilities as well as deflation risk. Our work here generally contributes to the literature proposing and estimating models for GDP growth and inflation that admit conditional non-Gaussianities, starting with the regime switching models of Hamilton (1990) for GDP growth and Evans and Wachtel (1993) for inflation. We show formally that our model fits the data better than regime-switching and asymmetric GARCH models. Our results are consistent with the quantile regression results in Adrian, Boyarchenko, and Giannone (2019), showing an important and time-varying left tail in US GDP growth,[4] and Jensen et al. (2020) showing that GDP growth skewness has decreased over the past three decades, ascribing it to increased leverage of households and firms. However, our focus is broader as we consider the joint distribution of GDP growth and inflation and show how it varies across AD and AS environments.

Second, our econometric model does not impose unrealistic restrictions on the correlation between volatility shocks and shocks to the levels of macroeconomic data. Carriero, Clark and Marcellino (CCM, 2018) point out that more often than not the estimation of uncertainty measures is not embedded in the econometric model used to identify shocks and the uncertainty measures are therefore inefficiently and/or inconsistently estimated (e.g. using a homoskedastic VAR to identify shocks). In illustrating the importance of this shortcoming within the context of a Bayesian VAR, CCM (2018) demonstrate that uncertainty indices produce significantly negative output effects, but ultimately uncertainty shocks are not as important as the shocks to the levels of the variables in the VARs themselves. In doing so, CCM (2018) make the important assumption that volatility and level shocks are independent. Alessandri and Mumtaz (2019) create an uncertainty index from 4 macro series, making the same independence assumption. However, level shocks may be naturally correlated with volatility shocks, with negative economic activity shocks being associated with higher volatility, mimicking the asymmetric volatility

[4]Salgado, Guvenen, and Bloom (2019) show related results for micro-dynamics, that is, the skewness of the growth rates of employment, sales, and productivity at the firm level is time-varying and procyclical.

effect in equities (see, e.g., Engle and Ng, 1993). Our general model does not impose a particular correlation structure on level versus volatility shocks. While the model admits a very flexible time-varying level-volatility correlation, it focuses on two structural shocks and can be estimated from just a few macro series. It is therefore complementary to the reduced form methodology of Gorodnichenko and Ng (2017), who infer volatility shocks from a large panel of data, without imposing correlation restrictions, using a factor model approach. We find strong positive (negative) correlation between demand shocks and shocks to Gaussian (bad) uncertainty for demand. In contrast, supply shocks appear to be largely independent of shocks to Gaussian uncertainty, but negatively correlated with shocks to bad supply uncertainty. Thus, the data support the notion that overall volatility shocks are negatively correlated with level shocks, although these correlations are time-varying and can even switch signs. Bloom et al. (2018) show that empirical impulse responses in a macro VAR can only be fit if they allow negative level shocks to be correlated with uncertainty shocks. This is consistent with our finding that volatility and level shocks are not independent and thus often occur simultaneously.

Third, we use the estimated conditional volatilities and their Gaussian and negatively skewed components to revisit the Great Moderation - a reduction in the volatility of many macroeconomic variables since the mid-1980s. We find it is attributed largely to strong decreases in the volatility of the Gaussian components of both AS and AD shocks. Meanwhile, the volatility of bad shocks has not experienced a significant decline. As a result, the frequency and severity of recessions, which are mostly associated with elevated bad volatility over the last 40 years, have not changed much over our sample. These results offer a refinement to the work of Jurado, Ludvigson and Ng (JLN, 2015), who find a strong counter-cyclical component to aggregate volatility. Our formal break tests confirm the observation in Jensen et al. (2021) that the recessions since the Great Moderation are in fact deeper that the pre-1984 ones, that is, the skewness of real GDP growth has significantly decreased over time. The same is true, but to a lesser extent, for inflation.

2 Modeling and Estimating AS/AD Shocks

We provide a brief overview on how we identify and estimate AS/AD shocks using an identification scheme relying on higher order moments of macro data. The approach is spelled out in Bekaert, Engstrom, and Ermolov (2022), but the actual model, underlying data, and estimation are different here.

2.1 Defining and identifying aggregate supply and demand shocks

Consider a bivariate system in real GDP Growth (g_t) and inflation (π_t):

$$\begin{aligned} g_t &= E_{t-1}[g_t] + u_t^g, \\ \pi_t &= E_{t-1}[\pi_t] + u_t^\pi, \end{aligned} \tag{1}$$

where E_{t-1} denotes the expectation operator conditional on information available at time $t-1$. The variables u_t^g and u_t^π are reduced-form shocks. We model the reduced-form shocks as linear combinations of two structural shocks, labeled supply and demand, and denoted u_t^s and u_t^d, respectively:

$$\begin{aligned} u_t^\pi &= -\sigma_{\pi s} u_t^s + \sigma_{\pi d} u_t^d, \\ u_t^g &= \sigma_{gs} u_t^s + \sigma_{gd} u_t^d. \end{aligned} \tag{2}$$

The σ parameters are the loadings of the reduced-form shocks onto the supply and demand shocks. We assume the σ parameters are all positive to make clear the sign restrictions that we are imposing. In this sense, our use of sign restrictions is different from the common methodology in macroeconomics, pioneered by Faust (1998), Canova and De Nicolo (2002) and Uhlig (2005), to impose sign restrictions on impulse responses to aid identification. The first fundamental economic shock, u_t^s, is an aggregate supply shock, defined so that it moves GDP growth and inflation in opposite directions, as happens, for instance, in episodes of stagflation. The second fundamental shock, u_t^d, is an aggregate demand shock, defined so that it moves GDP growth and inflation in the same direction as would be the case in a typical economic boom or recession from the past few decades. Supply and demand shocks are assumed to be uncorrelated and, without loss of generality, to have unit variance.

While more complex shock structures can be entertained, this minimal structure encompasses many important economic shocks. For example, standard monetary policy shocks can be viewed as demand shocks (see, e.g., Ireland, 2011).

Note that the sample covariance matrix of the reduced-form shocks from the bivariate system in equation (1) only yields three unique moments, but we need to identify four σ coefficients in equation (2) to extract the supply and demand shocks. In particular, the unconditional covariance matrix for inflation and growth shocks is:

$$\begin{bmatrix} \sigma_{\pi s}^2 + \sigma_{\pi d}^2 & -\sigma_{\pi s}\sigma_{gs} + \sigma_{\pi d}\sigma_{gd} \\ -\sigma_{\pi s}\sigma_{gs} + \sigma_{\pi d}\sigma_{gd} & \sigma_{gs}^2 + \sigma_{gd}^2 \end{bmatrix}. \quad (3)$$

Hence, absent additional assumptions, a system with Gaussian shocks would be unidentified.

This set-up is analogous to the identification problem on supply and demand data carefully laid out in Uhlig (2017). To achieve identification, we exploit the well-established result that distributions of macroeconomic data exhibit substantial non-Gaussian features; thus, demand and supply shocks may have non-zero unconditional skewness and excess kurtosis. Specifically, we use the four available unconditional skewness and co-skewness moments and five excess kurtosis and co-kurtosis moments for GDP growth and inflation. To aid identification, we assume that co-skewness (e.g., $E[(u_t^s)(u_t^d)^2]$) and asymmetric excess co-kurtosis (e.g., $E[(u_t^s)(u_t^d)^3]$) are zero, which is consistent with modeling the supply and demand shocks as independent, an assumption Uhlig (2017) also imposes. However, we weaken the assumption of independence to allow for supply and demand to have excess co-kurtosis, that is, $E[(u_t^s)^2(u_t^d)^2 - 1]$ may be nonzero. This unconditional moment captures that the volatilities of supply and demand shocks may be correlated. In sum, this moment together with accommodating non-zero skewness and excess kurtosis for the supply and demand shocks requires the estimation of five additional parameters associated with higher-order moments. Thus, to achieve identification econometrically through non-Gaussianity we must make use of at least 6 higher-order moments (see also Lanne, Meitz, and Saikkonen, 2017, and Lanne and Luoto, 2021, for theoretical work on obtaining identification through higher-order moments in a VAR).

The main advantage of the definition for supply and demand shocks above is that it carries minimal theoretical restrictions (only a sign restriction). Moreover, once we have estimated the σ parameters in equation (2), we can simply invert the supply and demand shocks without further assumptions:

$$\begin{aligned} u_t^s &= \frac{\sigma_{\pi d} u_t^g - \sigma_{gd} u_t^\pi}{\sigma_{\pi d}\sigma_{gs} + \sigma_{\pi s}\sigma_{gd}}, \\ u_t^d &= \frac{\sigma_{\pi s} u_t^g + \sigma_{gs} u_t^\pi}{\sigma_{\pi d}\sigma_{gs} + \sigma_{\pi s}\sigma_{gd}}. \end{aligned} \tag{4}$$

2.2 Measuring macro shocks

We use forecast revisions from survey data to operationalize equation (1), obviating the need for model selection in estimating reduced-form macro shocks. Not having to model conditional means helps mitigate the criticisms of CCM (2018) on the modeling of volatility shocks within an inconsistent econometric framework. The survey data are from the Survey of Professional Forecasters (SPF). Our sample is quarterly from 1968:Q4 to 2019:Q2 (203 quarters). The number of respondents in the SPF varies over time and across macro variables being forecasted but a typical number of respondents is about 40.

To identify inflation shocks using the survey data, we use:

$$u_t^\pi = \hat{\pi}_t - \hat{\pi}_{t,t-1}, \tag{5}$$

where $\hat{\pi}_t$ is the forecast in quarter t for the percentage change in the GDP deflator in quarter t (π_t), and $\hat{\pi}_{t,t-1}$ the forecast for π_t in the previous quarter. Therefore, u_t^π represents the revision to the expectation for π_t between periods $t-1$ and t. Note that published data for π_t is generally not fully available until many quarters after (at least until $t+1$ for an advance release), so $\hat{\pi}_t$ need not equal the eventually published official value for π_t. The SPF survey is typically published around the 10^{th} day of the month in February, May, August and November of each year. As a concrete example, our measured revision to inflation for the period 2018:Q1 is equal to the average SPF forecast (as of early February 2018) for inflation for Q1 inflation minus the expectation for Q1 inflation that was measured in the previous SPF survey, published in early November of 2017. The inflation data forecasted in the survey corresponds to the percentage change in the GDP price deflator over the first (calendar) quarter of 2018; this data is first published (through an advance release) by the U.S. Bureau of Economic Analysis (BEA) in April.

Our use of survey revisions to measure economic shocks is perhaps uncommon, but we believe it is well justified. First, the true pace of economic activity is never directly observed, only estimated. One estimate of economic activity is the BEA advance release that is published one month after the quarter end. Another estimate is the BEA final (revised) release, which is published many quarters (and often years) after the fact. The latter measure is the one which is perhaps most often used in academic papers, but it is the least plausible candidate for being in the minds of economic agents due to the lag in publishing. For example, Ghysels, Horan, and Moench (2018) show that the use of real time data substantially reduces the predictive power of macro variables for bond returns, suggesting that investors do not anticipate future data revisions. In addition, GDP and inflation data most certainly are plagued by measurement error (see, e.g., Aruoba et al., 2016, for GDP and Lebow and Rudd, 2006, for inflation), which renders the structural modeling of shocks more difficult. In contrast, current-quarter nowcasts from survey data also offer viable estimates of economic activity - those of the survey respondents, and they have the advantage of being available in real time, and are therefore plausibly reflected in the beliefs of economic agents. Moreover, they should be less subject to measurement error noise. Second and importantly, these revisions do correspond to a difference in realized value from its conditional expectation as in equation (1). Because of the law of iterated expectations, $\hat{\pi}_{t,t-1}$ is the conditional expectation for $\hat{\pi}_t$ in the previous quarter.

Finally, Ang, Bekaert and Wei (2007) show that inflation expectations from the SPF provide more accurate forecasts of future inflation than statistical, Phillips curve and term structure models. Coibion and Gorodnichenko (2012, 2015) show that the predictability of forecast errors from SPF inflation forecasts (including predictability coming from forecast revisions) is consistent with models of information rigidities and cast doubt on full rational expectations models. Our estimates are therefore more consistent with actual expectation formation than econometric models estimated on revised data would be. Similarly, we measure shocks to the outlook for real activity as forecast revisions for the percentage change in real GDP growth.[5]

Figure 1 depicts the resulting real GDP and inflation shocks, expressed as a percentage change at an annual rate. Shocks to real GDP shocks are generally larger earlier in the sample, and deeply negative spikes occur during recessions throughout the sample. Similarly, inflation variability is higher earlier in the sample and large positive and negative spikes are evident during recessions that occur early in the sample period. Later in the sample period, the overall variability of inflation decreases, and the shocks during recessions are notably negative.

In Appendix I we verify that these patterns appear largely consistent with patterns in macro shocks defined in a more standard VAR fashion. In particular, we extract GDP growth and inflation shocks from a bivariate VAR, which also uses SPF data as predictors (Ang, Bekaert, and Wei, 2007, emphasize the importance of survey forecasts in such regressions, and they considerably improve explanatory power in our setting). For example, we find that the correlation between GDP growth VAR shocks and forecast revisions is 0.59.

2.3 Estimating supply and demand shocks

To estimate the σ coefficients in equation (2), we use information in all available 2$^{\text{nd}}$, 3$^{\text{rd}}$, and 4$^{\text{th}}$ order unconditional moments of the reduced-form macroeconomic shocks in a classical minimum distance (CMD) estimation framework (see, e.g., Wooldridge, 2002, pp. 445-446). Specifically, we calculate 12 statistics based on the two series of shocks measured

[5]Mechanically, survey respondents fill out forecasts for nominal GDP and the GDP deflator separately, with their forecasts for real GDP being calculated as the ratio between the two.

in the survey data. These are the unconditional sample standard deviations (2 statistics), the correlation (1 statistic), univariate (scaled) skewness and excess kurtosis (4 statistics), co-skewness (2 statistics), and co-excess kurtosis (3 statistics). The parameters we use to match these moments include the loadings of inflation and real activity onto supply and demand shocks, ($\sigma_{\pi d}$, $\sigma_{\pi s}$,σ_{gd}, and σ_{gs}), the unconditional skewness, $E[(u_t^d)^3]$ and $E[(u_t^s)^3]$, and excess kurtosis of supply and demand shocks, ($E[(u_t^d)^4] - 3$) and ($E[(u_t^s)^4 - 3]$), and the excess co-kurtosis of supply and demand shocks, ($E[(u_t^d)^2(u_t^s)^2] - 1$).

With 12 moments to match and 9 parameters to estimate, our system is overidentified, thus requiring a weighting matrix. To generate a weighting matrix, we use the inverse of the covariance matrix of the sampling error for the statistics, consistent with asymptotic theory suggesting that this choice leads to efficient estimates. We use a block bootstrapping routine to calculate the covariance matrix. Specifically, we sample, with replacement, blocks of length 12 quarters of the two survey-based macroeconomic shocks, to build up a synthetic sample of length equal to that of our data. We calculate the same set of 2$^{\text{nd}}$, 3$^{\text{rd}}$, and 4$^{\text{th}}$ order statistics for each of 10,000 synthetic samples. We then calculate the covariance matrix of these statistics across bootstrap samples.

Table 1 reports the sample statistics that we use for the estimation. The volatility statistics have rather tight standard errors, but the unconditional correlation of inflation revisions and revisions to real growth is insignificantly different from zero with a point estimate of -0.13. Real growth shocks are significantly negatively skewed with estimated skewness of -1.23, and the co-skewness moment involving inflation revisions squared times real growth revisions is significantly negative. Together, these two moments suggest that real growth is, on average, more negative when inflation volatility is high and when real growth volatility is high. The excess kurtosis of real growth is significantly positive with a value of 4.71, as is the fourth moment involving squared inflation revisions times squared growth revisions. The latter indicates that the volatilities of inflation and real growth tend to move together. The p-value for the joint significance of all the 3^{rd} and 4^{th} order moments is 0.26 percent, strongly rejecting the hypothesis that the data are distributed unconditionally according to a multivariate Gaussian distribution and providing strong support for our identification scheme.

In Table 2, Panel A, we report the supply and demand loadings for GDP growth and inflation. These are generally quite precisely estimated. Our estimates suggest that supply and demand shocks contribute roughly equally to the unconditional variance of inflation shocks over this sample period: the inflation supply and demand loadings are -0.48 and 0.51, respectively. Unconditionally, supply shocks, contribute somewhat more than demand shocks to the overall variance of real growth shocks: the real GDP growth supply and demand loadings are 1.18 and 0.60 respectively (in variance terms the relative contribution is approximately 80% versus 20%). This is perhaps surprising at first glance, but our sample includes the stagflation of the 70's and the Great Recession, which many argue had a significant supply component (see, e.g., Ireland, 2011, or Mulligan, 2012). In addition, we find that the resulting supply and demand shocks are more than 99% correlated with out estimated shocks employing the standard generalized method of moments (GMM). Appendix II provides further detail.

Returning to Table 1, in square brackets we report the fitted values for all statistics. Recall that because the system is overidentified by 3 degrees of freedom, not all moments can be fit perfectly. Nonetheless the overall fit is quite good. All second and third-order moments are within a one standard error band of the point estimate, and all the fourth order moments are within a two standard error band. We also report a standard overidentification test for the CMD model fit. The corresponding p-value is 38.74 percent implying that the model is not rejected.

In Panel B of Table 2, we report the estimated skewness and kurtosis of the supply and demand shocks. Both shocks are negatively skewed and leptokurtic (though only for supply shocks are these estimates statistically significant). Interestingly, we find little evidence for excess co-kurtosis between supply and demand shocks, suggesting that the variances of supply and demand shocks may not covary strongly. The top panels of Figures 2 and 3 depict the supply and demand shocks that we recover from this exercise. Both sets of shocks exhibit greater overall variability early in the sample period, followed by a secular decline in variability that perhaps reflects the so-called "Great Moderation", although deeply negative shocks occur during recessions throughout the entire sample.

In Appendix I, we again verify that these patterns are consistent with patterns in demand and supply shocks inverted from VAR GDP growth and inflation shocks. For example, the correlation between both demand and supply shocks extracted using the VAR and forecast revisions is 0.50.

Our supply and demand shocks definitions do not necessarily comport with demand and supply shocks in, say, a New Keynesian framework (see, e.g., Woodford, 2003) or structural VARs in the Sims (1980) tradition.[6] However, Appendix III shows that the short and long-term effects of the AS/AD shocks thus identified are consistent with the standard Keynesian interpretation (e.g., Blanchard, 1989, or Blanchard and Quah, 1989). Furthermore, when used to characterize recessions as supply or demand driven, our results are consistent with Gali (1992) for the first five recessions in the sample.

3 Modeling and Estimating Time-varying AS and AD Macro Risk Factors

Having identified supply and demand shocks, we now examine how the magnitudes of uncertainty associated with supply and demand shocks evolve over time, and how those time-varying volatilities affect various forms of economic activity.

3.1 Defining macro risk factors

We define macro risk factors as the variables that capture the time-variation in the second and higher-order moments of supply and demand shocks. Statistically, we generalize the "bad environment-good environment" (BEGE) framework of Bekaert and Engstrom (2017) to accommodate potentially independent innovations to the level and volatility of supply and demand shocks.

Consider a generic shock, u_{t+1} (e.g., a supply or demand shock) to occur at time

[6] Furthermore, in some models the "supply" shocks might move real activity and inflation in the same direction: see, for instance, news shocks in Cochrane (1994).

$(t+1)$. We model u_{t+1} as having two components:

$$u_{t+1} = \sigma_{up}\omega_{p,t+1} - \sigma_{un}\omega_{n,t+1}, \tag{6}$$

where $\omega_{p,t+1}$ and $\omega_{n,t+1}$ are individual component shocks. The volatility parameters σ_{up} and σ_{un} are restricted to be positive. The component shocks are independent and distributed as centered-gamma:

$$\begin{aligned} \omega_{p,t+1} &\sim \tilde{\Gamma}(p_t, 1), \\ \omega_{n,t+1} &\sim \tilde{\Gamma}(n_t, 1), \end{aligned} \tag{7}$$

where the expression $\omega_{p,t+1} \sim \tilde{\Gamma}(p_t, 1)$ denotes that the random variable $\omega_{p,t+1}$ follows a centered gamma distribution with shape parameter p_t and a unit scale parameter.[7] Consider the first term on the right hand side of equation (6), $\sigma_{up}\omega_{p,t+1}$. Because the $\omega_{p,t+1}$ shock is right skewed, we refer to it as a "good" shock (though it has zero mean and it may, of course, have negative realizations). The variance of this component of u_{t+1} is $\sigma_{up}^2 p_t$, which is a well-known feature of the gamma distribution, and its (unscaled) third moment is $2\sigma_{up}^3 p_t$. When p_t is time-varying, we refer to p_t as the "good variance" state variable. Similarly, the second term in u_{t+1}, $-\sigma_{un}\omega_{n,t+1}$, is negatively skewed, with variance $\sigma_{un}^2 n_t$ and third moment of $-2\sigma_{un}^3 n_t$. We thus refer to n_t as the "bad variance" state variable. The standard (scaled) skewness coefficient decreases (increases) in p_t (n_t) and the demeaned gamma distribution converges to a Gaussian distribution for large p_t and n_t. In practice, a demeaned gamma distribution becomes indistinguishable from a Gaussian distribution when its shape parameter is over 20.

For an illustration of the density implied by equation (6), the upper part of Panel A in Figure 4 illustrates that the probability density function of $\sigma_{up}\omega_{p,t+1}$ (the "good" component) is bounded from the left and has an unbounded right tail. Similarly, the middle part of Panel A in Figure 4 shows that the probability density function of $-\sigma_{un}\omega_{n,t+1}$

[7] The probability density function is $\phi(\omega_{p,t+1}) = \frac{1}{\Gamma(p_t)}(\omega_{p,t+1} + p_t)^{p_t - 1} e^{-\omega_{p,t+1} - p_t}$ with $\omega_{p,t+1} > p_t$ and $\Gamma(p_t)$ representing the gamma function. This distribution has zero mean, unlike the standard gamma distribution.

(the “bad” component) is bounded from the right and has an unbounded left tail. Panel B of Figure 4 illustrates possible conditional distributions of u_t which could arise as a result of time variation in the shape parameters p_t and n_t. In particular, the probability density function at the top of Panel B in Figure 4 characterizes the situation where good volatility (as governed by p_t) is relatively large and the distribution has a pronounced right tail, while the probability density function in the bottom corresponds to the case where bad volatility is relatively large (i.e., a large value for n_t) with the distribution exhibiting a pronounced left tail.

To model the dynamics of macroeconomic uncertainty, we assume that the risk factors, p_t and n_t, follow simple autoregressive processes:

$$\begin{aligned} p_{t+1} &= \bar{p}(1-\rho_p) + \rho_p p_t + \sigma_{pp}\nu_{p,t+1}, \\ n_{t+1} &= \bar{n}(1-\rho_n) + \rho_n n_t + \sigma_{nn}\nu_{n,t+1}, \end{aligned} \tag{8}$$

where $\bar{p}$ and $\bar{n}$ are the unconditional means of the variables, and ρ_p and ρ_n govern their autocorrelation. The volatility parameters, σ_{pp} and σ_{nn} are restricted to be positive. The shocks to good and bad variance, $\nu_{p,t+1}$ and $\nu_{n,t+1}$ in equation (8), are independent gamma-distributed shocks and they also use p_t and n_t as their shape parameters:

$$\begin{aligned} \nu_{p,t+1} &\sim \tilde{\Gamma}(p_t, 1), \\ \nu_{n,t+1} &\sim \tilde{\Gamma}(n_t, 1). \end{aligned} \tag{9}$$

This is similar to other common “square root volatility” specifications in which the conditional volatility of the shock is proportional to the square root of the level of the series. For example, the conditional volatility of p_{t+1} is $\sigma_{pp}\sqrt{p_t}$. In the general model presented above, the following conditional moments for u_t follow from the properties of the gamma distribution:

$$
\begin{aligned}
&E_t[u_{t+1}] = 0, \\
&E_t[u_{t+1}^2] = Var_t[u_{t+1}] = \sigma_{up}^2 p_t + \sigma_{un}^2 n_t, \\
&E_t[u_{t+1}^3] = 2\sigma_{up}^2 p_t - 2\sigma_{un}^2 n_t, \\
&E_t[u_{t+1}^4] - 3(E_t[u_{t+1}^2])^2 = 6\sigma_{up}^4 p_t + 6\sigma_{un}^4 n_t.
\end{aligned}
\tag{10}
$$

Thus, the BEGE structure implies that the conditional variance of macro shocks varies through time, and that the shocks may be conditionally non-Gaussian, with time variation in the higher order moments all driven by p_t and n_t. Moreover, the conditional variances of GDP and inflation vary through time as functions of these structural macroeconomic risk factors, $[p_t^s, n_t^s, p_t^d, n_t^d]'$, with the "$s$" superscript denoting supply variables and "d" denoting demand. In addition, the model also implies that the conditional covariance between inflation and GDP growth shocks is time-varying and can switch signs:

$$
Cov_{t-1}[u_t^g, u_t^\pi] = -\sigma_{\pi s}\sigma_{gs} Var_{t-1}[u_t^s] + \sigma_{\pi d}\sigma_{gd} Var_{t-1}[u_t^d], \tag{11}
$$

where the subscripts on the Cov and Var operators denote that they may vary over time. As can be seen in equation (11), when demand variance dominates the covariance is positive but when supply variance dominates it is negative.

To close the model, we must make assumptions regarding the correlation between the "level" shocks, $\omega_{n,t}$ and $\omega_{p,t}$, and the "volatility" or "uncertainty" shocks, $\nu_{n,t}$ and $\nu_{p,t}$. As we noted above, such shocks are often assumed to be independent. Let's start by assuming that the two types of shocks (ω's and ν's) are independent. Then, to allow a flexible correlation between the level and volatility shocks, we replace equation (6) with:

$$
u_{t+1} = \sigma_{up}((1-\lambda_p^2)^{\frac{1}{2}}\omega_{p,t+1} + \lambda_p \nu_{p,t+1}) - \sigma_{un}((1-\lambda_n^2)^{\frac{1}{2}}\omega_{n,t+1} + \lambda_n \nu_{n,t+1}), \tag{12}
$$

where λ_p and λ_n are between 0 and 1. Although the formulation looks complex, it is simply structured to imply that the conditional correlation between the good component of u_{t+1} and p_{t+1} is equal to λ_p. Analogously, the conditional correlation between the bad component of u_{t+1} and n_{t+1} is $-\lambda_n$. Note that despite the complexity of the model in equation (12), the conditional variance, Var_t, of u_t is still $\sigma_{up}^2 p_t + \sigma_{un}^2 n_t$. Moreover, we have:

$$Cov_t[u_{t+1}, Var_{t+1}] = \lambda_p \sigma_{up}^3 \sigma_{pp} p_t - \lambda_n \sigma_{un}^3 \sigma_{nn} n_t. \tag{13}$$

Naturally, bad (good) variance shocks lower (increase) the conditional covariance between level shocks and uncertainty shocks. When the bad variance state variable dominates, the model generates the macroeconomic counterpart of asymmetric volatility in finance (e.g., Heston, 1993): negative shocks are associated with higher conditional volatility. When $\lambda_p = \lambda_n = 0$, the covariance between the level shocks and the conditional variance is zero, as is the case in a standard Gaussian GARCH model.

Because the specification in equation (12) requires two additional parameters and involves 4 latent shock variables, we also consider more parsimonious special cases. In one case we set the λ parameters equal to zero. Under this specification, $\nu_{p,t+1}$ and $\nu_{n,t+1}$ are "pure" volatility shocks, with no effect on the level of the overall macro shock. At the other extreme, the λ's equal 1. In this case, the level and risk factor shocks coincide. For example, when $\lambda_p = 1$, the good component of u_t is perfectly correlated with the shock to p_t. It is worth noting that even in this seemingly restrictive case, there is still independent variation between the observed macro shock, u_{t+1}, and the risk factors. To see this, note that when $\omega_{p,t+1} = \nu_{p,t+1}$, the conditional correlation between u_{t+1} and p_{t+1} is $Corr_t(u_{t+1}, p_{t+1}) = \frac{\sigma_{up}\sigma_{pp}p_t}{(\sigma_{up}^2 p_t + \sigma_{un}^2 n_t)^{\frac{1}{2}} (\sigma_{pp}^2 p_t)^{\frac{1}{2}}}$ which in general varies from 0 ($n_t \gg p_t$) to 1 ($p_t \gg n_t$).[8]

The model is also rich enough to approximate the popular "disaster risk" model (e.g.,

[8]There is also a large body of work in finance on the importance of "volatility of volatility" shocks, see, e.g., Bollerslev, Tauchen and Zhou (2009). The factor structure that we build in for the higher order moments of the structural shocks implies that the variance of the variance of u_t is also an affine function of p_t and n_t. Denoting, the variance of the variance of u_t by q_t, we have: $q_t = \sigma_{up}^4 \sigma_{pp}^2 p_t + \sigma_{un}^4 \sigma_{nn}^2 n_t$. Note that q_t is not perfectly correlated with the conditional variance of u_t, but the model does imply an intuitive positive correlation between the variance of u_t and its variance of variance, q_t.

Gabaix, 2012, or Wachter, 2013), where Gaussian shocks are combined with a jump process delivering occasional severe negative shocks (e.g., following a Poisson distribution). Such a model emerges when the $\omega_{p,t}$ shock is (nearly) Gaussian, and the $\omega_{n,t}$ shock is very skewed (with n_t having a very low mean).

3.2 Estimating the macro risk model

With the pre-estimated supply and demand shocks in hand from Section 2, we identify the BEGE model parameters separately for the supply and demand shocks. We use an estimation and filtering apparatus due to Bates (2006). The methodology is similar in spirit to that of the Kalman filter, but the Bates routine accommodates non-Gaussian shocks. We relegate a technical discussion to Appendix IV.

The estimation of the BEGE model uses pre-estimated coefficients and time series estimates for supply and demand shocks that are subject to sampling error. Therefore, we also execute the entire estimation process including the identification of the AS/AD shocks, using synthetic data, bootstrapped from our raw survey data observations to conduct statistical inference.

Because the identification scheme for structural shocks in Section 2 is model-free, we can employ any statistical model that can accommodate non-Gaussian unconditional moments in the data. The BEGE model is flexible and has been demonstrated to fit macroeconomic and financial data quite well, out-performing alternative non-linear models in several cases (see also Bekaert, Engstrom and Ermolov, 2015). For our macro data, we compare our BEGE specifications against regime-switching models of the Hamilton (1989)-type with 2-4 regimes and the GJR-GARCH model of Glosten, Jagannathan, and Runkle (1993), invariably finding the best BEGE models to outperform these alternative models (see also below).

3.3 Model Selection

The general model in equations (6)-(9) and (12) is quite highly parameterized. Therefore, we estimate a number of variations on the basic BEGE model for both supply and demand shocks to identify the most parsimonious specifications that are supported by the data. We are particularly interested in determining whether the shocks to the macro series, u_{t+1}, are correlated or independent of shocks to the good or bad risk factors. The various specifications that we investigate include:

1. $\omega_{p,t+1}$ and $\nu_{p,t+1}$ are (i) independent ($\lambda_p = 0$), (ii) coincide ($\lambda_p = 1$), (iii) partially correlated (λ_p free)

2. $\omega_{n,t+1}$ and $\nu_{n,t+1}$ are (i) independent ($\lambda_n = 0$), (ii) coincide ($\lambda_n = 1$), (iii) partially correlated (λ_n free)

3. p_t is time-varying or constant

4. n_t is time-varying or constant

5. $\omega_{p,t+1}$ is demeaned gamma or Gaussian

6. $\omega_{n,t+1}$ is demeaned gamma or Gaussian

Variations 1 and 2 impose different degrees of dependence between the good and bad components of u_{t+1}, and the shocks to the risk factors, $\nu_{p,t+1}$ and $\nu_{n,t+1}$. The partial correlation model involves the additional λ parameter as in equation (12). One variation in 3 and 4 restricts the good and/or bad variance risk factors to be constant. For instance, p_t being constant imposes $\rho_p = \sigma_{pp} = 0$, reducing the number of parameters, but also reducing the flexibility of the model, potentially to the detriment of the model fit of the data.

Variations 5 and 6 potentially replace the gamma distribution with the Gaussian distribution for $\omega_{p,t+1}$ and/or $\omega_{n,t+1}$. The Gaussian distribution requires one fewer parameter relative to the gamma distribution, but the Gaussian distribution cannot accommodate conditional skewness or other higher-order moments. Because the $\nu_{p,t+1}$ and $\nu_{n,t+1}$ shocks drive positive volatility processes, they are always assumed to follow a demeaned gamma distribution. Under mild restrictions on the parameter space, this specification keeps the volatility everywhere positive (see Gourieroux and Jasiak, 2006).

We use the small sample corrected Akaike information criterion (AICc) as the basis for model selection. In all, we estimate 64 different models for both supply and demand shocks, with full results reported in Table 3. The AIC criterion does indeed indicate that a relatively parsimonious specification is optimal. When we estimate the full model, the λ parameter is not precisely identified for either demand or supply shocks and the full model does not rank very high on the AIC criterion relative to other models. In fact, the fully flexible specification is dominated by more parsimonious ones, for both supply and demand.

For the supply shock, the optimal AICc model is:

$$\begin{aligned}
u_t^s &= \sigma_{sp}\omega_{p,t}^s - \sigma_{sn}\omega_{n,t}^s, \\
p_t^s &= \bar{p}^s(1-\rho_p^s) + \rho_p^s p_{t-1}^s + \sigma_{pp}^s \nu_{p,t}^s, \\
n_t^s &= \bar{n}^s(1-\rho_n^s) + \rho_n^s n_{t-1}^s + \sigma_{nn}^s \omega_{n,t}^s, \\
\omega_{p,t+1}^s &\sim \mathcal{N}(0, p_t^s), \\
\nu_{p,t+1}^s &\sim \mathcal{N}(0, p_t^s), \\
\omega_{n,t+1}^s &\sim \tilde{\Gamma}(n_t^s, 1).
\end{aligned} \tag{14}$$

One important finding is that the data support two volatility factors for supply shocks. That is, all of the models that use only one volatility factor are dominated by the two-factor model in equation (14) in which good and bad variance both evolve over time (as opposed to being constant). In addition, the "good" environment component of supply shocks, $\omega_{p,t}^s$, is well-modeled using a Gaussian distribution as opposed to a gamma distribution (the latter requires an additional parameter). Also, the good component of the supply shocks, $\omega_{p,t}^s$, and the good variance shock for supply, $\nu_{p,t}^s$, are independent under the chosen specification. Finally, the bad component of supply shocks is gamma-distributed, and a single shock, $\omega_{n,t}^s$, affects both the level of the supply shock and the shock to bad variance. This generates negative correlation (but not perfect negative correlation) between the overall supply shock and the bad variance shock.

For demand shocks, the optimal specification under AICc is:

$$\begin{aligned}
u_t^d &= \sigma_{dp}\omega_{p,t}^d - \sigma_{dn}\omega_{n,t}^d, \\
p_t^d &= \bar{p}^d(1-\rho_p^d) + \rho_p^d p_{t-1}^d + \sigma_{pp}^d \omega_{p,t}^d, \\
n_t^d &= \bar{n}^d(1-\rho_n^d) + \rho_n^d n_{t-1}^d + \sigma_{nn}^d \omega_{n,t}^d, \\
\omega_{p,t+1}^d &\sim \mathcal{N}(0, p_t^d), \\
\omega_{n,t+1}^d &\sim \tilde{\Gamma}(n_t^d, 1).
\end{aligned} \tag{15}$$

We see that for demand shocks, as with supply shocks, the AICc selects a specification with two volatility factors. As was the case for supply shocks, the good component of demand shocks is distributed as Gaussian and the bad component is gamma-distributed. Moreover, for demand, as for supply, the same shock, $\omega_{n,t}^d$, affects both the level of u_t^d as well as the bad variance, n_t^d. The only difference between the specification chosen for demand shocks versus supply shocks is that AIC does not select independent variation between the good component of the shock to demand and the level of good variance. That is, for demand $\omega_{p,t}^d$ is a common shock for both u_t^d and p_t^d.

Because for both supply and demand shocks AICc selects a Gaussian "good" component, which of course has zero skewness, we refer to these components below as the Gaussian volatility or Gaussian component rather than "good." Of course, both supply and demand shocks also feature a negatively skewed bad component of volatility that varies over time as well, which we continue to refer to as bad volatility.

We also compare our BEGE specifications against regime-switching models of the Hamilton (1989)-type with 2-4 regimes. Based on AICc, the best models for both demand and supply shocks are 2 state regime-switching models. For the demand shock AICc is 507.4, and for the supply shock AICc is 494.7. Both are larger than the best BEGE specifications in Table 3, indicating a worse model fit. The GJR-GARCH model of Glosten, Jagannathan, and Runkle (1993) results in AICc:s of 509.1 and 495.4 for demand and supply shocks, respectively, both larger than the best BEGE specifications.

3.4 BEGE Estimation

The parameter estimates for the BEGE model are reported in Table 4.[9] The parameter σ_p represents the unconditional volatility of the supply and demand shocks due to the Gaussian component. This parameter is similar across supply and demand shocks (0.84 for supply; 0.95 for demand). Because the unconditional variance of the shocks is one under our normalization assumption, most of the variation in both series is due to the Gaussian component (of course, during adverse times, the "bad" component can, and does, dominate). As discussed, p_t follows a Gaussian stochastic volatility model for both supply and demand shocks.[10] For demand, this variable is very persistent with an autocorrelation parameter of about 0.98 and a very low innovation standard deviation (σ_{pp} is 0.08). The properties of the gamma-distributed bad environment state variable for demand shocks, n_t^d, contrasts sharply with those of p_t^d. First, its mean is only 0.04, implying that its unconditional skewness is about 10. This generates substantial negative skewness for demand shocks. The bad environment shape parameter is also less persistent than the good environment variable, with an autocorrelation of only 0.50. Therefore, n_t captures short-lived periods of risk characterized by potentially deeply negative shocks.

The BEGE parameter estimates for supply shocks are broadly similar to those for demand shocks. The Gaussian component has very high persistence (0.99); however, the volatility of the volatility shocks (σ_{pp}) is much larger at 0.31 compared to 0.08 for demand shocks. Taken together, this implies that the variance of the variance is much larger for the Gaussian component of supply shocks than it is for the Gaussian component of demand shocks. The supply bad-environment distribution is substantially non-Gaussian with the unconditional mean of the shape parameter equal to 2.45. This implies unconditional skewness of -1.28. The bad environment risk factor for supply has somewhat

[9]The bootstrapped standard errors appear unrealistically high for $\bar{n}$ for both supply and demand shocks. This is due to a number of outlier runs in which $\bar{n}$ was poorly identified either because the CMD step did not correctly identify supply versus demand shocks, or the Bates estimation got stuck in an unidentified region for $\bar{n}$.

[10]Note that because the selected processes for p_t are Gaussian for both supply and demand shocks, the parameter $\bar{p}$, the unconditional value of the shape parameter, is not identified. Without loss of generality, we set $\bar{p}$=1. This implies that, for example, σ_{dp}, represents the unconditional standard deviation of the Gaussian shock component for demand shocks.

higher persistence than for demand shocks, suggesting that supply driven recessions may have greater duration than demand driven recessions. The volatility shock parameter (σ_{nn}) is 0.57, which is considerably larger than the corresponding coefficient for demand shocks, implying that bad environment supply variances are also more variable than bad environment demand variances.

The Bates estimation also yields filtered estimates for the shocks to the Gaussian and bad components for supply and demand. The components of the demand shocks are shown in the middle and bottom panels of Figure 2. The Gaussian shock is relatively variable early in the sample, but less so later on. The bad demand shock is mostly near zero, but spikes down during some recessions. Reminiscent of a "jump" shock, the extreme non-Gaussianity of this shock is clearly evident. An analogous decomposition for supply shocks is shown in the bottom three panels of Figure 3. Here too, the variance of the Gaussian shock decreases notably in the late 1980s, and remains near zero thereafter. In contrast, the bad environment shock retains its variability throughout the sample, showing downward sharp spikes during most recessions. Finally, the pure volatility shock for supply, which is independent of the level shocks, is also mostly variable in the seventies, and its variance becomes minuscule starting in the 1990s. Hence, pure uncertainty shocks appear to no longer play a large role in driving supply shock volatility. This implies that level and uncertainty shocks are likely quite correlated, contradicting the independence assumption often maintained in the literature (e.g., Alessandri and Mumtaz, 2019). As a robustness check, we also verify that our key results on the Great Moderation in section 4 continue to hold using the second best specifications for demand and supply shock dynamics in Table 3.

4 Characterizing the History of Macroeconomic Volatility in the US

Having used the AIC to select optimal two-factor models of volatility for demand and supply shocks, and having estimated the optimal BEGE model parameters, we can now use the BEGE model as a lens to interpret the history of U.S. macroeconomic uncertainty over the sample period.

4.1 Time series estimates of uncertainty

The Bates estimation procedure allows us to filter time series estimates of the risk factors governing Gaussian and bad variances for supply and demand shocks. These are plotted in Figure 5. Starting with the demand variances, in Panels A and B, the Gaussian component of demand variance was relatively high in the 70s and the early 80s, and subsequently decreased to lower levels. The bad demand variance shows much less pronounced low frequency variation but increases in most recessions after 1980. Figure 5, Panels C and D, performs the same exercise for supply variances. The good variance level is elevated through the mid-1980s after which it trends down. The bad supply variance increases in most recessions throughout the sample period. Figure 5, Panel E, graphs both demand and supply variances. The conditional variance of supply shocks is largest in the early part of the same, dominating the conditional variance of demand shocks, consistent with stagflation incidences during that period. In the second half of the sample, the supply and demand conditional variances seem often indistinguishable, but the conditional demand variance peaks more sharply in the last 2 recessions.

The Gaussian and bad components of supply and demand variances map linearly into the conditional variances of inflation and real GDP growth. The latter are graphed in the top two panels of Figure 6. Both GDP growth and inflation variances were relatively high in the early part of the sample, and both trended down dramatically through the 1980s. For both inflation and real GDP growth, variance continued to spike up during recessions through the end of the sample. As shown in the top two panels to the right, the secular decline in volatility for both real GDP and inflation owes largely to a decline in the Gaussian component of supply variance. A reduction in the Gaussian component of demand variance also contributed to the decline in inflation variability in the 1980s. Spikes in volatility during recessions owe to the bad variance components of both supply and demand.

From equation (13), it is evident that in an environment where demand (supply) variances dominate, the conditional covariance between inflation and real activity is positive (negative). The bottom panels of Figure 6 graph the conditional covariance between inflation and real GDP growth shocks, and its components. The covariance is mostly very negative up until around 1990 due to the high variance of the Gaussian supply shock, but afterwards, this source of variation dramatically declines, resulting in a covariance closer to zero with short dips into positive territory during demand-driven recessions. We also show the good and bad supply and demand covariance components of the total covariance. For example, the near-zero correlation between real GDP and inflation during the 2000-2001 and 2007-2008 recessions reflects increases in both bad supply and bad demand variance, with offsetting effects on the covariance.

Figure 7 plots the conditional correlations between shocks to the level and uncertainty for supply and demand. That is, for supply shocks, we graph the correlation between u^s_{t+1} and $Var_{t+1}[u^s_{t+2}]$ (see equation (13)), and analogously for demand shocks. The top panel illustrates that for supply, the correlation between shocks to the level and total uncertainty is always negative. This is because the optimal specification has independence between shocks to good variance and the level, while bad variance shocks are negatively correlated with the level shocks. When bad volatility rises (relative to good volatility), the correlation becomes more negative. As shown in the bottom panel, for demand, shocks to good volatility are positively correlated with level shocks whereas shocks to bad volatility are negatively correlated with the level. When good volatility is large (relative to bad volatility) the correlation between level shocks and total uncertainty is positive, and vice versa. There is substantial time-variation as well as sign-switching in this correlation. This finding casts doubt on models that, ex ante, impose independence between level and uncertainty shocks.

4.2 The Great Moderation

Our estimated time series for volatilities can contribute to the debate on the Great Moderation. The literature has mostly focused on overall output volatility and puts a "break point" for output volatility in the first quarter of 1984 (see McConnell and Perez-Quiros, 2000; Stock et al., 2002). For inflation, Baele et al. (2015) suggest a later date, the first quarter of 1990. Whereas most of the literature attributes the decreased volatility to either good luck or improvements in monetary policy (see e.g. Cogley and Sargent, 2005, Benati and Surico, 2009, Sims and Zha, 2006, and Baele et al., 2015, and the references therein), we decompose the overall changes in volatility into changes in demand versus supply variances and bad versus Gaussian variability. We also address a more recent question asking whether the Great Moderation is over. Baele et. al. (2015) suggest that the Great Moderation for both inflation and output has ended, even before (for inflation) or just with the onset (for output) of the Great Recession. In contrast, Gadea, Gomez-Loscos and Perez-Quiros (2015) argue that the Great Moderation is alive and well, despite the Great Recession experience.

To test these various hypotheses, Table 5 reports simple dummy variable regressions where the dependent variables are the estimated conditional variances for inflation and GDP growth, as well as their AS versus AD and Gaussian versus bad components. The columns report the constant and the coefficients for two dummy variables. The first dummy variable is equal to 1 in the post-1985 (for GDP growth) or post-1990 (for inflation) part of the sample and is designed to identify changes in volatility associated with the onset of the Great Moderation. The second dummy is equal to 1 after 2006 and is designed to capture any possible reversal of the Great Moderation in the period beginning with the 2007-2008 Great Recession.

The results for the inflation variance are shown in the top panel. The overall level of the variance declined by about $\frac{3}{4}$ of its prior level in the post-1990 period, consistent with the Great Moderation. Looking at the components of inflation variance, the Gaussian components of both supply and demand shocks exhibit substantial declines as well. Notably, bad supply and demand variances exhibit no such decline. Thus, for inflation, the Great Moderation reflects a decline in relatively benign Gaussian volatility, but no decline due to the more pernicious bad volatility, which is strongly associated with recessions. Turning to the third column, there is no statistical evidence of any change in the

Great Moderation for any volatility series in the post-2006 period. Thus, recessionary inflation (deflation) risk for AS (AD) driven recessions has not waned. These results are not inconsistent with the view that monetary policy changes played an important role in the Great Moderation but they do suggest that monetary policy had little effect on "bad" component risk (see Coibion and Gorodnichenko, 2011, for an elaborate discussion and evidence that monetary policy restored macroeconomic stability through its effect on the level of trend inflation).

The results for real GDP growth, shown in the bottom panel, tell a similar story. There is a dramatic decline in the overall level of volatility for real GDP growth and its Gaussian subcomponents in the post-1985 period, consistent with the Great Moderation. However, there is no evidence of a decline due to changes in bad demand volatility during the Great Moderation, nor is there any evidence that volatility changed in the post-2006 period. We do observe a small but significant decrease in the bad supply variance post 1985. Note that this decrease disappears, both economically and statistically, if the regression specification only accommodates one post-1985 dummy instead of featuring also a second dummy after the Great Recession. Thus, the decrease in the bad supply variance shown in Table 5 may simply result from the lack of large negative supply shocks between 1985 and 2007. Overall, we should not expect recessions to be less variable in the future than they were in the past, even though the Great Moderation appears to still apply for overall volatility.

Table 6 repeats the same exercise for skewness. The skewness for both inflation (Panel A) and real GDP growth (Panel B) decreased substantially in the past three decades. Interestingly, the unscaled (by the standard deviation to the third power) centered third moment has not changed for inflation and even slightly increased for real GDP growth. Thus, the decline in conditional skewness in both cases is due to the decline in the conditional variance, which, as documented in Table 5, is due to the decline of Gaussian demand and supply variances. The effect is stronger for GDP growth than for inflation; and GDP growth skewness further significantly decreases after 2007. This confirms intuition of the "deepening" of recessions explored in Jensen et al. (2020), who suggest to normalize the severity of the recession (the fall in GDP growth per unit of time) by the standard deviation.

4.3 The conditional distribution of growth and inflation

One novel feature of our model is that it provides estimates of the non-Gaussian conditional distribution of shocks to real growth and inflation. Figure 8 graphs the standard (scaled) conditional skewness coefficient for our macro variables. As shown in the top panel, the skewness of real GDP growth has become more negative over time. This owes to the downward trend in the Gaussian component of supply and demand variance, leaving the more pernicious and negatively skewed bad variance to dominate recessions. An important implication is that risk is in fact higher than before, notwithstanding the low volatility observed in normal periods. Despite being produced in a framework with only AS and AD shocks, these results are reminiscent of the "volatility paradox" generated in models with credit frictions (Adrian and Boyarchenko, 2012; Brunnermeier and Sannikov, 2014), where periods of low volatility of output growth may foreshadow future crises. Adrian, Boyarchenko and Giannone (2019) also show distinct negative conditional skewness of GDP growth, using quantile regressions. They show, consistent with our results, more variation in the left tail than in the right tail and explore how financial conditions affect "growth vulnerability." Jensen et al. (2021), focusing simply on GDP growth itself, show that skewness is lower over the 1984-2016 period than over the 1947-1984 period; and the ratio of downside over upside volatility higher.

As shown in the bottom panel, for inflation a similar picture emerges. Skewness was near zero for the early portion of the sample when the Gaussian variance components were high. In the latter portion of the sample, inflation skewness declines and occasionally moves down sharply, particularly during recessions when the bad variance spikes in the absence of a large Gaussian component. An important conclusion is that deflation risk is more pronounced than the risk of high positive inflation in recent recessions in the U.S.

The conditional distributions for supply and demand shocks, as governed by the estimated BEGE model, together with the loadings of reduced-form shocks onto the supply and demand shocks as in equation (2), imply a non-Gaussian bivariate conditional distribution for real GDP and inflation shocks that varies notably over time. To illustrate this, Figure 9 plots the conditional bivariate distribution for inflation and real GDP growth from four periods in our sample. Each panel shows the bivariate density, illustrated using iso-density contours. The total probability mass inside each contour is labeled in blue. In the upper left panel, we plot the distribution as of 1972Q4. This is an expansionary period according to the NBER. As discussed above, during expansionary periods in the 1970s the Gaussian components of both supply and demand shocks dominate the distributions. As a result, an ellipsoid distribution emerges, consistent with a nearly Gaussian bivariate relationship. In 1974Q4, a recession was underway, and as shown by the upper right panel, the distribution expands notably as both supply and demand volatility increased. Moreover, due to very high Gaussian supply variance, a more negative correlation between real GDP and inflation is evident. Still, the overall distribution appears to be mostly Gaussian. The bottom two plots show the distribution from two representative quarters in the 2000s. The lower left panel shows the distribution from the expansionary period 2006Q4. The distribution is narrower in all directions compared to the top two panels, consistent with the Great Moderation. However, it is also evidently less Gaussian, with a much less prevalent ellipsoid shape, and, in contrast with the upper panels, two distinct "bad" tails over low GDP growth outcomes: one with low growth and high inflation, and another with low growth and low inflation. This is a manifestation of the "bad" AS and AD risks being more prevalent even in an expansionary period. Finally, the lower right panel shows the distribution after the onset of the Great Recession in 2008Q4. Due to a surge in bad demand volatility, the distribution appears wide and highly non-Gaussian with a vastly expanded heavy tail towards outcomes characterized by low inflation and low growth, suggesting dominant AD risks.

5 Conclusion

In this article, we develop a new dynamic model for real GDP growth and inflation using forecast revisions from the SPF obviating any complex modeling of conditional mean macro-dynamics. We model the shocks using BEGE dynamics, which accommodates time-varying non-Gaussian features with good and bad volatility. We extract bad and good volatilities for aggregate demand and supply shocks. We find both aggregate demand and supply shocks to be negatively skewed and leptokurtic, but their "good" components are Gaussian. We use the model to provide several contributions to the literature on macroeconomic uncertainty.

First, we differentiate models with various degrees of correlation between level and volatility shocks, finding that in the best model volatility and level shocks are on average negatively correlated and the correlation has decreased over time. Second, we characterize the time-variation in these supply and demand macro risks and their resulting effect on the conditional variances of inflation and GDP growth. We show that the Great Moderation largely reflects secular and large declines in the Gaussian demand and supply variances. However, there is little evidence that the bad variances have decreased over time, and these variances almost invariably peak in recessions. Third, the conditional skewness of both GDP growth and inflation has decreased over time, heightening macro vulnerabilities and it appears that the prevailing macro risk with regard to price movements is one of deflation, not inflation.

Our work provides alternative measurement of macro uncertainty. In a leading paper in this literature, JLN, the uncertainty index is based on 114 different time series, comprising price and output indices, and some financial time series as well. While our uncertainty measures are based on only two macro series, they admit an easy economic interpretation of various types of uncertainty and are not contaminated by financial data.

Our model and findings contribute to both the empirical and theoretical literature on uncertainty and business cycles in various ways. For example, a number of articles (e.g., Caggiano, Castelnuovo and Groshenny, 2014, or Alessandri and Mumtaz, 2019) show that the effects of uncertainty shocks are much larger in recessions. This prompts researchers to favor models with regime dependent exposures to economic shocks, but our model incorporates such effects endogenously, as the relative importance of bad volatility varies over time.

Our parametric model may also prove useful in theoretical real business cycle models. As Fernández-Villaverde and Guerron-Quintana (2020) discuss, equilibrium models often generate overly small effects of uncertainty shocks. Accounting for the strong conditional non-Gaussianities in the macro data in a tractable fashion as our model does, can be helpful. While there are alternative models that may fit non-Gaussianities and time-varying volatilities in macro data (e.g., the rare disaster models in Gabaix, 2012, and Wachter, 2013), the case for the BEGE model was recently bolstered by Bakshi and Chabi-Yo (2012), Chabi-Yo and Liu (2020), and Chabi-Yo and Loudis (2020), who show that a tractable representative agent model with BEGE dynamics for the macro fundamentals outperforms alternative non-Gaussian models in fitting asset prices.

Our research comes at an opportune time with the COVID crisis representing a huge economic shock, accompanied by much economic uncertainty. Obviously, the efficacy of various policy responses may well depend on whether AS or AD shocks dominate and even which type of uncertainty dominates, if uncertainty has indeed a causal effect on business cycles, as asserted in the new business cycle literature. Bekaert, Engstrom and Ermolov (2020) find that an AS shock accounts for 57% of the hugely negative COVID shock in the second quarter of 2020. Although it is plausible that the COVID shock is fundamentally different from AS and AD shocks seen in the past, policy makers therefore operate in an economic environment with dramatic "bad" uncertainty of both demand and supply shocks.

References

Adrian, Tobias, and Nina Boyarchenko, 2012, Intermediary Leverage Cycles and Financial Stability, Working paper, International Monetary Fund and Federal Reserve Bank of New York.

Adrian, Tobias, Boyarchenko, Nina, and Domenico Giannone, 2019, Vulnerable Growth, *American Economic Review*, 109, pp. 1263-1289.

Alessandri, Piergiorgio, and Haroon Mumtaz, 2019, Financial Regimes and Uncertainty Shocks, *Journal of Monetary Economics*, 101, pp. 31-46.

Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?, *Journal of Monetary Economics*, 4, 2007, pp. 1163-1212.

Aruoba, Boragan, Diebold, Francis, Nalewaik, Jeremy, Schorfheide, Frank, and Dongho Song, 2016, Improving GDP Measurement: A Measurement-error Perspective, *Journal of Econometrics*, 191, pp. 384-397.

Baele, Lieven, Bekaert, Geert, Cho, Seonghoon, Inghelbrecht, Koen, and Antonio Moreno, Macroeconomic Regimes, *Journal of Monetary Economics*, 70, pp. 51-71.

Bakshi, Gurdip, and Fousseni Chabi-Yo, 2012, Variance Bounds on the Permanent and Transitory Components of Stochastic Discount Factors, *Journal of Financial Economics*, 105, pp. 191-208.

Bates, David, 2006, Maximum Likelihood Estimation of Latent Affine Processes, *Review of Financial Studies*, 19, pp. 909-965.

Bekaert, Geert, and Eric Engstrom, 2017, Asset Return Dynamics under Habits and Bad Environment-Good Environment Fundamentals, *Journal of Political Economy*, 125, pp. 713-760.

Bekaert, Geert, Engstrom, Eric, and Andrey Ermolov, 2015, Bad Environments, Good Environments: A Non-Gaussian Asymmetric Volatility Model, *Journal of Econometrics*, 186, pp. 258-275.

Bekaert, Geert, Engstrom, Eric, and Andrey Ermolov, 2020, Aggregate Demand and Aggregate Supply Effects of COVID-19: A Real-time Analysis, *Covid Economics*, 25, pp. 141-168.

Bekaert, Geert, Engstrom, Eric, and Andrey Ermolov, 2022, Identifying Aggregate Demand and Supply Shocks Using Sign Restrictions and Higher-Order Moments, Working paper, Columbia University, Federal Reserve Board of Governors, and Fordham University.

Benati, Luca, and Paulo Surico, 2009, VAR Analysis and the Great Moderation, *American Economic Review*, 99, pp. 1636-1652.

Bils, Mark, Peter Klenow, and Benjamin Malin, 2012, Reset Price Inflation and the Impact of Monetary Policy Shocks, *American Economic Review*, 102, pp. 2798-2825.

Blanchard, Olivier, 1989, A Traditional Interpretation of Macroeconomic Fluctuations, *American Economic Review*, 79, pp. 1146-1164.

Blanchard, Olivier, and Danny Quah, 1989, The Dynamic Effects of Aggregate Demand and Supply Disturbances, *American Economic Review*, 79, pp. 655-673.

Bloom, Nicholas, 2009, The Impact of Uncertainty Shocks, *Econometrica*, 77, pp. 623-685.

Bloom, Nicholas, Floetotto, Max, Jaimovich, Nir, Saporta-Eksten, Itay, and Stephen Terry, 2018, Really Uncertain Business Cycles, *Econometrica*, 86, pp. 1031-1065.

Bollerslev, Tim, Tauchen, George, and Hao Zhou, 2009, Expected Stock Returns and Variance Risk Premia, *Review of Financial Studies*, 22, pp. 4463-4492.

Brunnermeier, Markus, and Yuliy Sannikov, 2014, A Macroeconomic Model with a Financial Sector, *American Economic Review*, 104, pp. 379-421.

Caggiano, Giovanni, Castelnuovo, Efrem, and Nicolas Groshenny, 2014, Uncertainty Shocks and Unemployment Dynamics in U.S. Recessions, *Journal of Monetary Economics*, 67, pp. 78-92.

Canova, Fabio, and Gianni De Nicoló, 2002, Monetary Disturbances Matter for Business Fluctuations in the G-7, *Journal of Monetary Economics*, 49, pp. 1131-1159.

Carriero, Andrea, Clark, Todd, and Massimiliano Marcellino, 2018, Measuring Uncertainty and Its Impact on the Economy, *Review of Economics and Statistics*, 100, pp. 799-815.

Chabi-Yo, Fousseni, and Yan Liu, 2020, Maxing Out Entropy: A Conditioning Approach, Working Paper, University of Massachusetts Amherst and Purdue University.

Chabi-Yo, Fousseni, and Johnathan Loudis, 2020, A Decomposition of Conditional Risk Premia and Implications for Representative Agent Models, Working paper, University of Massachusetts Amherst and University of Notre Dame.

Cochrane, John, 1994, Shocks, *Journal of Monetary Economics*, 41, pp. 295-364.

Cogley, Timothy, and Thomas Sargent, 2005, The Conquest of US inflation: Learning and Robustness to Model Uncertainty, *Review of Economic Dynamics*, 8, pp. 528-563.

Coibion, Olivier, and Yuriy Gordnichenko, 2012, What Can Survey Forecasts Tell Us About Information Rigidities?, *Journal of Political Economy*, 120, pp. 116-159.

Coibion, Olivier, and Yuriy Gordnichenko, 2015, Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts, *American Economic Review*, 105, pp. 2644-2678.

Engle, Robert, and Victor, Ng, 1993, Measuring and Testing the Impact of News on Volatility, *Journal of Finance*, 48, pp. 1749-1778.

Evans, Martin, and Paul Wachtel, 1993, Were Price Changes during the Great Depression Anticipated?: Evidence from Nominal Interest Rates, *Journal of Monetary Economics*, 32, pp. 3-34.

Faust, Jon, 1998, The Robustness of Identified VAR Conclusions about Money, *Carnegie-Rochester Conference Series on Public Policy*, 49, pp. 207-244.

Fernández-Villaverde, Jesus, and Juan Rubio-Ramírez, 2013, Macroeconomics and Volatility: Data, Models, and Methods, in *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, Cambridge University Press.

Fernández-Villaverde, Jesus, Guerrón, Pablo, Kuster, Keith, and Juan Rubio-Ramírez, 2015, Fiscal Volatility Shocks and Economic Activity, *American Economic Review*, 105, pp. 3352-3384.

Fernández-Villaverde, Jesus, and Pablo Guerrón, 2020, Uncertainty Shocks and Business Cycle Research, *Review of Economic Dynamics*, 37, pp. S118-S146.

Gabaix, Xavier, 2012, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance, *The Quarterly Journal of Economics*, 127, pp. 645-700.

Gadea, Maria, Gomes-Loscos, Ana, and Gabriel Perez-Quiros, 2015, The Great Moderation in Historical Perspective: Is It That Great?, Working paper, Bank of Spain.

Gali, Jordi, 1992, How Well Does the IS-LM Model Fit Postwar U.S. Data?, *Quarterly Journal of Economics*, 107, pp. 709-738.

Ghysels, Eric, Horan, Casidhe, and Emanuel Moench, 2018, Forecasting Through the Rearview Mirror: Data Revisions and Bond Return Predictability, *Review of Financial Studies*, 31, pp. 678-714.

Glosten, Lawrence, Jagannathan, Ravi, and David Runkle, 1993, On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, *Journal of Finance*, 48, pp. 1779-1801.

Gorodnichenko, Yuriy, and Serena Ng, 2017, Level and Volatility Factors in Macroeconomic Data, *Journal of Monetary Economics*, 91, pp. 52-68.

Gourieroux, Christian, and Joann Jasiak, 2006, Autoregressive Gamma Processes, *Journal of Forecasting*, 25, pp. 129-152.

Hamilton, James, 1989, A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle, *Econometrica*, 57, pp. 357-384.

Heston, Steven, 1993, A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies*, 2, pp. 327-343.

Hodrick, Robert, and Edward Prescott, 1997, Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit, and Banking*, 29, pp. 1-16.

Ireland, Peter, 2011, A New Keynesian Perspective on the Great Recession, *Journal of Money, Credit, and Banking*, 43, pp. 31-54.

Jensen, Henrik, Petrella, Ivan, Søren Ravn, and Emiliano Santoro, 2020, Leverage and Deepening Business-Cycle Skewness, *American Economic Journal: Macroeconomics*, 12, pp. 245-281.

Jurado, Kyle, Sydney Ludvigson, and Serena Ng, 2015, Measuring Uncertainty, *American Economic Review*, 105, pp. 1177-1216.

Justiniano, Alejandro, and Giorgio Primiceri, 2008, The Time-varying Volatility of Macroeconomic Fluctuations, *American Economic Review*, 98, pp. 604-641.

Kilic, Mete, and Ivan Shalistovich, 2019, Good and Bad Variance Premia and Expected Returns, *Management Science*, 65, pp. 2445-2945.

Kozeniauskas, Nicholas, Orlik, Anna, and Laura Veldkamp, 2018, What Are Uncertainty Shocks?, *Journal of Monetary Economics*, 100, pp. 1-15.

Lanne, Markku, Meitz, Miika, and Pentti Saikkonen, 2017, Identification and Estimation of Non-Gaussian Structural Vector Autoregressions, *Journal of Econometrics*, 196, pp. 288-304.

Lanne, Markku, and Jani Luoto, 2021, GMM Estimation of Non-Gaussian Structural Vector Autoregression, *Journal of Business and Economic & Economic Statistics*, 39, pp. 69-81.

Lebow, David, and Jeremy Rudd, 2006, Inflation Measurement, Working paper, Board of Governors of the Federal Reserve System.

McConnell, Margaret, and Gabriel Perez-Quiros, 2000, Output Fluctuations in the United States: What Has Changed since the Early 1980's?, *American Economic Review*, 90, pp. 1464-1476.

Morley, James, and Jeremy Piger, 2012, The Asymmetric Business Cycle, *Review of Economics and Statistics*, 94, pp. 208-221.

Mulligan, Casey, 2012, The Redistribution Recession: How Labor Market Distortions Contracted the Economy, Oxford University Press.

Newey, Whitney K., and Kenneth D. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703-08.

Patton, Andrew, and Kevin Sheppard, 2015, Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility, *Review of Economics and Statistics*, 97, pp. 683-697.

Salgado, Sergio, Guvenen, Fatih, and Nicholas Bloom, 2019, Skewed Business Cycles, Working paper, University of Pennsylvania, University of Minnesota, and Stanford University.

Shapiro, Matthew, and Mark Watson, 1988, Sources of Business Cycle Fluctuations, *NBER Macroeconomics Annual*, 3, pp. 111-156.

Sichel, Daniel, 1993, Business Cycle Asymmetry: A Deeper Look, *Economic Inquiry*, 31, pp. 224-236.

Sims, Christopher, 1980, Macroeconomics and Reality, *Econometrica*, 48, pp. 1-48.

Sims, Christopher, and Tao Zha, 2006, Were There Regime Switches in U.S. Monetary Policy?, *American Economic Review*, 96, pp. 54-81.

Stock, James, and Mark Watson with comments by Gali, Jordi, and Robert Hall, 2002, Has the Business Cycle Changed and Why?, *NBER Macroeconomics Annual*, pp. 159-218.

Uhlig, Harald, 2005, What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure, *Journal of Monetary Economics*, 52, pp. 381-419.

Uhlig, Harald, 2017, Shocks, Sign Restrictions and Identification, Chapter 4 in Honoré, Bo, Pakes, Ariel, Piazzesi, Monika, and Larry Samuelson (eds.), *Advances in Economics and Econometrics: Eleventh World Congress* (Econometric Society Monographs), Cambridge University Press, pp. 95-127.

Wachter, Jessica, Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?, *Journal of Finance*, 68, pp. 987-1035.

Woodford, Michael, 2003, Interest and Prices: Foundations of a Theory of Monetary Policy, Princeton University Press.

Wooldridge, Jeffrey, 2002, Econometric Analysis of Cross Section and Panel Data, MIT Press.

Figure 1: Real GDP Growth and Inflation Shocks. Shocks are expressed as a percentage change at an annual rate. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

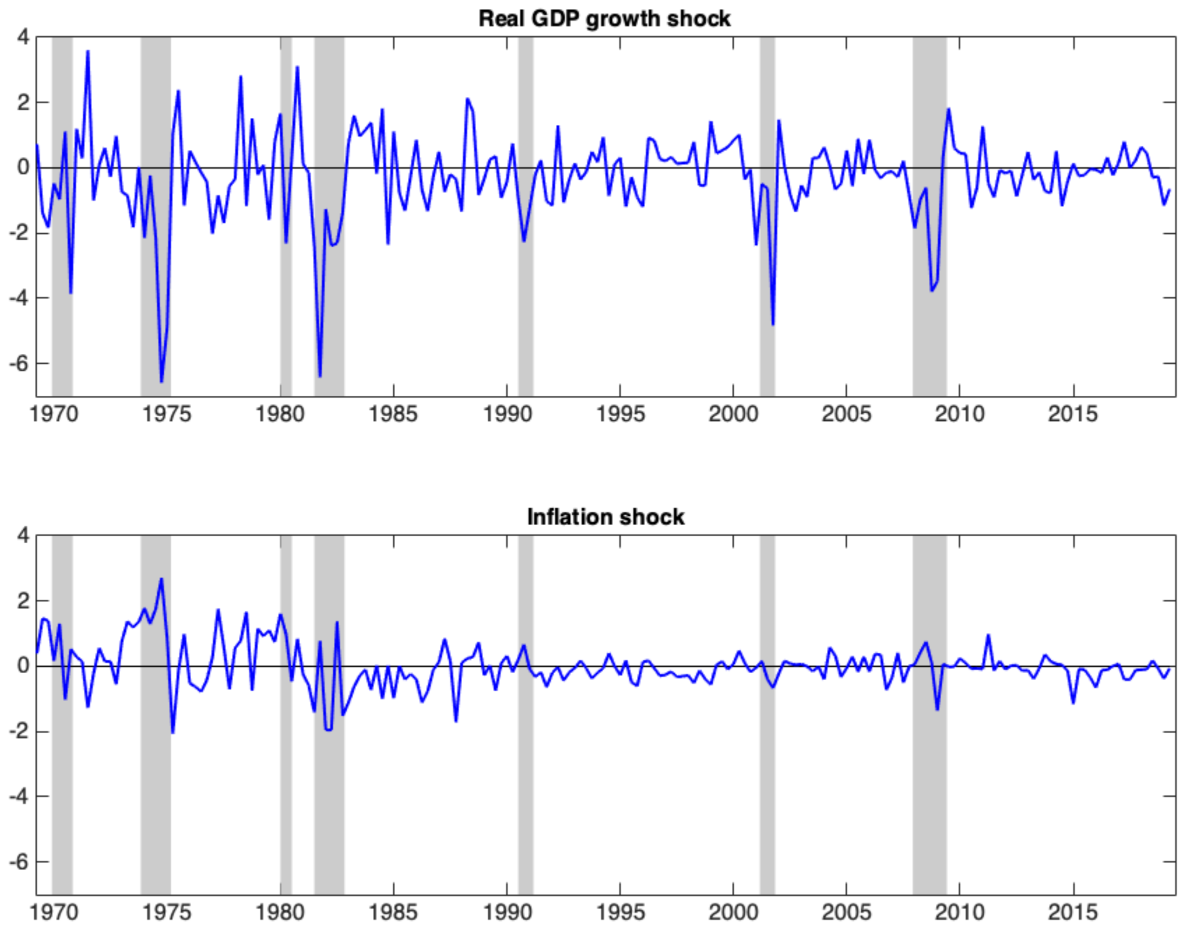

Figure 2: Demand Shock Decomposition. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions. Demand shocks, u_t^d, are extracted from the following system: $u_t^\pi = -\sigma_{\pi s} u_t^s + \sigma_{\pi d} u_t^d$ and $u_t^g = \sigma_{gs} u_t^s + \sigma_{gd} u_t^d$, where $\sigma_{\pi s}$, $\sigma_{\pi d}$, σ_{gs}, and σ_{gd} are positive constants, u_t^s supply shocks, and u_t^π and u_t^g are inflation and GDP growth shocks, respectively, extracted from survey forecast revisions as $u_t^\pi = E_t[\pi_t] - E_{t-1}[\pi_t]$ and $u_t^g = E_t[g_t] - E_{t-1}[g_t]$. Demand shock dynamics is $u_t^d = \sigma_{dp}\omega_{p,t}^d - \sigma_{dn}\omega_{n,t}^d$ with $\omega_{p,t}^d \sim \mathcal{N}(p_{t-1}^d)$ and $\omega_{n,t}^d \sim \Gamma(n_{t-1}^d, 1)$. Furthermore, $p_t^d = \bar{p}^d + \rho_p^d(p_{t-1}^d - \bar{p}^d) + \sigma_{pp}^d \omega_{p,t}^d$ and $n_t^d = \bar{n}^d + \rho_n^d(n_{t-1}^d - \bar{n}^d) + \sigma_{nn}^d \omega_{n,t}^d$. $\mathcal{N}(p_t)$ denotes a zero-mean Gaussian distribution with variance p_t. $\Gamma(n_t, 1)$ denotes a centered gamma distribution with shape parameter n_t and a unit scale parameter.

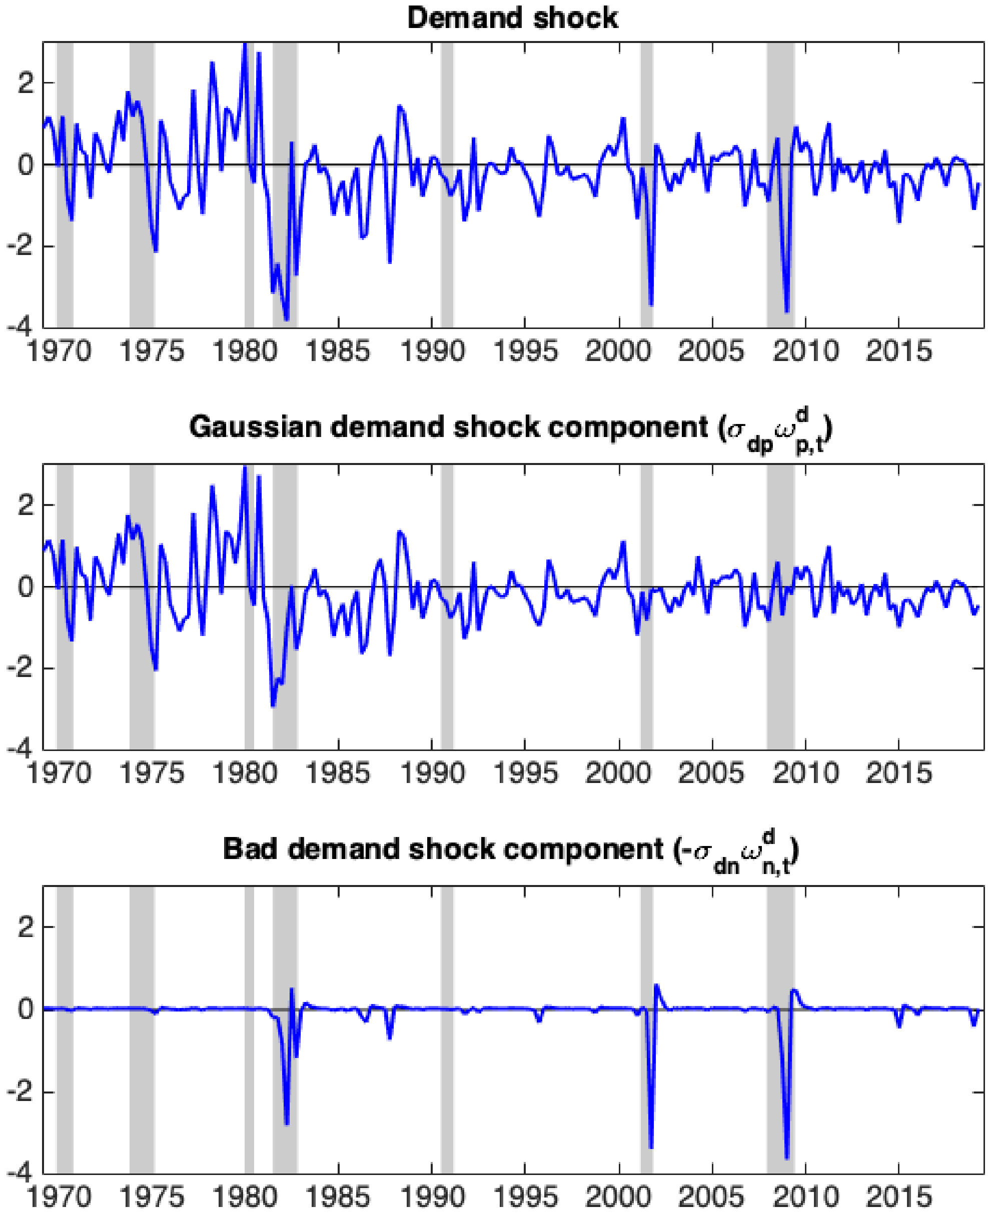

Figure 3: Supply Shock Decomposition. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions. Supply shocks, u_t^s, are extracted from the following system: $u_t^\pi = -\sigma_{\pi s} u_t^s + \sigma_{\pi d} u_t^d$ and $u_t^g = \sigma_{gs} u_t^s + \sigma_{gd} u_t^d$, where $\sigma_{\pi s}$, $\sigma_{\pi d}$, σ_{gs}, and σ_{gd} are positive constants, u_t^d supply shocks, and u_t^π and u_t^g are inflation and GDP growth shocks, respectively, extracted from survey forecast revisions as $u_t^\pi = E_t[\pi_t] - E_{t-1}[\pi_t]$ and $u_t^g = E_t[g_t] - E_{t-1}[g_t]$. Supply shock dynamics is $u_t^s = \sigma_{sp}\omega_{p,t}^s - \sigma_{sn}\omega_{n,t}^s$ with $\omega_{p,t}^s \sim \mathcal{N}(p_{t-1}^s)$ and $\omega_{n,t}^s \sim \Gamma(n_{t-1}^s, 1)$. Furthermore, $p_t^s = \bar{p}^s + \rho_p^s(p_{t-1}^s - \bar{p}^s) + \sigma_{pp}^s \nu_{p,t}^s$, $\nu_{p,t}^s \sim \Gamma(p_{t-1}^s)$, and $n_t^s = \bar{n}^s + \rho_n^s(n_{t-1}^s - \bar{n}^s) + \sigma_{nn}^s \omega_{n,t}^s$. $\mathcal{N}(p_t)$ denotes a zero-mean Gaussian distribution with variance p_t. $\Gamma(n_t, 1)$ denotes a centered gamma distribution with shape parameter n_t and a unit scale parameter.

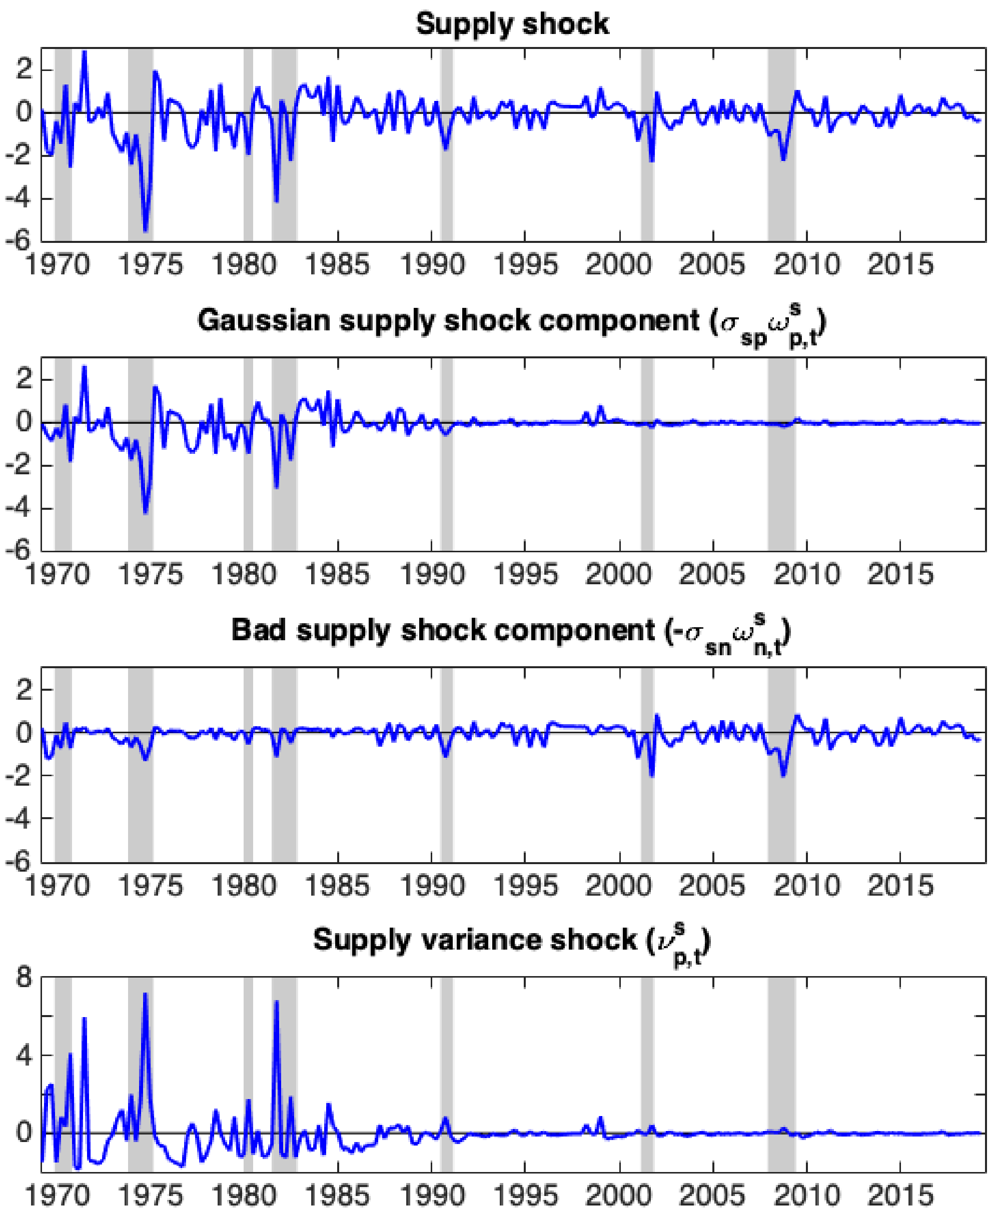

Figure 4: Components of Bad Environment - Good Environment Distribution.

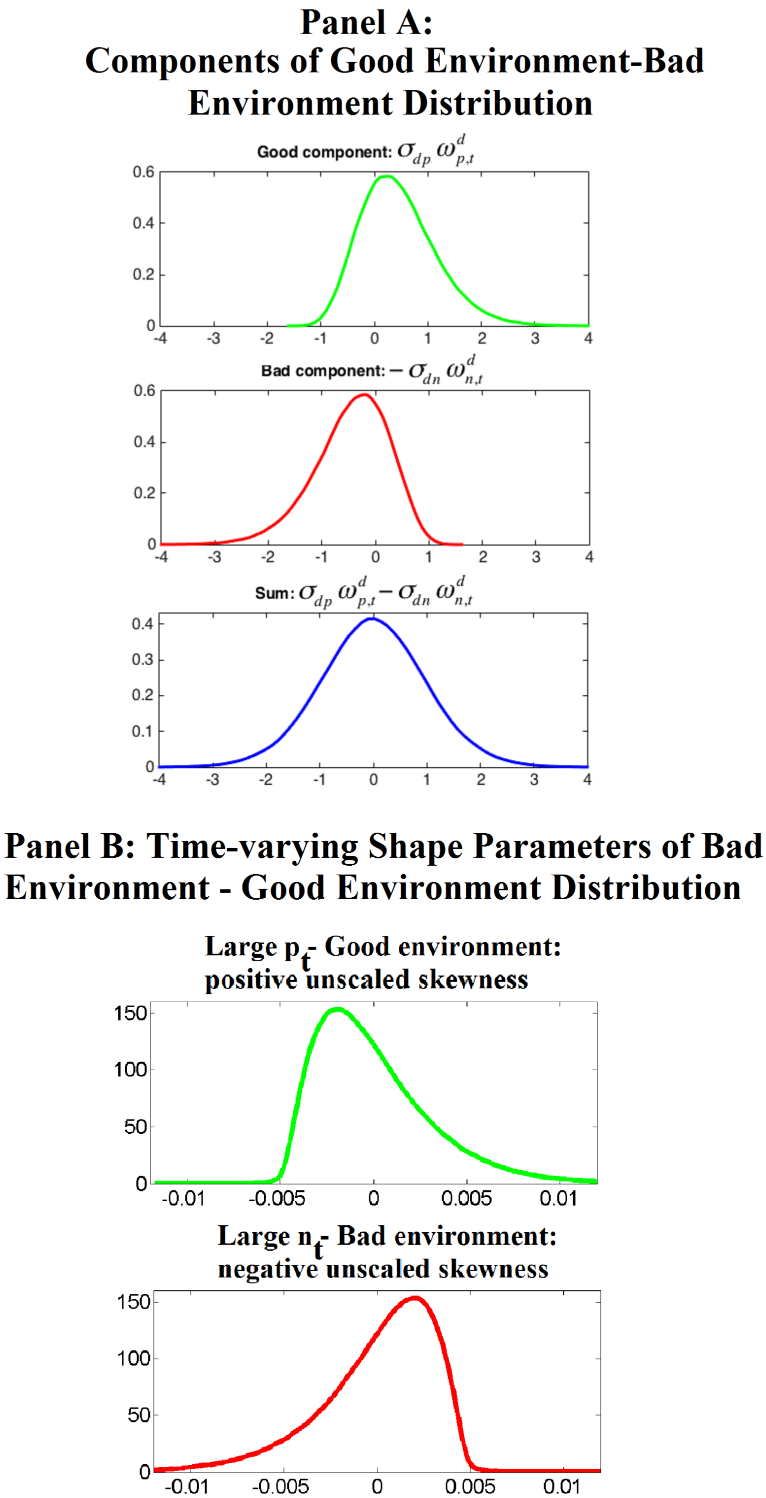

Figure 5: Conditional Demand and Supply Variances. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

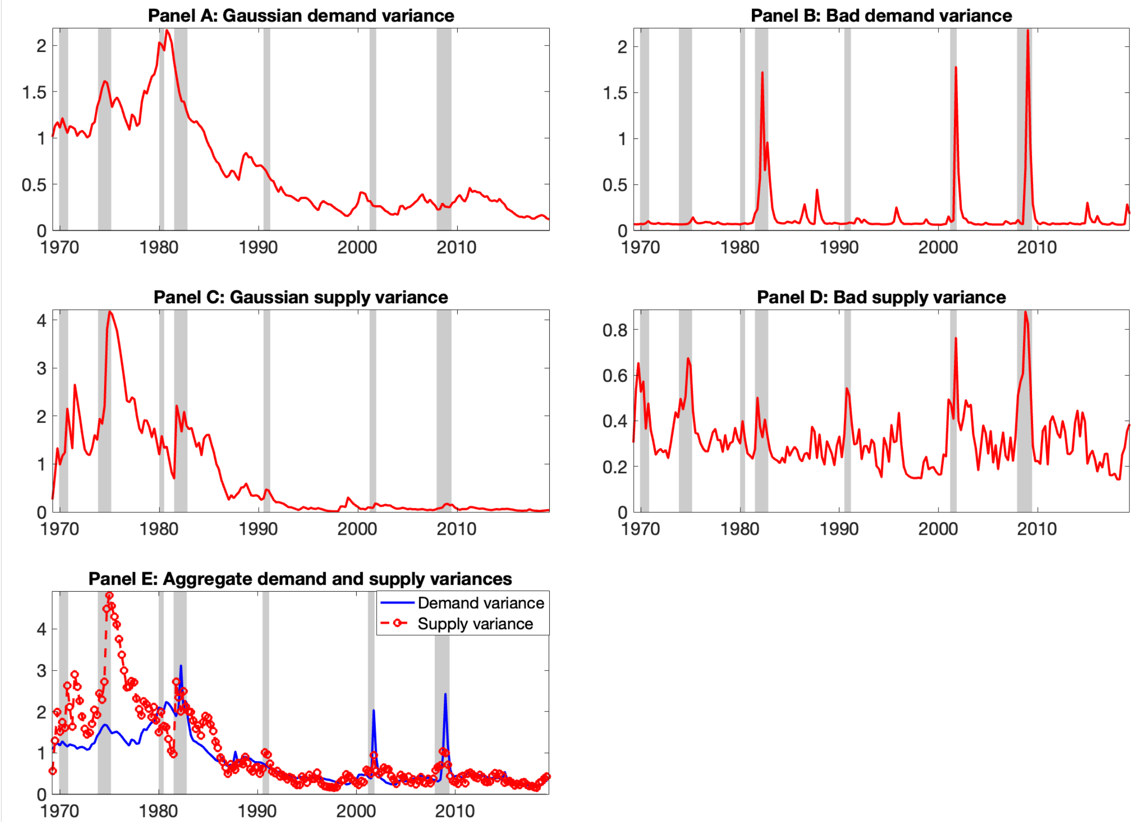

Figure 6: Conditional Second Moments of Real GDP Growth and Inflation. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

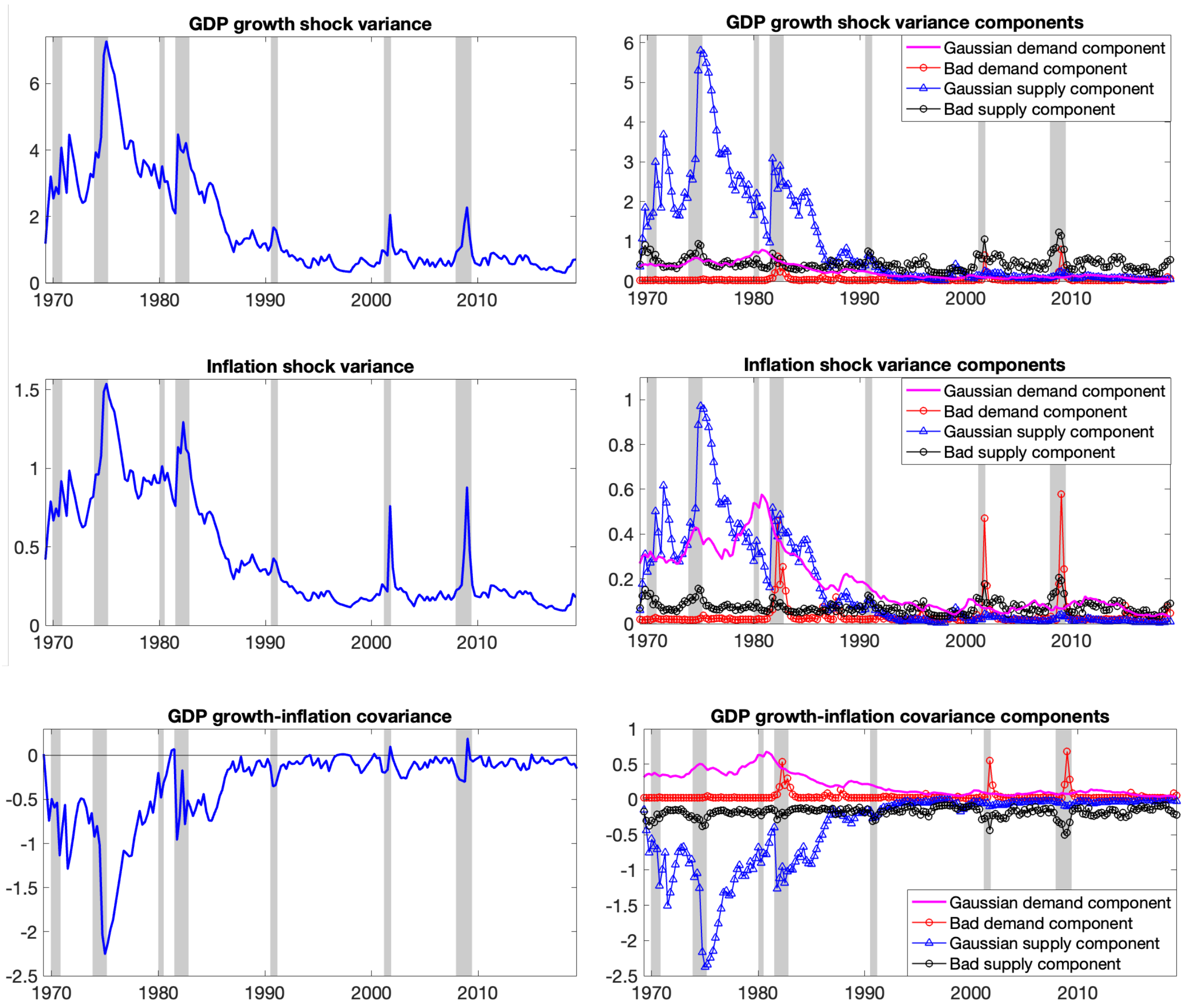

Figure 7: Conditional Correlations between Level and Variance Shocks. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

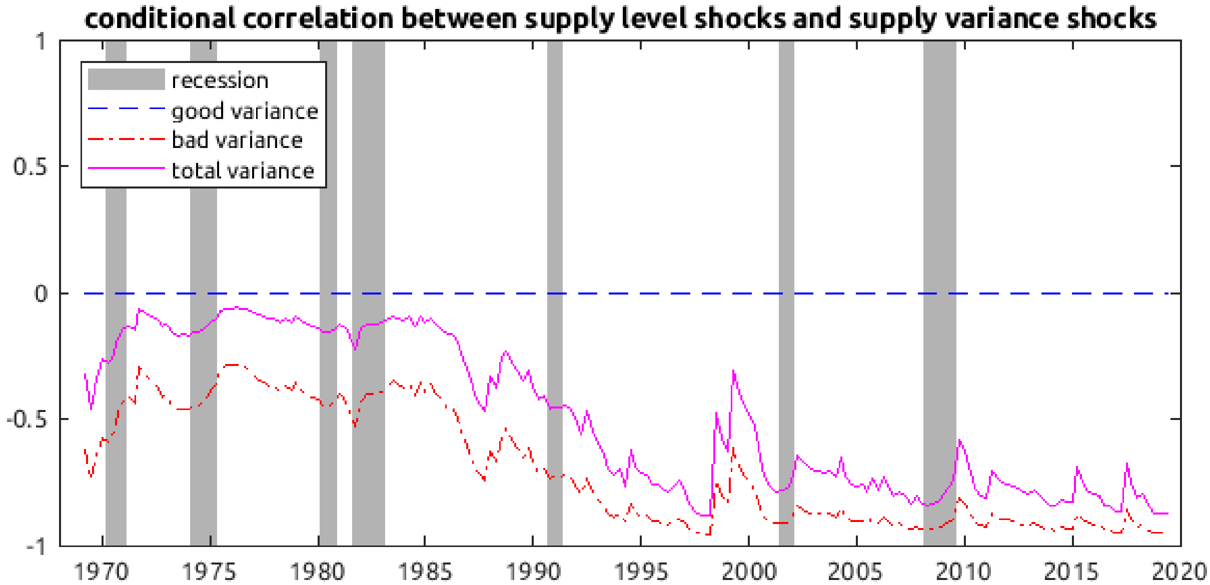

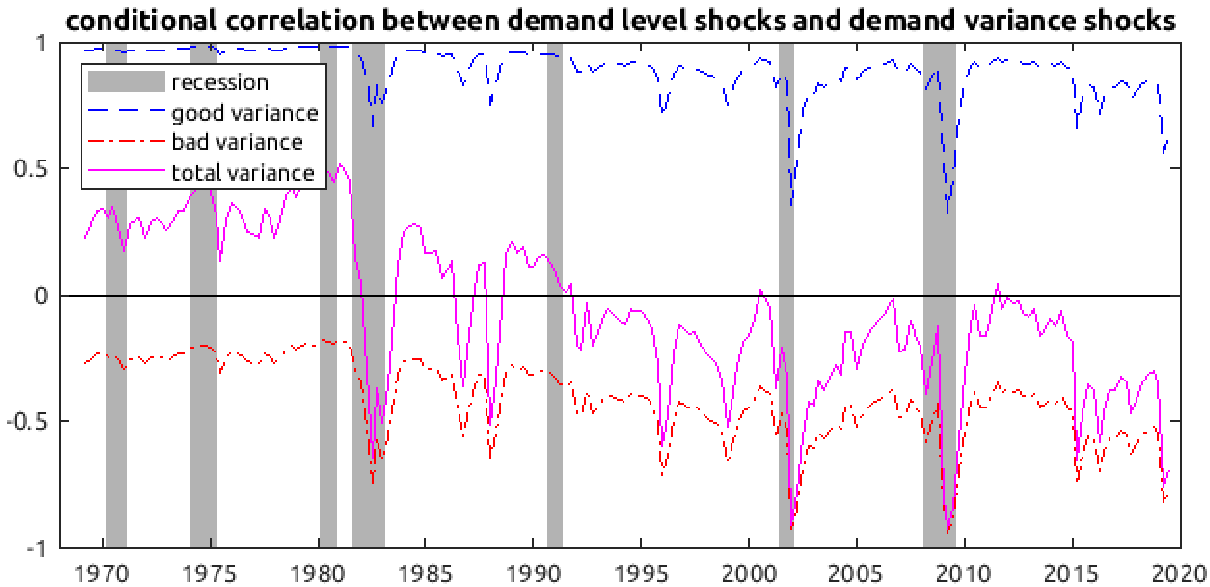

Figure 8: Conditional Third Moments of Real GDP Growth and Inflation. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

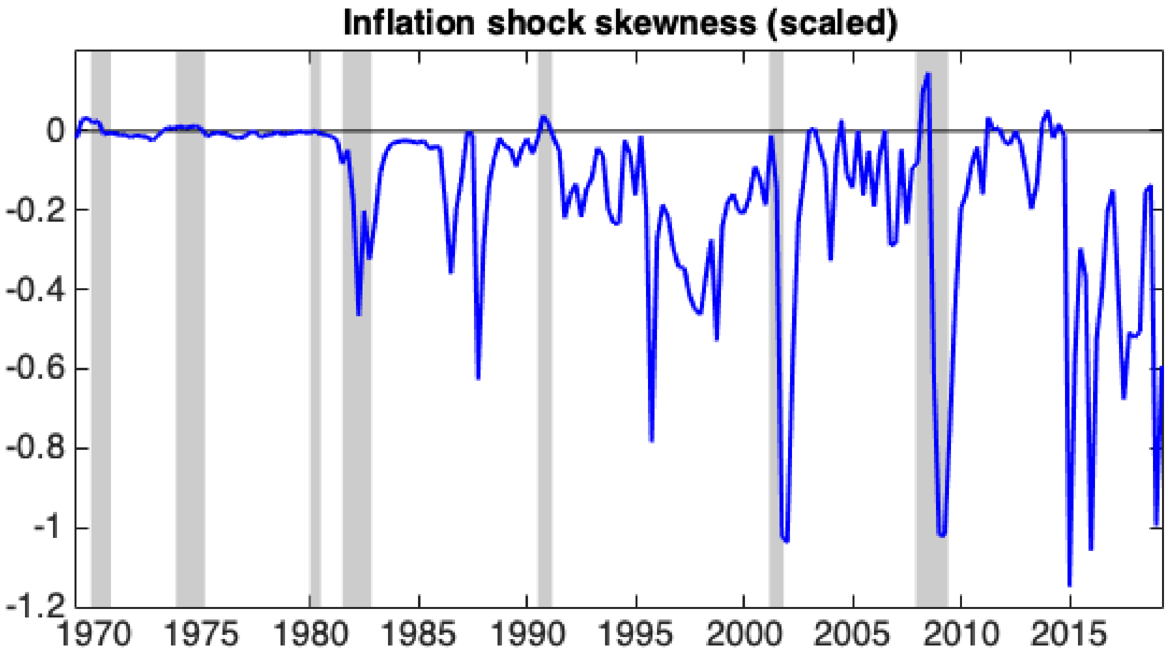

Figure 9: Conditional Contour Plots of Joint Real GDP Growth - Inflation Distributions. Numbers correspond to percentiles. Values are annualized. Plots are constructed by simulation.

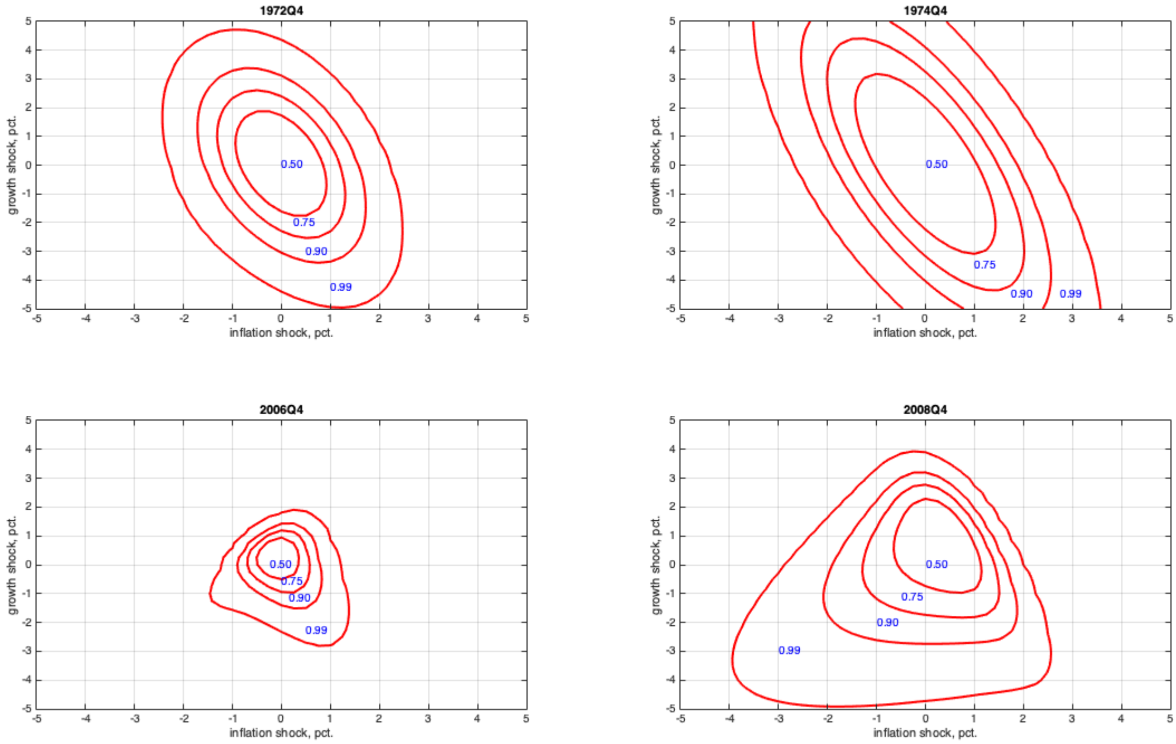

Table 1: Unconditional Moments of Macroeconomic Revisions: Classical Minimum Distance Fit. u_t^π and u_t^g are inflation and GDP growth shocks, respectively, extracted from survey forecast revisions as $u_t^\pi = E_t[\pi_t] - E_{t-1}[\pi_t]$ and $u_t^g = E_t[g_t] - E_{t-1}[g_t]$. The sample is quarterly 1968Q4-2019Q2. Asterisks *, **, and *** correspond to the statistical significance at the 10, 5, and 1 percent levels, respectively.

	Volatility		Correlation		
	u_t^π	u_t^g	$u_t^\pi u_t^g$		
Data	0.6361***	1.1885***	-0.1344		
Standard error	(0.0913)	(0.1448)	(0.1555)		
Fitted value	[0.7083]	[1.3295]	[-0.2776]		
	Skewness		Coskewness		
	u_t^π	u_t^g	$(u_t^\pi)^2 u_t^g$	$u_t^\pi (u_t^g)^2$	
Data	0.2005	-1.2343***	-0.7873***	0.4309	
Standard error	(0.3712)	(0.3890)	(0.2674)	(0.4884)	
Fitted value	[0.3663]	[-1.4465]	[-0.9808]	[0.4874]	
	Excess kurtosis		Excess cokurtosis		
	u_t^π	u_t^g	$(u_t^\pi)^2 (u_t^g)^2$	$(u_t^\pi)^3 u_t^g$	$u_t^\pi (u_t^g)^3$
Data	1.7280*	4.7138***	1.9239**	-0.5464	-1.6186
Standard error	(0.9813)	(1.3877)	(0.8979)	(1.1467)	(1.5647)
Fitted value	[1.7502]	[4.3216]	[2.6462]	[-1.7761]	[-3.2401]
Test for joint significance of 3rd and 4th order moments					
J-stat	25.3618				
p-value	0.26%***				
Overidentification test					
J-stat	2.9781				
p-value	38.74%				

Table 2: CMD Parameter Estimates. u_t^π and u_t^g are inflation and GDP growth shocks, respectively, extracted from survey forecast revisions as $u_t^\pi = E_t[\pi_t] - E_{t-1}[\pi_t]$ and $u_t^g = E_t[g_t] - E_{t-1}[g_t]$. u_t^d and u_t^s are demand and supply shocks, respectively. The system is $u_t^\pi = -\sigma_{\pi s} u_t^s + \sigma_{\pi d} u_t^d$ and $u_t^g = \sigma_{gs} u_t^s + \sigma_{gd} u_t^d$, where $\sigma_{\pi s}$, $\sigma_{\pi d}$, σ_{gs}, and σ_{gd} are positive constants. The sample is quarterly 1968Q4-2019Q2. Asymptotic standard errors are in parentheses.

Panel A: Loadings of Reduced-form Shocks onto Supply and Demand Shocks		
	u_t^π	u_t^g
u_t^s	-0.4829	1.1802
	(0.0566)	(0.1129)
u_t^d	0.5141	0.6035
	(0.0685)	(0.1064)
Panel B: Higher-order Moments of Supply and Demand Shocks		
	Skewness	Excess kurtosis
u_t^s	-1.9563	6.8535
	(0.3873)	(1.5692)
u_t^d	-0.6896	1.0062
	(0.5413)	(1.6825)
Co-excess kurtosis	-0.0095	
	(0.2843)	

Table 3: Demand and Supply Model Comparison. Models with the best (lowest) AICc criteria are shaded.

$\omega_{p,t}$			$\omega_{n,t}$			AICc-supply	AICc-demand
Distribution	Time-variation	λ	Distribution	Time-variation	λ		
Gaussian	constant		Gaussian	constant		575.4	574.2
Gaussian	constant		Gaussian	time-varying	0	494.6	508.8
Gaussian	constant		Gaussian	time-varying	1	519.4	567.2
Gaussian	constant		Gamma	time-varying	free	494.6	510.9
Gaussian	constant		Gamma	constant		538.9	556.0
Gaussian	constant		Gamma	time-varying	0	490.3	510.6
Gaussian	constant		Gamma	time-varying	1	507.8	550.1
Gaussian	constant		Gamma	time-varying	free	492.0	512.8
Gaussian	time-varying	0	Gaussian	constant		493.0	508.8
Gaussian	time-varying	0	Gaussian	time-varying	0	497.3	513.5
Gaussian	time-varying	0	Gaussian	time-varying	1	493.8	513.3
Gaussian	time-varying	0	Gamma	time-varying	free	498.2	515.5
Gaussian	time-varying	0	Gamma	constant		493.4	510.7
Gaussian	time-varying	0	Gamma	time-varying	0	497.6	513.6
Gaussian	time-varying	0	Gamma	time-varying	1	488.0	515.7
Gaussian	time-varying	0	Gamma	time-varying	free	496.5	515.8
Gaussian	time-varying	1	Gaussian	constant		512.2	524.0
Gaussian	time-varying	1	Gaussian	time-varying	0	497.2	504.3
Gaussian	time-varying	1	Gaussian	time-varying	1	523.4	526.2
Gaussian	time-varying	1	Gamma	time-varying	free	498.2	506.2
Gaussian	time-varying	1	Gamma	constant		501.7	504.2
Gaussian	time-varying	1	Gamma	time-varying	0	494.8	504.1
Gaussian	time-varying	1	Gamma	time-varying	1	494.3	503.2
Gaussian	time-varying	1	Gamma	time-varying	free	496.2	505.3
Gaussian	time-varying	free	Gaussian	constant		495.2	508.4
Gaussian	time-varying	free	Gaussian	time-varying	0	499.4	506.4
Gaussian	time-varying	free	Gaussian	time-varying	1	495.8	511.7
Gaussian	time-varying	free	Gamma	time-varying	free	500.4	508.3
Gaussian	time-varying	free	Gamma	constant		495.8	510.5
Gaussian	time-varying	free	Gamma	time-varying	0	496.9	506.2
Gaussian	time-varying	free	Gamma	time-varying	1	490.3	515.0
Gaussian	time-varying	free	Gamma	time-varying	free	498.7	507.4
Gamma	constant		Gaussian	constant		577.7	576.3
Gamma	constant		Gaussian	time-varying	0	493.4	511.2
Gamma	constant		Gaussian	time-varying	1	521.0	554.0
Gamma	constant		Gamma	time-varying	free	496.3	513.3
Gamma	constant		Gamma	constant		530.2	537.1
Gamma	constant		Gamma	time-varying	0	490.2	511.8
Gamma	constant		Gamma	time-varying	1	507.7	530.8
Gamma	constant		Gamma	time-varying	free	492.9	514.0
Gamma	time-varying	0	Gaussian	constant		505.2	525.3
Gamma	time-varying	0	Gaussian	time-varying	0	500.4	515.5
Gamma	time-varying	0	Gaussian	time-varying	1	500.8	530.1
Gamma	time-varying	0	Gamma	time-varying	free	501.4	517.8
Gamma	time-varying	0	Gamma	constant		503.7	520.1
Gamma	time-varying	0	Gamma	time-varying	0	497.4	516.0
Gamma	time-varying	0	Gamma	time-varying	1	490.4	524.2
Gamma	time-varying	0	Gamma	time-varying	free	499.3	518.1
Gamma	time-varying	1	Gaussian	constant		581.6	540.4
Gamma	time-varying	1	Gaussian	time-varying	0	500.3	508.4
Gamma	time-varying	1	Gaussian	time-varying	1	505.5	526.3
Gamma	time-varying	1	Gamma	time-varying	free	501.3	510.4
Gamma	time-varying	1	Gamma	constant		510.2	507.5
Gamma	time-varying	1	Gamma	time-varying	0	497.4	516.4
Gamma	time-varying	1	Gamma	time-varying	1	498.0	507.5
Gamma	time-varying	1	Gamma	time-varying	free	499.3	510.3
Gamma	time-varying	free	Gaussian	constant		507.3	523.6
Gamma	time-varying	free	Gaussian	time-varying	0	502.2	510.3
Gamma	time-varying	free	Gaussian	time-varying	1	502.3	525.1
Gamma	time-varying	free	Gamma	time-varying	free	503.1	513.0
Gamma	time-varying	free	Gamma	constant		500.9	514.6
Gamma	time-varying	free	Gamma	time-varying	0	499.4	510.3
Gamma	time-varying	free	Gamma	time-varying	1	492.9	510.3
Gamma	time-varying	free	Gamma	time-varying	free	501.3	511.0

Table 4: Demand and Supply Shock Dynamics Parameter Estimates. Asymptotic standard errors are in parentheses. Bootstrap standard errors controlling for uncertainty in both CMD and maximum likelihood steps are in square brackets. Note that $\bar{p}$ is missing, because "good" components of both demand and supply shocks are Gaussian.

	u_t^s	u_t^d
σ_p	0.8441	0.9515
	(0.0307)	(0.0427)
	[0.1272]	[0.0479]
ρ_p	0.9922	0.9847
	(0.0082)	(0.0066)
	[0.1159]	[0.1457]
σ_{pp}	0.3129	0.0806
	(0.1408)	(0.0177)
	[0.2309]	[0.1066]
σ_n	0.3423	1.5126
	(0.1085)	(1.6994)
	[0.4169]	[0.6686]
$\bar{n}$	2.4547	0.0413
	(1.3388)	(0.0714)
	[19.4438]	[2.1465]
ρ_n	0.5952	0.5044
	(0.2058)	(0.4283)
	[0.1982]	[0.1098]
σ_{nn}	0.5739	0.3242
	(0.3234)	(0.3219)
	[0.2119]	[0.1664]

Table 5: Decomposing the Great Moderation into Changes in Demand and Supply Variances. The sample is quarterly 1968Q4-2019Q2. Coefficients are OLS regression coefficients from regressing the dependent variable on a constant and both dummies. Standard errors in parentheses are Newey-West (1987) standard errors computed with 40 lags. The asterisks, ** and ***, correspond to statistical significance at the 5 and 1 percent levels, respectively.

Panel A: Aggregate Inflation			
Dependent variable	Constant	Dummy 1991Q1-	Dummy 2007Q1-
Aggregate variance	0.7753***	-0.5697***	0.0009
	(0.1153)	(0.1194)	(0.0338)
Supply variance	0.4268***	-0.3336***	-0.0022
	(0.0834)	(0.0850)	(0.0130)
Gaussian supply variance	0.3497***	-0.3245***	-0.0106
	(0.0789)	(0.0803)	(0.0074)
Bad supply variance	0.0771***	-0.0091	0.0084
	(0.0063)	(0.0079)	(0.0096)
Demand variance	0.3486***	-0.2361***	0.0031
	(0.0399)	(0.0424)	(0.0212)
Gaussian demand variance	0.3116***	-0.2313***	-0.0076
	(0.0383)	(0.0425)	(0.0122)
Bad demand variance	0.0369**	-0.0047	0.0107
	(0.0182)	(0.0425)	(0.0122)
Panel B: Real GDP Growth			
Dependent variable	Constant	Dummy 1985Q1-	Dummy 2007Q1-
Aggregate variance	3.6030***	-2.6769***	-0.2239
	(0.2486)	(0.3432)	(0.2161)
Supply variance	3.0572***	-2.3256***	-0.1882
	(0.2549)	(0.3333)	(0.1722)
Gaussian supply variance	2.5723***	-2.2438***	-0.2412
	(0.2423)	(0.3175)	(0.1494)
Bad supply variance	0.4850***	-0.0818**	0.0530
	(0.0357)	(0.0388)	(0.0559)
Demand variance	0.5458***	-0.3513***	-0.0356
	(0.0497)	(0.0525)	(0.0457)
Gaussian demand variance	0.4920***	-0.3413***	-0.0508
	(0.0347)	(0.0464)	(0.0380)
Bad demand variance	0.0538***	-0.0100	0.0151
	(0.0196)	(0.0205)	(0.0181)

Table 6: The Great Moderation and Macroeconomic Skewness. The sample is quarterly 1968Q4-2019Q2. Reported coefficients are OLS regression coefficients from regressing the dependent variable on a constant and both dummies. Standard errors in parentheses are Newey-West (1987) standard errors computed with 40 lags. The asterisks, *, **, and ***, correspond to statistical significance at the 10, 5, and 1 percent levels, respectively. Note that good demand and supply components are Gaussian and, thus, have 0 skewness.

Panel A: Aggregate Inflation			
Dependent variable	Constant	Dummy 1991:Q1-	Dummy 2007:Q1-
Skewness	-0.0513*	-0.1583***	-0.0610
	(0.0269)	(0.0442)	(0.0476)
Unscaled centered 3^{rd} moment	-0.0322*	0.0044	-0.0139
	(0.0170)	(0.0246)	(0.0179)
Supply component (bad)	0.0254**	-0.0030	0.0028
	(0.0021)	(0.0026)	(0.0032)
Demand component (bad)	-0.0576***	0.0074	0.0167
	(0.0154)	(0.0176)	(0.0210)
Panel B: Real GDP Growth			
Dependent variable	Constant	Dummy 1991:Q1-	Dummy 2007:Q1-
Skewness	-0.0826***	-0.4950***	-0.2469*
	(0.0106)	(0.1056)	(0.1330)
Unscaled centered 3^{rd} moment	-0.4897***	0.0843**	-0.0705
	(0.0210)	(0.0371)	(0.0758)
Supply component (bad)	-0.3914***	0.0660**	-0.0428
	(0.0288)	(0.0313)	(0.0451)
Demand component (bad)	-0.0984***	0.0182	-0.0277
	(0.0357)	(0.0375)	(0.0331)

Appendix I: Survey Revisions versus VAR shocks

Real GDP Growth and Inflation Shocks. Shocks are expressed as a percentage change at an annual rate. VAR shocks are shocks to realized GDP growth and inflation from VAR(1) with realized real GDP growth and inflation and their Survey of Professional Forecasters expectations as variables. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

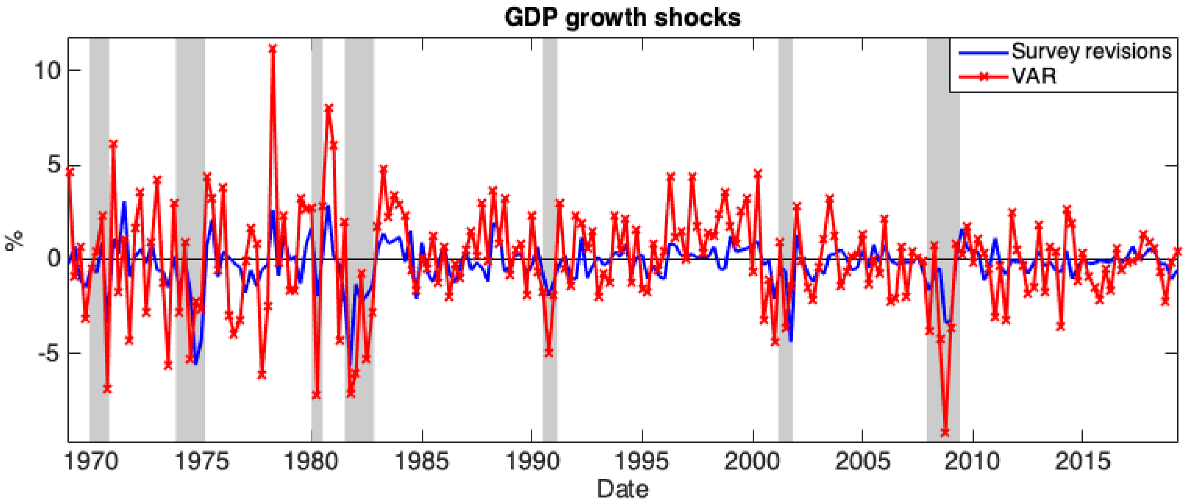

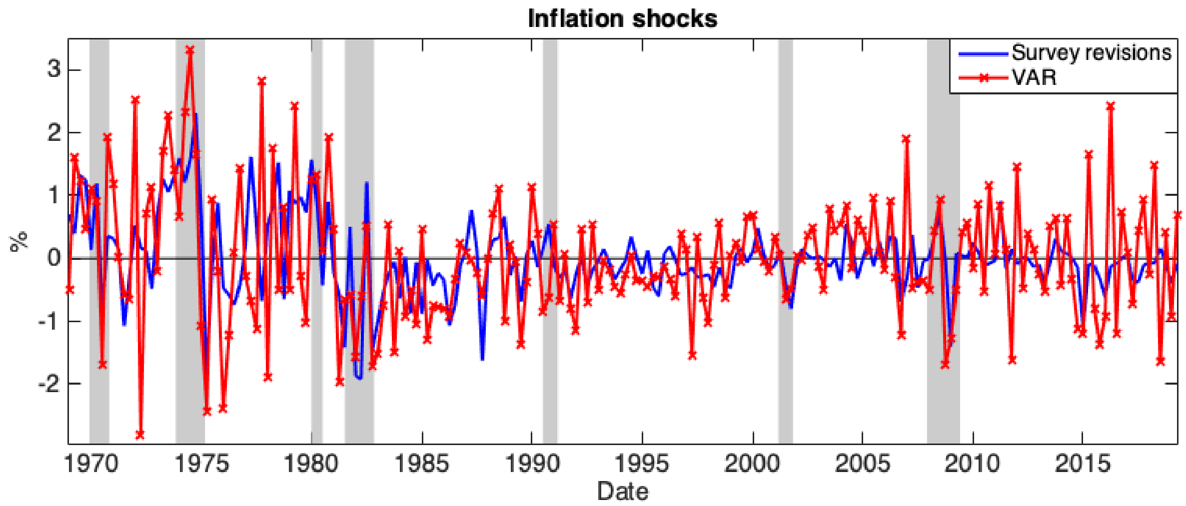

Demand and Supply Shocks. Shocks are inverted from real GDP growth and inflation shocks. VAR shocks are shocks to realized GDP growth and inflation from VAR(1) with realized real GDP growth and inflation and their Survey of Professional Forecasters expectations as variables. The sample is quarterly 1968:Q4-2019:Q2. Shading corresponds to NBER Recessions.

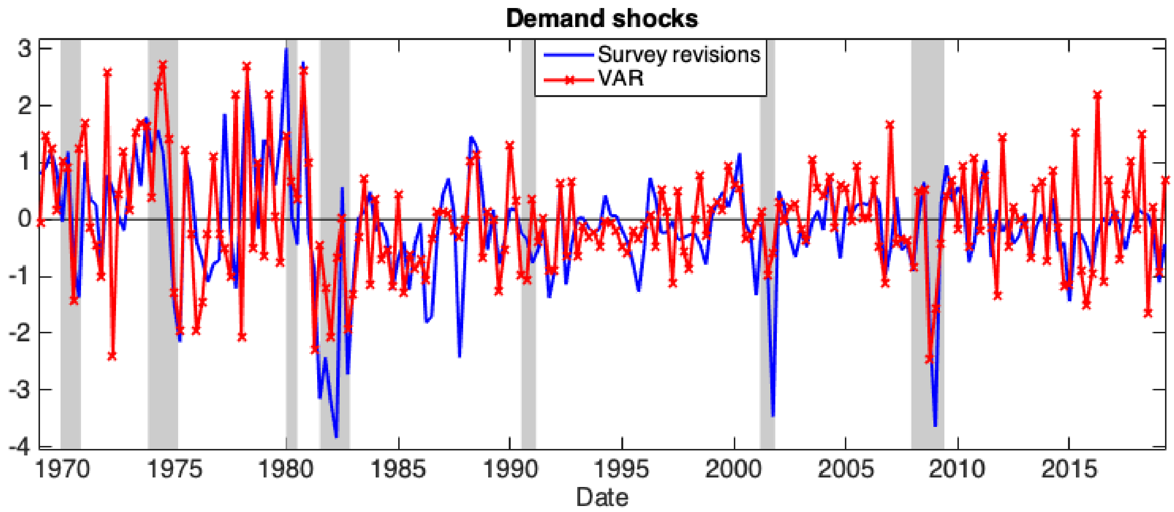

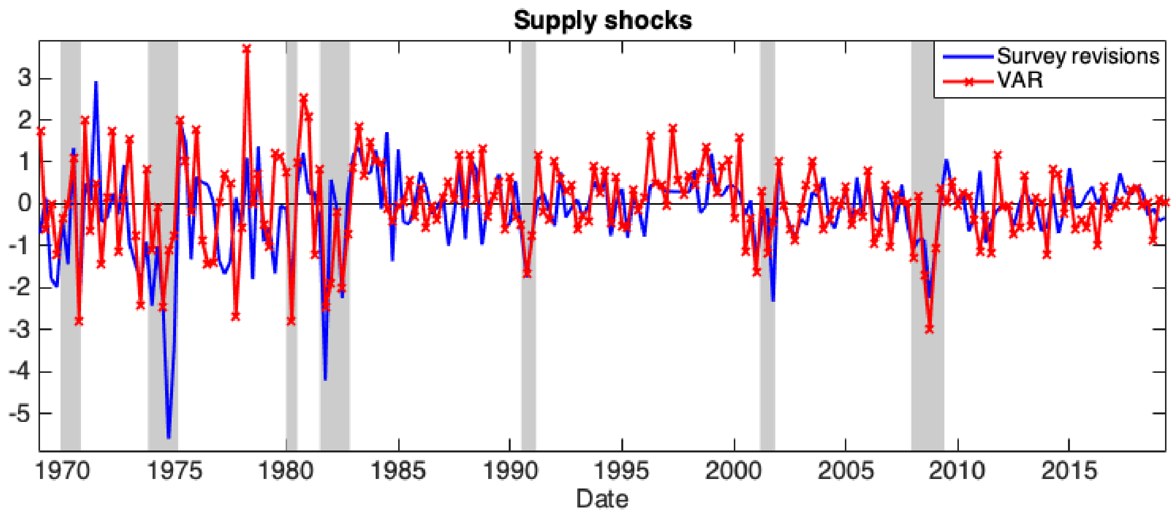

Appendix II: GMM Estimation of Demand and Supply Shocks

The CMD methodology employs statistics rather than moments, and here we verify that our results are robust to using standard GMM on regular moments of inflation and GDP growth. To this end, we use the following raw moments: $(u_t^\pi)^2$, $(u_t^g)^2$, $u_t^\pi u_t^g$, $(u_t^\pi)^3$, $(u_t^g)^3$, $(u_t^\pi)^4$, $(u_t^g)^4$, $(u_t^\pi)^2 u_t^g$, $u_t^\pi (u_t^g)^2$, $(u_t^\pi)^2 (u_t^g)^2$, $(u_t^\pi)^3 u_t^g$, and $u_t^\pi (u_t^g)^3$.

Of course, the statistics we use for the CMD procedure (volatilities, correlation, and scaled higher-order moments) are simple functions of these unscaled moments, so we expect results to be similar. We have the same 9 parameters to fit these 12 moments in a GMM system. Another difference is that we now can use a standard weighting matrix, based on the spectral density at frequency zero of the orthogonality conditions (instead of the bootstrap procedure used with the CMD methodology). To compute an initial GMM weighting matrix, we use our usual CMD parameters (from Table 2) as starting values, obtain residuals, and compute the spectral density. We use the optimal GMM weighting matrix using 30 Newey-West (1987) lags. Below are the results. In a nutshell, the loadings in Panel A are very similar to the ones reported in Table 2, confirming GDP growth shocks loading more on supply than on demand shocks. Moreover, the inverted supply and demand series are virtually indistinguishable, as the regressions in Panel B show.

GMM Estimation of Demand and Supply Shocks.

Panel A: Loadings of Reduced-form Shocks onto Supply and Demand Shocks			
	u_t^π	u_t^g	
u_t^s	-0.4022	1.0125	
	(0.0509)	(0.1647)	
u_t^d	0.4804	0.5461	
	(0.0471)	(0.0721)	
Panel B: OLS Regression Coefficients of CMD extracted shocks on GMM extracted shocks			
	Constant	Slope	R^2
u_t^s	0.0017	1.0000	0.9998
u_t^d	-0.0029	0.9999	0.9997

Appendix III: VAR Impulse Responses to Aggregate Demand and Aggregate Supply Shocks

Cumulative 20 Quarter VAR Impulse Responses of Real GDP and Aggregate Price Level to One Standard Deviation Demand and Supply Shocks. The data are 1968:Q4-2019:Q2 quarterly. The VAR model is: $Y_t = A_0 + A_1 Y_{t-1} + S[u_t^s, u_t^d]' + \epsilon_t$, where Y_t is the vector of final, revised real GDP growth and inflation, $[u_t^s, u_t^d]'$ are pre-estimated structural shocks from the SPF, and ϵ_t is a residual noise vector. The cumulative impulse responses include the quarter 0 (where the shocks happened) responses. Numbers in parentheses are probabilities that the impulse response is less than 0 obtained from 10,000 block-bootstrap samples of historical length with the block size of 8 quarters. The asterisks, ***, correspond to statistical significance at the 1 percent level.

Shock	Real GDP level	Price level
Demand	0.00%	1.17%***
	(52.25%)	(0.00%)
Supply	0.66%***	-0.45%
	(0.00%)	(93.95%)

Appendix IV: Maximum Likelihood Estimation of Demand and Supply Shocks Parameters

The estimation procedure is a version of Bates (2006) algorithm for the component model of two gamma distributed variables. The step-by-step estimation strategy for the demand shock is described below. The estimation for the supply shock is identical.

The methodology below is an approximation, because, in order to facilitate the computation, at each time point the conditional distribution of state variables p_t^d and n_t^d is assumed to be gamma, although the distribution does not have a closed form solution. The choice of the approximating distributions is discussed in details in section 1.3 of Bates (2006). Here the gamma distributions are used, because they are bounded from the left at 0, which ensures that the shape parameters of the gamma distribution in the model (p_t^d and n_t^d) will always stay positive, like they should.

The system to estimate is:

$$u^d_{t+1} = \sigma_{dp}\omega^d_{p,t+1} - \sigma_{dn}\omega^d_{n,t+1},$$

$$\omega^d_{p,t+1} \sim \Gamma(p^d_t, 1) - p^d_t,$$

$$\omega^d_{n,t+1} \sim \Gamma(n^d_t, 1) - n^d_t,$$

$$p^d_{t+1} = \bar{p}^d(1-\rho^d_p) + \rho^d_p p^d_t + \sigma^d_{pp}\omega^d_{p,t+1},$$

$$n^d_{t+1} = \bar{n}^d(1-\rho^d_n) + \rho^d_n n^d_t + \sigma^d_{nn}\omega^d_{n,t+1}.$$

The following notation is defined:

$U^d_t \equiv \{u^d_1, ..., u^d_t\}$ is the sequence of observations up to time t.

$F(i\phi, i\psi^1, i\psi^2|U^d_t) \equiv E(e^{i\phi u^d_{t+1} + i\psi^1 p^d_{t+1} + i\psi^2 n^d_{t+1}}|U^d_t)$ is the next period's joint conditional characteristic function of the observation and the state variables.

$G_{t|s}(i\psi^1, i\psi^2) \equiv E(e^{i\psi^1 p^d_t + i\psi^2 n^d_t}|U^d_s)$ is the characteristic function of the time t state variables conditioned on observing data up to time s.

At time 0, the characteristic function of the state variables $G_{0|0}(i\psi^1, i\psi^2)$ is initialized. As mentioned above, the distribution of p^d_0 and n^d_0 is approximated with gamma distributions. Note that the unconditional mean and variance of p^d_t are $E(p^d_t) = \bar{p}^d$ and $Var(p^d_t) = \frac{\sigma^2_{pp}}{1-\rho^{d2}_p}\bar{p}^d$, respectively. The approximation by the gamma distribution with the shape parameter k_0 and the scale parameter σ^p_0 is done by matching the first two unconditional moments. Using the properties of the gamma distribution, $k^p_0 = \frac{E^2 p^d_t}{Var(p^d_t)}$ and $\theta^p_0 = \frac{Var(p^d_t)}{E(p^d_t)}$. Thus, p^d_0 is assumed to follow $\Gamma(k^p_0, \theta^p_0)$ and n^d_0 is assumed to follow $\Gamma(k^n_0, \theta^n_0)$, where k^n_0 and θ^n_0 are computed in the same way. Using the properties of the expectations of the gamma variables, $G_{0|0}(i\psi^1, i\psi^2) = e^{-k^p_0 \ln(1-\theta^p_0 i\psi^1) - k^n_0 \ln(1-\theta^n_0 i\psi^2)}$. Given $G_{0|0}(i\psi^1, i\psi^2)$, computing the likelihood of U^d_T is performed by repeating the steps 1-3 below for all subsequent values of t.

Step 1. Computing the next period's joint conditional characteristic function of the observation and the state variables:

$$
\begin{aligned}
&F(i\Phi, i\psi^1, i\psi^2|U_t^d) = E(E(e^{i\Phi(\sigma_{dp}\omega_{p,t+1}^d-\sigma_{dn}\omega_{n,t+1}^d)+i\psi^1(\bar{p}^d+\rho_p^d p_t^d+\sigma_{pp}^d\omega_{p,t+1}^d)+i\psi^2(\bar{n}^d(1-\rho_n^d)+\rho_n^d n_t^d+\sigma_{nn}^d\omega_{n,t+1}^d)}|U_t^d) \\
&= E(e^{i\psi^1\bar{p}^d(1-\rho_p^d)+i\psi^2\bar{n}^d(1-\rho_n^d)+(i\psi^1\rho_p^d-\ln(1-i\Phi\sigma_{dp}-i\psi^1\sigma_{pp}^d)-i\Phi\sigma_{dp}-i\psi^1\sigma_{pp}^d)p_t^d+(i\psi^2\rho_n^d-\ln(1+i\Phi\sigma_{dn}-i\psi^2\sigma_{nn}^d)+i\Phi\sigma_{dn}-i\psi^2\sigma_{nn}^d)n_t^d}|U_t^d) \\
&= e^{i\psi^1\bar{p}^d(1-\rho_p^d)+i\psi^2\bar{n}^d(1-\rho_n^d)}G_{t|t}(i\psi^1\rho_p^d-\ln(1-i\Phi\sigma_{dp}-i\psi^1\sigma_{pp}^d)-i\Phi\sigma_{dp}-i\psi^1\sigma_{pp}^d, i\psi^2\rho_n^d-\ln(1+i\Phi\sigma_{dn}-i\psi^2\sigma_{nn}^d)+i\Phi\sigma_{dn}-i\psi^2\sigma_{nn}^d).
\end{aligned}
$$

Step 2. Evaluating the conditional likelihood of the time $t+1$ observation:

$$
p(u_{t+1}^d|U_t^d) = \frac{1}{2\pi}\int_{-\infty}^{\infty} F(i\Phi, 0, 0|U_t^d)e^{-i\Phi u_{t+1}^d)}d\Phi,
$$

where the function F is defined in step 1 and the integral is evaluated numerically.

Step 3. Computing the conditional characteristic function for the next period, $G_{t+1|t+1}(i\psi^1, i\psi^2)$:

$$
G_{t+1|t+1}(i\psi^1, i\psi^2) = \frac{\frac{1}{2\pi}\int_{-\infty}^{\infty} F(i\Phi, i\psi^1, i\psi^2|U_t^d)e^{-i\Phi u_{t+1}^d}d\Phi}{p(u_{t+1}^d|U_t^d)}.
$$

As above, the function $G_{t+1|t+1}(i\psi^1, i\psi^2)$ is also approximated with the gamma distribution via matching the first two moments of the distribution. The moments are obtained by taking the first and second partial derivatives of the joint characteristic function:

$$
\begin{aligned}
E_{t+1}p_{t+1}^d &= \frac{1}{2\pi p(u_{t+1}^d|U_t^d)}\int_{-\infty}^{\infty} F_{\psi^1}(i\Phi, 0, 0|U_t^d)e^{-i\Phi u_{t+1}^d}d\Phi, \\
Var_{t+1}p_{t+1}^d &= \frac{1}{2\pi p(u_{t+1}^d|U_t^d)}\int_{-\infty}^{\infty} F_{\psi^1\psi^1}(i\Phi, 0, 0|U_t^d)e^{-i\Phi u_{t+1}^d}d\Phi - E_{t+1}^2 p_{t+1}^d, \\
E_{t+1}n_{t+1}^d &= \frac{1}{2\pi p(u_{t+1}^d|U_t^d)}\int_{-\infty}^{\infty} F_{\psi^2}(i\Phi, 0, 0|U_t^d)e^{-i\Phi u_{t+1}^d}d\Phi, \\
Var_{t+1}n_{t+1}^d &= \frac{1}{2\pi p(u_{t+1}^d|U_t^d)}\int_{-\infty}^{\infty} F_{\psi^2\psi^2}(i\Phi, 0, 0|U_t^d)e^{-i\Phi u_{t+1}^d}d\Phi - E_{t+1}^2 n_{t+1}^d,
\end{aligned}
$$

where F_{ψ^i} denotes the derivative of F with respect to ψ^i. The expressions inside the integral are obtained in closed form by taking the derivative of the function $F(i\Phi, i\psi^1, i\psi^2|U_t^d)$ in step 1, and integrals are evaluated numerically. Using the properties of the gamma distribution, the values of the shape and the scale parameters are $k_{t+1}^p = \frac{E_{t+1}^2 p_{t+1}^d}{Var_{t+1}p_{t+1}^d}$ and $\theta_{t+1}^p = \frac{Var_{t+1}p_{t+1}^d}{E_{t+1}p_{t+1}^d}$, respectively. The expressions for k_{t+1}^n and θ_{t+1}^n are similar.

The total likelihood of the time series is the sum of individual likelihoods from step 2: $L(Y_T) = \ln p(u_1^d | k_0^p, \theta_0^p) + \sum_{t=2}^{T} \ln p(u_{t+1}^d | U_t^d)$.