The Zero Lower Bound and Endogenous Uncertainty*

Michael Plante Alexander W. Richter Nathaniel A. Throckmorton

First Draft: May 21, 2014
This Draft: September 10, 2016

ABSTRACT

This paper examines the correlation between uncertainty and real GDP growth. We use the volatility of real GDP growth from a VAR, stock market volatility, survey-based forecast dispersion, and the index from Jurado et al. (2015) as proxies for uncertainty. In each case, a stronger negative correlation emerged in 2008. We contend the zero lower bound (ZLB) on the federal funds rate contributed to our finding. To test our theory, we estimate a New Keynesian model with a ZLB constraint to generate a data-driven, forward-looking uncertainty measure. The correlations between that measure and real GDP growth are close to the values in the data.

Keywords: Bayesian Estimation; Monetary Policy; Uncertainty; Zero Lower Bound
JEL Classifications: C11; E43; E58

*Plante, Federal Reserve Bank of Dallas, 2200 N Pearl St, Dallas, TX 75201 (michael.plante@dal.frb.org); Richter, Federal Reserve Bank of Dallas and Auburn University, 2200 N Pearl St, Dallas, TX 75201 (alex.richter@dal.frb.org); Throckmorton, College of William & Mary, P.O. Box 8795, Williamsburg, VA 23187 (nathrockmorton@wm.edu).
We thank our discussants, Thiago Ferreira and Fernando Martin, our editor, Morten Ravn, as well as Rudi Bachmann, Roberto Billi, Mario Crucini, Ferre De Graeve, Martin Gervais, Francois Gourio, Rob Hicks, Lance Kent, Keith Kuester, Thomas Lubik, Christian Matthes, Tai Nakata, Chris Vickers, Nic Ziebarth, and four referees for useful comments. Throckmorton thanks the Reves Center for travel support and Eric Walter for helping with the SciClone cluster. Richter thanks Bradley Morgan for assisting with the Hopper cluster. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 INTRODUCTION

There is significant interest in understanding the relationship between uncertainty and economic activity. Several papers find a negative relationship in the data using various measures of uncertainty. For example, Bloom (2009) shows unexpected increases in uncertainty, given by stock market volatility, are associated with declines in industrial production. Bekaert et al. (2013), Bloom et al. (2014), and Pinter et al. (2013) find similar relationships in the data. While those papers focus on financial market volatility, Jurado et al. (2015) develop a broader measure of uncertainty and find that spikes in uncertainty are more infrequent but more persistent and more negatively correlated with hours, employment, and industrial production than the literature has typically reported.[1] Leduc and Liu (2014) estimate that uncertainty shocks led to a larger increase in the unemployment rate and a much slower recovery from the Great Recession than in past recessions. They speculate that their finding is attributable to the zero lower bound (ZLB) constraint on the federal funds rate.

There is also an established literature that explores how uncertainty affects real activity in structural models. That literature typically examines how endogenous variables respond to volatility shocks that exogenously increase uncertainty about one or more variables in the model.[2] In this setup, higher volatility reduces economic activity, but the quantitative effects are sometimes small.

While the literature has primarily focused on how economic activity responds to changes in specific types of uncertainty, this paper examines how recent events affected macroeconomic uncertainty and its correlation with real GDP growth. Our analysis reveals that the correlation was significantly more negative during the ZLB period. We first estimate a time-varying parameter VAR with stochastic volatility (SV) and calculate the correlation between real GDP growth and its predicted volatility. Then we use stock market volatility, survey-based forecast dispersion, and the index developed by Jurado et al. (2015) as measures of uncertainty and correlate those series with real GDP growth. In each case, a much stronger and statistically more negative correlation emerged in late 2008. Before then the correlation was weaker and sometimes not statistically below zero.

A major difference between the recent economic downturn and previous recessions is that the Fed has been constrained by the ZLB on the federal funds rate. We contend that the ZLB constraint contributed to the stronger negative correlation that emerged in late 2008 because it restricts the ability of the central bank to stabilise the economy. That situation makes real GDP more responsive to shocks hitting the economy and, therefore, increases the uncertainty surrounding future growth.

To test our theory, we estimate a nonlinear New Keynsian model with a ZLB constraint on the nominal interest rate using data from 1986Q1 to 2014Q2. A major benefit of estimating the model is that it allows us to generate a data-driven, forward-looking estimate of macroeconomic uncertainty, which equals the expected volatility of the forecast error for real GDP growth. Another benefit is that it allows us to make empirical predictions about how the level of uncertainty and its correlation with real GDP growth changed in the ZLB period. We find our uncertainty series has a strong and statistically significant negative correlation with real GDP growth in the ZLB period, just like we found in the data. We also show there is a dramatic increase in our uncertainty series at the start of the ZLB period, similar to what occurred with all of the other uncertainty measures.

[1] Others find a negative relationship between fiscal uncertainty and economic activity [Fernández-Villaverde et al. (2015), Born and Pfeifer (2014)] and oil price uncertainty and economic activity [Elder and Serletis (2010), Jo (2014)].

[2] A few recent examples include Bachmann and Bayer (2013), Bloom (2009), Bloom et al. (2014), Christiano et al. (2014), Fernández-Villaverde et al. (2011), Gilchrist et al. (2014), Justiniano and Primiceri (2008), and Mumtaz and Zanetti (2013). There is also an older literature that examines similar research questions to those posed in the stochastic volatility literature. See, for example, Leland (1968), Levhari and Srinivasan (1969), and Sandmo (1970).

In our baseline model, the ZLB is the only mechanism that generates meaningful movements in macroeconomic uncertainty, which makes it impossible to determine its correlation with real GDP growth when the ZLB does not bind. Furthermore, there is no way to assess whether exogenous or endogenous sources of uncertainty play a larger role in the ZLB period. To overcome those two drawbacks, we follow the SV literature and add an exogenous source of time-varying uncertainty about real GDP growth to our model. We then compute correlations in three settings: (1) At the ZLB without SV; (2) at the ZLB with SV; and (3) away from the ZLB with SV. In the model with SV, uncertainty about real GDP growth regularly fluctuates at and away from the ZLB, the correlations in the ZLB period are stronger and statistically more negative than in the pre-ZLB period, and most of the uncertainty in the ZLB period is due to the constraint instead of SV shocks.

We also examine the correlations between the uncertainty measures, the correlations at leads and lags of real GDP growth, and the correlations between inflation uncertainty and real GDP growth. Nearly all of the uncertainty measures are positively correlated, and the correlations are especially pronounced in the ZLB period. The strongest correlation is between the series from our structural model and the Jurado et al. (2015) macro uncertainty index, which is significant because their model imposes no structure and it uses a broad information set. The cross-correlations indicate that exogenous uncertainty shocks have persistent effects on real GDP growth (i.e., negative correlations with leads of real GDP growth), while endogenous uncertainty arises due to what is happening in the economy (i.e., negative correlations with lags of real GDP growth). The correlations with inflation uncertainty are very similar to our results with real GDP growth uncertainty.

Our paper is related to the literature that examines the effects of ZLB constraint. Gust et al. (2013) estimate a New Keynesian model with a ZLB constraint and find the constraint accounts for 20% of the drop in U.S. real GDP in 2008. Nakov (2008) finds the optimal discretionary monetary policy leads to a more negative output gap at the ZLB when households are uncertain about the real interest rate. Nakata (2012) finds higher uncertainty about discount factor shocks increases the slope of the policy function for output, so positive shocks lead to a larger decline in output when the ZLB binds. Basu and Bundick (2014) show cost and demand uncertainty shocks generate larger business cycles when the ZLB binds. They calculate that demand uncertainty shocks can account for one-fourth of the drop in output in 2008. Basu and Bundick (2015) show that the ZLB amplifies the adverse effects of exogenous volatility shocks through precautionary saving and contractionary bias channels, but the central bank can attenuate that outcome with a history-dependent policy rule.

Other papers study endogenous uncertainty in different contexts. Bachmann and Moscarini (2012) set up a model where uncertainty increases in recessions because it is less costly for firms to experiment with price changes to learn about their market power. Van Nieuwerburgh and Veldkamp (2006) argue that lower production during a recession leads to noisy forecasts that impede learning and slow the recovery. In a related paper, Fajgelbaum et al. (2014) set up a model where information moves slowly during recessions, which discourages already low investment and results in an uncertainty trap. Gourio (2014) argues the volatility of output is countercyclical because people expect larger losses when a shock raises the probability of default. Navarro (2014) develops a model where financial shocks endogenously generate the volatility during the Great Recession.

The paper proceeds as follows. Section 2 computes correlations between our uncertainty measures in the data and economic activity. Section 3 describes our model and measure of endogenous uncertainty. Section 4 reports the correlations between output growth and our estimated uncertainty series. Section 5 shows the correlations between our uncertainty measures, correlations at leads and lags of real GDP growth, and the correlations with inflation uncertainty. Section 6 concludes.

2 Relationship Between Economic Activity and Uncertainty

Using various measures of uncertainty, this section documents that a much stronger negative correlation between economic activity and macroeconomic uncertainty emerged during the ZLB period.

2.1 Time-Varying Parameter VAR with Stochastic Volatility We begin by computing correlations based on a time-varying measure of uncertainty in a reduced-form model of the U.S. economy. Following Primiceri (2005), we use a time-varying parameter VAR with SV to estimate the volatility of real GDP growth and then show how those estimates are correlated with economic activity. One difference between Primiceri (2005) and our VAR is that we use real GDP growth instead of the unemployment rate, so we can draw comparisons to equivalent statistics from our structural model. Another difference is we use data from 1975Q3 to 2014Q2, where the first ten years train the prior distributions of the parameters. The reduced-form VAR model is given by

$$y_t = b_t + B_{1,t}y_{t-1} + B_{2,t}y_{t-2} + A_t^{-1}\Sigma_t\varepsilon_t, \quad t = 1, \dots, T.$$

Since the model contains a 2-quarter lag, our estimates are based on data from 1986Q1 to 2014Q2.[3] y_t is a 3×1 vector that includes per capita real GDP growth, the inflation rate, and the federal funds rate, b_t is a 3×1 vector of time varying intercepts, $B_{1,t}$ and $B_{2,t}$ are 3×3 matrices of time varying coefficients, and ε_t is 3×1 vector of shocks with a multivariate standard normal distribution. The shock coefficients are the result of a triangular reduction of the variance-covariance matrix, where

$$A_t = \begin{bmatrix} 1 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 \end{bmatrix}, \qquad \Sigma_t = \begin{bmatrix} \sigma_{1,t} & 0 & 0 \\ 0 & \sigma_{2,t} & 0 \\ 0 & 0 & \sigma_{3,t} \end{bmatrix}.$$

Define $\beta_t = \text{vec}(b_t, B_{1,t}, B_{2,t})$ and $X_{t-1} = I_3 \otimes [1, y'_{t-1}, y'_{t-2}]$. Then the model can be written as

$$y_t = X_{t-1}\beta_t + A_t^{-1}\Sigma_t\varepsilon_t.$$

Now let $\alpha_t = (\alpha_{21,t}, \alpha_{31,t}, \alpha_{32,t})$ and $h_t = (h_{1,t}, h_{2,t}, h_{3,t})$, with $h_{j,t} = \log\sigma_{j,t}$ for $j = 1,2,3$. All of the time-varying parameters evolve according to first-order random walk processes, given by,

$$\beta_t = \beta_{t-1} + \nu_t, \qquad \alpha_t = \alpha_{t-1} + \zeta_t, \qquad h_t = h_{t-1} + \eta_t,$$

where the shocks, ν_t, ζ_t, and η_t, are jointly normally distributed with mean zero and variance covariance matrix V. Since the log volatilities follow a random walk, the realised volatility of each observable is nearly equal to its expected future volatility (i.e., $\sigma_{j,t}\exp(0.5\sigma^2_{\eta_j}) = E_t\sigma_{j,t+1}$).

The model is estimated with Bayesian MCMC methods using code that accompanies Koop and Korobilis (2010). The code implements a correction to the algorithm outlined in Appendix A of Primiceri (2005), as explained by Del Negro and Primiceri (2015). We calculate correlations between per capita real GDP growth and its estimated volatility in both the pre-ZLB and ZLB samples. Our estimates are based on $100{,}000$ draws from the posterior distribution, keeping every 100th draw. For more details about the model and the estimation procedure see the online appendix.

Figure 1 compares real GDP volatility from the VAR to per capita real GDP growth in the data. Volatility regularly fluctuates at and away from the ZLB, it peaks when the ZLB first binds, and it appears that a stronger negative correlation with real GDP growth emerged during the ZLB period.

[3] We also estimated models with 1- and 3-quarter lags, but those specifications had little effect on our main results.

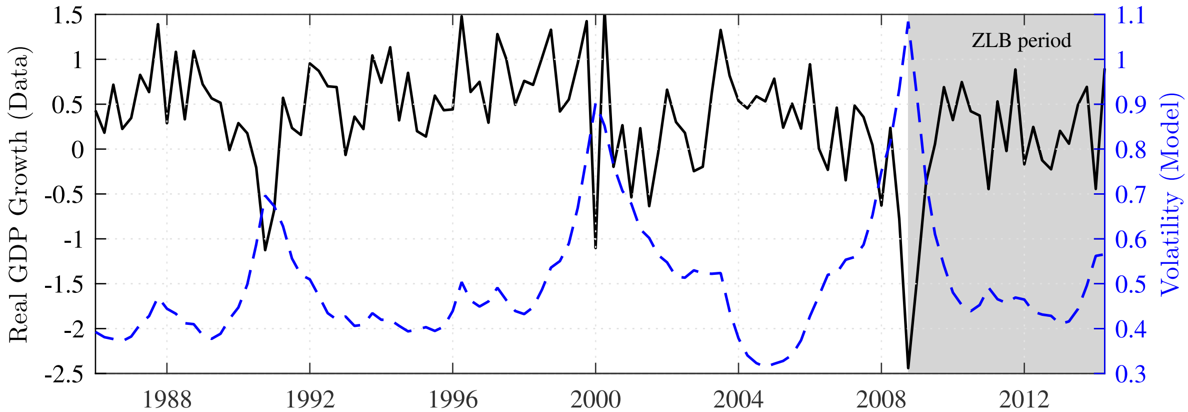

Figure 1: Time series of real GDP growth (solid line) and the median path of real GDP growth volatility (dashed line)

The first column of table 1 shows the correlation between real GDP growth and its estimated volatility in the baseline VAR. We test whether the correlations in the pre-ZLB (1986Q1-2008Q3) and ZLB (2008Q4-2014Q4) samples are statistically significant by calculating the fraction of draws where each correlation is positive. In the pre-ZLB sample, the median correlation is -0.24 and 4.5% of draws are positive. In contrast, the median correlation in the ZLB sample is -0.60 and only 0.6% of draws are positive. Although our results indicate that the correlations in both samples are statistically below zero, the correlation in the ZLB sample is stronger. To provide statistical evidence, we subtract each correlation in the pre-ZLB sample from the correlation in the ZLB sample. The differences are positive in only 7% of draws and the median difference is -0.35.

Another potential source of the stronger correlation in the ZLB sample is the rise in financial uncertainty. To control for that outcome, we add the index from Ludvigson et al. (2015) to our VAR. The index includes 147 financial variables, and it is constructed the same way as the macro uncertainty index in Jurado et al. (2015). We use this index because it has a different relationship with real activity. Specifically, heightened financial uncertainty is an exogenous impulse that causes recessions, while macro uncertainty is an endogenous response to economic fundamentals.

The second column of table 1 shows our results are robust to controlling for financial uncertainty. The median correlation between real GDP growth and its predicted volatility is -0.21 in the pre-ZLB sample and -0.57 in the ZLB sample. The difference between the ZLB and pre-ZLB correlations is positive in only 10% of draws and the median difference is -0.34. The fact that the ZLB correlation remains statistically more negative than the pre-ZLB correlation supports our finding that the ZLB was the main source of the stronger correlation and not financial uncertainty.

Our VAR allows parameters to change over time, but it is possible the random walk assumption makes it too slow to respond to the changes that occur in the data when the ZLB first binds. To show our results are robust, we estimate the model with interest rates that were not constrained by the ZLB. The last three columns of table 1 show the correlations between real GDP growth and its estimated uncertainty series with three substitutes for the federal funds rate, including Wu and Xia's (2016) shadow rate, the 10-year T-Bond rate, and an *ex-ante* real rate computed with the median one-quarter ahead forecast of the GDP deflator from the Survey of Professional Forecasters (SPF). In each case, the ZLB correlation is more negative than the pre-ZLB correlation, and the differences between the correlations are positive in less than 10% of draws. Those results provide evidence that our VAR estimates are robust to concerns about a structural break in the coefficients.

	Fed Funds Rate	Fed Funds Rate & Fin. Uncertainty	Shadow Rate	10-Year Treasury	Real Rate
Pre-ZLB Sample (1986Q1-2008Q3)	-0.24**	-0.21*	-0.24**	-0.25**	-0.23*
ZLB Sample (2008Q4-2014Q4)	-0.60***	-0.57**	-0.60***	-0.61***	-0.59***
Differences (ZLB - Pre-ZLB)	-0.35*	-0.34*	-0.35*	-0.35*	-0.36*

Table 1: Correlations between per capita real GDP growth and the median path of real GDP growth volatility. Column 1 is the baseline VAR, which uses the federal funds rate. Column 2 is the baseline VAR, augmented with the financial uncertainty index from Ludvigson et al. (2015). Columns 3-5 replace the federal funds rate in the baseline VAR with alternative interest rates. An asterisk indicates a correlation in the pre-ZLB or ZLB sample is statistically less than 0 or the difference between the correlations in the two samples is significant at a ***1%, **5%, and *10% level.

2.2 ALTERNATIVE UNCERTAINTY MEASURES We consider four different proxies for macroeconomic uncertainty that have been used in the literature: (1) the Chicago Board Options Exchange S&P 100 Volatility Index (VXO), (2) the dispersion in large manufacturers' forecasts of general business activity from the Business Outlook Survey (BOS), (3) the dispersion in forecasts of real GDP growth 1-quarter ahead from the Survey of Professional Forecasters (SPF), and (4) the macro uncertainty index from Jurado et al. (2015) (JLN). Figure 2 plots the time series of each measure.

The VXO measures the risk-neutral expected volatility in the S&P 100 stock market index over the next 30 days at an annualised rate. For example, if the value on the vertical axis is $x\%$, then people expect there is a 68% chance the index will change by at most $\pm x/\sqrt{12}\%$ over the next 30 days. We average the daily series each quarter so it matches the frequency of real GDP releases.

The BOS asks large manufacturing firms to forecast whether business activity will increase, decrease, or remain unchanged over the next six months. The forecast dispersion (FD) is given by

$$\text{BOS FD}_t = \sqrt{\text{Frac}_t^+ + \text{Frac}_t^- - \left(\text{Frac}_t^+ - \text{Frac}_t^-\right)^2},$$

where an increase (decrease) in business activity is labelled as $+1$ (-1) and Frac^+ (Frac^-) is the fraction of firms who forecast that outcome. Thus, the BOS FD is the standard deviation of the responses in each month. We average the monthly BOS FD series each quarter and then standardise the values so the vertical axis displays the number of standard deviations from the mean response.

The SPF asks people who regularly make forecasts to predict macroeconomic aggregates up to 4 quarters ahead. We use the inter-quartile range of real GDP growth ($\hat{y}$) forecasts 1 quarter ahead,

$$\text{SPF FD}_t = \hat{y}^{75}_{t+1|t-1} - \hat{y}^{25}_{t+1|t-1}.$$

This value is the percent difference between the 75th and 25th percentiles of the quarter t forecasts of real GDP growth in quarter $t+1$, given all observations in quarter $t-1$ and earlier.[4] We use the inter-quartile range rather than more extreme percentiles because on average only 41 firms complete the survey, which means the tails of the distribution would have too small of a sample.

It is possible that the three uncertainty measures described above, as well as other commonly used measures, are only weakly related to actual uncertainty. For example, forecasters can disagree even when they are confident in their projections and stock market volatility may reflect changes in

[4] Our results are based on the D2 measure of forecast dispersion, but the D3 measure produces very similar results.

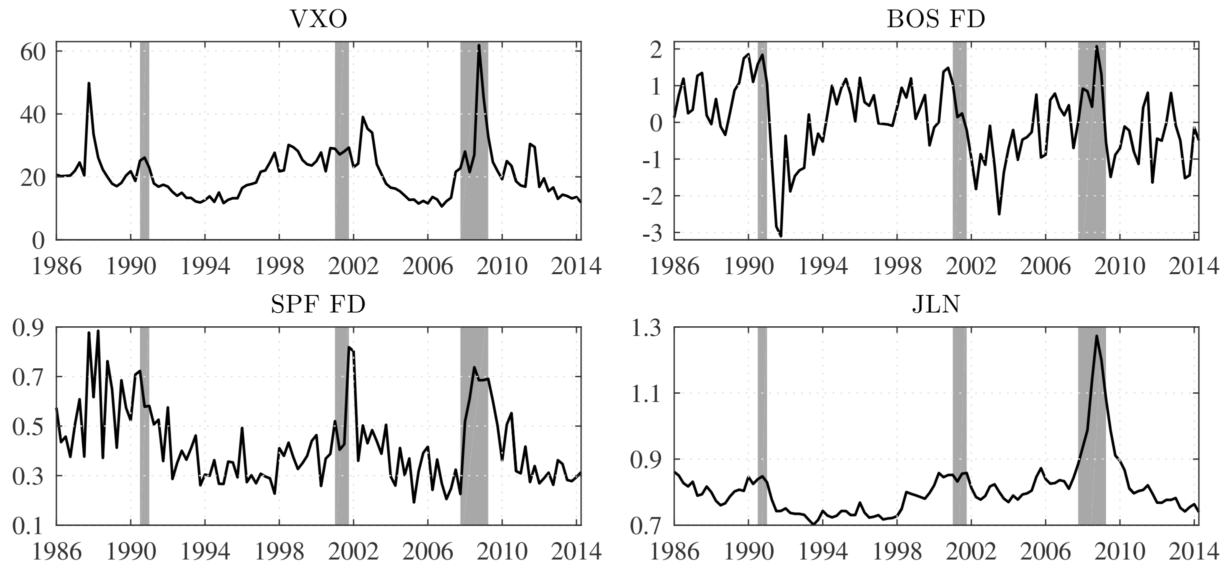

Figure 2: Measures of uncertainty. (1) Chicago Board Options Exchange Volatility Index (VXO); (2) Business Outlook Survey Forecast Dispersion (BOS FD); (3) Survey of Professional Forecasters real GDP Forecast Dispersion (SPF FD); (4) Macro Uncertainty Index developed by Jurado et al. (2015) (JLN). The shaded regions are NBER recessions.

leverage or risk instead of uncertainty. Therefore, Jurado et al. (2015) develop a macro uncertainty index based on a data rich time series model that avoids these concerns. They define macro uncertainty as the conditional volatility of the unforecastable component of a future variable. There are two major advantages of their index over other measures—it distinguishes between uncertainty and conditional volatility and measures the common variation in 134 macroeconomic time series.

Table 2a shows correlations between per capita real GDP growth and the four uncertainty measures.[5] The top and middle rows are based on the pre-ZLB (1986Q1-2008Q3) and ZLB (2008Q4-2014Q2) samples we used to calculate the correlations with our VAR model. We use a one-tailed t-test to determine whether the correlations are statistically less than zero. While some of the correlations in the pre-ZLB sample are not statistically below zero (VXO, SPF FD), three of the correlations in the ZLB sample are below zero at a 1% level. More importantly, all of the results indicate that a much stronger negative correlation emerged in the ZLB sample. We use the Fisher z-transformation to test whether the correlations in the two samples are statistically different. Those tests reveal that the correlations are statistically different at a 1% level when calculated with the VXO and BOS FD and at a 5% level with the SPF FD and JLN uncertainty measures (bottom row).

We use per capita real GDP growth as a measure of economic activity, but other work often uses industrial production (IP) [e.g., Bloom (2009), Bekaert et al. (2013), and Jurado et al. (2015)]. To draw comparisons with the literature, the first four columns of table 2b reproduce the correlations in table 2a with IP growth as a measure of economic activity. Consistent with our previous results, all of the correlations are stronger in the ZLB sample than the pre-ZLB sample. The Fisher z-transformation test shows the correlations are different at a 1% significance level when calculated with the VXO, a 5% level with the BOS FD and SPF FD, and a 12% level with the JLN measure.[6]

[5]We also examined realised stock price volatility, but the results are similar to the VXO so they are not reported.

[6]Given the results in Bloom (2009), the weak correlations with the VXO in our pre-ZLB sample may be surprising. Choi (2013) shows that Bloom's results are sample-dependent. Although the uncertainty shocks have a significant impact when the sample begins in 1962, they no longer have a negative effect on IP when the sample begins in 1983.

	VXO	BOS FD	SPF FD	JLN
Pre-ZLB Sample (1986Q1-2008Q3)	-0.09	-0.19**	-0.08	-0.45***
ZLB Sample (2008Q4-2014Q2)	-0.72***	-0.70***	-0.47**	-0.73***
Difference (ZLB-Pre-ZLB)	-0.64***	-0.51***	-0.39**	-0.29**

(a) Correlations with per capita real GDP growth

	Quarterly Data				Monthly Data		
	VXO	BOS FD	SPF FD	JLN	VXO	BOS FD	JLN
Pre-ZLB Sample (1986Q1-2008Q3)	-0.11	-0.19**	-0.26***	-0.64***	-0.10*	-0.13**	-0.41***
ZLB Sample (2008Q4-2014Q2)	-0.74***	-0.59***	-0.61***	-0.78***	-0.50***	-0.38***	-0.54***
Difference (ZLB-Pre-ZLB)	-0.62***	-0.40**	-0.34**	-0.14	-0.40***	-0.25**	-0.13

(b) Correlations with industrial production growth

Table 2: Correlations between economic activity and various measures of macroeconomic uncertainty before and during the ZLB period. An asterisk indicates a correlation in the pre-ZLB or ZLB sample is statistically less than 0 or the difference between the correlations in the two samples is significant at a ***1%, **5%, and *10% level.

A benefit of using industrial production is that it is released every month. Since the BOS FD and JLN time series are monthly and the VXO series is daily, we can calculate the same correlations at a monthly frequency, which increases our sample size. The last three columns of table 2b show the correlations with monthly data are qualitatively similar to the correlations with quarterly data. All of the correlations in the pre-ZLB sample are weaker than the correlations in the ZLB sample, and the correlations in the two samples are significantly different at a 1% level when calculated with the VXO, a 5% level with the BOS FD, and a 11% level with the JLN uncertainty measure.[7]

Interestingly, the correlations in the pre-ZLB sample are generally more negative than when we use real GDP, which suggests that there is something different about the manufacturing sector that sets it apart from the overall economy. In our paper, we are interested in the connection between uncertainty and real GDP growth. It is beyond the scope of this paper and our structural estimation to formally explore why the relationship is stronger with manufacturing related variables, but we speculate that the difference may have to do with the fact that manufacturing activity is more volatile and more sensitive to changes in business conditions than other segments of the economy.

3 Structural Model and Measure of Uncertainty

3.1 Model There are three actors in the model: (1) a representative household that has access to a one-period nominal bond, (2) a representative firm that bundles a continuum of intermediate inputs to produce a final good, and (3) a central bank that sets the short-term nominal interest rate.

Households A representative household chooses $\{c_t, n_t, b_t\}_{t=0}^{\infty}$ to maximise expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \tilde{\beta}_t [\log(c_t - hc^a_{t-1}) - \chi n_t^{1+\eta}/(1+\eta)]$, where $\chi > 0$, $1/\eta$ is the Frisch elasticity

[7] In this paper, we focus on survey-based forecasts over a one-quarter horizon. The SPF also asks forecasters to predict various macro variables over longer horizons. The qualitative results in table 2 are robust to longer horizons.

of labour supply, c is consumption, c^a is aggregate consumption, h is the degree of external habit persistence, n is labour hours, b is the real value of a privately-issued 1-period nominal bond, E_0 is an expectation operator conditional on information in period 0, $\tilde{\beta}_0 \equiv 1$, and $\tilde{\beta}_t = \prod_{j=1}^{t>0} \beta_j$. Following Eggertsson and Woodford (2003), β is a time-varying discount factor that follows

$$\log \beta_t = (1 - \rho_\beta) \log \bar{\beta} + \rho_\beta \log \beta_{t-1} + \sigma_\upsilon \upsilon_t, \tag{1}$$

where $\bar{\beta}$ is the discount factor along the steady state growth path, $0 \leq \rho_\beta < 1$ is the persistence of the discount factor, and υ is a standard normal shock to the discount factor. These choices are constrained by $c_t + b_t = w_t n_t + i_{t-1} b_{t-1}/\pi_t + d_t$, where π is the gross inflation rate, w is the real wage rate, i is the gross nominal interest rate set by the central bank, and d is a real dividend received from ownership of intermediate firms. The optimality conditions to the household's problem imply

$$\begin{aligned} w_t &= \chi n_t^\eta (c_t - h c_{t-1}^a), \\ 1 &= i_t E_t[q_{t,t+1}/\pi_{t+1}], \end{aligned}$$

where $q_{t,t+1} \equiv \beta_{t+1}(c_t - h c_{t-1}^a)/(c_{t+1} - h c_t^a)$ is the pricing kernel between periods t and $t+1$.

Firms The production sector consists of a continuum of monopolistically competitive intermediate goods firms owned by households and a final goods firm. Intermediate firm $f \in [0,1]$ produces a differentiated good, $y_t(f)$, according to $y_t(f) = z_t n_t(f)$, where $n(f)$ is the labour hired by firm f. z is labour productivity, which is common across all firms and evolves according to a random walk, $\log z_t = \log z_{t-1} + \log g_t$. The deviations from the steady state growth rate, $\bar{g}$, follow

$$\log g_t = (1 - \rho_g) \log \bar{g} + \rho_g \log g_{t-1} + \sigma_\varepsilon \varepsilon_t, \tag{2}$$

where $0 \leq \rho_g < 1$ is the persistence and ε is a standard normal shock to the growth rate. Each intermediate firm chooses its labour to minimise costs, $w_t n_t(f)$, subject to its production function.

The representative final goods firm purchases $y_t(f)$ units from each intermediate goods firm to produce the final good, $y_t \equiv [\int_0^1 y_t(f)^{(\theta-1)/\theta} df]^{\theta/(\theta-1)}$, according to a Dixit and Stiglitz (1977) aggregator, where $\theta > 1$ measures the elasticity of substitution between the intermediate goods. The final goods firm then maximises dividends to determine its demand function for intermediate good f, $y_t(f) = (p_t(f)/p_t)^{-\theta} y_t$, where $p_t = [\int_0^1 p_t(f)^{1-\theta} df]^{1/(1-\theta)}$ is the price of the final good.

Following Rotemberg (1982), each intermediate firm faces a cost to adjusting its price, $adj_t(f)$, which emphasises the negative effect that price changes can have on customer-firm relationships. Using the functional form in Ireland (1997), $adj_t(f) = \varphi[p_t(f)/(\bar{\pi} p_{t-1}(f)) - 1]^2 y_t/2$, where $\varphi \geq 0$ scales the size of the adjustment cost and $\bar{\pi}$ is the gross inflation rate along the steady state growth path. Real dividends are then given by $d_t(f) = (p_t(f)/p_t) y_t(f) - w_t n_t(f) - adj_t(f)$. Intermediate firm f chooses its price, $p_t(f)$, to maximise the expected discounted present value of real dividends, given by, $E_t \sum_{k=t}^\infty q_{t,k} d_k(f)$, where $q_{t,t} \equiv 1$ and $q_{t,k} \equiv \prod_{j=t+1}^k q_{j-1,j}$. In a symmetric equilibrium, all firms make identical decisions and the optimality condition reduces to

$$\varphi \left(\frac{\pi_t}{\bar{\pi}} - 1 \right) \frac{\pi_t}{\bar{\pi}} = (1 - \theta) + \theta (w_t/z_t) + \varphi E_t \left[q_{t,t+1} \left(\frac{\pi_{t+1}}{\bar{\pi}} - 1 \right) \frac{\pi_{t+1}}{\bar{\pi}} \frac{y_{t+1}}{y_t} \right].$$

Without price adjustment costs (i.e., $\varphi = 0$), the real marginal cost of producing a unit of output (w_t/z_t) equals $(\theta - 1)/\theta$, which is the inverse of a firm's markup of price over marginal cost, μ.

Monetary Policy The central bank sets the gross nominal interest rate according to

$$i_t = \max\{\underline{\imath}, i_t^*\}, \quad i_t^* = (i_{t-1}^*)^{\rho_i}(\bar{\imath}(\pi_t/\bar{\pi})^{\phi_\pi}(y_t^{gdp}/(\bar{g}y_{t-1}^{gdp}))^{\phi_y})^{1-\rho_i}\exp(\sigma_\nu \nu_t),$$

where $\underline{\imath}$ is the effective lower bound, y^{gdp} is real GDP (i.e., the level of output minus the resources lost due to price adjustment costs), i^* is the notional interest rate (i.e., the rate the central bank would set if it was not constrained), $0 \leq \rho_i < 1$ is the persistence of the notional rate, ϕ_π and ϕ_y are the responses to deviations of inflation from target and deviations of real GDP growth from the steady state growth rate, $\bar{\imath}$ and $\bar{\pi}$ are the inflation and interest rate targets, which equal their values along the steady state growth path, and ν is a serially uncorrelated standard normal policy shock.

Competitive Equilibrium The aggregate resource constraint is given by $c_t = y_t - adj_t \equiv y_t^{gdp}$. Given the unit root in labour productivity, the model does not possess a steady-state. In order to make the model stationary, we redefine all of the variables that grow along the steady state growth path in terms of labour productivity (i.e., $\tilde{x}_t \equiv x_t/z_t$). The detrended equilibrium system contains

$$\tilde{w}_t = \chi \tilde{y}_t^\eta(\tilde{c}_t - h\tilde{c}_{t-1}/g_t), \tag{3}$$

$$1 = i_t E_t\left[\frac{\beta_{t+1}}{g_{t+1}\pi_{t+1}}\frac{\tilde{c}_t - h\tilde{c}_{t-1}/g_t}{\tilde{c}_{t+1} - h\tilde{c}_t/g_{t+1}}\right], \tag{4}$$

$$\varphi\left(\frac{\pi_t}{\bar{\pi}} - 1\right)\frac{\pi_t}{\bar{\pi}} = (1-\theta) + \theta\tilde{w}_t + \varphi E_t\left[\beta_{t+1}\frac{\tilde{c}_t - h\tilde{c}_{t-1}/g_t}{\tilde{c}_{t+1} - h\tilde{c}_t/g_{t+1}}\left(\frac{\pi_{t+1}}{\bar{\pi}} - 1\right)\frac{\pi_{t+1}}{\bar{\pi}}\frac{\tilde{y}_{t+1}}{\tilde{y}_t}\right], \tag{5}$$

$$\tilde{c}_t = [1 - \varphi(\pi_t/\bar{\pi} - 1)^2/2]\tilde{y}_t, \tag{6}$$

$$i_t^* = (i_{t-1}^*)^{\rho_i}(\bar{\imath}(\pi_t/\bar{\pi})^{\phi_\pi}(g_t\tilde{y}_t^{gdp}/(\bar{g}\tilde{y}_{t-1}^{gdp}))^{\phi_y})^{1-\rho_i}\exp(\sigma_\nu\nu_t), \tag{7}$$

the ZLB constraint, and the stochastic processes, which impose the bond market clearing condition, $b_t = 0$, and the aggregation rule, $\tilde{c}_t = \tilde{c}_t^a$. A competitive equilibrium includes sequences of quantities, $\{\tilde{c}_t, \tilde{y}_t\}_{t=0}^\infty$, prices, $\{w_t, i_t, i_t^*, \hat{\pi}_t\}_{t=0}^\infty$, and exogenous variables, $\{\beta_t, g_t\}_{t=0}^\infty$, that satisfy the detrended system, given the initial conditions, $\{\tilde{c}_{-1}, i_{-1}^*, \beta_0, g_0, \nu_0\}$, and the shocks, $\{\varepsilon_t, \upsilon_t, \nu_t\}_{t=1}^\infty$.

3.2 MEASURE OF ENDOGENOUS UNCERTAINTY

A recent segment of the literature adds SV to dynamic stochastic general equilibrium (DSGE) models. Our work is different in that we primarily focus on how uncertainty about future variables endogenously responds to the state of the economy.

We first describe how to measure uncertainty in a model with SV. Suppose a model includes an exogenous variable, x, that evolves according to $x_t = (1-\rho_x)\bar{x} + \rho_x x_{t-1} + \sigma_t\varepsilon_t$, where $\bar{x}$ is the mean, $0 \leq \rho_x < 1$ is the persistence, and ε is a standard normal shock. SV is introduced by assuming the standard deviation of the shock is time-varying, which relaxes the common assumption of homoskedastic innovations. If ε and σ are uncorrelated, then the mean forecast error, FE_x, equals

$$E_t[FE_{x,t+1}] = E_t[x_{t+1} - E_t x_{t+1}] = 0.$$

Although the forecast error is mean zero, there is uncertainty about its future value. One measure of that uncertainty is the expected standard deviation of the forecast error for x, which is given by

$$\sqrt{E_t[FE_{x,t+1}^2]} = \sqrt{E_t[(x_{t+1} - E_t x_{t+1})^2]} = \sqrt{E_t\sigma_{t+1}^2}.$$

Models with SV can match features of the data that models with homoskedastic errors cannot. However, even in models without SV there is uncertainty about future endogenous variables due to the presence of first moment shocks. We refer to this type of uncertainty as endogenous time-varying uncertainty, and we quantify it by following the logic of the SV literature. Specifically, the endogenous uncertainty surrounding real GDP growth, $\hat{y}^{gdp}$, 1 quarter in the future, is given by

$$\sigma_{\hat{y}^{gdp},t} \equiv \sqrt{E_t[(\hat{y}^{gdp}_{t+1} - E_t\hat{y}^{gdp}_{t+1})^2]},$$

which can vary over time in response to the state of the economy, even when SV shocks are not included in the model. For example, a large discount factor shock can cause the ZLB to bind, which restricts the central bank's ability to stabilise the economy. In that case, the constraint makes real GDP more sensitive to shocks. Therefore, the uncertainty surrounding real GDP growth is a function of both the underlying shock volatilities and the state of the economy. We focus on real GDP growth uncertainty, but we also calculate this measure of uncertainty for the inflation rate.

By itself, the DSGE model is useful for investigating the theoretical properties of our time-varying measure of endogenous uncertainty. By estimating the model, we go a step further and generate a data-driven estimate of endogenous uncertainty that we correlate with real GDP growth.

3.3 SOLUTION METHOD We solve the model using the policy function iteration algorithm described in Richter et al. (2014), which is based on the theoretical work on monotone operators in Coleman (1991). This method discretises the state space and iteratively solves for updated policy functions until a tolerance is met. We use linear interpolation to approximate future variables, since it accurately captures the kink in the policy functions, and discretise the stochastic processes following Rouwenhorst (1995) to numerically integrate. See Appendix A for a formal description.

There are two main advantages of using a global solution method that depends on a continuum of future shocks. One, it allows for variation in endogenous variables when the ZLB binds, which occurs in the data and is necessary to calculate correlations in the ZLB period. Two, it permits recurring ZLB events and accounts for the expectational effects of going to and leaving the ZLB, which is important for understanding changes in uncertainty. A simplification commonly used in the literature is a two-state Markov chain on the discount factor. With that setup, however, the expectational effect of going to the ZLB is fixed and endogenous variables do not vary at the ZLB.

Benhabib et al.'s (2001) finding that models with a ZLB constraint have two steady-state equilibria has generated discussion about whether there are conditions in which a unique minimum state variable solution exists. Richter and Throckmorton (2015) show the algorithm used in this paper converges to the inflationary equilibrium as long as there is a sufficient expectation of returning to a policy rule that satisfies the Taylor principle, and it never converges to the deflationary equilibrium. Addressing questions related to the deflationary steady state is an interesting topic for future research, but our analysis, like most research on the ZLB, focuses on the solution centered around the inflationary steady state. One exception is Aruoba et al. (2016), but they do not find evidence that the U.S. switched to the deflationary equilibrium following the Great Recession.

3.4 ESTIMATION PROCEDURE We estimate the model with quarterly data on per capita real GDP ($RGDP/CNP$), the GDP deflator (DEF), and the federal funds rate (FFR) from 1986Q1 to 2014Q2. A description of our data sources is given in Appendix B. The matrix of observables is

$$\hat{\mathbf{x}}^{data} \equiv [\log(RGDP_t/CNP_t) - \log(RGDP_{t-1}/CNP_{t-1}), \log(DEF_t/DEF_{t-1}), \log((1+FFR_t/100)^{1/4}].$$

We calibrate five poorly identified parameters (table 3). The steady-state discount factor, $\bar{\beta}$, is set to 0.9984, which equals $(1/T)\sum_{t=1}^{T}(1+G_t/400)(1+\Pi_t)/(1+FFR_t/100)^{1/4}$ where T is the sample size, G_k is the annual utilisation-adjusted growth rate of technology from Fernald (2012) and $\Pi_k = \log(DEF_k/DEF_{k-1})$. The leisure preference parameter, χ, is set so that steady-state labour equals $1/3$ of the available time. The elasticity of substitution between intermediate goods, θ, is set to 6, which corresponds to an average markup over marginal cost equal to 20%. The lower bound on the nominal interest rate, $\underline{\imath}$, is calibrated to 1.00017, which is the minimum federal funds rate. The Frisch labour supply elasticity, $1/\eta$, is set to 3, to match the estimate in Peterman (2016).

Steady-State Discount Factor	$\bar{\beta}$	0.9984	Nominal Interest Rate Lower Bound	$\underline{\imath}$	1.00017
Frisch Elasticity of Labour Supply	$1/\eta$	3	Real GDP Growth Rate Measurement Error SD	$\sigma_{me,\hat{y}}$	0.00194
Elasticity of Substitution between Goods	θ	6	Inflation Rate Measurement Error SD	$\sigma_{me,\pi}$	0.00075
Steady-State Labour	$\bar{n}$	0.33	Federal Funds Rate Measurement Error SD	$\sigma_{me,i}$	0.00206

Table 3: Calibrated parameters and measurement error standard deviations

The other parameters are estimated with a random walk Metropolis-Hastings algorithm that uses a particle filter to evaluate the likelihood of the posterior distribution. Following Herbst and Schorfheide (2016), we use a version of the particle filter that incorporates the information contained in the current observation, which helps the model better match outliers in the data. The filter uses $40{,}000$ particles and systematic resampling with replacement following Kitagawa (1996). Given the simulated paths from the model, we transform the predictions for real GDP, inflation, and the policy rate according to $\hat{\mathbf{x}}_t^{model} = [\log(g_t\tilde{y}_t^{gdp}/\tilde{y}_{t-1}^{gdp}), \log(\pi_t), \log(i_t)]$. The observables contain measurement error, so $\hat{\mathbf{x}}_t^{data} = \hat{\mathbf{x}}_t^{model} + \xi_t$, where $\xi \sim \mathbb{N}(0, \Sigma)$ is a vector of measurement errors and $\Sigma = \mathrm{diag}([\sigma^2_{me,\hat{y}}, \sigma^2_{me,\pi}, \sigma^2_{me,i}])$. Unlike linear models estimated with a Kalman filter, the particle filter requires some measurement error. Without any error, the particle filter suffers from degeneracy—a situation when all but a few weights on the particles are near zero. We set the variance of each measurement error to 10% of the variance of each observable (table 3), but we also estimated the model with alternative Σ's and found that it had little effect on our estimates.

We obtain $100{,}000$ draws from the joint posterior distribution and then keep every 100th draw. The remaining $1{,}000$ draws are used to calculate the posterior means, standard deviations, and $(5\%, 95\%)$ credible sets. The online appendix provides estimation diagnostics, including the trace plots, a comparison of the prior and posterior distributions, and the filtered shocks and observables.

The prior and posterior distributions are shown in table 4. Most of our priors are based on Gust et al. (2013), who also estimate a model with a ZLB constraint. In many cases, our priors are also similar to those used in Justiniano et al. (2011). We opted to use diffuse priors for most parameters. Three exceptions are the priors for the standard deviations of the structural shocks, which are less diffuse than in An and Schorfheide (2007) and Smets and Wouters (2007) because our constrained nonlinear model generates more volatility than linear models that do not impose a ZLB constraint.

Our parameter estimates are similar to the values in Gust et al. (2013). The Rotemberg price adjustment cost coefficient is equivalent to a Calvo price duration of roughly five quarters in a linear model. The persistence and standard deviation of the discount factor are large enough to explain movements in real GDP growth and inflation that drive the policy rate to its ZLB. The policy rate is highly persistent and the monetary response to the inflation and output gaps is stronger than most estimates obtained using data prior to 2008. We attribute that result to the long ZLB episode that occurred even though output growth and inflation remained relatively close to their targets from 2010 onward. Shocks to the steady state growth rate are not very persistent, but the standard deviation is large relative to the discount factor. The mean estimates of the annualised growth rate

	Prior			Posterior			
Parameter	Distribution	Mean	SD	Mean	SD	5%	95%
φ	Gamma	80.000	20.000	96.80137	20.03770	67.71867	131.85091
h	Beta	0.500	0.200	0.44428	0.08186	0.30733	0.57745
ρ_g	Beta	0.500	0.200	0.20064	0.09857	0.06547	0.37805
ρ_β	Beta	0.500	0.200	0.90245	0.01884	0.87001	0.92958
ρ_i	Beta	0.500	0.200	0.81158	0.03339	0.75375	0.86060
σ_ε	InvGamma	0.010	0.010	0.00968	0.00151	0.00738	0.01241
σ_υ	InvGamma	0.010	0.010	0.00215	0.00039	0.00159	0.00286
σ_ν	InvGamma	0.010	0.010	0.00199	0.00035	0.00148	0.00261
ϕ_π	Normal	2.500	1.000	4.06383	0.47424	3.33170	4.90267
ϕ_y	Normal	1.000	0.400	1.49057	0.23504	1.12702	1.87727
$\bar{g}$	Normal	1.004	0.001	1.00376	0.00068	1.00260	1.00489
$\bar{\pi}$	Normal	1.006	0.001	1.00622	0.00038	1.00556	1.00683

Table 4: Prior and posterior distributions of the estimated parameters. The last two columns report the 5th and 95th percentiles of the posterior distribution. The model is estimated using quarterly data from 1986Q1 to 2014Q2.

and inflation rate are 1.51% and 2.51%, respectively, which are slightly higher than the values in the data because they are unconditional and under-represent the effects of the long ZLB episode.[8]

4 MONETARY POLICY, THE ZLB CONSTRAINT, AND UNCERTAINTY

This section first explains why the ZLB constraint affects the relationship between real GDP growth and uncertainty and then shows whether the correlations in the model are consistent with the data.

4.1 MECHANISM We begin by showing the relationship between real GDP growth and its uncertainty with a generalised impulse response function (GIRF). The advantage of a GIRF is that it is based on an average of model simulations where the shocks are consistent with households' expectations over time. To compute the GIRF, we follow the procedure in Koop et al. (1996). Specifically, we first calculate the mean of $10{,}000$ model simulations, conditional on random shocks in every quarter (i.e., the baseline path). We then calculate a second mean from another set of $10{,}000$ simulations, but this time the shock in the first quarter is replaced with the shock of interest. The GIRF reports the difference between the two mean paths. See the online appendix for more details.

Figure 3 plots the responses to a 2 standard deviation positive discount factor (left panel), productivity growth (middle panel), and monetary policy (right panel) shock. The parameters are set to their posterior means, and the simulations are initialised at two different states. As a benchmark, the steady-state simulations (solid lines) are initialised at the stochastic steady state. We compare the benchmark responses to the responses when the ZLB first binds in the data (dashed lines) by initializing each of the simulations at the filtered state vector corresponding to 2008Q4.

A discount factor shock is the main mechanism that influences demand and causes ZLB events in our model. It is a proxy for a demand shock because it determines households' degree of patience. In either initial state, a discount factor shock causes households to postpone consumption, which reduces real GDP growth on impact. When the nominal interest rate is far from its ZLB,

[8]The online appendix compares unconditional moments and the duration of ZLB events in the model to the data.

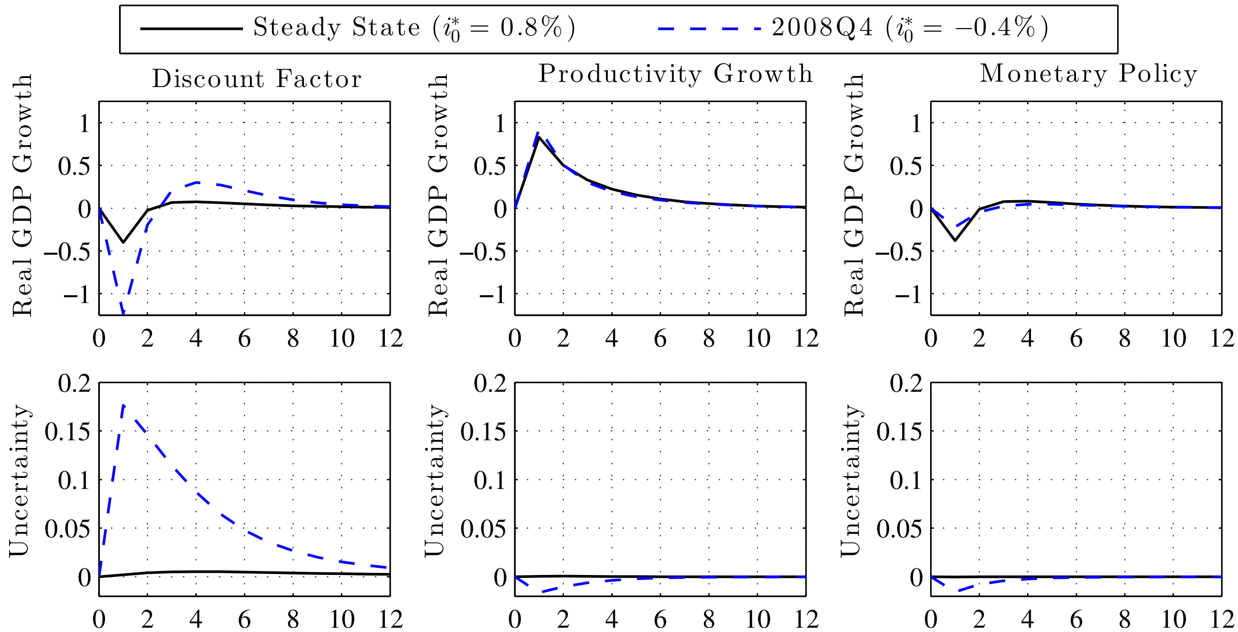

Figure 3: Generalised impulse responses to a 2 standard deviation positive shock at and away from the ZLB. The steady-state simulation (solid line) is initialised at the stochastic steady state. The other simulation (dashed line) is initialised at the filtered state corresponding to 2008Q4. The vertical axis is the percentage change in real GDP growth (or difference in uncertainty) from the baseline simulation. The horizontal axis displays the time period in quarters.

the drop in real GDP growth is damped by the monetary policy response. There is little change in uncertainty because households expect future shocks will have the same effect on real GDP regardless of the state of the economy. When the ZLB binds, however, the central bank cannot respond by lowering its policy rate, which leads to larger declines in real GDP. In that case, uncertainty sharply increases since households expect a wider range of future realisations of real GDP growth.

A positive productivity growth shock lowers the marginal cost of production, which causes firms to increase output, regardless of whether the ZLB binds. The central bank responds to the higher real GDP growth by increasing its policy rate. The shock has little effect on uncertainty when the economy begins in steady state but slightly decreases in the deep ZLB state, since the higher policy rate reduces the constraint on monetary policy. Similarly, a positive monetary policy shock raises the policy rate and reduces the likelihood of remaining at the ZLB. Therefore, there is a small decrease in uncertainty in the deep ZLB state, but in this case real GDP growth declines.

A comparison of the responses of real GDP growth suggests that when the ZLB does not bind productivity shocks have the largest influence on economic activity. They not only produce the biggest impact but are also the most persistent. In a deep recession that drives the policy rate to zero, however, discount factor shocks play a much more important role than in the pre-ZLB period.

4.2 ESTIMATES OF UNCERTAINTY

We generate the path of real GDP growth uncertainty predicted by the model by filtering the data at the posterior mean. Figure 4a plots that estimated path on the right axis (dashed line) along with a time series of per capita real GDP growth on the left axis (solid line). The two time series provide an estimate of the periods when uncertainty was elevated and an illustration of their historical relationship. The shaded area shows the ZLB period.

There are two key takeaways from our estimates. One, uncertainty about future real GDP

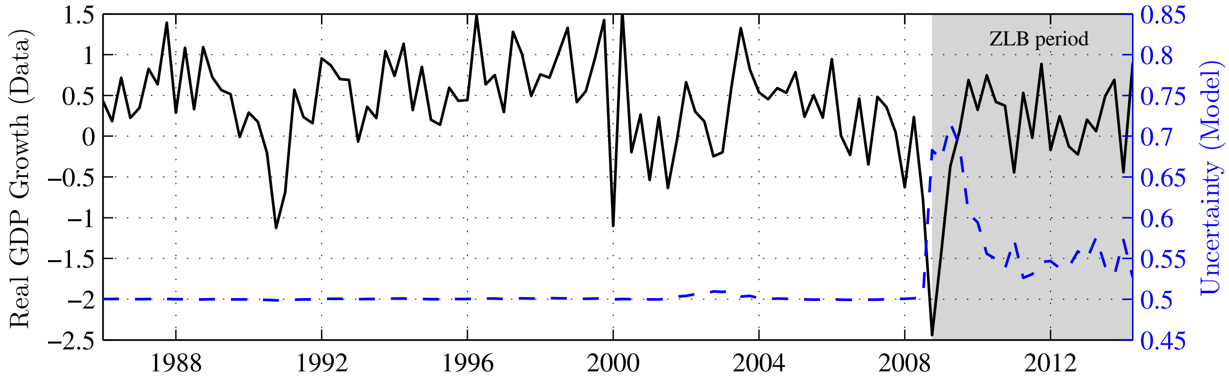

(a) Model without stochastic volatility

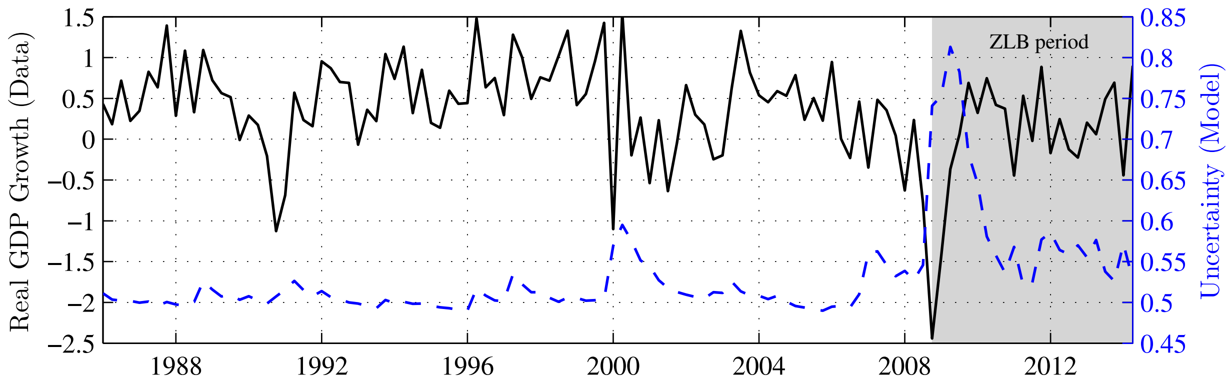

(b) Model with stochastic volatility

Figure 4: Time series of real GDP growth (solid line) and the mean path of real GDP growth uncertainty (dashed line).

growth is state-dependent. When the ZLB does not bind, uncertainty is essentially constant, except in quarters when the nominal interest rate is near its ZLB. In those situations, the high probability of hitting the ZLB next period leads to persistently higher uncertainty. The closer the nominal interest rate is to the ZLB, the higher the uncertainty, which underscores the importance of expectational effects. When the ZLB binds, uncertainty is as much as 42% higher than its value outside the ZLB. The amount of uncertainty depends on the notional interest rate. The lower the notional interest rate, the less likely it is the nominal interest rate will exit the ZLB and the higher the uncertainty.

Two, it is possible for real GDP growth uncertainty to decline while the ZLB binds, which implies that uncertainty is time-varying at the ZLB. That outcome occurs whenever the central bank is constrained and there is a negative discount factor shock. A lower discount factor means households are more optimistic about the future economy, which increases the expected policy rate and reduces real GDP growth uncertainty. This feature of the model is important because macroeconomic uncertainty continued to fluctuate in the data even after many central banks reduced their policy rates to zero. In other words, our theory does not claim that uncertainty is always elevated when the ZLB binds, rather it is more strongly correlated with real GDP growth and the strength of that correlation depends on the severity of the recession and the likelihood of staying at the ZLB.

In the data, uncertainty is time-varying in the pre-ZLB period. However, it is virtually constant in figure 4a when there is no expectation of hitting the ZLB, since the ZLB constraint is the most important nonlinearity and by far the strongest source of time-varying uncertainty in the model.

Therefore, it is hard to assess the correlation between real GDP growth and uncertainty when the ZLB does not bind. There is also no way to determine the relative importance of exogenous and endogenous uncertainty. To overcome those drawbacks, we follow the SV literature and add an exogenous source of time-varying uncertainty about real GDP growth. There are several types of SV shocks we could introduce. We assume the standard deviation of productivity growth follows

$$\log \sigma_{\varepsilon,t} = (1-\rho_\sigma)\log \bar{\sigma}_\varepsilon + \rho_\sigma \log \sigma_{\varepsilon,t-1} + \sigma_x x_t, \tag{8}$$

where $\bar{\sigma}_\varepsilon$ is the average standard deviation, x is a standard normal shock, and $0 \leq \rho_\sigma < 1$ is the persistence of the shock. We believe this specification gives the model the best chance to take into account the numerous factors that affect the amount of uncertainty surrounding real GDP growth.

We would prefer to estimate the model with SV. Unfortunately, it adds another state variable and complicates the numerical integration, which makes it too costly to estimate. Therefore, we estimate the parameters in (8) by searching for the posterior mode of (ρ_σ, σ_x), conditional on the posterior mean from the baseline model. We then solve the SV model for each posterior draw from the baseline model given the mode for the SV parameters, $(\rho_\sigma, \sigma_x) = (0.737, 0.128)$, and filter the data with those solutions. The online appendix provides more information about the mode search.

Figure 4b reproduces the series in figure 4a for the model with SV. In this case, uncertainty about real GDP growth fluctuates at and away from the ZLB, but the level and volatility of uncertainty are still higher in the ZLB sample than the pre-ZLB sample. Comparing the two figures, it is apparent that adding an exogenous uncertainty shock is not necessary to generate significant movement in uncertainty at the ZLB. Although SV slightly magnifies the increase in uncertainty, most of the uncertainty at the ZLB appears to be endogenous rather than from exogenous factors.

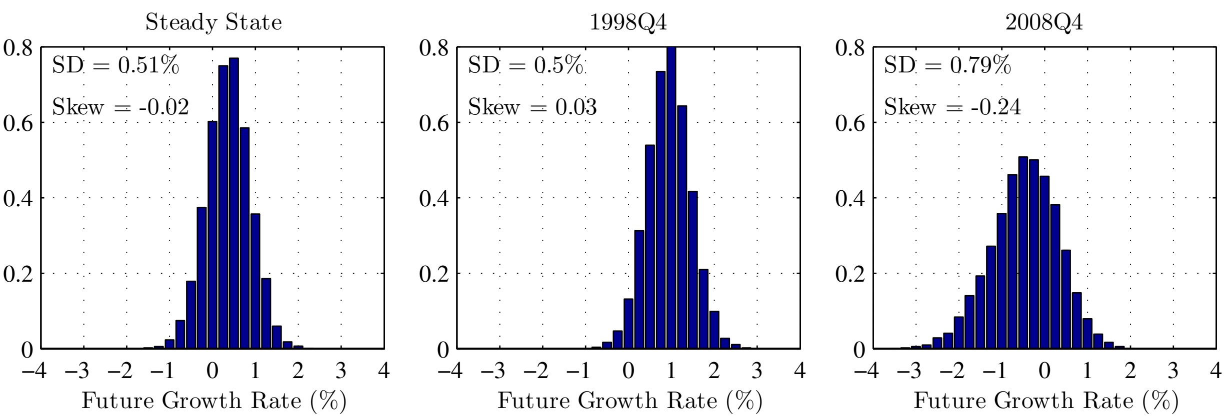

Figure 5: Histograms of 1-quarter ahead real GDP growth rate forecasts, based on 10,000 Monte Carlo simulations of our model with SV. Each simulation is initialised at the indicated state using the posterior mean parameterization.

Another way to visualise the time-variation in uncertainty is to look at the standard deviation and skewness of the distribution of future real GDP growth at different points in time. Figure 5 plots the distributions of the 1-quarter ahead forecasts of real GDP growth in our model with SV. To generate those distributions, we simulate the model $10{,}000$ times conditional on the posterior mean and initialise each simulation at the stochastic steady state (left panel) and the filtered states corresponding to 1998Q4 (middle panel) and 2008Q4 (right panel). In steady state, future real GDP forecasts have a standard deviation of 0.54% and a skewness of -0.02. A quarter that is representative of the forecasts made in steady state is 1998Q4. In that quarter, the policy rate was far from the ZLB, and real GDP growth forecasts also had a standard deviation of 0.54% with a skewness near zero. At that time, the Fed was not constrained and, if necessary, it could have responded to lower inflation or real GDP growth by cutting the federal funds rate to dampen the effects of a fall in demand. Therefore, the distribution of real GDP growth forecasts was fairly symmetric and uncertainty was low. In 2008Q4, however, sufficiently large shocks, primarily to the discount factor, caused large enough declines in demand that the Fed was compelled to cut its policy rate to zero. Since the Fed was unable to further reduce its policy rate, the volatility of the forecasts rose to 0.81% and the skewness was -0.18. In other words, the economy became more sensitive to further declines in demand, which dramatically increased macroeconomic uncertainty.

	Pre-ZLB Sample (1986Q1-2008Q3)	ZLB Sample (2008Q4-2014Q4)	Differences (ZLB - Pre-ZLB)
Without SV	-	-0.48***	-
With SV	-0.19***	-0.54***	-0.34***

Table 5: Correlations between per capita real GDP growth and our estimate of real GDP growth uncertainty. Row 1 is our baseline model with constant shock volatilities. Row 2 introduces time-varying volatility about productivity growth in the baseline model. An asterisk indicates a correlation in the pre-ZLB or ZLB sample is statistically less than 0 or the difference between the correlations in the two samples is significant at a ***1%, **5%, and *10% level.

4.3 CORRELATIONS The empirical measures introduced earlier indicate that a stronger correlation between real GDP growth and uncertainty emerged in late 2008. To determine the range of correlations from our structural model, we filter the data and generate a time series of real GDP growth uncertainty for each draw from our posterior distribution. We then calculate the correlation between per capita real GDP growth and our filtered uncertainty series. Table 5 reports the median correlations from our baseline model without SV (row 1) and our model that adds SV (row 2). In our baseline model, the median correlation in the ZLB sample is -0.48, and every draw is less than zero, so SV is not necessary to generate a negative correlation in the ZLB period. We do not report the pre-ZLB correlation since it is uninformative. In our model with SV, the median correlation in the ZLB period is -0.54, which is slightly more negative than in the model without SV, and every correlation is still negative. The pre-ZLB median correlation is -0.19, which is smaller but still statistically less than zero. To determine whether the differences between the two correlations are statistically significant, we subtract the pre-ZLB correlations from the ZLB correlations just like in the VAR model. The median difference is -0.34 and the distribution has a negative support, which is consistent with our finding that a stronger negative correlation emerged during the ZLB period.

Using the same parameters and filtered states that take into account the entire sample, we also compute correlations in our model with SV using two counterfactual ZLB samples—one that removes 2008Q4 and one that omits 2008Q4 and 2009Q1. Those quarters are important because they represent a period when there were significant declines in real GDP growth in conjunction with a sharp decline in the notional rate that generated a dramatic increase in uncertainty. However, several other events were happening at the beginning of the ZLB period. For example, Lehman Brothers filed for bankruptcy on September 15, 2008, which temporarily paralysed financial markets and increased uncertainty. In spite of removing these quarters, the median correlations are only slightly less negative (-0.45 from 2009Q1-2014Q2 and -0.26 from 2009Q2-2014Q2) and none of the correlations are positive. Furthermore, the ZLB correlation is significantly less than the pre-ZLB correlation at a 1% level when 2008Q4 is removed and a 10% level when both 2008Q4 and 2009Q1 are omitted. Those results provide evidence that the strong negative correlation in the ZLB sample is not solely the result of factors that are unique to the beginning of the ZLB period.

5 Additional Results and Supporting Evidence

This section first compares our measures of uncertainty in the data to our structural estimates. It then tests whether a stronger correlation with inflation uncertainty also emerged in the ZLB period.

Pre-ZLB Sample / ZLB Sample	DSGE	DSGE SV	VAR	VXO	BOS FD	SPF FD	JLN
DSGE	–	0.18	0.07	0.40	−0.35	0.08	0.05
DSGE SV	0.99	–	0.70	0.29	−0.03	−0.02	0.41
VAR	0.77	0.76	–	0.43	0.09	0.17	0.57
VXO	0.67	0.67	0.86	–	0.11	0.46	0.30
BOS FD	0.22	0.20	0.54	0.60	–	0.13	0.25
SPF FD	0.83	0.83	0.75	0.74	0.34	–	0.35
JLN	0.87	0.87	0.92	0.91	0.52	0.88	–

Table 6: Correlations between the various uncertainty measures. The pre-ZLB sample is based on data from 1986Q1 to 2008Q3 and the ZLB sample is based on data from 2008Q4 to 2014Q2 for all of the uncertainty measures. DSGE, DSGE SV, and VAR are direct measures of real GDP growth uncertainty, while the others are indirect measures.

5.1 Uncertainty Measures

One way to validate our uncertainty measures is to look at their relationships. Table 6 plots the correlations between the various measures, including the estimates from our DSGE model with and without SV, our time-varying parameter VAR with SV, the VXO, survey-based forecast dispersion, and the JLN macro uncertainty index. In the upper triangle are correlations in the pre-ZLB sample and the lower triangle shows correlations in the ZLB sample.

There are several important takeaways. First, most of the correlations are positive, including every value in the ZLB sample. Second, all of the correlations are stronger in the ZLB sample than the pre-ZLB sample, which indicates that there is a unique feature of the data that affected every measure of uncertainty in that sample. Third, in the ZLB sample all of the uncertainty measures, except the BOS FD, are strongly correlated with the measures from our structural models. The strongest correlation is with the JLN measure, which is significant because it is based on the broadest information set. Fourth, the correlations in both samples are stronger when our model includes exogenous uncertainty shocks. For example, in the pre-ZLB sample the correlation between the VAR and our baseline model is weak, but it is much stronger when the model includes SV. That result is not surprising because ZLB is by far the most important mechanism for generating time-varying uncertainty in our baseline model. Besides 2004, there were essentially no instances when the economy was near the ZLB in the pre-ZLB sample. However, the fact that our DSGE model with SV is positively correlated with both the VAR and the JLN uncertainty measures indicates that the SV process captures many of the movements in uncertainty that occur in the pre-ZLB sample.

	-3	-2	-1	0	1	2	3
DSGE	-0.17*	-0.14*	-0.07	0.02	0.07	0.14*	0.16*
DSGE SV	-0.13	-0.21**	-0.24**	-0.20**	-0.13	-0.26***	-0.29***
VAR	-0.18**	-0.28***	-0.29***	-0.35***	-0.31***	-0.29***	-0.26***
VXO	-0.07	0.03	-0.09	-0.09	-0.11	0.00	0.04
BOS FD	0.34***	0.17*	-0.11	-0.19**	-0.14*	-0.15*	-0.14*
SPF FD	-0.11	-0.17*	-0.24**	-0.12	-0.12	-0.09	-0.06
JLN	-0.16*	-0.31***	-0.35***	-0.45***	-0.41***	-0.43***	-0.45***

(a) Pre-ZLB sample (1986Q1-2008Q3)

	-3	-2	-1	0	1	2	3
DSGE	-0.83***	-0.84***	-0.62***	-0.62***	-0.19	0.04	0.27
DSGE SV	-0.85***	-0.82***	-0.60***	-0.59***	-0.20	0.03	0.24
VAR	-0.54***	-0.67***	-0.79***	-0.79***	-0.55***	-0.16	0.13
VXO	-0.33*	-0.48**	-0.61***	-0.72***	-0.55***	-0.32*	-0.02
BOS	0.42**	0.07	-0.33*	-0.70***	-0.39**	-0.27	-0.15
SPF	-0.53***	-0.64***	-0.64***	-0.47***	-0.32*	0.03	0.12
JLN	-0.72***	-0.79***	-0.74***	-0.73***	-0.47**	-0.09	0.15

(b) ZLB sample (2008Q4-2014Q2)

Table 7: Correlations between real GDP growth j periods ahead and current real GDP growth uncertainty (i.e., $\mathrm{corr}(\sigma_{\hat{y}^{gdp},t}, \hat{y}_{t+j})$). An asterisk indicates a correlation in the pre-ZLB or ZLB sample is statistically less than 0 or the difference between the correlations in the two samples is significant at a ***1%, **5%, and *10% level.

5.2 CROSS-CORRELATIONS

Throughout our analysis, we have focused on contemporaneous correlations. Our theory predicts that changes in real GDP growth will affect the expected volatility of the one-period ahead forecast error for real GDP growth (i.e., our measure of uncertainty) in the same period that the nominal interest rate is near or at its ZLB. The correlations between uncertainty and leads and lags of economic activity provide an additional check on the model's predictions. In our baseline model without SV, the primary driver of the correlation is the ZLB constraint. If there is a large shock that drives the economy deep into the ZLB, such as the one that occurred during the Great Recession, and the economy remains in that state for some time, a strong negative correlation between contemporaneous uncertainty and lags of real GDP growth is likely to occur. In our model with SV, volatility shocks to productivity growth are the primary driver of the correlation between uncertainty and real GDP growth in the pre-ZLB sample. That model still predicts a negative contemporaneous correlation, but it is probable that a negative correlation with future real GDP growth would also occur since positive volatility shocks reduce economic activity.

Table 7 shows the contemporaneous correlations between real GDP growth and various measures of macroeconomic uncertainty as well as the same correlations at leads and lags of real GDP growth. The results from our structural models provide evidence for our intuition. In the pre-ZLB period, the model without SV predicts a weak correlation between contemporaneous uncertainty and both leads and lags of real GDP growth. However, once we allow for the possibility of SV shocks, we find negative correlations with leads of real GDP growth, which are generally much stronger than with the lags. During the ZLB period, we find a very strong correlation between contemporaneous uncertainty and both contemporaneous real GDP growth and lags of that variable. The patterns of the correlations with the JLN measure of uncertainty at the various leads and lags are very similar to our DSGE model with SV. With the exception of the BOS FD, which is based on a narrower measure of economic activity, all of the measures have similar qualitative features.

	DSGE	DSGE SV	VAR	SPF FD CPI
Pre-ZLB Sample (1986Q1-2008Q3)	-0.23***	-0.28***	-0.19**	-0.25***
ZLB Sample (2008Q4-2014Q2)	-0.53***	-0.55***	-0.30	-0.55***
Difference (ZLB-Pre-ZLB)	-0.30***	-0.27***	-0.13	-0.31*

Table 8: Correlations between real GDP growth and various measures of inflation uncertainty before and during the ZLB period. An asterisk indicates a correlation in the pre-ZLB or ZLB sample is statistically less than 0 or the difference between the correlations in the two samples is significant at a ***1%, **5%, and *10% level.

5.3 CORRELATION WITH INFLATION UNCERTAINTY

We have focused on the correlation between real GDP growth and its uncertainty. This section examines the correlation between real GDP growth and inflation uncertainty. Table 8 shows estimates from our structural models, our time-varying parameter VAR with SV, and the dispersion in forecasts of inflation. The uncertainty series from our DSGE models are based on the expected volatility of the inflation rate forecast error 1-quarter ahead from the particle filter. The series from the VAR is based on the estimated volatility of the inflation rate from the covariance matrix, $A_t^{-1}\Sigma_t$, and the forecast dispersion equals the inter-quartile range of the individual SPF forecasts of the CPI inflation rate 1-quarter ahead.

In our DSGE models, there is a much stronger correlation between real GDP growth and inflation uncertainty in the ZLB sample than the pre-ZLB sample. The correlations in the pre-ZLB sample are statistically below zero, but the difference between the ZLB and pre-ZLB correlations is positive in less than 1% of draws in the model with and without SV. In the VAR model, the median correlation in the pre-ZLB sample is also statistically below zero, but it is less negative than the correlation in the ZLB sample. The correlations based on SPF forecast dispersion are significantly negative in both the pre-ZLB and ZLB samples. A z-transformation test, however, shows the correlations in the two samples are significantly different at a 10% level. These results provide evidence that a stronger correlation between per capita real GDP growth and inflation uncertainty emerged when the federal funds rate hit its ZLB in 2008Q4, just like with real GDP growth uncertainty.

6 CONCLUSION

This paper finds a stronger negative correlation between macroeconomic uncertainty and real GDP growth emerged in late 2008. Before then the correlation was weak and sometimes not statistically below zero. We provide empirical evidence using estimates from a time-varying parameter VAR with SV as well as several other measures, including stock market volatility, survey-based forecast dispersion, and the uncertainty index from Jurado et al. (2015). We contend the ZLB constraint on the federal funds rate contributed to the stronger correlation. To test our theory, we estimate a nonlinear New Keynesian model with an occasionally binding ZLB constraint and generate a data-driven, forward looking measure of uncertainty from the model. The correlations in the ZLB period are stronger and statistically more negative than the values in the pre-ZLB period, just like we found in the data. We also find that all of the uncertainty measure are positively correlated and most of them developed a strong negative correlation with lags of real GDP growth. There is also evidence that a strong negative correlation with inflation uncertainty emerged in the ZLB period.

Our work presents avenues for future research. One, it is possible to calculate the measure of uncertainty from our DSGE model for other variables in closed and open economy models, which could be used to explain the behaviour of uncertainty in the data. Second, the ZLB generates uncertainty because it creates a kink in the policy functions. It stands to reason that other nonlinearities, such as irreversible investment, may have similar impacts on uncertainty that can explain data.

REFERENCES

AN, S. AND F. SCHORFHEIDE (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113–172.

ARUOBA, S., P. CUBA-BORDA, AND F. SCHORFHEIDE (2016): "Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries," Manuscript, University of Maryland.

BACHMANN, R. AND C. BAYER (2013): "'Wait-and-See' Business Cycles?" *Journal of Monetary Economics*, 60, 704–719.

BACHMANN, R. AND G. MOSCARINI (2012): "Business Cycles and Endogenous Uncertainty," Manuscript, RWTH Aachen University.

BASU, S. AND B. BUNDICK (2014): "Uncertainty Shocks in a Model of Effective Demand," FRB Kansas City Working Paper 14-15.

——— (2015): "Endogenous Volatility at the Zero Lower Bound: Implications for Stabilization Policy," FRB Kansas City Working Paper 15-1.

BEKAERT, G., M. HOEROVA, AND M. LO DUCA (2013): "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics*, 60, 771–788.

BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): "The Perils of Taylor Rules," *Journal of Economic Theory*, 96, 40–69.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2014): "Really Uncertain Business Cycles," Center for Economic Studies Working Paper 14-18.

BORN, B. AND J. PFEIFER (2014): "Policy Risk and the Business Cycle," *Journal of Monetary Economics*, 68, 68–85.

CHOI, S. (2013): "Are the Effects of Blooms Uncertainty Shocks Robust?" *Economics Letters*, 119, 216–220.

CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

COLEMAN, II, W. J. (1991): "Equilibrium in a Production Economy with an Income Tax," *Econometrica*, 59, 1091–1104.

CUBA-BORDA, P. (2014): "What Explains the Great Recession and the Slow Recovery?" Manuscript, University of Maryland.

DEL NEGRO, M. AND G. E. PRIMICERI (2015): "Time Varing Structural Vector Autoregressions and Monetary Policy: A Corrigendum," *Review of Economic Studies*, 82, 1342–1345.

DIXIT, A. K. AND J. E. STIGLITZ (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308.

EGGERTSSON, G. B. AND M. WOODFORD (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34(1), 139–235.

ELDER, J. AND A. SERLETIS (2010): "Oil Price Uncertainty," *Journal of Money, Credit and Banking*, 42, 1137–1159.

FAJGELBAUM, P., M. TASCHEREAU-DUMOUCHEL, AND E. SCHAAL (2014): "Uncertainty Traps," NBER Working Paper 19973.

FERNALD, J. (2012): "Productivity and Potential Output before, during, and after the Great Recession," Federal Reserve Bank of San Francisco Working Paper 2012-18.

FERNÁNDEZ-VILLAVERDE, J., G. GORDON, P. GUERRÓN-QUINTANA, AND J. F. RUBIO-RAMÍREZ (2015): "Nonlinear Adventures at the Zero Lower Bound," *Journal of Economic Dynamics and Control*, 57, 182–204.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–61.

GILCHRIST, S., J. W. SIM, AND E. ZAKRAJŠEK (2014): "Uncertainty, Financial Frictions, and Investment Dynamics," NBER Working Paper 20038.

GOURIO, F. (2014): "Financial Distress and Endogenous Uncertainty," Manuscript, Federal Reserve Bank of Chicago.

GUST, C., D. LÓPEZ-SALIDO, AND M. E. SMITH (2013): "The Empirical Implications of the Interest-Rate Lower Bound," CEPR Discussion Paper 9214.

HERBST, E. P. AND F. SCHORFHEIDE (2016): *Bayesian Estimation of DSGE Models*, Princeton University Press.

IRELAND, P. N. (1997): "A Small, Structural, Quarterly Model for Monetary Policy Evaluation," *Carnegie-Rochester Conference Series on Public Policy*, 47, 83–108.

JO, S. (2014): "The Effects of Oil Price Uncertainty on Global Real Economic Activity," *Journal of Money, Credit and Banking*, 46, 1113–1135.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 14, 101–121.

JUSTINIANO, A. AND G. E. PRIMICERI (2008): "The Time-Varying Volatility of Macroeconomic Fluctuations," *American Economic Review*, 98, 604–41.

KITAGAWA, G. (1996): "Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models," *Journal of Computational and Graphical Statistics*, 5, pp. 1–25.

KOOP, G. AND D. KOROBILIS (2010): "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends(R) in Econometrics*, 3, 267–358.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 74, 119–147.

KOPECKY, K. AND R. SUEN (2010): "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 13, 701–714.

LEDUC, S. AND Z. LIU (2014): "Uncertainty Shocks are Aggregate Demand Shocks," FRB San Francisco Working Paper 2012-10.

LELAND, H. E. (1968): "Saving and Uncertainty: The Precautionary Demand for Saving," *The Quarterly Journal of Economics*, 83, 465–473.

LEVHARI, D. AND T. N. SRINIVASAN (1969): "Optimal Savings under Uncertainty," *Review of Economic Studies*, 36, 153–63.

LUDVIGSON, S. C., S. MA, AND S. NG (2015): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" NBER Working Paper 21803.

MUMTAZ, H. AND F. ZANETTI (2013): "The Impact of the Volatility of Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 45, 535–558.

NAKATA, T. (2012): "Uncertainty at the Zero Lower Bound," Finance and Economics Discussion Series, 2013-09.

NAKOV, A. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4, 73–127.

NAVARRO, G. (2014): "Financial Crises and Endogenous Volatility," Manuscript, New York University.

PETERMAN, W. B. (2016): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," *Economic Inquiry*, 54, 100–120.

PINTER, G., K. THEODORIDIS, AND T. YATES (2013): "Risk News Shocks and the Business Cycle," Bank of England Working Paper 483.

PRIMICERI, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852.

RICHTER, A. W. AND N. A. THROCKMORTON (2015): "The Zero Lower Bound: Frequency, Duration, and Numerical Convergence," *B.E. Journal of Macroeconomics*, 15, 157–182.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

ROUWENHORST, K. G. (1995): "Asset Pricing Inplications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton University Press, 294–330.

SANDMO, A. (1970): "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies*, 37, 353–60.

SIMS, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

VAN NIEUWERBURGH, S. AND L. VELDKAMP (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53, 753–772.

WU, J. C. AND F. D. XIA (2016): "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48, 253–291.

A SOLUTION METHOD

A formal description of the numerical algorithm begins by writing the model compactly as

$$\mathbb{E}[f(\mathbf{v}_{t+1}, \mathbf{w}_{t+1}, \mathbf{v}_t, \mathbf{w}_t)|\Omega_t] = 0,$$

where f is a vector-valued function, $\mathbf{v}$ are the exogenous variables, $\mathbf{w} = (\tilde{w}, i, i^*, \pi, \tilde{c}, \tilde{y})$ are the endogenous variables, and $\Omega = \{S, P, \mathbf{z}\}$ is the information set, which contains the structural model, S, its parameters, P, and the state vector, $\mathbf{z}$. In the baseline model, $\mathbf{v} = (g, \beta)$ and $\mathbf{z}_t =$

$(\tilde{c}_{t-1}, i^*_{t-1}, g_t, \beta_t, \nu_t)$. We discretise the state variables into $(7, 7, 8, 12, 8)$ points. The bounds on the endogenous state variables, $\tilde{c}_{t-1}$ and i^*_{t-1}, are set to $\pm 4\%$ and $\pm 2.5\%$ of steady state. In the model with SV, $\mathbf{v} = (g, \beta, \sigma_\varepsilon)$ and $\mathbf{z}_t = (\tilde{c}_{t-1}, i^*_{t-1}, g_t, \beta_t, \nu_t, \sigma_{\varepsilon,t})$. We use the same number of points as the basline model, except we increase the number on g_t to 11 and put 8 points on $\sigma_{\varepsilon,t}$.

The following steps outline our policy function iteration algorithm:

1. Obtain initial conjectures for $\tilde{c}^A_0$ and π^A_0 on each node from the log-linear model without the ZLB constraint imposed. We obtain those values with Sims's (2002) `gensys` algorithm.
2. For iteration $i \in \{1, \ldots, I\}$ and node $d \in \{1, \ldots, D\}$, implement the following steps:
   (a) Solve for $\{\tilde{w}_t, \tilde{y}_t, i^*_t, i_t\}$ given $\tilde{c}^A_{i-1}(\mathbf{z}^d_t)$ and $\pi^A_{i-1}(\mathbf{z}^d_t)$ with the ZLB constraint imposed.
   (b) Use linear interpolation to update consumption and inflation, $\{\tilde{c}_{t+1}, \pi_{t+1}\}^M_{m=1}$, given the updated state vector, $\mathbf{z}_{t+1}$. In our baseline model, we calculate $\{g^m_{t+1}, \beta^m_{t+1}, \nu^m_{t+1}\}^M_{m=1}$ following Rouwenhorst (1995), which Kopecky and Suen (2010) show outperforms other methods. In the model with SV, we discretise the productivity growth process with Gauss-Hermite quadrature since it depends on the standard deviation of the shock.
   (c) Given $\{\tilde{c}_{t+1}, \pi_{t+1}\}^M_{m=1}$, solve for the rest of $\mathbf{x}^m_{t+1}$ and compute
   $$\mathbb{E}\left[f(\mathbf{x}^m_{t+1}, \mathbf{x}^d_t)|\Omega^d_t\right] \approx \textstyle\sum_{m=1}^M \phi(g^m_{t+1}, \beta^m_{t+1}, \nu^m_{t+1}) f(\mathbf{x}^m_{t+1}, \mathbf{x}^d_t),$$
   where $\mathbf{x} \equiv (\mathbf{v}, \mathbf{w})$, and ϕ are the respective Rouwenhorst weights. The superscripts on $\mathbf{x}$ indicate which realisations of the state variables are used to compute expectations.
   (d) Use Chris Sims' `csolve` to find policy functions that satisfy $\mathbb{E}[f(\cdot)|\Omega_t] = 0$.
3. Define $\text{maxdist}_i \equiv \max\{|\tilde{c}^A_i - \tilde{c}^A_{i-1}|, |\pi^A_i - \pi^A_{i-1}|\}$. Repeat step 2 until $\text{maxdist}_i < 10^{-7}$ for all d. When that condition occurs, the algorithm has converged to an approximate solution.

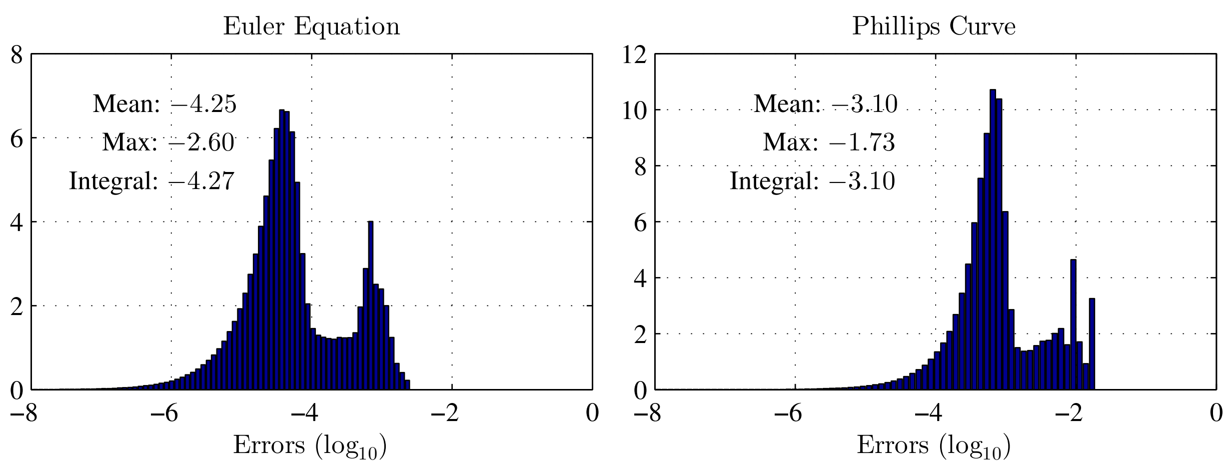

Figure 6: Distribution of Euler equation errors in base 10 logarithms

Richter et al. (2014) use Euler Equation errors to show the accuracy of the solution to a model with a ZLB constraint. To measure errors between nodes, we use Gauss-Hermite quadrature instead of the Rouwenhorst method, which forces exogenous variables to be on the grid. Figure 6 shows the distribution of the absolute value of the errors in base 10 logarithms for the consumption

Euler equation and Phillips curve in the baseline model. For example, an Euler equation error of -3 means the household makes an error equal to one of every $1{,}000$ consumption goods. We also report the mean, maximum, and integral, which is a weighted average of the errors at each state vector using its probability distribution. Both the mean and maximum errors are similar to those in Cuba-Borda (2014), which are based on a model with a ZLB constraint and capital accumulation.

B DATA SOURCES

The following series are available on the Federal Reserve Economic Database (FRED):

U.S. Real GDP: Chained 2009 dollars, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 1.1.6. (FRED ID: GDPC1).

U.S. Population: Quarterly average of monthly values in thousands of persons. Source: U.S. Bureau of Labour Statistics, Current Population Survey. (FRED ID: CNP16OV).

U.S. GDP Deflator: Index 2009=100, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 1.1.9. (FRED ID: GDPDEF).

Federal Funds Rate: Quarterly average of monthly values. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates (Monthly), H.15. (FRED ID: FEDFUNDS).

U.S. Industrial Production: Index 2007 = 100, Total index, seasonally adjusted quarterly averages of monthly values. Source: Board of Governors of the Federal Reserve System, Industrial Production and Capacity Utilisation, G.17. (FRED ID: INDPRO).

U.S. 10-Year Treasury Rate: Quarterly average of monthly values. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates (Monthly), H.15. (FRED ID: GS10).

The following series are available from other sources:

U.S. VXO: Expected volatility in the S&P 100 over the next 30 days at an annualised rate. We calculate a quarterly average of the daily observations. Source: Chicago Board Options Exchange, VIX Historical Price Data ("old methodology").

U.S. BOS: Future general activity; percent of manufacturers that forecast a decrease (GAFDSA), an increase (GAFISA) and no change (GAFNSA). Source: Federal Reserve Bank of Philadelphia, Business Outlook Survey, revised monthly data.

U.S. SPF: Cross-sectional forecasts dispersion of Real GDP Growth, IP Growth, and CPI. Source: Federal Reserve Bank of Philadelphia, Dispersion measure D2 (D1 for CPI).

Fernald (2012) TFP: Quarterly values. Source: Federal Reserve Bank of San Francisco, Total Factor Productivity, TFP measure dtfp_util.

Jurado et al. (2015) Macro Uncertainty Index: Quarterly average of monthly values. Source: Authors' personal websites, Uncertainty measure $h = 3$ (one quarter forecast horizon).

Ludvigson et al. (2015) Financial Uncertainty: Quarterly average of monthly values. Source: Authors' personal websites, Uncertainty measure $h = 3$ (one quarter forecast horizon).

Wu and Xia (2016) Shadow Rate: Quarterly average of monthly values. Source: Federal Reserve Bank of Atlanta, Wu-Xia Rate.

All of the data are available for our full sample from 1986Q1 to 2014Q2.