

Uncertainty shocks as second-moment news shocks

David Berger, Ian Dew-Becker, and Stefano Giglio*

November 20, 2018

Abstract

We provide evidence on the relationship between aggregate uncertainty and the macroeconomy. Identifying uncertainty shocks using methods from the news shocks literature, the analysis finds that innovations in realized stock market volatility are robustly followed by contractions, while shocks to forward-looking uncertainty have no significant effect on the economy. Moreover, investors have historically paid large premia to hedge shocks to realized but not implied volatility. A model in which fundamental shocks are skewed left can match those facts. Aggregate volatility matters, but it is the *realization* of volatility, rather than uncertainty about the future, that has been associated with declines.

*Berger: Northwestern University and NBER. Dew-Becker: Northwestern University and NBER. Giglio: Yale University and NBER. This paper was previously circulated and presented under the title “Contractionary volatility or volatile contractions?” We appreciate helpful comments and discussions from Robert Barsky, Nick Bloom, Gideon Bornstein, Brent Bundick, Nicolas Crouzet, Larry Christiano, Guido Lorenzoni, Aytek Malkhozov, Stephen Terry (who also generously provided simulation results), Amir Yaron, and seminar participants at Booth, CITE, SITE, the Federal Reserve Banks of Boston, Chicago, Richmond, and San Francisco, Northwestern, the BI-SHoF finance conference, LBS, the NBER Summer Institute, UCL Workshop on Uncertainty, the SED, AFA, and WFA, Arizona State, Yale, Texas A&M, and the SFS Cavalcade.

1 Introduction

This paper aims to estimate the magnitude of the effects of aggregate uncertainty shocks on economic activity. Uncertainty about the economy varies substantially over time, countercyclically, and often driven by policy choices. A growing literature has developed a range of theoretical channels through which uncertainty shocks might cause recessions, though in most of those models the sign of the effect of uncertainty shocks is actually ambiguous.¹ Our analysis directly applies to those theories by quantifying how the economy responds to identified shocks to uncertainty. On a more practical level, it lets us ask whether events like Brexit or elections can have negative impacts on the economy purely through their effects on uncertainty, regardless of any direct first-moment effects they may have.

The major identification problem we face is that uncertainty about the future tends to be related to current economic conditions. The arrival of large shocks – i.e. the *realization* of volatility – is correlated with increases in expected future volatility, that is, uncertainty about the future (sometimes called GARCH effects, after the models of Engle (1982) and Bollerslev (1986)). Our key distinction from past work is to construct shocks to uncertainty that are orthogonal to current volatility. Current volatility measures how large are the shocks that have just occurred, whereas uncertainty is about how large agents expect *future* shocks to be. Models of the effects of uncertainty, such as those with wait-and-see or precautionary saving effects, are driven by variation in agents’ subjective distributions of future shocks, as opposed to the variance of the distribution from which today’s shocks were drawn.

Past work, such as Bloom (2009), found that shocks to stock market volatility, measured as a mixture of realized and expected future volatility, have negative effects on the economy. We show that the distinction between current volatility and news about the future is critical for understanding the effects of uncertainty shocks. When uncertainty shocks are properly defined as forward looking, they actually have no effects on the economy. Rather it is current realized volatility that is associated with downturns and drives the previous results. Moreover, an identification scheme separating realized from expected volatility accurately uncovers the true structural uncertainty shocks in two recent state-of-the-art models of the effects of aggregate uncertainty from Basu and Bundick (2017) and Bloom et al. (2017).

Empirically, the identified uncertainty shocks are not small, accounting for 30–60 percent of the variation in total uncertainty, and they are associated with economically plausible uncertainty events, such as the debt ceiling debates and the Euro crisis in 2010–2012. Those events are particularly useful for identification because they increased uncertainty but did not have any direct effect on the economy on impact, unlike, for example, the Financial Crisis or 9/11. In a vector autoregression (VAR), the identified uncertainty shocks are found to have no significant effect on real activity, with confidence bands that we argue are meaningfully tight. Uncertainty about the

¹For theoretical models, among many others, see Basu and Bundick (2017), Bloom et al. (2016), and Leduc and Liu (2016). Bloom et al. (2016) discuss how in their model uncertainty shocks can have both expansionary and contractionary effects, and similar forces are present in other settings.

economy driven by major political and economic events does not appear, by itself, to harm the economy.

The uncertainty shocks are primarily identified empirically in our analysis by option prices (the VIX), so one possible explanation for our results is that even though investors in option markets have extra information about future volatility, the average person might not (though note that average investors do have access to option markets). The simplest way to check that is to examine alternative measures of uncertainty. We show that alternative uncertainty indexes used in the related literature are also able to predict future volatility above and beyond what is contained in the past history of realized volatility, showing people do not simply extrapolate past volatility when forecasting the future. Furthermore, when those indexes are used in our VAR, the finding that uncertainty shocks are not contractionary is unchanged. However, we do show theoretically that if agents use biased forecasts of future volatility, while our identification scheme still can correctly identify the sign of the effect of uncertainty shocks, the magnitude may be biased down.

While there are many models explaining how uncertainty shocks might drive the economy, there are no current models available matching our findings that uncertainty shocks are not contractionary while realized volatility shocks are. The last section of the paper presents a simple purely rational model that quantitatively matches our evidence. Its key mechanism is that shocks to technology are negatively skewed.² Negative skewness means that large shocks – which cause high realized volatility – also tend to be negative shocks, immediately generating the observed negative correlation between realized volatility and output. Our finding that realized volatility is associated with contractions does not therefore imply that realized volatility per se is contractionary – it can be simply capturing the occurrence of a large negative shock (in our model, a downward jump in technology). In simulations, the model matches the empirical left skewness in real activity and also the relative sizes of the financial risk premia on shocks to realized and implied volatility. Most importantly, when we use our VAR identification on simulations of the model, the IRFs are quantitatively similar to the data and the identified shocks are almost perfectly correlated with the true structural shocks, providing further theoretical support for our identification scheme.

In addition to the macroeconomic studies discussed above, our work is also closely related to an important strand of research in finance studying the differences between expected and realized volatility, e.g. Andersen, Bollerslev, and Diebold (2007). A jump in stock prices, such as a crash or the response to a particularly bad macro data announcement, mechanically generates high realized volatility. On the other hand, news about future uncertainty, such as an approaching presidential election, increases expected volatility (Kelly, Pastor, and Veronesi (2016)). This paper also contributes to the literature on the pricing of shocks to realized and implied volatility: we show that historically, investors have paid large premia to hedge shocks to realized volatility, but the premia associated with shocks to implied volatility are close to zero, implying that investors do not view uncertainty shocks as unfavorable.

²See also Morley and Piger (2012), McKay and Reis (2008), and Salgado, Guvenen, and Bloom (2016) for recent work on skewness. While our skewness is exogenous for simplicity, a richer model could endogenize it.

Our VAR estimation method is closely related to the recent literature on news shocks in macroeconomics.³ Past work has focused on *first-moment* news shocks: news about the average future path of the economy (e.g. future productivity growth). This paper extends the estimation to *second-moment* news shocks – changes in expected future *squared* growth rates. News about the expectation of future squared growth rates represents an uncertainty shock in that it is a change in the conditional variance (see Beaudry and Portier (2014)). In the same way that a shock to the current level of TFP is not the same as a shock to expected future TFP growth, a surprisingly large squared shock today – realized volatility – is not the same as an increase in conditional variance, or uncertainty.

Our work is also related to a large empirical literature that studies the relationship between aggregate volatility and the macroeconomy. A wide range of measures of volatility in financial markets and the real economy have been found to be countercyclical.⁴ Data used in past work has generally only measured either uncertainty (e.g. the surveys studied in Bachmann, Elstner, and Sims (2013)) or realized volatility (e.g. the cross-sectional variance of income growth in Storesletten, Telmer, and Yaron (2007)). Part of our contribution is to measure both.

To identify causal effects, a number of papers use VARs, often with recursive identification, to measure the effects of volatility shocks on the economy.⁵ Ludvigson, Ma, and Ng (2015), like us, distinguish between different types of uncertainty and examine reverse causation from activity to uncertainty.⁶ Caldara et al. (2016) use a penalty-function based identification scheme to distinguish between the effects of uncertainty and financial conditions.

Many different concepts of uncertainty have been studied in the literature. This paper aims to understand the effects of uncertainty about the aggregate economy, and, similar to past work, we focus on uncertainty about the aggregate stock market. Our negative results therefore do not rule out potential effects of other factors, such as uncertainty about idiosyncratic shocks (e.g. Christiano, Motto, and Rostagno (2014)) or policies with cross-sectional effects.

³E.g. Beaudry and Portier (2006), Barsky and Sims (2011), and Barsky, Basu, and Lee (2015).

⁴See, among others, Gilchrist, Sim, and Zakrajsek (2014), Campbell et al. (2001), Storesletten, Telmer, and Yaron (2004), Guvenen, Ozkan, and Song (2014), Eisfeldt and Rampini (2006), Alexopoulos and Cohen (2009), Baker, Bloom, and Davis (2015), and Caggiano, Castelnuovo, and Groshenny (2014).

⁵See Bloom (2009) and Basu and Bundick (2015), who study the VIX; and Baker, Bloom, and Davis (2015) and Alexopoulos and Cohen (2009), who study news-based measures of uncertainty. Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2015) measure uncertainty based on squared forecast errors for a large panel of macroeconomic time series (using a two-sided filter to extract a latent volatility factor). Baker and Bloom (2013) use cross-country evidence to argue that there is causal and negative relationship between uncertainty and growth. See also Born and Pfeifer (2014), Caldara et al. (2016), Fernandez-Villaverde et al. (2011), Fernandez-Villaverde et al. (2013), Leduc and Liu (2016).

⁶Other papers arguing that causality could run from real activity to volatility and uncertainty include Decker, D’erasmo and Boedo (2016), Berger and Vavra (2013), Ilut, Kehrig and Schneider (2015), Kozlowski, Veldkamp, and Venkateswaran (2016), Creal and Wu (2017), and Cesa-Bianchi, Pesaran, and Rebucci (2018). See also Carriero, Clark, and Marcellino (2016) and Popiel (2017) for alternative identification schemes.

2 Preliminary evidence

We begin by briefly presenting simple regressions that illustrate the basic results of the paper. Table 1 reports results of regressions of monthly growth rates of employment and industrial production on the current value and four lags of realized S&P 500 volatility and option-implied volatility. We discuss the data in more detail below, but realized volatility (RV) is calculated as the sum of daily squared returns on the S&P 500 index each month, and option-implied volatility (V_1) is nearly identical to the VIX and it is calculated based on option prices at the end of each month. Realized volatility therefore measures current conditions, while implied volatility is related to uncertainty about the future. All variables are converted to z-scores to aid interpretation of the coefficients. We report only the average of the coefficients on realized and implied volatility (they are relatively stable across lags).

The first column in the two sections of table 1 shows results from regressions that include only option-implied volatility (V_1):

$$y_t = a + b(L) V_{1,t} + \varepsilon_t \quad (1)$$

where y_t corresponds to the different macroeconomic measures, and the table reports the average of the coefficients in the fourth-order lag polynomial b . The results show that high option-implied volatility is correlated with declines in employment and industrial production.

The second column of each section adds realized volatility – the sum of daily squared stock returns during month t (RV_t) – and its lags to the regression:

$$y_t = a + c(L) RV_t + b(L) V_{1,t} + e_t \quad (2)$$

Once realized volatility is included, implied volatility no longer carries negative coefficients – the coefficients are actually *positive* on average in both regressions. That results implies that implied volatility is not associated with downturns because of the fact that it represents uncertainty about the future, but rather because it is correlated with how volatile returns are in month t . But contemporaneous volatility is conceptually distinct from uncertainty about the future. Table 1 gives the simplest way to see the paper’s basic result that uncertainty about the future is not associated with declines in real activity.

2.1 Identification and expectation biases

The regressions in table 1, and our main empirical analysis below, use option-implied volatility to help measure uncertainty, and interpreting them requires making assumptions about the relationship between option-implied uncertainty and that of the average person. Suppose the true statistical measure of uncertainty, U_t , and the uncertainty of the average consumer in the economy,

U^* , are (ignoring constants throughout)

$$U_t = u_1 RV_t + u_2 V_{1,t} \quad (3)$$

$$U_t^* = u_1^* RV_t + u_2^* V_{1,t} \quad (4)$$

If agents extrapolate excessively, we would say that $u_1^* > u_1$ (they put too much weight on RV) and perhaps $u_2^* < u_2$ (they ignore some of the information in V_1), while $u_1^* = u_1$ and $u_2^* = u_2$ represents rationality.⁷

The relationship between output and volatility has two components. First, agents' uncertainty might have a direct impact on output Y_t . Second, it is possible that output and RV share a common shock. We show in our structural model in section 8 that that will happen if shocks to output are skewed left – then relatively large shocks, with high RV , will tend to be negative. So we have the system,

$$Y_t = -b_U U_t^* - J_t \quad (5)$$

$$RV_t = U_{t-1} + b_{RV} J_t \quad (6)$$

where $-J_t$ is a binary variable representing a downward jump that reduces output and increases RV . Note that while it is the agents' perceived uncertainty U_t^* that determines the macroeconomic response in (5), the conditional mean of RV is by definition the statistical measure of uncertainty, U_t , in (6).

Suppose we regress Y_t on RV_t and $V_{1,t}$ and their lags, as above. Combining the above equations,

$$Y_t = b_{RV}^{-1} U_{t-1} - (b_U u_1^* + b_{RV}^{-1}) RV_t - b_U u_2^* V_{1,t} \quad (7)$$

RV can have a negative coefficient in these regressions either because uncertainty causes declines in output or because RV and output share a common shock, J . Observing a negative relation between output and RV therefore does not prove the existence of negative effects of uncertainty. Moreover, the negative relationship between output and RV_t is not causal; RV_t just helps measure a level shock to output.

Most importantly, the coefficient on implied volatility shows that as long as $u_2^* > 0$, i.e. as long as there is some relationship between agents' uncertainty and option-implied volatility, then if uncertainty has contractionary effects, at least the sign of this relationship will be identified from the coefficient on $V_{1,t}$.

⁷This specification tries to capture in reduced form different potential sources of extrapolative expectations and other behavioral biases, studied in a vast literature. See for example Barberis and Shleifer (2003), Greenwood and Shleifer (2014), Greenwood and Hanson (2014), Barberis et al. (2015, 2016), Gennaioli, Ma, and Shleifer (2015), and Bordalo, Gennaioli and Shleifer (2017).

2.2 Challenges

The paper’s basic result is in table 1: realized volatility drives out implied volatility, indicating that uncertainty shocks are not contractionary. Table 1 is insufficient by itself to convincingly make that point, though. First, the simple linear regressions allow only very minimal dynamics, they do not allow us to estimate impulse response functions or identify structural shocks, and they do not yield variance decompositions to help understand the sources of variation in output. Second, one might naturally ask whether the identification scheme used here would work in leading models of uncertainty shocks. Third, even if we find empirically with more sophisticated methods that identified uncertainty shocks do not affect real activity, the model above shows that could be because $u_2^* = 0$ rather than $u_2 = 0$ – agents forecast volatility based only on its past realizations. Fourth, what kind of model would be required to rationalize the result that realized volatility predicts contractions but uncertainty and implied volatility do not?

The remainder of the paper addresses those four issues. We examine shocks to realized and implied volatility in a VAR, show that the results are robust to a wide range of specification and identification choices, show that the identification scheme successfully uncovers structural shocks in three different theoretical models, and provide evidence on u_2^* from uncertainty surveys and returns on stock options. Finally, we develop a novel model with skewed shocks that can rationalize the empirical results.

3 General empirical framework

This section describes how we define, measure, and identify uncertainty shocks in the data.

3.1 Conditional variances

Denote the log of the total return stock index as s_t . Uncertainty about the future value of the stock market relative to its value today is measured in this paper as the conditional variance, $Var_t [s_{t+n}]$.

The one-period log stock return is $r_t \equiv s_t - s_{t-1}$. If returns are unpredictable and time periods are sufficiently short that $E_t r_{t+1} \approx 0$, we have:

$$Var_t [s_{t+n}] \approx E_t \left[\sum_{j=1}^{n-1} r_{t+j}^2 \right] \quad (8)$$

The conditional variance of stock prices on some future date is equivalent to the expected total variance of returns over that same period. Under standard conditions on the returns process, as the length of a time period approaches zero, (8) becomes an equality. This equation therefore says that uncertainty about future returns is equal to the expectation of future squared returns.

So a shock to uncertainty is a shock to expected future volatility. That is conceptually distinct from a shock to current squared returns, r_t^2 . Realized and expected volatility may be correlated, but

they are not the same object, and the results in table 1 suggest that they have distinct relationships with real activity.

Since the data on real activity is measured at the monthly frequency, but stock returns are available at a daily frequency, realized volatility is measured as the sum of daily squared returns over each month,⁸

$$RV_t \equiv \sum_{days \in t} r_{daily}^2 \quad (9)$$

We then have

$$Var_t [s_{t+n}] \approx E_t \left[\sum_{j=1}^n RV_{t+j} \right] \quad (10)$$

3.2 Vector autoregressions

To identify uncertainty shocks and estimate their effects, we estimate VARs of the form

$$\begin{bmatrix} RV_t \\ Y_t \end{bmatrix} = C + F(L) \begin{bmatrix} RV_{t-1} \\ Y_{t-1} \end{bmatrix} + A\varepsilon_t \quad (11)$$

where RV_t is realized volatility from (9), Y_t is a vector including measures of real activity, variables that help forecast future values of realized volatility, and other controls, C is a vector of constants, and $F(L)$ is a matrix polynomial in the lag operator, L . ε_t is a vector of uncorrelated innovations with unit variances. (11) can be estimated by ordinary least squares. The VAR has a moving average representation,

$$\begin{bmatrix} RV_t \\ Y_t \end{bmatrix} = (I - F(1))^{-1} C + B(L) A\varepsilon_t \quad (12)$$

$$\text{where } B(L) = \sum_{j=0}^{\infty} B_j L^j = (I - F(L))^{-1} \quad (13)$$

The shocks to realized volatility and uncertainty are identified under a timing restriction (which is relaxed in sections 6.1.2 and A.3). The realized volatility shock is treated as moving first, so that its impulse response functions (IRFs) are defined as

$$\frac{\partial E [Y_{t+j} \mid RV_t, RV^{t-1}, Y^{t-1}]}{\partial RV_t} = B_j A_{(:,1)} \quad (14)$$

where $A_{(:,1)}$ denotes the first column of A . E denotes the expectation operator conditional on the VAR (11) and RV^{t-1} and Y^{t-1} are the histories of RV and Y up to date $t - 1$. As in Hamilton

⁸Calculating RV with just the monthly return – or with monthly output or TFP growth – would yield a much noisier measure of actual volatility. If returns are i.i.d., for example, and there are on average 21 trading days in a month, then the variance of the sum of daily squared returns is 21 times smaller than the variance of the squared monthly return. So daily data can give a much more accurate estimate of the true underlying current volatility (as the sampling frequency becomes infinite, volatility can potentially be estimated without error, in fact).

(1994), the impulse response is defined as the average change in expectations for the future following a unit innovation in realized volatility. We do not view the realized volatility shock as a structural shock – obviously volatility in the stock market depends on many different underlying shocks. Instead, the IRF tells us *on average* how the economy (output, employment, etc.) changes when realized volatility changes.

The second estimated shock is the residual innovation in uncertainty over the following n periods, $Var_t[s_{t+n}]$ or, equivalently, the residual innovation in expectations of future volatility. It is the component of $E_t \sum_{j=1}^n RV_{t+j} - E_{t-1} \sum_{j=1}^n RV_{t+j}$ that is orthogonal to $RV_t - E_{t-1}RV_t$. In other words, the uncertainty shock is ordered second, following realized volatility.

The identifying assumption is that on average, structural uncertainty/volatility news shocks do not cause changes in realized volatility on impact, while realized volatility can potentially affect uncertainty. The latter effect, as discussed above, is sometimes called a GARCH effect – high volatility today predicts high volatility in the future. The motivation for the assumption that uncertainty shocks do not affect realized volatility on impact is that while a large increase in uncertainty should drive stock prices down, generating a large squared return, a large reduction in uncertainty drives stock prices up, thus *also* generating a large squared return. That is, the causal effect of uncertainty shocks on realized volatility is approximately quadratic: sometimes it is positive, sometimes negative. On average, the effect of uncertainty on RV on impact is zero. Section 6.1.1 shows that this intuition in fact holds in three theoretical models, and that the timing assumption successfully identifies structural uncertainty shocks in those cases.

To construct the news shock, we use the total change in cumulative expected volatility up to time $t + n$, which is

$$E_t \sum_{j=1}^n RV_{t+j} - E_{t-1} \sum_{j=1}^n RV_{t+j} = \left(e_1 \sum_{j=1}^n B_j \right) A \varepsilon_t \quad (15)$$

where $e_1 = [1, 0, \dots]$. The parameter n determines the horizon over which the news shock is calculated. Cumulative expected volatility depends on the sum of the first rows of the MA matrices up to lag n . The innovation to expectations over horizon n , i.e. the news about future volatility, is then the linear combination of shocks represented by $e_1 \sum_{j=1}^n B_j A$. As in Barsky, Basu, and Lee (2014) and Barsky and Sims (2011), we orthogonalize that linear combination with respect to the innovation to RV_t so that the impact shock to RV_t is uncorrelated with the news shock.⁹

The uncertainty shock, denoted u_t , is equal to the component of $\left(e_1 \sum_{j=1}^n B_j \right) A \varepsilon_t$ that is orthogonal to the RV shock, $e_1 A \varepsilon_t$, and its impulse responses are defined as

$$\partial E [Y_{t+j} \mid u_t, RV_t, RV^{t-1}, Y^{t-1}] / \partial u_t \quad (16)$$

Again, the VAR yields the average behavior of the economy following a unit increase in uncertainty

⁹The orthogonalized innovation is simply $e_1 \sum_{j=1}^n B_j \bar{A} \varepsilon_t$, where \bar{A} is equal to A but with the first column set to zero.

(conditional on RV_t), as measured by u_t .

The two impulse responses are only identified up to some normalization. We scale them so that they have the same effect on uncertainty. The uncertainty/news shock is reported as a unit standard deviation impulse, while the realized volatility shock is rescaled so that its cumulative effect on uncertainty (i.e. its IRF for realized volatility over the next 24 months) is the same as that for the news shock. The shocks then have equal effects on uncertainty about the future and differ only in their effect on realized volatility on impact.

We set $n = 24$ months. Past work finds that volatility shocks have half-lives of 6–12 months, so 24 months represents the point at which the average shock has dissipated by 75 percent or more. Choosing n between 12 and 60 months yields similar results.

4 Data

4.1 Macroeconomic data

As in other work in this literature (Bloom (2009), Leduc and Liu (2016)), we use monthly data to maximize statistical power, since fluctuations in both expected and realized volatility are very short-lived. Real activity is measured by the Federal Reserve’s measure of industrial production (IP) for the manufacturing sector and total non-farm private employment.

4.2 Financial data

We obtain data on daily stock returns of the S&P 500 index from the CRSP database and use it to construct RV_t at the monthly frequency. Option-implied volatilities, $V_{n,t}$, are included in the state vector to help forecast future volatility (the assumption is not that they represent statistical forecasts of future volatility). They are constructed using prices of S&P 500 options obtained from the Chicago Mercantile Exchange (CME), with traded maturities from one to at least six months since 1983. We thus have a substantially longer sample of implied volatility than has been used in past work. Given that shocks to stock market volatility have half lives estimated to be approximately six to nine months (see Bloom (2009) and Drechsler and Yaron (2011)), one- to six-month options will contain information about the dominant shocks to stock market uncertainty. Technically, $V_{n,t}$ is defined as the option-implied variance of the stock index on date $t + n$ under the pricing measure Q ,¹⁰

$$V_{n,t} \equiv Var_t^Q [s_{t+n}] \tag{17}$$

The appendix provides the details of the calculation. We use the measure of $Var_t^Q [s_{t+n}]$ in (17) because it maps directly to our object of interest, the conditional variance of log stock prices. In practice, it is extremely close to the VIX calculated by the CBOE and other related model-free

¹⁰The pricing measure, Q , is equal to the true (or physical) pricing measure multiplied by $M_{t+1}/E_t [M_{t+1}]$, where M_{t+1} is the pricing kernel.

implied volatility measures. In the remainder of the paper we focus on the logs of realized and option-implied volatility ($rv_t \equiv \log RV_t$, $v_{n,t} \equiv V_{n,t}$) due to their high skewness.

Figure 1 plots the history of realized volatility along with 1-month option-implied volatility in annualized standard deviation terms. Both realized and option-implied volatility vary considerably over the sample. The two most notable jumps in volatility are the financial crisis in 2008 and the 1987 market crash, which both involved realized volatility above 75 annualized percentage points and rises of $V_{1,t}$ to 65 percent. More generally, realized volatility tends to be highest in periods when the market falls, and that relationship is strongest around extreme events. In particular, the monthly correlation between realized volatility and S&P 500 stock returns is -0.28. The correlation between the cubes of those variables, which emphasizes extreme returns, is -0.71. On the other hand, the correlation between the cube roots of realized volatility and S&P 500 returns is only -0.23 (these correlations are all significant at the 1-percent level). The months with the most extreme realized volatility also tend to have the most extremely negative returns – the four highest realizations of realized volatility all have negative returns for the S&P 500, including the two lowest over our sample.

At lower frequencies, the periods 1997–2003 and 2008–2012 are associated with persistently high uncertainty, while it is lower in other periods, especially the early 1980’s, early 1990’s, and mid-2000’s. There are also distinct spikes in uncertainty in the summers of 2010 and 2011, likely due to concerns about the stability of the Euro zone and the willingness of the United States government to continue to pay its debts. Those events will be particularly important for our identification.

5 VAR results

All VARs below use four lags, as suggested by the Akaike information criterion for the main specification. The vector of variables in the benchmark specification is $[rv_t, v_{1,t}, FFR_t, ip_t, emp_t]$, where the latter three variables are the Fed Funds rate, log industrial production, and log employment, respectively. The benchmark specification imposes the restriction that the coefficients on the lags of FFR , ip , and emp in the equations for rv and v_1 are equal to zero, which helps reduce overfitting and is consistent with predictive regressions discussed in section 6.3.1. A χ^2 test of the validity of those restrictions yields a p-value of 0.61, implying that they are consistent with the data. Section 6.3.2 examines the issue of overfitting the volatility equations in more detail.

5.1 Benchmark specification

Figure 2 has three columns for the responses of rv , employment, and industrial production to the shocks. As discussed above, these impulse response functions (IRFs) are scaled so that the two shocks – current rv and the identified uncertainty shock – have the same total effect on volatility expectations (uncertainty) 2–24 months in the future (i.e. not counting the impact period). Each panel shows the point estimate for the IRFs along with 1-standard deviation (68-percent) and

90-percent bootstrapped confidence intervals.

The first row shows the response of the economy to the rv shock. The shock to realized volatility is very short-lived: the IRF falls by half within two months, and by three-fourths within five months, showing that realized volatility has a highly transitory component. Those transitory increases in realized volatility are associated with statistically and economically significant declines in both employment and industrial production. So, consistent with past work, we find a significant negative relationship between *volatility* and real activity. But this result does not imply that an uncertainty shock is contractionary, since the realized volatility shock is a combination of shocks to both *current* realized volatility and also expectations about the future. In the language of section 2.1, the rv shock may include the shock to the level of output, J .

The second row of panels in figure 2 plots IRFs for the identified pure uncertainty shock, which has no contemporaneous effect on rv . The news shock forecasts high realized volatility in the future at a high level of statistical significance – p-values testing whether the IRF is positive range between 0.001 and 0.05 – indicating that the identified news shock contains statistically significant news about uncertainty, even after controlling for the contemporaneous rv shock.

Surprisingly, though, the second-moment news shocks are associated with no significant change in either employment or industrial production; the IRFs both stay very close to zero at all horizons. And the confidence bands are sufficiently narrow that the point estimates for the responses of employment and industrial production to the rv shock are outside the 90-percent confidence bands for the uncertainty shock.

To further examine the relative magnitudes, the bottom row of panels in figure 2 reports the difference in the IRFs for the uncertainty and realized volatility shocks along with confidence bands. The two shocks have the same cumulative impacts on the future path of realized volatility (by construction, due to the scaling of the IRFs), and they have different effects on rv on impact due to the identifying assumptions. The difference in the IRFs is statistically significant for employment, weakly so for industrial production. So innovations in rv are followed by statistically significant declines in real activity, while uncertainty shocks are not, and that difference itself is statistically significant in one case.

An alternative way to interpret the bottom row of figure 2 is that it represents the response of the economy to a pure shock to realized volatility that has no net effect on forward-looking uncertainty. By construction, the shock in the third row has a positive effect on rv on impact, but the IRF for rv over the following 24 periods sums to zero. The figure shows that such a shock has negative effects on the economy.

Overall, then, figure 2 shows that under our baseline identification scheme, shocks to rv are associated with statistically significant subsequent declines in output, while uncertainty shocks are not, and the difference between those two results is itself statistically significant. That result is notable given that past work (e.g. Bloom (2009), Basu and Bundick (2017), Leduc and Liu (2016)) has found that increases in option-implied volatility are followed by declines in output (a result we also obtain in our data; see appendix figure A.5). The results here show that it is actually

realizations of volatility that seem to drive such effects, as opposed to shocks to uncertainty about the future. These results are essentially the same as those in table 1, simply with more controls, a flexible lag structure, and a formal identification scheme.

5.2 Variation in volatility and uncertainty

To further understand the importance of the uncertainty and rv shocks, figure 3 reports forecast error variance decompositions (FEVDs). As in figure 2, we report the effect of the rv shock, the uncertainty shock, and their difference.

The first column reports variance decompositions for uncertainty, again defined as $E_t \left[\sum_{j=1}^{24} rv_{t+j} \right]$. At the point estimates the rv shock accounts for 65 percent of the variance of uncertainty, while the uncertainty shock accounts for the remaining 35 percent. The confidence bands are wide enough, though, that at most horizons we cannot reject the hypothesis that the rv and news shocks account for equal fractions of the variance of uncertainty at even the 32-percent significance level.

To further examine the relative importance of the two shocks for uncertainty, figure 4 plots fitted uncertainty from the VAR along with the parts driven by the rv and identified uncertainty shocks.¹¹ Ignoring constants, total uncertainty is equal to the sum of those two parts (which follows from the linearity of the model). The component from the uncertainty shocks does not represent total uncertainty because shocks that affect rv can also affect uncertainty (e.g. through GARCH effects). The component from the uncertainty shocks is the part generated by the pure news shocks.

The standard deviation of the part of uncertainty coming from news is only 22 percent smaller than that coming from rv , which is visible from the similar overall variation in the two series. The periods with the largest increase in news-driven uncertainty are associated with events that are clearly associated with uncertainty: the 1987 and 2008 market crashes and the third quarter of 2011 (debt ceiling and Euro crisis). At lower frequencies, the news-driven component of uncertainty is high around the 1990 recession (i.e. around the First Gulf War), in the late 1990's (Asian financial crisis, Russian default, LTCM), and between 2010 and 2012 (debt ceiling debates, Euro crisis). These are all periods that are plausibly related to high forward-looking uncertainty about real events.

The financial crisis and 9/11, on the other hand, are periods of high uncertainty, but driven by rv rather than pure uncertainty shocks. In the context of the model, they would be driven by the jump shock, J . Intuitively, those were real events that had first-moment effects on the economy, while simultaneously affecting uncertainty.

¹¹In the VAR, there exists a vector B such that

$$E_t \left[\sum_{j=1}^{24} rv_{t+j} \right] = B \left[\tilde{Y}'_t, \tilde{Y}'_{t-1}, \dots, \tilde{Y}'_{t-4} \right]' \quad (18)$$

where $\tilde{Y}_t \equiv [Y'_t, RV'_t]'$. The time series \tilde{Y}^{rv} is constructed by using the VAR structure and setting all the estimated shocks to zero except for the rv shock. Similarly, \tilde{Y}^{news} is constructed by setting all estimated shocks except the news shock to zero. The parts of uncertainty coming from rv and news are then $B\tilde{Y}^{rv}$ and $B\tilde{Y}^{news}$, respectively.

The identified shocks are also both strongly correlated with returns on the S&P 500. Specifically, the correlations of the identified rv and uncertainty shocks with S&P 500 returns are:

Correlations with S&P 500 returns

rv shock: -0.412

Uncertainty shock: -0.405

The rv and uncertainty shocks (which are uncorrelated with each other by construction) have substantial negative and nearly identical correlations with stock returns. So not only is the uncertainty shock significantly associated with future volatility, but it is also associated with substantial declines in stock returns.¹² These facts make it all the more surprising that the uncertainty shock is not associated with any significant change in real activity.

5.3 Forecast error variance decompositions for employment and IP

In addition to the FEVDs for uncertainty, figure 3 also reports FEVDs for the two real outcomes, employment and industrial production. The realized volatility shock explains 25 percent of the variance of employment and 10 percent of the variance of industrial production at the two-year horizon, while the point estimates for the fraction of the variance accounted for by second-moment news are two percent or less, and the upper ends of the 90-percent confidence intervals are near 5 percent for the first year.

Note also that the lack of importance of the uncertainty shock for real outcomes is not simply due to its size. Again, the pure uncertainty shock in the second row accounts for 35 percent of the total variance of uncertainty. So if we scale up the variance decompositions for employment and industrial production by multiplying them by 3, the point estimates still say that all uncertainty variation accounts for less than 5 percent of the total variance of employment and industrial production.

Looking at the third row, the behavior of the rv and news shocks is again significantly different. For employment, the variance accounted for by the rv shock is larger at the 90-percent level, while for IP, they differ at only a 1-standard-deviation level.

6 Identification challenges

This section addresses numerous potential concerns about the VAR results, including threats to the identification scheme, the interpretation of the identified shocks, and sensitivity to specification choices.

¹²Note that this result is entirely consistent with our identification scheme. The assumption is not that uncertainty shocks do not cause on average changes in returns, but changes in realized volatility, which is a quadratic function of returns.

6.1 The identification scheme

This subsection first shows that the baseline identification scheme successfully identifies structural uncertainty shocks in three theoretical models, then shows that the empirical results also hold under milder restrictions than the strict timing assumption used in the baseline case.

6.1.1 Identification in structural models

The benchmark identification scheme involves a timing restriction: it identifies uncertainty shocks that do not affect realized volatility contemporaneously. In section 2.1, that is imposed by the assumption that RV_t does not depend on uncertainty on date t . Obviously there are data generating processes that will violate that assumption. We ask, though, whether it should be expected to be violated in theoretical models.

A natural worry about the timing assumption is that a shock that increases uncertainty about the future would be expected to reduce stock prices and thus create realized volatility. But a shock that *decreases* uncertainty also changes stock prices, just in the opposite direction, also creating realized volatility. Since there is a quadratic relationship between shocks and realized volatility – both large positive and large negative shocks create high RV – the average (i.e., expected) effect of uncertainty shocks on RV can be close to zero, which is what is required for our identification scheme to work.¹³

To test the identification scheme and its performance, we examine state-of-the-art models of uncertainty shocks from Basu and Bundick (2017; BB) and Bloom et al. (2017; RUBC).¹⁴ The BB and RUBC models capture the two most prominent channels through which uncertainty has been proposed to affect the economy: precautionary saving and wait-and-see effects, respectively. We also develop a model below that can rationalize the empirical results through left skewness in technology shocks, so we refer to it as the RSBC model – “really skewed business cycles” – to highlight the key difference with Bloom et al.’s (2017) RUBC – “really uncertain business cycles”.

To test the identification in these models, we estimate the VAR specification in simulations.¹⁵ The correlations of the identified uncertainty shock with the true shock to uncertainty in the three models are:

Correlation of VAR-identified uncertainty shock with structural uncertainty shock:

¹³As a formal example, if shocks to uncertainty are symmetrically distributed and they have a linear relationship with stock returns, then the average effect of an uncertainty shock onto squared stock returns is exactly zero. Empirically, monthly changes in the log of S&P 500 implied volatility are indeed close to symmetric: the 10th, 25th, 75th, and 90th percentiles are -0.59, -0.34, 0.25, and 0.62 (see also Amengual and Xiu (2014)).

¹⁴The simulations of the BB model are carried out with the replication files available from *Econometrica*’s website. Stephen Terry generously provided simulations of the RUBC model.

¹⁵The three models are all specified in discrete time with time periods equal to a quarter in the case of BB and RUBC and a month in the case of RSBC, so there is not an exact mapping in the models to the daily realized volatility that we study empirically. We approximate it here as the square of the total realized return in the model in each period, after subtracting mean returns as a normalization.

BB	RUBC	RSBC
0.94	0.94	0.91

The identification scheme robustly identifies the true structural innovation to uncertainty across three very different structural models. That happens because in all three cases the identifying assumptions hold. The uncertainty shocks in the models all increase the future volatility of stock returns, and they are uncorrelated with realized volatility on impact:

Contemporaneous correlation of structural uncertainty shocks with realized volatility:

BB	RUBC	RSBC
-0.01	-0.10	0.00

These correlations confirm, for the three models, the identifying assumption discussed in section 3.2 (which also appeared in the model of section 2.1).

Figures A.6 and A.7 in the appendix report the IRFs estimated using the VAR in the BB and RUBC models. They show that in both models, the estimated uncertainty shocks drive output down significantly (which is the effect of structural uncertainty shocks in the models), while realized volatility shocks have no effect, the opposite of what we find empirically. Section 8 and figure 9 below show that the RSBC model, on the other hand, qualitatively and quantitatively matches the empirical VAR results.

A reason that it is important to control for RV first in the VAR is that RV itself may raise future uncertainty. That effect is a standard result in the finance literature on volatility, often referred to as a GARCH (generalized autoregressive conditional heteroskedasticity) effect (Engle (1982), Bollerslev (1986)). We can also use the simulations of the models to explore what happens if the assumed timing of the shocks is reversed, failing to control for realized volatility in identifying the uncertainty shock. That is, what if we assume that uncertainty shocks move first and realized volatility second? In that case, the identifying assumption is that the change in uncertainty is affected by no other shocks in the economy. In the three simulated models, we then find that the VAR-identified uncertainty shock is less strongly correlated with the true uncertainty shock. Specifically, in the BB, RUBC, and RSBC models, the correlations between the VAR-identified and the true uncertainty shock become 0.94, 0.66, and 0.83, respectively. In the BB model, the timing is not relevant because in the VAR in that model innovations to realized volatility and uncertainty are uncorrelated (which is strongly at odds with the data). In the other two models, though, realized volatility and uncertainty are more correlated, and the ordering of the variables does in fact matter.

To summarize, then, our identification scheme does extremely well in capturing the true uncertainty shock across the three models; conversely, none of the shocks in the reverse ordering correctly captures the uncertainty shock in general.

6.1.2 Weakening the timing assumption

Section A.3 of the appendix reports results from a Bayesian version of the estimation that uses a somewhat different identification scheme. Rather than making a strict timing assumption to identify the shocks, it sets a prior that the RV shock has a larger effect on RV on impact than in the future, while the news shock has a larger effect in the future than on impact (i.e. the effect is hump-shaped). There are also some more minor technical differences from the main specification.

The results of the Bayesian estimation are in figures A.17 and A.18. The first figure reports the IRFs for employment, the second for IP. The RV shock again is associated with significant declines in employment and industrial production while the news shock is not. Moreover, that zero appears to be well identified, at least for employment, in the sense that the posterior probability bands are smaller than the priors by 60 percent, showing that the data is informative. For IP, the confidence band is wider, but the point estimate is in fact substantially positive. As in the main results, the difference between the IRFs of the RV and news shocks for employment and IP is significant at the 10-percent level (though in this case that difference is not reported in the bottom row of panels, instead we report the third shock so that the figure summarizes the full set of parameter priors and estimates). Figures A.17 and A.18 therefore show that the results do not rely on the strict timing assumption, nor do they necessarily rely on the exact VAR structure used in the baseline analysis.

6.1.3 Time aggregation

Since the main analysis involves a timing assumption at the monthly level, it is subject to a standard time-aggregation critique (Christiano and Eichenbaum (1986); Bansal, Kiku, and Yaron (2016)). That is, if the true joint process for real activity, implied volatility, and realized volatility is determined at, say, the daily frequency, but we estimate the model at the monthly frequency, then the identifying assumptions may be invalid, even if they are actually valid at the daily frequency. Specifically, the potential bias in the present case is that there might be changes in implied volatility within month t , say on day 1, which then also affect realized volatility in month t (starting on day 2), making monthly uncertainty shocks no longer contemporaneously uncorrelated with realized volatility.

The question of the correct frequency at which to estimate models of this sort depends on the decision frequency of agents and on the dynamic process for volatility. If agents respond on a daily basis to every change in implied volatility, then there could be a bias. If their decision frequency is at the monthly or quarterly level, though, as is assumed in general in the empirical literature and in the models that we are focused on testing, then a monthly identification scheme will be valid. This issue has been widely examined in the asset pricing literature, and Bansal, Kiku, and Yaron (2016) find that a monthly decision frequency fits the data best.¹⁶

¹⁶The bias will be most important if uncertainty varies primarily at frequencies greater than a month. Among others, Drechsler and Yaron (2011) argue that volatility shocks with half-lives of six months fit the data well, suggesting that the within-month bias may not be important quantitatively.

An important feature of this study that distinguishes it from most past work is that we have high-frequency data on implied and realized volatility. Much of the past work on uncertainty has used timing assumptions subject to the time-aggregation critique, and the data involved has been quarterly or monthly. Using daily data allows us to evaluate the importance of the time aggregation problem.

To account for time aggregation, we estimate a daily VAR with implied and realized volatility and use it to construct an alternative time series of realized volatility and uncertainty shocks that is purged of the within-month time aggregation bias and satisfies the timing assumption at the monthly level by construction. Appendix A.4 describes the details of the procedure. Figure A.19 shows that when using daily data to construct estimates robust to time-aggregation bias, the main results are highly similar to the baseline, implying that time aggregation is not a serious problem in practice in our setting. We also find that table 1 is essentially entirely unchanged (we report the version purged of within-month time-aggregation bias as table A.1 in the appendix).

6.2 Validating the identified uncertainty shocks

The key question in section 2.1 is how agents form expectations about future volatility. In particular, to what extent do their expectations match statistical forecasts (i.e. does $u_2 = u_2^*$?). One way to evaluate that question is to examine how our statistical forecasts relate to survey-based measures. Table 2 reports results from regressions of four alternative uncertainty indexes on the VAR-implied rv and news components of uncertainty. The first three measures are from Baker, Bloom, and Davis (2015); they are their total uncertainty index and the components from monetary policy and the stock market. These are not surveys of, say, household beliefs, so we do not take them as somehow more direct measures of uncertainty. Rather, they let us understand what our measure is related to. The fourth measure is from the Michigan Consumer Survey and used by Leduc and Liu (2016) and is a more direct measure, but it only asks about whether uncertainty has affected a household’s recent car purchases, so it is rather narrow.

Across the four regressions, the coefficients are all positive, but typically only significantly so for the rv -driven component of uncertainty. However, we can never reject the hypothesis, at even marginal significance levels, that the coefficients on the rv - and news-driven components are equal (the third column reports the p-value for the test of their difference). So while the t-statistics suggest that the relationship of these alternatives may primarily be with realized volatility rather than expectations of the future, a formal test of a difference in the relationships shows no support. In other words, these regressions do not reject the hypothesis that $u_2 = u_2^*$, but they also do not provide strong affirmative evidence.

Figures A.14, A.15, and A.16 report results replacing implied volatility with the alternative indexes in the VAR. The responses of IP and employment to uncertainty shocks constructed with the other indexes are actually *positive*. Furthermore, the response of rv to these news shocks remain statistically significantly positive, showing that the alternative uncertainty indexes contain

information about uncertainty beyond what is contained in current realized volatility. That is, people do not appear to simply extrapolate past realized volatility when forming expectations about the future. In the context of the model in section 2.1, this means that $u_2^* > 0$, which is necessary for us to have identification. The figures therefore show that whether we construct volatility expectations using implied volatility or other recent uncertainty measures, there is no evidence that uncertainty shocks are contractionary, and in many cases realized volatility continues to have statistically significantly more negative effects.

6.3 Predicting future volatility

Uncertainty is identified in the VAR as expected future realized volatility. The baseline case, and the model of section 2.1, uses current realized and option-implied volatility as the only forecasting variables. This section examines the justification for that choice and then allows more variables to potentially forecast volatility in the VAR.

6.3.1 Volatility forecasting regressions

Table 3 reports results of regressions of future realized volatility, $\sum_{j=1}^6 rv_{t+j}$, on various predictors. The first column uses rv_t and $v_{1,t}$, as in the main VAR. $v_{1,t}$ has a substantially larger t-statistic, indicating that it has greater explanatory power. The marginal R^2 of $v_{1,t}$ is 0.0554, while that of rv_t is smaller by a factor of eight at 0.0069. rv_t is in fact only significant at the ten-percent level. Given that asset prices are expected to aggregate information efficiently, it is not surprising that option-implied volatility is a substantially stronger predictor of future volatility.

The second column of table 3 shows that the results are unaffected by including a lagged value of rv . In fact, if we include 6 lags of rv , their combined marginal R^2 is still only half the marginal R^2 of $v_{1,t}$. The third column of table 3 adds the six-month implied volatility, $v_{6,t}$, and finds that it adds no incremental information. We obtain similar results using a principal component of the term structure of implied volatilities.

The fourth and fifth columns of table 3 show that adding various macroeconomic and financial indicators also has no effect, and that they are not significant, which is why they are excluded from the analysis. Table A.2 in the appendix shows that when rv and v_1 are excluded from the regression, a number of the macroeconomic and financial variables become significant predictors of future volatility. In other words, the macro and financial time series on their own can help predict future volatility, but their forecasting power is subsumed by current realized and option-implied volatility.

Table A.3 in the appendix shows, in regressions analogous to those in table 3, that v_1 is predicted only by its own lags and those of rv . Duffee (2011) shows that in standard linear term structure models, except for in knife-edge cases (even allowing for time-varying risk premia), investor expectations of the future can be extracted from observed asset prices. In our setting, that result corresponds to the view that the current state of the term structure of option-implied volatilities

and rv should, together, encode all available information about future values of rv and option-implied volatility. The fact that macro variables are driven out of these regressions is therefore not surprising.

6.3.2 Allowing more variables to predict future volatility

Lasso While the assumption that current v_1 and rv contain all available information about future volatility is consistent with economic theory, the regressions in table 3, and a χ^2 test in the VAR, it is still important to examine robustness to relaxing it. We do that here with lasso (Tibshirani (1996)), which adds to the objective a penalty on the sum of the absolute values of the coefficients. The advantage of using lasso for our purposes is that it is economically agnostic and driven by statistical considerations. It thus yields restrictions on the VAR to help reduce overfitting similarly to our benchmark specification, but without imposing the set of restrictions based purely on theory.

Section A.2 describes the details of our implementation. In the estimation, lasso restricts many of the coefficients on the macro variables in the equations for rv and v_1 to zero, but not all of them. It also restricts the coefficients on some of the lags of rv and v_1 to be zero (they were not restricted in the benchmark). So it reduces the restrictions on the macro variables but increases them on the lags of the volatility measures.

The third row in figures 5 and 6 reports results using lasso for the estimation instead of the benchmark restrictions. The results remain highly similar to the benchmark case, showing that the restrictions applied in the benchmark model do not drive the results alone.

Completely unrestricted model The bottom rows in figures 5 and 6 report results that put no restrictions on the VAR coefficients and include $v_{6,t}$ as an additional potential predictor of volatility to maximize explanatory power. These are thus our most general and least constrained results. Shocks to realized volatility continue to be contractionary, and we again find no evidence that the news shock has negative effects on output, but the confidence bands for the news shock in this case become so wide as to be uninformative. The difference between the IRFs for rv and news is now statistically insignificant. We obtain similar results when we include other predictors of future volatility instead of $v_{6,t}$ such as the default spread.

The unrestricted specification has low power for identifying the effect of news shocks because the news shock itself is difficult to identify. There are, in this case, 24 potential predictors, and noise in the estimated coefficients on them is inherited by the identified news shock. That is why the specifications with restrictions and using lasso, which both have fewer coefficients to estimate, are much more stable.

Beyond power, though, the unrestricted model is potentially problematic because the news shock may be overfit. That can be seen partly from the fact that the news shock in this case actually accounts for a substantially *larger* amount of the variance of uncertainty than the rv shock does – up to 60 percent of the total variance at short horizons.

As further evidence of overfitting, the bottom panel of figure 4 plots total uncertainty, $E_t \left[\sum_{j=1}^{24} rv_{t+j} \right]$, from the benchmark and completely unrestricted specifications. In the unrestricted specification, uncertainty varies much more – its standard deviation is higher by a factor of 1.6. Moreover, it rises well in advance of the 2008 financial crisis, even though realized and option-implied volatility were low. That is, the unrestricted model seems to have seen the crisis coming before investors, consistent with overfitting. Uncertainty in the restricted model leads the benchmark specification in a number of other episodes. This lead-lag relationship can also be confirmed by examining cross-correlations. Finally, whereas section 6.2 shows that uncertainty in the restricted model is positively correlated with alternative measures of uncertainty, the Baker–Bloom–Davis index and the Michigan index, in the unrestricted model uncertainty is actually substantially *negatively* correlated with those series (with coefficients of -0.25 and -0.30), again emphasizing the implausibility of this specification.

6.4 Alternative specification choices

This section shows that the VAR results are robust to a number of changes to the details of the specification.

Figures 5 and 6 report impulse responses for employment and IP to the two identified shocks and their difference across a number of perturbations of the benchmark specification.

First, we consider alternative orderings of the variables in the VAR. The effects of the ordering depend ultimately on the correlation matrix of the innovations, which we report below:

	rv	v_1	Fed Funds	Empl.	IP
rv	1				
v_1	0.73	1			
Fed Funds	0.01	-0.06	1		
Empl.	0.01	0.05	0.05	1	
IP	-0.04	0.01	0.04	0.54	1

The shocks to rv and v_1 are correlated with each other, but almost completely uncorrelated with those to the other variables, implying that if the news shock were orthogonalized not just to the shock to rv_t but to all the macro variables also, its IRF would remain unchanged, which the top panels of figures 5 and 6 confirm by ordering rv and v_1 last.

The second row of panels in figures 5 and 6 reports results when we substitute v_6 for v_1 . The six-month option-implied volatility seems like potentially a more natural variable to include since it might represent a more realistic economic decision horizon. As the figure shows, in that case we find results even less favorable to uncertainty shocks, with these shocks now having slightly expansionary effects on employment and industrial production.

The goal of the main analysis is to identify a pure shock to uncertainty that has no contemporaneous impact on realized volatility. A natural question is what happens if we reverse the ordering

of the identification so that the first shock is the *entire* shock to uncertainty, while the second captures the residual variation in rv . In this case, then, the first shock is a combined shock to uncertainty and rv , while the second shock is a pure shock to realized volatility that has no net effect on uncertainty. Figure A.8 reports results from such a specification. It shows that both these shocks have essentially the same effect on employment and industrial production, exactly as should be expected from our main analysis. When two shocks have the same initial effect on rv , they have the same effects on output (figure A.8) whereas when they have different effects on rv , they have different effects on output (figure 2). In other words, it is the impact of a shock on contemporaneous realized volatility, not its effect on uncertainty, that determines how it affects output.

Figures A.9 and A.10 report results from four additional robustness tests:

1. Replacing our volatility measure with the VIX
2. Limiting the sample to 1988–2006 to remove the jumps in volatility in 1987 and 2008.
3. Controlling for the S&P 500 in the VAR before the identified shocks.
4. Using RV and V_1 – i.e. levels rather than logs.

The results of those robustness tests are qualitatively and quantitatively consistent with our baseline results.

Figures A.11 and A.12 report results from a quarterly VAR similar to what is estimated in Basu and Bundick (2017) with the set of variables now rv , v_1 , GDP, aggregate consumption, aggregate investment (all in real terms), the GDP deflator, and the Fed funds rate. The large number of variables and smaller number of time series observations gives this VAR specification much lower power overall than our benchmark monthly specification. Furthermore, since uncertainty shocks typically die out over a period of only a few months (see figure 2, and similar results are obtained in Bloom (2009) and elsewhere), even with a long time series the quarterly specification would have trouble identifying the effects of such shocks.

As in the benchmark results, output, consumption, and investment all fall following an increase in rv in the quarterly specification. However, in this case the confidence bands around the uncertainty shock are wide enough to render the estimates essentially uninformative. When we use v_1 to help measure uncertainty, the shock appears slightly contractionary, while using v_6 makes it appear actually expansionary. While the quarterly data does not yield sufficient power to draw firm conclusions, it in no way conflicts with our main results, and it reinforces the robustness of the finding of contractionary effects for realized volatility shocks.

Figure A.13 in the appendix reports results from a factor-augmented VAR using the first three principal components from the set of macroeconomic time series studied by Ludvigson and Ng (2007). The advantage of the FAVAR specification is that it incorporates information from an extremely broad range of variables, instead of just using employment and industrial production (we use the setup of Bernanke, Boivin, and Eliasziw (2005)). The results are highly similar to the benchmark.

7 Evidence from risk premia

We now show that investors have historically paid large premia for insurance against increases in realized volatility, but not for insurance against increases in market-implied uncertainty, which suggests that investors do not view periods in which uncertainty rises as having high marginal utility – i.e. as being bad times – consistent with our VAR results. This fact has been established in past work; here we review the evidence and extend the time series further back than in other analyses.¹⁷

A one-month variance swap is an asset whose final payoff is the sum of daily squared log returns of the underlying index (the S&P 500, in our case) over the next month. That asset gives the buyer protection against a surprise in equity return volatility (rv) over the next month. If investors are averse to periods of high realized volatility, then, we would expect to see negative average returns on one-month variance swaps, reflecting the cost of buying that insurance. A simple way to see that is to note that the Sharpe ratio of an asset, the ratio of its expected excess return to its standard deviation, is

$$\frac{E_t [R_{t+1} - R_{f,t+1}]}{SD_t [R_{t+1}]} = -corr_t (R_{t+1}, MU_{t+1}) \times std (MU_{t+1}) \quad (19)$$

for any return R_{t+1} , where $R_{f,t+1}$ is the risk-free rate and MU_{t+1} denotes the marginal utility of consumption on date $t + 1$. Assets that covary positively with marginal utility, and hence are hedges, earn negative average returns. So if realized volatility is high in high marginal utility states (in most models, bad times), then one-month variance swaps will earn high Sharpe ratios.

The first point on the left in the left-hand panel of Figure 7 (which is drawn from Dew-Becker et al. (2017)) plots average annualized Sharpe ratios on 1-month S&P 500 variance swaps between 1996 and 2014.¹⁸ The average Sharpe ratio is -1.4 , approximately three times larger (with the opposite sign) than the Sharpe ratio on the aggregate equity market in that period. Investors have historically paid extraordinarily large premia for protection against periods of high realized volatility, suggesting that they view those times as particularly bad (or as having very high marginal utility).

Now consider a j -month variance forward, whose payoff, instead of being the sum of squared returns over the next month ($t + 1$), is the sum of squared returns in month $t + j$ (so then the one-month variance swap above can also be called a 1-month variance forward). If an investor buys a j -month variance forward and holds it for a single month, selling it in month $t + 1$, then the variance forward protects her over that period against *news* about volatility in month $t + j$. If between t and $t + 1$ investors receive news that volatility will be higher in the future – i.e. if uncertainty rises

¹⁷See for example Egloff, Leippold, and Wu (2010); Ait-Sahalia et al. (2015); Dew-Becker et al. (2017). A large literature in finance studies the pricing of realized and expected future volatility. See, among many others, Adrian and Rosenberg (2008), Bollerslev et al. (2009), Heston (1993), Ang et al. (2006), Carr and Wu (2009), Bakshi and Kapadia (2003), Egloff, Leippold, and Wu (2010), and Ait-Sahalia, Karaman, and Mancini (2013) (see Dew-Becker et al. (2017) for a review).

¹⁸The data is described in Dew-Becker et al. (2017); it is obtained from a large asset manager and Markit, but may be closely approximated by portfolios of options, for which prices are widely available (e.g. from Optionmetrics).

– the holding period return on that j -period forward will increase. The left-hand panel of figure 7 also plots one-month holding period Sharpe ratios for variance forwards with maturities from 2 to 12 months. For all maturities higher than 2 months, the Sharpe ratios are near zero, and in fact the sample point estimates are positive. The Sharpe ratios are also all statistically significantly closer to zero than the Sharpe ratio on the one-month variance swap.

The left panel of Figure 7 therefore shows that investors have paid large premia for protection against surprises in realized volatility, but news about future uncertainty has had a premium that is indistinguishable from zero, and may even be positive. Realized volatility thus appears to have a large positive correlation with marginal utility, while shocks to expected volatility have a correlation that is close to zero or negative.

Using the options data described above, it is possible to extend those results further, back to 1983. The right-hand panel of figure 7 reports the average shape of the term structure of variance forward prices constructed from data on S&P 500 options (we study the term structure with this data because it is estimated more accurately than returns). The variance forwards are constructed from synthetic variance swaps, a calculation almost identical to our calculation of $Var_t^Q [s_{t+n}]$. The term structure reported here is directly informative about risk premia. The average return on an n -month variance claim is:

$$\underbrace{E \left[\frac{F_{n-1,t} - F_{n,t-1}}{F_{n,t-1}} \right]}_{\text{Average return of the forward}} \approx \frac{\overbrace{E [F_{n-1}] - E [F_n]}^{\text{Average slope of the forward curve}}}{E [F_n]} \quad (20)$$

where $F_{n,t}$ is the price on date t of an n -month volatility forward. The slope of the average term structure thus indicates the average risk premium on news about volatility n months forward. If the term structure is upward sloping, then the prices of the variance claims fall on average as their maturities approach, indicating that they have negative average returns. If it slopes down, then average returns are positive.

The right-hand panel of Figure 7 plots the average term structure of variance forward prices for the period 1983–2013. The term structure is strongly upward sloping for the first two months, again indicating that investors have paid large premia for assets that are exposed to realized variance and expected variance one month in the future. But the curve quickly flattens, indicating that the risk premia for exposure to fluctuations in expected variance further in the future have been much smaller.

The asset return data says that investors appear to have been highly averse to news about high realized volatility, while shocks to expected volatility do not seem to have been related to marginal utility. Figure 7 therefore confirms and complements the results from our VAR, that show that shocks to rv are associated with recessions but uncertainty shocks are not.

8 Equilibrium model and skewness evidence

The paper thus far has provided empirical evidence on two basic points: first, surprises in realized volatility in the stock market are associated with future declines in real activity, while uncertainty shocks, identified as second-moment news, are not; second, investors have historically paid large premia to hedge shocks to realized volatility, but have paid premia that have averaged to nearly zero to hedge shocks to uncertainty.

In this section we present a simple structural model that is consistent with those features of the data. The key ingredient is asymmetry in fundamental shocks. Intuitively, when fundamental shocks are skewed to the left, large shocks, which are associated with high realized volatility, tend to be negative.¹⁹ That is simply the definition of left skewness: the squared innovation is negatively correlated with the level of the innovation. We first discuss evidence that there is left skewness in economic activity and then describe the model and show that estimated impulse responses to uncertainty and realized volatility shocks in the model match what we find in the data.

8.1 Skewness

A potential source of negative correlation between output and realized volatility is negatively skewed shocks. Specifically, if some shock ε is negatively skewed, then $E[\varepsilon^3] < 0 \Rightarrow \text{cov}(\varepsilon, \varepsilon^2) < 0$. There are large literatures studying skewness in both aggregate stock returns and economic growth. We therefore provide a brief overview of the literature and basic evidence.

Table 4 reports the skewness of monthly and quarterly changes in a range of measures of economic activity. It includes both the standard skewness based on the scaled third moment and also the Kelley skewness based on percentiles, with bootstrapped standard errors. Nearly all the variables are negatively skewed, at both the monthly and quarterly levels.

In addition to real variables, table 4 also reports realized and option-implied skewness for S&P 500 returns.²⁰ The implied and realized skewness of monthly stock returns is substantially negative and similar to the skewness of capacity utilization. The realized skewness of stock returns is less negative than option-implied skewness, which is consistent with investors demanding a risk premium on assets that have negative returns in periods when realized skewness is especially negative (i.e. that covary positively with skewness).

In addition to the basic evidence reported here, there is a large literature providing much more sophisticated analyses of asymmetries in the distributions of output and stock returns. Morley and Piger (2012) provide an extensive analysis of asymmetries in the business cycle and review the large literature (e.g. McKay and Reis (2008)). Estimating a wide range of models, they show that

¹⁹Skewness in equilibrium quantities could arise because the fundamental shocks are skewed, or because symmetrical shocks are transmitted to the economy asymmetrically (perhaps because constraints, such as financial frictions, bind more tightly in bad times; see Kocherlakota (2000), or because firms respond to shocks in a concave manner, as in Ilut, Kehrig, and Schneider (2016)).

²⁰We obtain option-implied skewness from the CBOE's time series of its SKEW index, which is defined as $SKEW = 100 - 10 \times \text{Skew}(R)$. We thus report $10 - SKEW/10$.

those that fit aggregate output best have explicit non-linearity and negative skewness. Even after averaging across models using a measure of posterior probability, they find that their measure of the business cycle is substantially skewed to the left, consistent with the results reported in table 4. More recently, Salgado, Guvenen, and Bloom (2016) provide evidence that left skewness is a robust feature of business cycles, at both the macro and micro levels and across many countries.

The finance literature has also long recognized that there is skewness in aggregate equity returns and in option-implied return distributions (Campbell and Hentschel (1992); Ait-Sahalia and Lo (1998); Bakshi, Kapadia, and Madan (2003)). The skewness that we measure here appears to be pervasive and has existed in returns reaching back even to the 19th century (Campbell and Hentschel (1992)).

Taken as a whole, then, across a range of data sources and estimation methods, there is a substantial body of evidence that fluctuations in the economy are negatively skewed. If negative skewness drives our results, then we might expect that it would be the downward component of returns that would be most relevant for the VAR analysis in the previous section. To see that, we define the downward component of realized volatility as

$$rv_t^{down} \equiv \sum_{days \in t} r_t^2 \times 1\{r_t < 0\} \quad (21)$$

where $1\{\cdot\}$ is the indicator function. rv_t^{down} therefore isolates the downward component of realized volatility. We construct the upward part of realized volatility not as the squared returns on positive days but rather as the part of rv_t that is orthogonal to rv_t^{down} , i.e.

$$rv_t^{up} \equiv rv_t - proj\left(rv_t \mid rv_t^{down}\right) \quad (22)$$

where $proj\left(rv_t \mid rv_t^{down}\right)$ is the projection of rv_t onto rv_t^{down} . In this way, rv_t^{up} represents the part of rv_t that is separate from downward movements.

Figures A.20 and A.21 in the appendix report results from estimates of the benchmark specification where we replace rv_t with rv_t^{down} and rv_t^{up} , respectively. The top panels show that shocks to rv^{down} are associated with declines in employment and IP while news shocks are not, as in the main results. In the middle panels, though, shocks to rv^{up} have no effect, while the news shocks are now associated with declines. This result follows from the fact that the news shock is still a linear function of the shock to v_1 , and we know from past work, e.g. Bloom (2009), that v_1 by itself, when we do not control for rv , is followed by declines. So figures A.20 and A.21 show that when controlling for rv , as in our main analysis, what is relevant is controlling for the downward component, rather than the upward part.²¹ Finally, the bottom panels of figures A.20 and A.21 show that if we replace rv with a measure of just jumps in stock prices,²² the results are again

²¹The experiment can be reversed, with rv^{up} being defined by rv on days when the market rises and rv^{down} as the residual part of rv , and the results are similar.

²²The jump series here is constructed as the difference between realized volatility and bipower variation, where bipower variation is a measure of the diffusive part of returns each month. We obtain the data on bipower variation

similar to the benchmark, suggesting that our results can be accounted for purely through large movements in stock prices.

8.2 Equilibrium model

We now present a simple extension to the RBC model consistent with our empirical findings that (1) shocks to realized volatility are associated with declines in real activity, while shocks to expected volatility are not; (2) Sharpe ratios on short-term claims to volatility are much more negative than those on longer-term claims; and (3) output growth and equity returns are negatively skewed.

8.3 Model structure

Firms produce output with technology, A_t , capital, K_t , and labor, N_t ,

$$Y_t = A_t^{1-\alpha} K_t^\alpha N_t^{1-\alpha} \quad (23)$$

We set $\alpha = 0.33$, consistent with capital's share of income. Capital is produced subject to adjustment costs according to the production function

$$K_t = \left[(1 - \delta) + \frac{I_{t-1}}{K_{t-1}^{agg}} - \frac{\zeta}{2} \left(\frac{I_{t-1}}{K_{t-1}^{agg}} - \overline{I/K} \right)^2 \right] K_{t-1} \exp(-J\nu_t) \quad (24)$$

where I_t is gross investment, K_t^{agg} is the aggregate capital stock (which is external to individual firm decisions), ζ is a parameter determining the magnitude of adjustment costs and $\overline{I/K}$ is the steady-state investment/capital ratio. The term $\exp(-J\nu_t)$ is associated with downward jumps in productivity that are described below – those jumps destroy both technology and capital, similar to Gourio (2012). We set $\delta = 0.08/12$ (corresponding to a monthly calibration) and $\zeta = 0.5$.²³ Given the structure for adjustment costs and production, the equilibrium price and return on a unit of installed capital are

$$P_{K,t} = \frac{1}{1 - \zeta \left(\frac{I_t/K_{t-1}}{\overline{I/K}} - 1 \right)} \quad (25)$$

$$R_{K,t} = \frac{\alpha A_t^{1-\alpha} K_{t-1}^{\alpha-1} N_t^{1-\alpha} + (1 - \delta) P_{K,t}}{P_{K,t-1}} \quad (26)$$

A representative agent maximizes Epstein–Zin (1991) preferences over consumption and leisure,

$$V_t = \arg \max_{C,N} \log \left(C_{t+j} - bC_{t+j-1}^{agg} \right) - \theta \frac{N_{t+j}^{1+\chi}}{1+\chi} + \frac{\beta}{1-\gamma} \log E_t \exp((1-\gamma)V_{t+1}) \quad (27)$$

from the Oxford-Man Institute Realized Library (<https://realized.oxford-man.ox.ac.uk/>).

²³See, e.g., Cummins, Hassett, and Hubbard (1994) for estimates of adjustment costs similar to this value. Jermann (1998) use a similar values. $\zeta = 0.5$ is on the lower end of estimates based on aggregate data and more consistent with micro evidence, but our results are not sensitive to the choice of this parameter.

where C^{agg} is aggregate consumption, subject to the budget constraint

$$C_t + I_t \leq Y_t \quad (28)$$

Agents have log utility over consumption minus an external habit. We set the magnitude of the habit to $b = 0.8$ to help generate smoothness in consumption. β is set to $0.99^{1/12}$, χ to $1/3$ for a Frisch elasticity of 3, and θ to generate steady-state employment of $1/3$ (e.g. working 8 of 24 hours). γ is set to 12 to generate a Sharpe ratio on a claim to capital of 0.35, similar to what is observed for US equities.

The model is closed by the Euler equation and the optimization condition for labor,

$$1 = E_t \left[\beta \frac{\exp((1-\gamma)V_{t+1})}{E_t \exp((1-\gamma)V_{t+1})} \frac{C_t - bC_{t-1}^{agg}}{C_{t+1} - bC_t^{agg}} R_{K,t+1} \right] \quad (29)$$

$$\theta N_t^\chi (C_t - bC_{t-1}^{agg}) = (1-\alpha) A_t^{1-\alpha} K_{t-1}^\alpha N_t^{-\alpha} \quad (30)$$

We model realized volatility as in the empirical analysis as the square of the levered excess return on capital,

$$RV_t = (\lambda(R_{K,t+1} - R_{f,t}) + R_{f,t})^2 \quad (31)$$

where $R_{f,t}$ is the risk-free rate and the parameter λ determines the leverage of equity. We set $\lambda = 9.2$ to match the volatility of the aggregate stock market.

The only exogenous variable in the model is technology, A_t , which follows the process

$$\Delta \log A_t = \sigma_{t-1} \bar{\sigma}_\varepsilon \varepsilon_t - J(\nu_t - \bar{p}) + \mu \quad (32)$$

$$\log \sigma_t = \phi_\sigma \log \sigma_{t-1} + \sigma_\sigma \eta_t + \kappa_{\sigma,\nu} \nu_t \quad (33)$$

$$\varepsilon_t, \eta_t \sim N(0, 1) \quad (34)$$

$$\nu_t \sim \text{Bernoulli}(\bar{p}) \quad (35)$$

Technology follows a random walk in logs with drift, μ , set to 2 percent per year. ε_t is a normally distributed innovation that affects technology in each period, while ν_t is a shock that is equal to zero in most periods but equal to 1 with probability \bar{p} . It induces downward jumps in technology, with J determining the size of the jump and \bar{p} the average frequency. σ_t determines the volatility of normally distributed shocks to technology. It is itself driven by two shocks: an independent shock η_t (volatility news) and also the jumps ν_t .²⁴ Downward jumps in technology can be associated with higher volatility, generating GARCH-type effects. The volatility process thus has two features that will be important in matching the data: it has news shocks, and it is countercyclical for $\kappa_{\sigma,\nu} < 0$.

ϕ_σ and σ_σ are calibrated so that $\log \sigma_t$ has a standard deviation of 0.35 and a one-month autocorrelation of 0.91, consistent with the behavior of the VIX. $\kappa_{\sigma,\nu}$ is set to -0.7, which implies that

²⁴There are various mechanisms through which skewness could arise, e.g. financial frictions (Kocherlakota (2000)), investment dynamics (Gilchrist and Williams (2000)), or employment dynamics (McKay and Reis (2008)). Here we use a simple setup in which skewness comes directly from technology shocks.

a jump in technology increases σ_t by 2 standard deviations, generating countercyclical volatility. $\bar{\sigma}_\varepsilon$ is set so that normally distributed shocks on average generate a standard deviation of output growth close to the value of 1.92 we observe empirically. Jumps on average reduce technology by 8 percent (which is 3.2 times $\bar{\sigma}_\varepsilon$, the average standard deviation of the Gaussian TFP shocks) and are calibrated to occur once every 10 years on average. We thus think of them as representing small disasters or large recessions (consistent Backus, Chernov, and Martin (2011) and with the view of skewed recessions in Salgado, Guvenen, and Bloom (2016)), rather than depression-type disasters.²⁵ While the size of the jump seems large initially, recall that the model is calibrated to match the standard deviation of output growth.

We solve the model by projecting the consumption and value functions on a set of Chebyshev polynomials up to the 6th order (a so-called global solution) to ensure accuracy not only for real dynamics but also for asset prices and realized volatility. Integrals are calculated using Gaussian quadrature with 20 points. Euler equation errors are less than $10^{-5.0}$ across the range of the state space that the simulation explores and have an average absolute value of $10^{-5.3}$. The use of a global solution method allows for high accuracy in the solution, but also makes it infeasible to search over many parameters or estimate the model, which is why we explore just a single calibration here.

8.4 Simulation results

We examine three sets of implications of the model: basic moments of real and financial variables, risk premia, and VAR estimates. All results are population statistics calculated from a simulation lasting 10,000 years.

Table 5 reports basic moments of returns on capital and growth rates of output, consumption, and investment. The model generates negative skewness in output, consumption, and investment, consistent with the data, but the skewness is much larger than is observed empirically. Mean growth rates of real variables are similar to the data. The standard deviation of output is very close to the data, while consumption is more and investment less volatile; the gap between the two is smaller than observed empirically, however. Table 5 thus suggests that the model generates moments that are broadly consistent with the data, in particular generating comovement among aggregate variables (the three growth rate series in table 5 have correlations between 0.53 and 0.96) and volatilities that are empirically reasonable.

The model also fits the financial variables well. It is calibrated in order to match the mean and volatility of stock returns (through the choice of leverage and risk aversion), and they are also skewed down (though again more than in the data). For the VIX – one-month implied volatility – the model is able to match the positive variance risk premium in the data: implied volatility is on average 1/3 higher than realized volatility. While the VIX in the model is less volatile than in the data, it does share the observed right skewness.

²⁵A realistic extension of the model would be to allow for jumps to be drawn from a distribution, rather than all having the same size. See, e.g., Barro and Jin (2011).

Figure 8 plots the Sharpe ratios of volatility claims in the model that correspond to the forward volatility claims examined in Figure 7. As in the left panel of Figure 7, the Sharpe ratio of the one-month asset, which is a claim to realized volatility, is far more negative than the Sharpe ratios for the claims with longer maturities. Intuitively, this is because shocks to volatility expectations, η_t , have relatively small effects on consumption and lifetime utility, hence earning a small risk premium. Shocks to realized volatility, on the other hand, tend to isolate the jumps, ν_t , (as we will show below), so they earn larger premia.

Finally, figure 9 plots IRFs from the estimation of our VAR in the simulation of the model. The gray shaded regions represent confidence bands from the VAR in the empirical data.²⁶ The red lines are the population IRFs from the VAR estimated in the simulated model. The IRFs here represent responses to unit standard deviation impulses to the identified shocks.²⁷

The model matches well in terms of how an uncertainty shock predicts future stock market volatility and output. Most importantly, it generates a large and empirically reasonable decline in output following the realized volatility shock and an economically small response of output to the identified uncertainty shock. Both IRFs lie within the confidence bands at most horizons. The model performs more poorly in matching the behavior of employment. Employment responds little to either of the shocks, so it matches the empirical IRF for the news shock, but not the realized volatility shock.

The VAR results are notable because they replicate the behavior of output observed empirically even though there is no structural “realized volatility shock” in the model. Rather, the identified RV shock comes from the jump in TFP in the model ($J\nu_t$). To see that, we report the correlations between the VAR-identified shocks and the structural shocks in the model in the bottom section of table 5. The RV shock is correlated nearly exclusively with $J\nu_t$, the jump shock in the model. So the VAR successfully identifies the jumps as realized volatility shocks, which are then structurally, but obviously not causally, related to declines in output.

The identified uncertainty shock, as we would hope, is almost purely correlated with η_t , the volatility news shock. So, as discussed above, and similar to the BB and RUBC models, our main VAR specification does a good job in this setting – a non-linear production model – of actually identifying true structural shocks and also fitting the qualitative behavior of our empirical VAR analysis.

To further illustrate the mechanisms driving the model, the dotted blue lines in figure 9 plot the dynamic responses of the variables to the true structural shocks in the model. In the top row the dotted lines plot the responses to the jump shock, ν_t , while in the bottom row they plot the responses to the uncertainty shock, η_t (again in both cases scaled to be unit standard deviations). In

²⁶These confidence bands are slightly different from what is reported in figure 2 because they use the levels of *RV* and option-implied volatility instead of their logs (since the model is in discrete time and realized volatility is the monthly squared return, it can approach zero, causing the log transform to break down).

²⁷Note that this differs from the empirical analysis above, which scaled the realized volatility shock to have the same effect on uncertainty (expected future volatility) as the news shock. The reason for the difference here is that realized volatility in the model is a relatively weak predictor of future volatility, making the scaling sometimes explode to infinity.

four of the cases the response to the structural shock is essentially identical to the estimate from the VAR. The responses of output are slightly different (since the VAR is not a perfect representation of the structural model) but still very close to the VAR.

The fact that the response of employment to the jump shock is small is a consequence of the use of preferences consistent with balanced growth following King, Plosser, and Rebelo (1988). Since the jumps in technology also involve destruction of capital, they represent essentially a shift along the balanced growth path. In the absence of habit formation, they would have precisely zero effect on employment; the habit causes a slight positive response. Obtaining more negative responses of employment would require adding frictions or changing the preferences, which we leave to future work.

The results in this section show that two features of the model explain the behavior of the VAR in the simulations. First, the negative relationship between realized volatility and output comes from the fact that realized volatility is high in months with jumps, and all jumps are negative. Second, the lack of a relationship between output and uncertainty comes from the fact that the true uncertainty shock has essentially zero effect on output.

In the end, then, this section shows that a simple production model can match the basic features of the data that we have estimated in this paper: output responds negatively to shocks to realized volatility but not to shocks to uncertainty, there is a much larger risk premium for realized than expected volatility, and economic activity and stock returns are both skewed to the left.

9 Conclusion

The key distinction that this paper draws is between realized volatility and uncertainty. Volatility matters for output, but it is the realized part that is robustly followed by downturns. Changes in expected volatility – uncertainty shocks – appear to have no significant negative effects. Evidence from asset prices and risk premia is consistent with these findings, and we develop a simple model that can rationalize the data and also justifies our identification scheme.

The empirical results are inconsistent with theories in which pure shocks to aggregate uncertainty play an important role in driving real activity. Moreover, the identification scheme used in this paper is shown to successfully identify true structural shocks to uncertainty in leading recent models.

More constructively, this paper aims to lay out a specific view of the joint behavior of stock market volatility and the real economy. There appear to be negative shocks to the stock market that occur at business cycle frequencies, are associated with high realized volatility and declines in output, and are priced strongly by investors. The simple idea that fundamentals are skewed left can explain our VAR evidence, the pricing of volatility risk, and the negative unconditional correlation between economic activity and volatility.

References

- Adrian, Tobias and Joshua Rosenberg**, “Stock Returns and Volatility: Pricing the Short-Run and Long-run Components of Market Risk,” *The Journal of Finance*, 2008, *63*(6), 2997–3030.
- Ait-Sahalia, Yacine and Andrew W Lo**, “Nonparametric estimation of state-price densities implicit in financial asset prices,” *The Journal of Finance*, 1998, *53* (2), 499–547.
- , **Mustafa Karaman, and Loriano Mancini**, “The Term Structure of Variance Swaps, Risk Premia and the Expectations Hypothesis,” 2015. Working paper.
- Alexopoulos, Michelle and Jon Cohen**, “Uncertain Times, Uncertain Measures,” 2009. Working paper.
- Amengual, Dante and Dacheng Xiu**, “Resolution of policy uncertainty and sudden declines in volatility,” 2014. Working paper.
- Andersen, Torben G. and Oleg Bondarenko**, “Construction and Interpretation of Model-Free Implied Volatility,” Working Paper, National Bureau of Economic Research September 2007.
- Andersen, Torben G, Tim Bollerslev, and Francis X Diebold**, “Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility,” *The Review of Economics and Statistics*, 2007, *89* (4), 701–720.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang**, “The Cross-Section of Volatility and Expected Returns,” *The Journal of Finance*, 2006, *61*(1), 259–299.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R Sims**, “Uncertainty and economic activity: Evidence from business survey data,” *American Economic Journal: Macroeconomics*, 2013, *5* (2), 217–249.
- Baker, Scott and Nicholas Bloom**, “The Effects of Uncertainty Shocks on International Economic Growth,” 2013. Working paper.
- Bakshi, Gurdip and Nikunj Kapadia**, “Delta-Hedge Gains and the Negative Market Volatility Risk Premium,” *The Review of Financial Studies*, 2003, *16*(2), 527–566.
- , —, and **Dilip Madan**, “Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options,” *Review of Financial Studies*, 2003, *16* (1), 101–143.
- Bansal, Ravi, Dana Kiku, and Amir Yaron**, “Risks for the long run: Estimation with time aggregation,” *Journal of Monetary Economics*, 2016, *82*, 52–69.
- Barberis, Nicholas and Andrei Shleifer**, “Style investing,” *Journal of financial Economics*, 2003, *68* (2), 161–199.
- , **Robin Greenwood, Lawrence Jin, and Andrei Shleifer**, “X-CAPM: An extrapolative capital asset pricing model,” *Journal of financial economics*, 2015, *115* (1), 1–24.
- , —, —, and —, “Extrapolation and bubbles,” Technical Report, National Bureau of Economic Research 2016.
- Barro, Robert J and Tao Jin**, “On the size distribution of macroeconomic disasters,” *Econometrica*, 2011, *79* (5), 1567–1589.
- Barsky, Robert B and Eric R Sims**, “News shocks and business cycles,” *Journal of Monetary Economics*, 2011, *58* (3), 273–289.


- , **Susanto Basu**, and **Keyoung Lee**, “Whither News Shocks?,” Technical Report, National Bureau of Economic Research 2014.
- Beaudry, Paul and Franck Portier**, “Stock Prices, News, and Economic Fluctuations,” *The American Economic Review*, 2006, *96* (4), 1293–1307.
- and — , “News-driven business cycles: Insights and challenges,” *Journal of Economic Literature*, 2014, *52* (4), 993–1074.
- Berger, David and Joseph Vavra**, “Volatility and Pass-through,” *NBER Working Paper Serie No. 19651*, 2013.
- Bernanke, Ben S., Jean Boivin, and Piotr Elias**, “Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach.,” *Quarterly Journal of Economics*, 2005, *120* (1), 387–422.
- Bloom, Nicholas**, “The Impact of Uncertainty Shocks,” *Econometrica*, 2009, *77*(3), 623–685.
- , **Scott Baker**, and **Steven J. Davis**, “Measuring Economic Policy Uncertainty,” 2015. Working paper.
- Bollerslev, Tim**, “Generalized autoregressive conditional heteroskedasticity,” *Journal of econometrics*, 1986, *31* (3), 307–327.
- , **George Tauchen**, and **Hao Zhou**, “Expected Stock Returns and Variance Risk Premia,” *Review of Financial Studies*, 2009, *22*(11), 4463–4492.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic expectations and credit cycles,” *The Journal of Finance*, 2017.
- Born, Benjamin and Johannes Pfeifer**, “Policy Risk and the Business Cycle,” *Journal of Monetary Economics*, 2014, *68*, 65–85.
- Broadie, Mark and Jerome Detemple**, “American option valuation: new bounds, approximations, and a comparison of existing methods,” *Review of Financial Studies*, 1996, *9* (4), 1211–1250.
- Caggiano, Giovanni, Efrem Castelnuovo, and Nicolas Groshenny**, “Uncertainty shocks and unemployment dynamics in US recessions,” *Journal of Monetary Economics*, 2014, *67*, 78–92.
- Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek**, “The macroeconomic impact of financial and uncertainty shocks,” *European Economic Review*, 2016.
- Campbell, John Y. and Ludger Hentschel**, “No news is good news: An asymmetric model of changing volatility in stock returns,” *Journal of Financial Economics*, 1992, *31* (3), 281–318.
- , **Martin Lettau, Burton G. Malkiel, and Yexiao Xu**, “Have Individual Stocks Become More Volatile? an Empirical Exploration of Idiosyncratic Risk,” *Journal of Finance*, 2001, *56* (1), 1–43.
- Carr, Peter and Liuren Wu**, “Variance Risk Premiums,” *Review of Financial Studies*, 2009, *22*(3), 1311–1341.
- Carriero, Andrea, Todd E. Clark, and Massimiliano Giuseppe Marcellino**, “Measuring uncertainty and its impact on the economy,” 2016. Working paper.

- Cesa-Bianchi, Ambrogio, M. Hashem Pesaran, and Alessandro Rebucci**, “Uncertainty and Economic Activity: A Multi-Country Perspective,” 2018. Working paper.
- Christiano, Lawrence J. and Martin S. Eichenbaum**, “Temporal Aggregation and Structural Inference in Macroeconomics,” 1986. Working paper.
- , **Roberto Motto, and Massimio Rostagno**, “Risk Shocks,” *American Economic Review*, 2014, *104*(1), 27–65.
- Creal, Drew D and Jing Cynthia Wu**, “Monetary policy uncertainty and economic fluctuations,” *International Economic Review*, 2017, *58* (4), 1317–1354.
- Cummins, Jason G., Kevin A. Hassett, and R. Glenn Hubbard**, “A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments,” *Brookings Papers on Economic Activity*, 1994, *1994* (2), 1–74.
- Decker, Ryan, Pablo N. D’Erasmus, and Hernan Moscoso Boedo**, “Market exposure and endogenous firm volatility over the business cycle,” *American Economic Journal: Macroeconomics*, 2016, *8* (1), 148–198.
- Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez**, “The Price of Variance Risk,” 2016. Working paper.
- Drechsler, Itamar and Amir Yaron**, “What’s Vol Got to Do with it?,” *The Review of Financial Studies*, 2011, *24*(1), 1–45.
- Duffee, Gregory R.**, “Information in (and not in) the Term Structure,” *Review of Financial Studies*, 2011, *24*(9), 2895–2934.
- Egloff, Daniel, Markus Leippold, and Liuren Wu**, “The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments,” *Journal of Financial and Quantitative Analysis*, 2010, *45*(5), 1279–1310.
- Eisfeldt, Andrea L. and Adriano A. Rampini**, “Capital Reallocation and Liquidity,” *Journal of Monetary Economics*, 2006, *53* (3), 369–399.
- Engle, Robert F**, “Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation,” *Econometrica: Journal of the Econometric Society*, 1982, pp. 987–1007.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *The Journal of Political Economy*, 1991, *99*(2), 263–286.
- Fernandez-Villaverde, Jesus, Pablo Guerron, Juan F. Rubio-Ramirez, and Martin Uribe**, “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 2011, *101*(6), 2530–2561.
- , **Pablo Guerron-Quintana, Keith Kuester, and Juan Rubio-Ramirez**, “Fiscal Volatility Shocks and Economic Activity,” 2013. Working paper.
- Gatheral, Jim and Antoine Jacquier**, “Convergence of Heston to SVI,” *Quantitative Finance*, 2011, *11* (8), 1129–1132.
- and —, “Arbitrage-free SVI volatility surfaces,” *Quantitative Finance*, 2014, *14* (1), 59–71.

- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer**, “Expectations and investment,” *NBER Macroeconomics Annual*, 2016, 30 (1), 379–431.
- Gilchrist, Simon and John C Williams**, “Putty-clay and investment: a business cycle analysis,” *Journal of Political Economy*, 2000, 108 (5), 928–960.
- , **Jae W. Sim, and Egon Zakrajsek**, “Uncertainty, Financial Frictions, and Investment Dynamics,” Working Paper 20038, National Bureau of Economic Research April 2014.
- Gourio, Francois**, “Disaster Risk and Business Cycles,” *American Economic Review*, 2012, 102(6), 2734–2766.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of returns and expected returns,” *The Review of Financial Studies*, 2014, 27 (3), 714–746.
- and **Samuel G Hanson**, “Waves in ship prices and investment,” *The Quarterly Journal of Economics*, 2014, 130 (1), 55–109.
- Güvener, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, 122 (3), 621–660.
- Homescu, Cristian**, “Implied Volatility Surface: Construction Methodologies and Characteristics,” 2011. Working paper.
- Ilut, Cosmin, Matthias Kehrig, and Martin Schneider**, “Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News,” *NBER Working Paper Series*, 2015.
- Jiang, George J. and Yisong S. Tian**, “Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index,” *The Journal of Derivatives*, 2007, 14 (3), 35–60.
- Jurado, Kyle, Sydney Ludvigson, and Serena Ng**, “Measuring Uncertainty,” *American Economic Review*, 2015, 105 (3), 1177–1216.
- Kelly, Bryan, Lúboš Pástor, and Pietro Veronesi**, “The price of political uncertainty: Theory and evidence from the option market,” *The Journal of Finance*, 2016.
- King, Robert G., Charles I. Plosser, and Sergio T. Rebelo**, “Production, Growth and Business Cycles : I. The Basic Neoclassical Model,” *Journal of Monetary Economics*, 1988, 21(2–3), 195–232.
- Kocherlakota, Narayana**, “Creating business cycles through credit constraints,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2000, 24 (3), 2–10.
- Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran**, “The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation,” 2016. Working paper.
- Leduc, Sylvain and Zheng Liu**, “Uncertainty Shocks are Aggregate Demand Shocks,” 2015. Working paper.
- Lee, Roger W.**, “The Moment Formula for Implied Volatility at Extreme Strikes,” *Mathematical Finance*, 2004, 14 (3), 469–480.
- Ludvigson, Sydney C. and Serena Ng**, “The Empirical Risk-Return Relation: A Factor Analysis Approach,” *Journal of Financial Economics*, 2007, 83(1), 171–222.

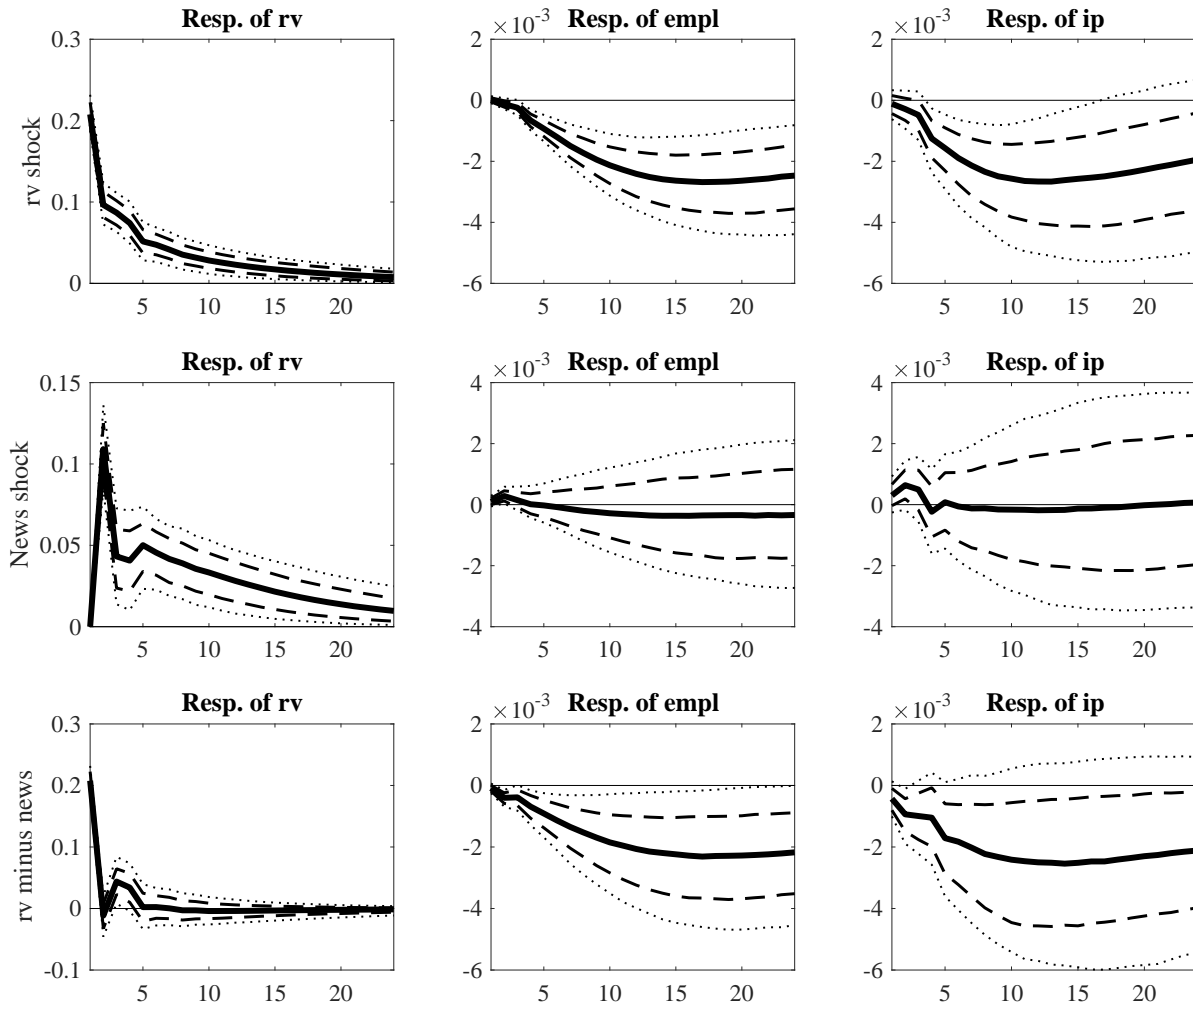
- Ludvigson, Sydney C, Sai Ma, and Serena Ng**, “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?,” Technical Report, National Bureau of Economic Research 2015.
- McKay, Alisdair and Ricardo Reis**, “The brevity and violence of contractions and expansions,” *Journal of Monetary Economics*, 2008, 55 (4), 738–751.
- Morley, James and Jeremy Piger**, “The asymmetric business cycle,” *Review of Economics and Statistics*, 2012, 94 (1), 208–221.
- Popiel, Michal**, “Fiscal policy uncertainty and US output,” 2017. Working paper.
- Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom**, “Skewed Business Cycles.” Working paper.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 2004, 112 (3), 695–717.
- , **Christopher I. Telmer, and Amir Yaron**, “Asset pricing with idiosyncratic risk and overlapping generations,” *Review of Economic Dynamics*, 2007, 10 (4), 519 – 548.
- Systems, Zeliade**, “Quasi-Explicit Calibration of Gatheral’s SVI model,” Technical Report, Zeliade Systems 2009.
- Taylor, Stephen J., Pradeep K. Yadav, and Yuanyuan Zhang**, “The information content of implied volatilities and model-free volatility expectations: Evidence from options written on individual stocks,” *Journal of Banking & Finance*, 2010, 34 (4), 871–881.
- Tibshirani, Robert**, “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, 58 (1), 267–288.

Figure 1: Time series of realized volatility and expectations



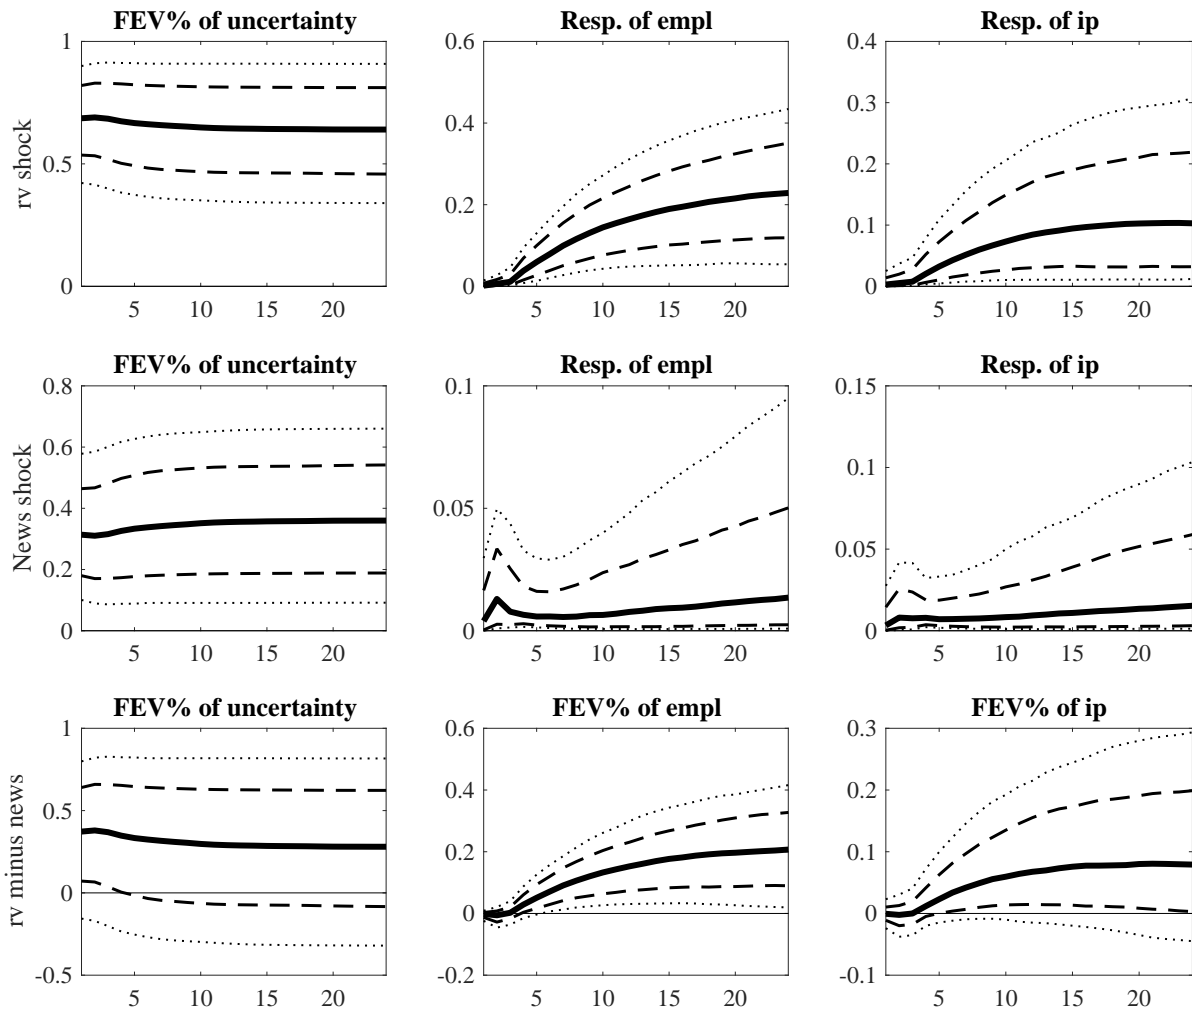
Note: Time series of realized volatility (RV), and 1-month option-implied volatility (V_1), in annualized standard deviation units. Grey bars indicate NBER recessions.

Figure 2: Impulse response functions from benchmark VAR




Note: Responses of rv , employment, and industrial production to shocks to rv and the identified uncertainty shock, in a VAR with rv , v_1 , federal funds rate, log employment, and log industrial production. The IRFs are scaled so that the two shocks have equal cumulative effects on rv over months 2–24 following the shock. The sample period is 1983–2014. The dotted lines are 68% and 90% confidence intervals.

Figure 3: Forecast error variance decomposition




Note: Fraction of the forecast error variance of uncertainty (the expected sum of *rv* over the next 24 months), employment, and industrial production to shocks to *rv* and uncertainty in the VAR of figure 2.

Figure 4: Fitted uncertainty




(a) Decomposition of total uncertainty in the benchmark specification




(b) Total uncertainty in benchmark vs. unrestricted specification

Note: The top panel reports a decomposition of total uncertainty (the conditional expectation of the sum of rv over the next 24 months) between the component driven by the rv shock and the component driven by the uncertainty shock in the benchmark model. The bottom panel reports the total uncertainty in the benchmark model and in the unrestricted specification.


Figure 5: Response of employment to rv and uncertainty shocks across specifications




(a) rv and v_1 ordered last



(b) Replacing v_1 with v_6




(c) Lasso




(d) Unrestricted

Note: Response of employment to RV shocks (left panels) and news shocks (middle panels) with the difference in the right panel and different model specifications in each row. Row (a) orders rv and v_1 last. Row (b) uses v_6 instead of v_1 . Row (c) uses lasso to estimate the VAR (see section A.2 for details). Row (d) estimates the benchmark VAR without any coefficient restrictions.


Figure 6: Response of IP to rv and uncertainty shocks across specifications




(a) rv and v_1 ordered last



(b) Replacing v_1 with v_6




(c) Lasso



(d) Unrestricted


Note: See figure 5.

Figure 7: Forward variance claims: returns and prices



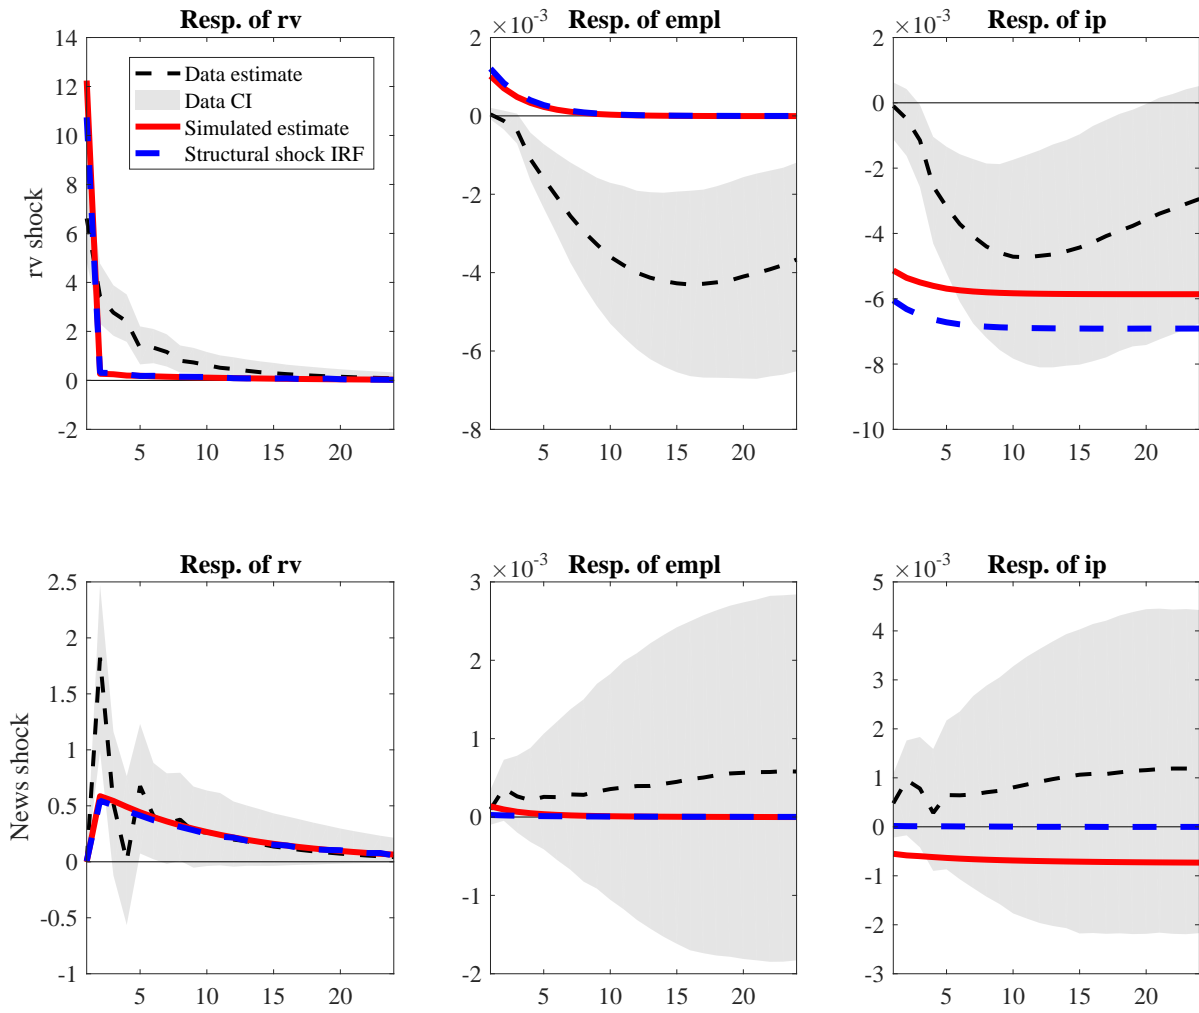
Note: Panel A shows the annualized Sharpe ratio for the forward variance claims, constructed using variance swaps. The returns are calculated assuming that the investment in an n -month variance claim is rolled over each month. Dotted lines represent 95% confidence intervals. All tests for the difference in Sharpe ratio between the 1-month variance swap and any other maturity confirm that they are statistically different with a p -value of 0.03 (for the second month) and < 0.01 (for all other maturities). The sample used is 1996–2013. For more information on the data sources, see Dew-Becker et al. (2017). Panel B shows the average prices across maturities of synthetic forward variance claims constructed from option prices for the period 1983–2014. All prices are reported in annualized standard deviation units. Maturity zero corresponds to average realized volatility.

Figure 8: Annual Sharpe ratios on forward claims (simulated structural model)



Note: Annual Sharpe ratios on forward variance claims in the simulated model of section 8. The Sharpe ratios are constructed as in Figure 7.

Figure 9: IRFs from structural model



Note: Impulse response functions from data simulated from the model in Section 8. Solid lines correspond to IRFs estimated using our VAR methodology as in Figure 2. Dashed lines correspond to IRFs for the two structural shocks $J\nu_t$ and η_t .

Table 1: Relationship between employment, industrial production and volatility

	Employment		Industrial Production	
v_1	-0.10**	0.05	-0.09*	0.16
	(0.04)	(0.07)	(0.05)	(0.12)
rv		-0.14**		-0.23*
		(0.07)		(0.13)

Note: Results from regressions of changes in log employment and industrial production on the current value and four lags of implied and realized volatility (in logs). The coefficients and standard errors in the table are for the average of the coefficients on each variable. Standard errors are calculated using the Newey–West (1987) method with 12 lags.

Table 2: VAR-implied and alternative measures of uncertainty

	RV	News	p(difference)
BBD	14.6***	17.1	0.28
	(2.3)	(6.8)	
BBD (monetary)	23.2***	10.5	0.20
	(7.1)	(8.9)	
BBD (stock market)	38.7***	19.7	0.26
	(14.4)	(18.1)	
Michigan	0.79*	0.82	0.96
	(0.44)	(0.59)	

Note: The table shows regressions of alternative measures of uncertainty on the components of VAR-implied uncertainty from rv and news shocks. BBD is the Baker, Bloom, and Davis (2015) measure, BBD (monetary) and BBD (stock market) are the corresponding subcomponents of that measure, and Michigan is a measure of uncertainty from the Michigan Consumer Survey (used by Leduc and Liu (2016)). The last column reports the p-value for a test of the difference between the two coefficients.

Table 3: Predictability of 6-month rv

Predictors	(1)	(2)	(3)	(4)	(5)
rv_t	0.16*	0.15	0.16*	0.13	0.09
	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)
$v_{1,t}$	0.57***	0.46***	0.60**	0.58***	0.60***
	(0.14)	(0.16)	(0.29)	(0.15)	(0.15)
$v_{6,t}$			-0.03		
			(0.30)		
rv_{t-1}		0.12**			
		(0.05)			
FFR				0.01	
				(0.01)	
$\Delta empl$				-4.98	
				(10.87)	
Δip				-4.81	
				(3.57)	
PC1					-0.010
					(0.009)
PC2					-0.017
					(0.012)
PC3					-0.011
					(0.007)
$R_{S\&P}$					0.12
					(0.41)
Default spread					0.07
					(0.09)
Adj. R^2	0.44	0.45	0.44	0.45	0.46

Note: Results of linear predictive regressions of realized volatility over the next six months on lagged rv , option-implied volatility, and various macroeconomic variables, with Hansen–Hodrick (1980) standard errors using a 6-month lag window. PC1–3 are principal components from the data set used in Ludvigson and Ng (2007). The default spread is the difference in yields on Baa and Aaa bonds. The sample is 1983–2014.

Table 4: Skewness

	Skewness		Kelley Skewness		Sample
	Monthly	Quarterly	Monthly	Quarterly	
<i>Panel A: real economic activity</i>					
Employment	-0.31 (0.33)	-0.87*** (0.19)	-0.14*** (0.04)	-0.23*** (0.07)	1952-2016
Capacity Utilization	-1.06*** (0.31)	-1.20*** (0.37)	-0.03 (0.05)	-0.15 (0.11)	1969-2016
IP	0.14 (0.50)	-0.63** (0.30)	-0.05 (0.04)	-0.09 (0.08)	1952-2016
IP, starting 1960	-0.87** (0.32)	-0.99** (0.35)	-0.07 (0.04)	-0.16** (0.08)	1960-2016
Output		-0.32 (0.24)		0.02 (0.07)	1952-2016
Non-durable consumption		-0.33 (0.34)		0.09 (0.08)	1952-2016
Durable consumption		-0.25 (0.43)		-0.14** (0.06)	1952-2016
Investment		-0.62** (0.26)		0.03 (0.10)	1952-2016
Residential investment		-0.41 (0.29)		-0.13* (0.07)	1952-2016
Non-residential investment		-0.76** (0.32)		-0.21*** (0.06)	1952-2016
<i>Panel B: skewness of S&P 500 monthly returns</i>					
	Estimate	Std. Err.			
Implied (since 1990)	-1.81***	(0.09)			
Realized (since 1926)	0.18	(0.55)			
Realized (since 1948)	-0.54***	(0.19)			
Realized (since 1990)	-0.67***	(0.20)			

Note: Panel A reports the skewness of log growth rates of measures of real activity. The first column reports the skewness of monthly changes, the second column the skewness of quarterly changes. Kelley skewness is computed as $KSK = \frac{P(90)+P(10)-2P(50)}{P(90)-P(10)} \in [-1, 1]$ where $P(*)$ denotes percentiles. Panel B reports the realized skewness of S&P 500 monthly returns in different periods, as well as the implied skewness computed by the CBOE using option prices. Standard errors are constructed by bootstrapping with 1000 replications.

Table 5: Model Calibration

<i>Panel A: Moments</i>	Model			Data		
	Mean	Std.	Skewness	Mean	Std.	Skewness
Output	1.99	2.02	-4.48	1.33	1.92	-0.11
Investment	1.99	3.64	-4.46	2.35	7.43	-0.03
Consumption	1.99	1.58	-3.41	1.20	1.06	-0.28
Returns	6.50	15.00	-4.92	7.46	14.77	-0.48
VIX	20.92	3.56	1.25	20.05	7.75	1.94

<i>Panel B: Corr. of VAR and structural shocks</i>	Structural shocks			
	Jv_t	η_t	ε_t	
	RV	0.88	0.00	0.07
VAR identified shocks	Uncertainty	0.11	0.91	0.04

Note: Panel A reports the mean, standard deviation, and skewness of financial and macroeconomic variables in the data and in the model (returns are reported in annualized terms). Panel B shows the correlation between the structural shocks in the model and the shocks identified in the VAR.

A.1 Construction of option-implied volatility, V_n

In this section we describe the details of the procedure we use to construct model implied uncertainty at different horizons, starting from our dataset of end-of-day prices for American options on S&P 500 futures from the CME.

Our implied volatility is written as a function of option prices,

$$V_{n,t} \equiv \text{Var}_t^Q [s_{t+n}] \tag{A.1}$$

$$= 2 \int_0^\infty \frac{1 - \log\left(\frac{K}{e^{rt} S_t}\right)}{B_t(n) K^2} O(K) dK - \left(e^{rt} \int_0^\infty \frac{O(K)}{B_t(n) K^2} dK \right)^2 \tag{A.2}$$

Note that this formula holds generally, requiring only the existence of a well-behaved pricing measure; there is no need to assume a particular specification for the returns process. $\text{Var}_t^Q [s_{t+n}]$ is calculated as an integral over option prices, where K denotes strikes, $O_t(n, K)$ is the price of an out-of-the-money option with strike K and maturity n , and $B_t(n)$ is the price at time t of a bond paying one dollar at time $t + n$. $V_{n,t}$ is equal to the option-implied variance of log stock prices n months in the future.

The result for $\text{Var}_t^Q [s_{t+n}]$ is obtained from equation 3 in Bakshi, Kapadia, and Madan (2003) by first setting $H(S) = \log(S)$ to obtain $E_t^Q [\log S_{t+n}]$ and then defining $G(S) = \left(\log(S) - E_t^Q [\log S_{t+n}] \right)^2$ and inserting it into equation 3 in place of H .

A.1.1 Main steps of construction of V_n

A first step in constructing the model-free implied volatility is to obtain implied volatilities corresponding to the observed option prices. We do so using a binomial model.¹ For the most recent years, CME itself provides the implied volatility together with the option price. For this part of the sample, the IV we estimate with the binomial model and the CME's IV have a correlation of 99%, which provides an external validation on our implementation of the binomial model.

Once we have estimated these implied volatilities, we could in theory simply invert them to yield implied prices of European options on forwards. These can then be used to compute V_n directly as described in equation (17).

In practice, however, an extra step is required before inverting for the European option prices and integrating to obtain the model-free implied volatility. The model-free implied volatility defined in equation (17) depends on the integral of option prices over *all strikes*, but option prices are only observed at discrete strikes. We are therefore forced to interpolate option prices between available strikes and also extrapolate beyond the bounds of observed strikes.² Following the literature, we fit a parametric model to the Black–Scholes implied volatilities of the options and use the model to then interpolate and extrapolate across all strikes (see, for example, Jiang and Tian (2007), Carr

¹See for example Broadie and Detemple (1996) and Bakshi, Kapadia, and Madan (2003), among others.

²See Jiang and Tian (2007) for a discussion of biases arising from the failure to interpolate and extrapolate.

and Wu (2009), Taylor, Yadav, and Zhang (2010), and references therein). Only after this extra interpolation-extrapolation step, the fitted implied volatilities are then inverted to yield option prices and compute V_n according to equation (17). To interpolate and extrapolate the implied volatility curve, we use the SVI (stochastic volatility inspired) model of Gatheral and Jacquier (2014).

In the next sections, we describe in more detail the interpolation-extrapolation step of the procedure (SVI fitting) as well as our construction of V_n after fitting the SVI curve. Finally, we report a description of the data we use and some examples and diagnostics on the SVI fitting method.

A.1.2 SVI interpolation: theory

There are numerous methods for fitting implied volatilities across strikes. Homescu (2011) provides a thorough review. We obtained the most success using Gatheral’s SVI model (see Gatheral and Jacquier 2014). SVI is widely used in financial institutions because it is parsimonious but also known to approximate well the behavior of implied volatility in fully specified option pricing models (e.g. Gatheral and Jacquier (2011)); SVI also satisfies the limiting results for implied volatilities at very high and low strikes in Lee (2004), and, importantly, ensures that no-arbitrage conditions are not violated.

The SVI model simply assumes a hyperbolic relationship between implied variance (the square of the Black–Scholes implied volatility) and the log moneyness of the option, k (log strike/forward price).

$$\sigma_{BS}^2(k) = a + b \left(\rho(k - m) + \sqrt{(k - m)^2 + \sigma^2} \right)$$

where $\sigma_{BS}^2(k)$ is the implied variance under the Black–Scholes model at log moneyness k . SVI has five parameters: a , b , ρ , m , and σ . The parameter ρ controls asymmetry in the variances across strikes. Because the behavior of options at high strikes has minimal impact on the calculation of model-free implied volatilities, and because we generally observe few strikes far above the spot, we set $\rho = 0$ (in simulations with calculating the VIX for the S&P 500 – for which we observe a wide range of options – we have found that including or excluding ρ has minimal impact on the result).

We fit the parameters of SVI by minimizing the sum of squared fitting errors for the observed implied volatilities. Because the fitted values are non-linear in the parameters, the optimization must be performed numerically. We follow the methodology in Zeliade (2009) to analytically concentrate a and b out of the optimization. We then only need to optimize numerically over σ and m (as mentioned above, we set $\rho = 0$). We optimize with a grid search over $\sigma \times m = [0.001, 10] \times [-1, 1]$ followed by the simplex algorithm.

For many date/firm/maturity triplets, we do not have a sufficient number of contract observations to fit the implied volatility curve (i.e. sometimes fewer than four). We therefore include strike/implied volatility data from the two neighboring maturities and dates in the estimation. The

parameters of SVI are obtained by minimizing squared fitting errors. We reweight the observations from the neighboring dates and maturities so that they carry the same amount of weight as the observations from the date and maturity of interest. Adding data in this way encourages smoothness in the estimates over time and across maturities but it does not induce a systematic upward or downward bias. We drop all date/firm/maturity triplets for which we have fewer than four total options with $k < 0$ or fewer than two options at the actual date/firm/maturity (i.e. ignoring the data from the neighboring dates and maturities).

When we estimate the parameters of the SVI model, we impose conditions that guarantee the absence of arbitrage. In particular, we assume that $b \leq \frac{4}{(1+|\rho)T}$, which when we assume $\rho = 0$, simplifies to $b \leq \frac{4}{T}$. We also assume that $\sigma > 0.0001$ in order to ensure that the estimation is well defined. Those conditions do not necessarily guarantee, though, that the integral determining the model-free implied volatility is convergent (the absence of arbitrage implies that a risk-neutral probability density exists – it does not guarantee that it has a finite variance). We therefore eliminate observations where the integral determining the model-free implied volatility fails to converge numerically. Specifically, we eliminate observations where the argument of the integral does not approach zero as the log strike rises above two standard deviations from the spot or falls more than five standard deviations below the strike (measured based on the at-the-money implied volatility).

A.1.3 Construction of V_n from the SVI fitted curve

After fitting the SVI curve for each date and maturity, we compute the integral in equation (17) numerically, over a range of strikes from -5 to +2 standard deviations away from the spot price.³ We then have V_n for every firm/date/maturity observation. The model-free implied volatilities are then interpolated (but not extrapolated) to construct V_n at maturities from 1–6 months for each firm/date pair.

A.1.4 Data description and diagnostics of SVI fitting

Our dataset consists of 2.3 million end-of-day prices for all American options on S&P 500 futures from the CME.

When more than one option (e.g. a call and a put) is available at any strike, we compute IV at that strike as the average of the observed IVs. We keep only IVs greater than zero, at maturities higher than 9 days and lower than 2 years, for a total of 1.9 million IVs. The number of available options has increased over time, as demonstrated by Figure A.2 (top panel), which plots the number of options available for V_n estimation in each year.

The maturity structure of observed options has also expanded over time, with options being introduced at higher maturities and for more intermediate maturities. Figure A.1 (top panel)

³In general this range of strikes is sufficient to calculate V_n . However, the model-free implied volatility technically involves an integral over the entire positive real line. Our calculation is thus literally a calculation of Andersen and Bondarenko’s (2007) corridor implied volatility. We use this fact also when calculating realized volatility.

reports the cross-sectional distribution of available maturities in each year to estimate the term structure of the model-free implied volatility. The average maturity of available options over our sample was 4 months, and was relatively stable. The maximum maturity observed ranged from 9 to 24 months and varied substantially over time.

Crucial to compute the model-free implied volatility is the availability of IVs at low strikes, since options with low strikes receive a high weight in the construction of V_n . The bottom panel of Figure A.1 reports the minimum observed strike year by year, in standard deviations below the spot price. In particular, for each day we computed the minimum available strike price, and the figure plots the average of these minimum strike price across all days in each year; this ensures that the number reported does not simply reflect outlier strikes that only appear for small parts of each year.

Figure A.1 shows that in the early part of our sample, we can typically observe options with strikes around 2 standard deviations below the spot price; this number increases to around 2.5 towards the end of the sample.

These figures show that while the number of options was significantly smaller at the beginning of the sample (1983), the maturities observed and the strikes observed did not change dramatically over time.

Figure A.3 shows an example of the SVI fitting procedure for a specific day in the early part of our sample (November 7th 1985). Each panel in the figure corresponds to a different maturity. On that day, we observe options at three different maturities, of approximately 1, 4, and 8 months. In each panel, the x's represent observed IVs at different values of log moneyness k . The line is the fitted SVI curve, that shows both the interpolation and the extrapolation obtained from the model.

Figure A.4 repeats the exercise in the later part of our sample (Nov. 1st 2006), where many more maturities and strikes are available.

Both figures show that the SVI model fits the observed variances extremely well. The bottom panel of Figure A.2 shows the average relative pricing error for the SVI model in absolute value. The graph shows that the typical pricing error for most of the sample is around 0.02, meaning that the SVI deviates from the observed IV by around 2% on average. Only in the very first years (up to 1985) pricing errors are larger, but still only around 10% of the observed IV.

Overall, the evidence in this section shows that our observed option sample since 1983 has been relatively stable along the main dimensions that matter for our analysis – maturity structure, strikes observed, and goodness of fit of the SVI model.

A.2 Lasso

Lasso is a regularization method for regressions that penalizes coefficients based on their absolute values. Specifically, the objective that is minimized under lasso is the sum of squared residuals plus a tuning parameter, which we denote λ , multiplied by the sum of the absolute values of the coefficients.

Lasso is not invariant to the scaling of the variables in the regression. We therefore rescale the variables as follows. rv , v_6 , and $slope$ are all translated into z-scores. The three macro variables (FFR , emp , and ip) are multiplied by constants so that their first differences have unit variances. We use that transformation because those three variables have approximate unit roots in our sample.

We examine two methods to select λ . The first is to use leave-one-out cross validation. We choose λ separately for the three volatility and three macro series. The cross-validation criterion implies setting $\lambda = 0.013$ for the volatility series and $\lambda = 0$ (i.e. no lasso) for the macro series. The results reported in the text use this choice of λ .

The second method is to choose the smallest (i.e. least restrictive) value of λ that causes the coefficients on all the lags of the macro variables in the rv equation to be zero, consistent with the benchmark specification. The motivation for this method is that it takes the restrictions that we impose on economic grounds and then essentially tries to impose similar restrictions on the other equations, for the sake of parity. In this case we find a value of λ of 0.055. The results with this value are not reported here but are consistent with our main findings. In this case we find slightly negative effects for uncertainty shocks, but they are still statistically significantly less negative than those for rv shocks, and the forecast error variance decompositions put an upper bound on the fraction driven by uncertainty shocks of 15 percent.

A.3 Bayesian estimation

This section examines an alternative estimation scheme for our main VAR. We study a Bayesian specification using the framework of Plagborg-Miçøeller (2017). In this specification, the model is studied in its vector moving average (VMA) form, i.e. directly in terms of IRFs. The priors, rather than being on coefficients in the VAR specification, which are sometimes difficult to interpret directly, are instead on the IRFs themselves. The drawback of this specification, though, is that one must have priors on all the IRFs, which becomes difficult when the number of variables in the model grows. We therefore consider a minimal specification in which the only included variables are rv , v_1 , and real activity (either employment or industrial production, but not both). Since we have the prior that the IRFs for rv and v_1 should be very similar (since v_1 is close to an expectation of rv), the three variables are actually entered in the model as v_1 , $rv - v_1$, and real activity.

The model is therefore written in the form

$$\begin{bmatrix} v_{1,t} \\ rv_t - v_{1,t} \\ \Delta y_t \end{bmatrix} = B(L) \varepsilon_t \quad (\text{A.3})$$

where $B(L)$ is the moving average lag polynomial. In what follows, though, we continue to report IRFs for rv_t and y_t directly. The IRF for rv_t is simply the sum of the IRFs for v_1 and $(rv_t - v_{1,t})$, while the IRF for y_t is the cumulative sum of the IRF for Δy_t .

The priors for the IRFs to the various shocks are plotted in figures A.17 and A.18 (along with

posteriors, which are discussed below). Figure A.17 has results for employment and A.18 for IP; the priors are the same in both cases. The solid thick lines represent prior modes and the gray areas are 90-percent probability bands. Some of the priors have zero standard deviations at certain points when they are normalized – shock volatilities are estimated as separate parameters.

The RV shock is assumed to increase RV on impact; about half the increase dies out after the first period and the remainder decays geometrically. The response of v_1 to the RV shock is constructed by treating v_1 as though it is the one-step-ahead forecast of rv_t . Empirically that will not be true, as it ignores variation in risk premia and the horizon of v_1 , but it is a simple prior approximation.

The news shock is identified separately from the RV shock as a shock that has a hump-shaped effect on RV. As with the RV shock, the prior for the effect on v_1 is equal to the one-step-ahead expectation of rv . The shocks therefore both increase v_1 on impact, and their difference is simply in how they affect rv on impact.

The priors for the effects of the news and RV shocks on activity both have zero mean and large standard deviations. The log volatilities of the innovations have Gaussian priors with means of 0.3 and standard deviations of 2.

Finally, to enforce a prior on smoothness, the IRFs are assumed to be correlated across time periods with a coefficient of 0.9.

Finally, the third shock represents a shock to real activity. We have a prior that it increases activity and assume that it has zero effect on rv and v_1 , with tight priors.

We implement the estimation using code available on Mikkel Plagborg-Mi \ddot{c} celler’s website. It uses a Hamiltonian Monte Carlo method to sample from the posterior distribution under the assumption of a Gaussian likelihood. Since the model is specified in logs, as in the main text, the distribution of the innovations is reasonably well approximated as being Normal.

A.3.1 Results

Figures A.17 and A.18 plot the posterior means and 95-percent probability intervals for the estimated responses to the three shocks. Across the various panels, the fact that the posterior probability intervals are substantially narrower than the priors shows that the data is informative for the IRFs. This is particularly important for showing that the news and RV shocks can be identified without the strong timing assumptions used in the benchmark analysis.

The results in figures A.17 and A.18 are highly similar to those in the benchmark. The RV shock produces a large initial increase in volatility followed by a smaller persistent component. It is associated with declines in both employment and industrial production that are both estimated with tight confidence bands.

The news shock again produces a hump-shaped response of realized volatility, and that response is significantly positive, indicating that it contains information about future realized volatility. That is, the news shock again represents an increase in uncertainty. The posteriors for the responses

of employment and industrial production show, though, that this uncertainty shock seems to have little or no effect on real activity in either direction. Furthermore, the posterior confidence bands on average 60-percent as wide as the priors, showing that the data is informative, providing affirmative evidence of zero effect for the uncertainty shocks.

As in the main text, it is also possible to examine the difference between the effects of the RV and news shocks. In results not reported here, we find that the effect of the RV shock is again significantly more negative than that of the news shock, with the difference significant at the 10-percent level across the 12-month horizon of the IRFs.

Overall, then, the Bayesian estimation here provides support for our main results in that we obtain qualitatively highly similar results while using a specification that differs from the main text in a number of important ways: the identifying assumptions, the functional form of the specification (a VMA(12) instead of a VAR(4)), and the set of state variables.

A.4 Accounting for within-month variation

This section describes how we estimate a daily VAR with realized and implied volatility in order to construct monthly shocks that are immune to time-aggregation concerns. The baseline analysis assumes that realized volatility is not affected by the contemporaneous shock to uncertainty at the monthly frequency. If volatility and uncertainty are driven by a simple daily model, though, that assumption can be violated. This section formalizes that intuition and shows how to use daily data to construct realized and implied volatility shocks that purge the effects of the within-month implied volatility shocks. The basic idea is to estimate a daily VAR and then shut down the shocks to uncertainty within each month, which are the basic source of the bias.

Empirically, returns, r_t , are the daily log return on the S&P 500 (from the close on day $t-1$ to the close on day t). Implied volatility, σ_t^2 , is measured by the VIX (the VXO in the period where the VIX is not available) calculated at the end of the day. These are consistent with the empirical choices in the paper. In what follows, the index t denotes days while s denotes months.

Suppose the true dynamics of volatility and returns at the daily level are

$$r_t^2 = \sigma_{t-1}^2 \varepsilon_t^2 \tag{A.4}$$

$$\sigma_t^2 = (1 - \rho - k) \bar{\sigma}^2 + \rho \sigma_{t-1}^2 + k r_t^2 + v_t \tag{A.5}$$

The model for volatility is a standard GARCH model (Bollerslev (1986)) plus an independent shock, v_t . We set $E\varepsilon_t^2 = 1$ as a normalization so that σ_t^2 represents the conditional variance of returns (with the conditional mean of r set to zero for simplicity).

The conditional means of r^2 and σ^2 are

$$E_{t-1} r_t^2 = \sigma_{t-1}^2 \tag{A.6}$$

$$E_{t-1} \sigma_t^2 = (1 - \rho - k) \bar{\sigma}^2 + (\rho + k) \sigma_{t-1}^2 \tag{A.7}$$

and the statistical innovations are

$$r_t^2 - E_{t-1}r_t^2 = \sigma_{t-1}^2 (\varepsilon_t^2 - 1) \quad (\text{A.8})$$

$$\sigma_t^2 - E_{t-1}\sigma_{t-1}^2 = k\sigma_{t-1}^2 (\varepsilon_t^2 - 1) + v_t \quad (\text{A.9})$$

One can immediately see the innovations satisfy the timing restriction from the main text and take a lower-triangular form. This model is simply a daily version of the lower-frequency specification used in Bloom (2009), Bloom et al. (2017), and Basu and Bundick (2017), among others.

The system has a VAR representation,

$$\begin{bmatrix} r_t^2 \\ \sigma_t^2 \end{bmatrix} = C + \Gamma \begin{bmatrix} r_{t-1}^2 \\ \sigma_{t-1}^2 \end{bmatrix} + \Sigma \mu_t \quad (\text{A.10})$$

where

$$C = \begin{bmatrix} 0 \\ (1 - \rho) \bar{\sigma}^2 \end{bmatrix} \quad (\text{A.11})$$

$$\Gamma = \begin{bmatrix} 0 & 1 \\ 0 & \rho + k \end{bmatrix} \quad (\text{A.12})$$

$$\Sigma = \begin{bmatrix} \text{var}(\sigma_{t-1}^2 (\varepsilon_t^2 - 1)) & 0 \\ k^2 \text{var}(\sigma_{t-1}^2 (\varepsilon_t^2 - 1)) & 1 \end{bmatrix} \quad (\text{A.13})$$

While the μ_t will be heteroskedastic and non-Gaussian, this is still a linear system that can be consistently estimated through OLS. The structural shocks, μ_t , are identified by premultiplying the reduced-form residuals by Σ^{-1} .

A.4.1 Monthly levels purged of contamination

We construct an alternative version of realized volatility within each month that sets the second element of μ_t – corresponding to the daily IV shock – to zero for all days t in the month. Specifically, define $\mu_t^{iv-purged}$ to be equal to μ_t but with the second element – the shock to *iv* – set to zero. We then construct $\left(r_t^{iv-purged}\right)^2$ by iterating on (11) in month s , replacing μ_t by $\mu_t^{iv-purged}$ for $t \in s$. $\left(r_t^{iv-purged}\right)^2$ then represents the history of realized volatility in month s in the absence of the implied volatility shocks. That is, it is what would have happened if realized volatility was purely driven by the daily RV shocks. $\left(r_t^{iv-purged}\right)^2$ is therefore constructed explicitly to be independent of any shocks to implied volatility – v_t – during month s , so that the identifying assumptions hold at the monthly frequency by construction.

In the main analysis, we construct monthly realized volatility as the sum of daily r_t^2 . That is,

true realized volatility in month s is

$$RV_s \equiv \sum_{t \in s} r_t^2 \quad (\text{A.14})$$

Similarly, one can construct a counterfactual variable $RV_s^{iv-purged}$ that represents the part of RV_s driven only by the uncertainty shocks during month s , and not by the IV shocks during the month,⁴

$$RV_s^{iv-purged} = \sum_{t \in s} \left(r_t^{iv-purged} \right)^2 \quad (\text{A.15})$$


$RV_s^{iv-purged}$ is then, by construction, purged of influence of the within-month uncertainty shocks. This variable can then be placed directly into our main VAR specification. Figure A.19 replicates the main VAR analysis in that case and shows that the results are similar.

We have also confirmed through a simulation study that this method successfully identifies structural shocks from monthly data when the true joint process for implied and realized volatility is determined at the daily frequency. The intuition is simple: $RV_s^{iv-purged}$ is, by construction, uncorrelated with uncertainty shocks during month s . So when it is ordered first, the VAR identified shock must be orthogonal to the uncertainty shock, v . The identified uncertainty shock, ordered second in the VAR, then accounts for the part of the variation in implied volatility that does not come from ARCH effects, which is v . Depending on how the daily data is aggregated, nonlinearity can arise, making correlations between the VAR-implied shocks and the true structural shocks less than 1, but the purged RV shock will always avoid any correlation with the IV shock, again by construction. The simulation study is available on request.

⁴More specifically, creating $RV_s^{iv-purged}$ for a particular value of s involves the following steps. Taking the VAR along with the estimates of μ_t as given:


1. Define $\hat{\mu}_t^s$ to be equal to μ_t except that for days $t \in \text{month } s$, set the first element of μ_t to zero.
2. Iterate on the VAR to construct a counterfactual history of $\left(r_t^{iv-purged} \right)^2$. This will differ from the true value of r_t^2 beginning in month s .
3. Sum the counterfactual $\left(r_t^{iv-purged} \right)^2$ within month s .

Figure A.1: Maturities and strikes in the CME dataset



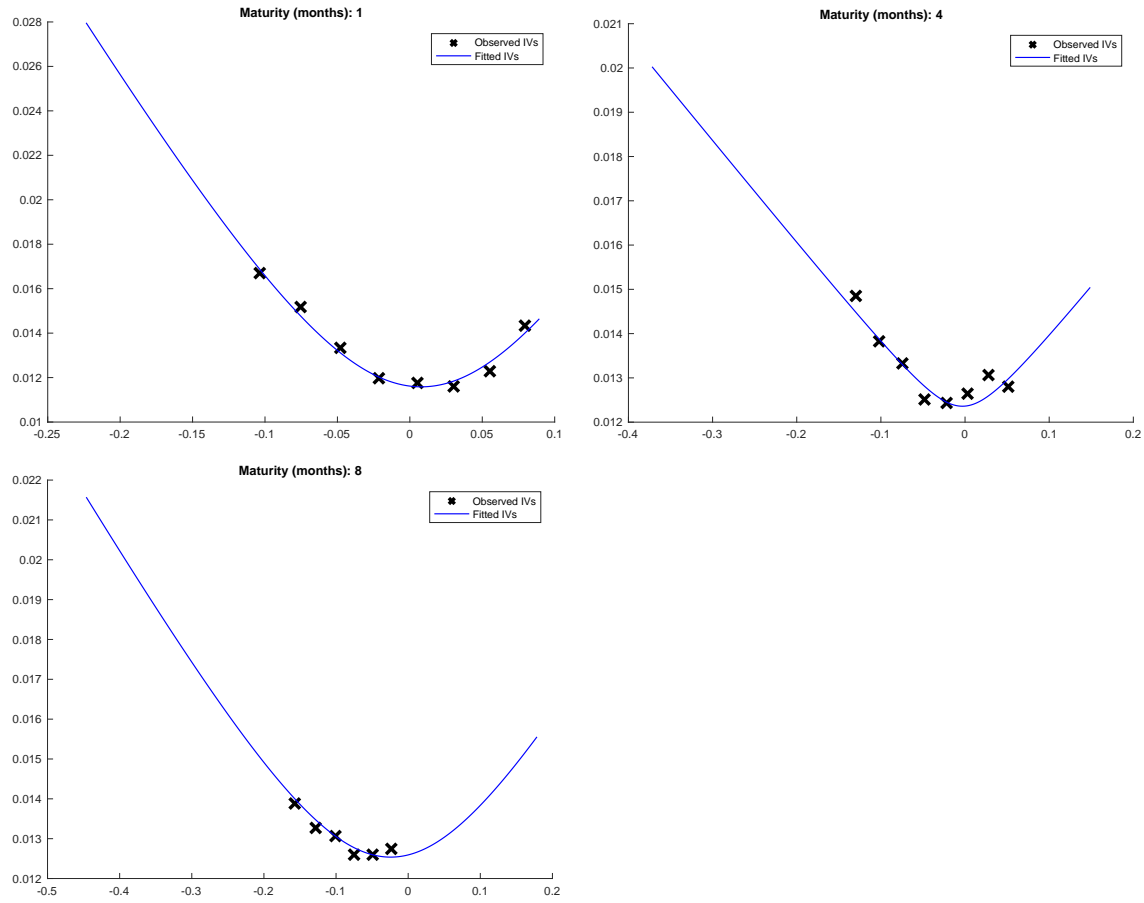
Note: The top panel reports the distribution of maturities of options used to compute implied volatility in each year, in months. The bottom panel reports the average minimum strike in each year, in standard deviations below the forward price. The number is obtained by computing the minimum observed strike in each date and at each maturity (in standard deviations below the forward price), and then averaging it within each year to minimize the effect of outliers.

Figure A.2: Number of options to construct implied volatility and pricing errors




Note: The top panel reports the number of options used to compute implied volatility in each year, in thousands. The bottom panel reports the average absolute value of the pricing error of the SVI fitted line relative to the observed implied variances, in proportional terms (i.e. 0.02 means absolute value of the pricing error is 2% of the observed implied variance).

Figure A.3: SVI fit: 11/7/1985




Note: Fitted implied variance curve on 11/7/1987, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve.

Figure A.4: SVI fit: 11/1/2006



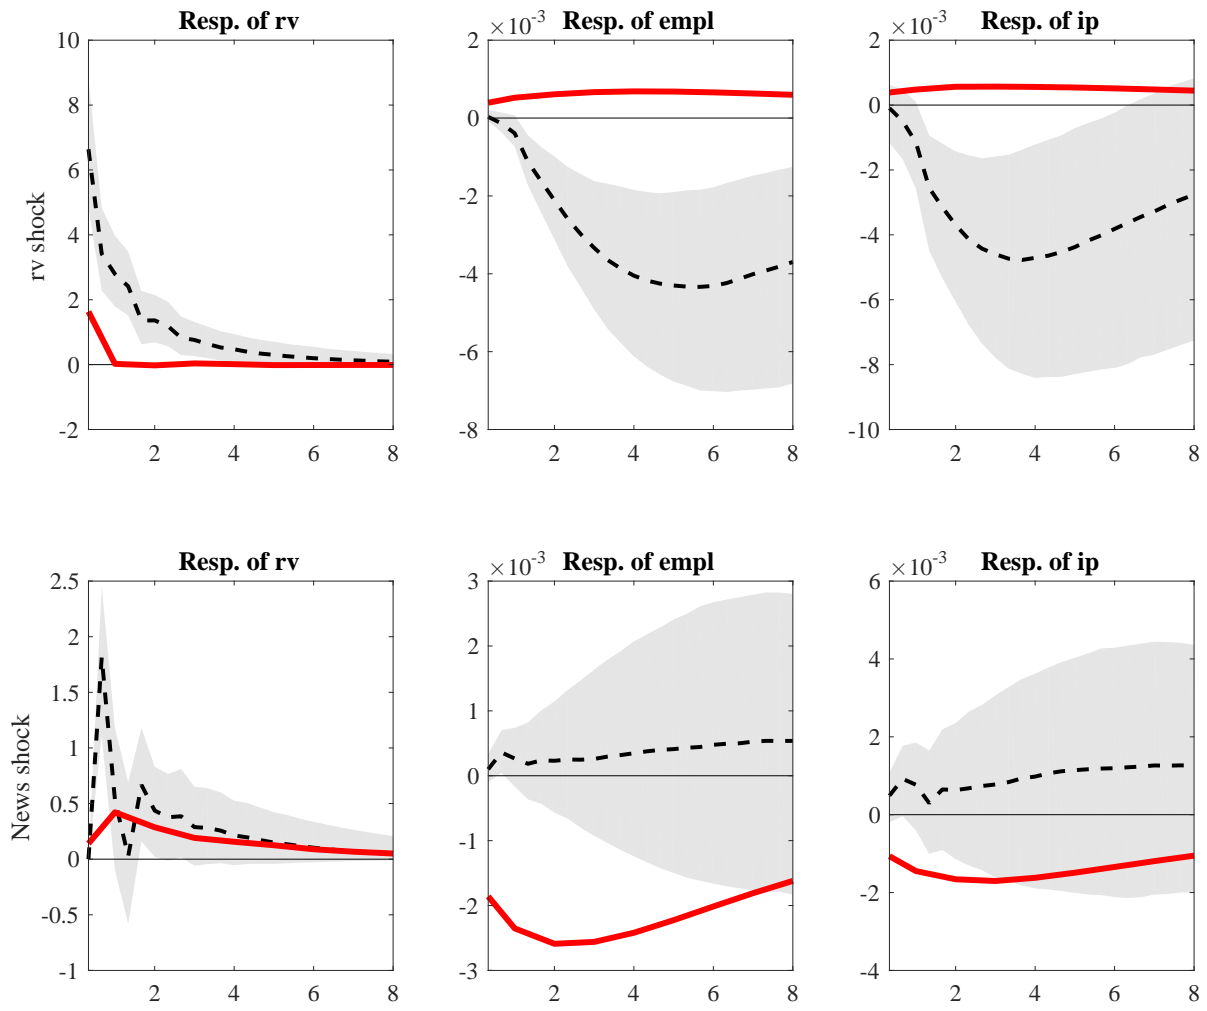
Note: Fitted implied variance curve on 11/1/2006, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve. On 11/1/2006 also a maturity of 5 months was available (not plotted for reasons of space).

Figure A.5: Impulse response functions from VAR with v_1 but not rv



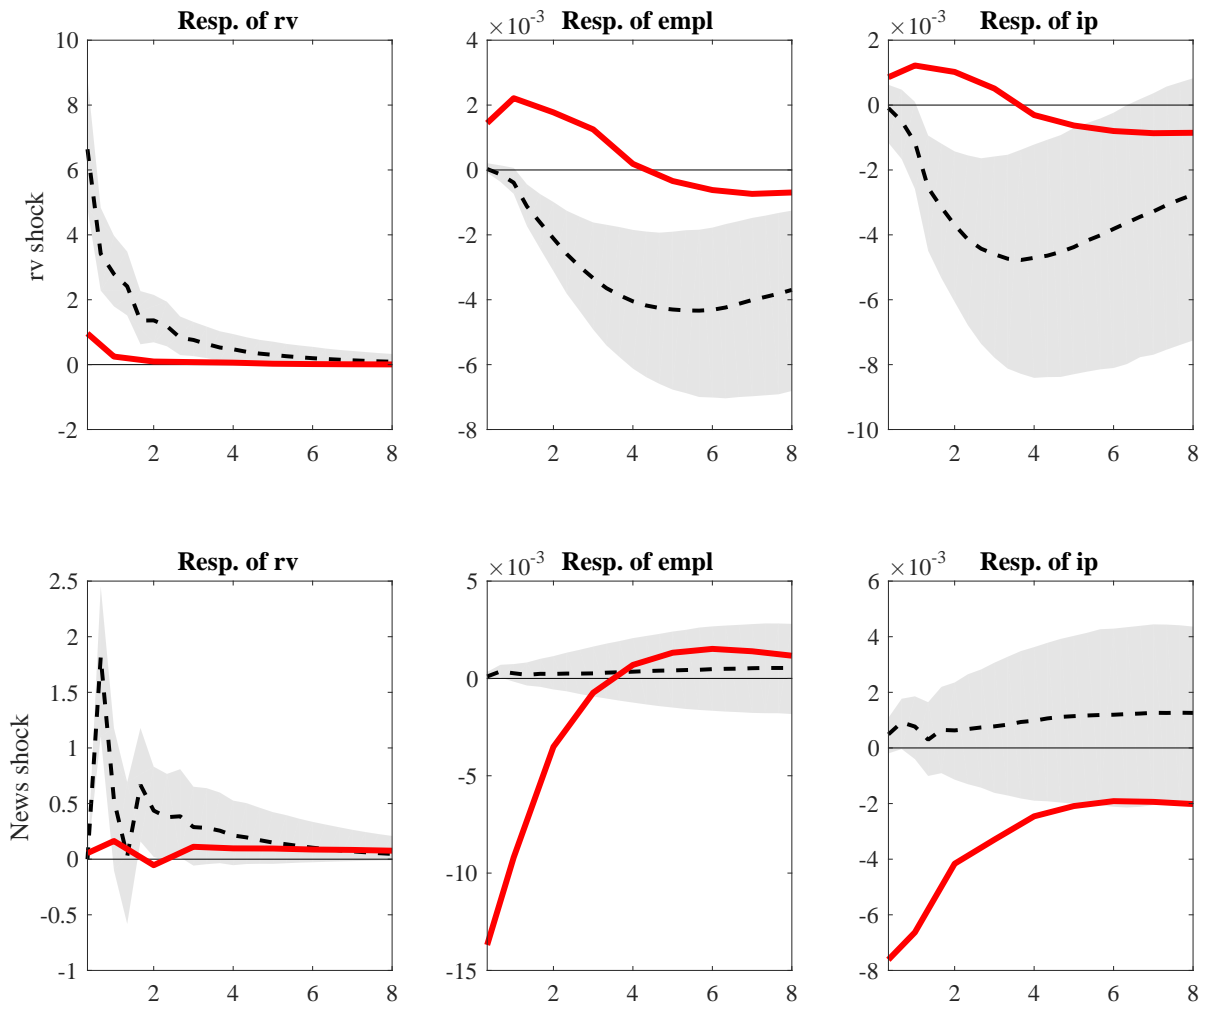
Note: The figure shows responses of volatility (measured by v_1), log employment, and log industrial production to a reduced-form shock to v_1 in a VAR with v_1 , the Fed funds rate, log employment, and log industrial production with 68% and 90% confidence intervals. Sample period 1986-2014.

Figure A.6: Impulse response functions in the Basu and Bundick (2017) model



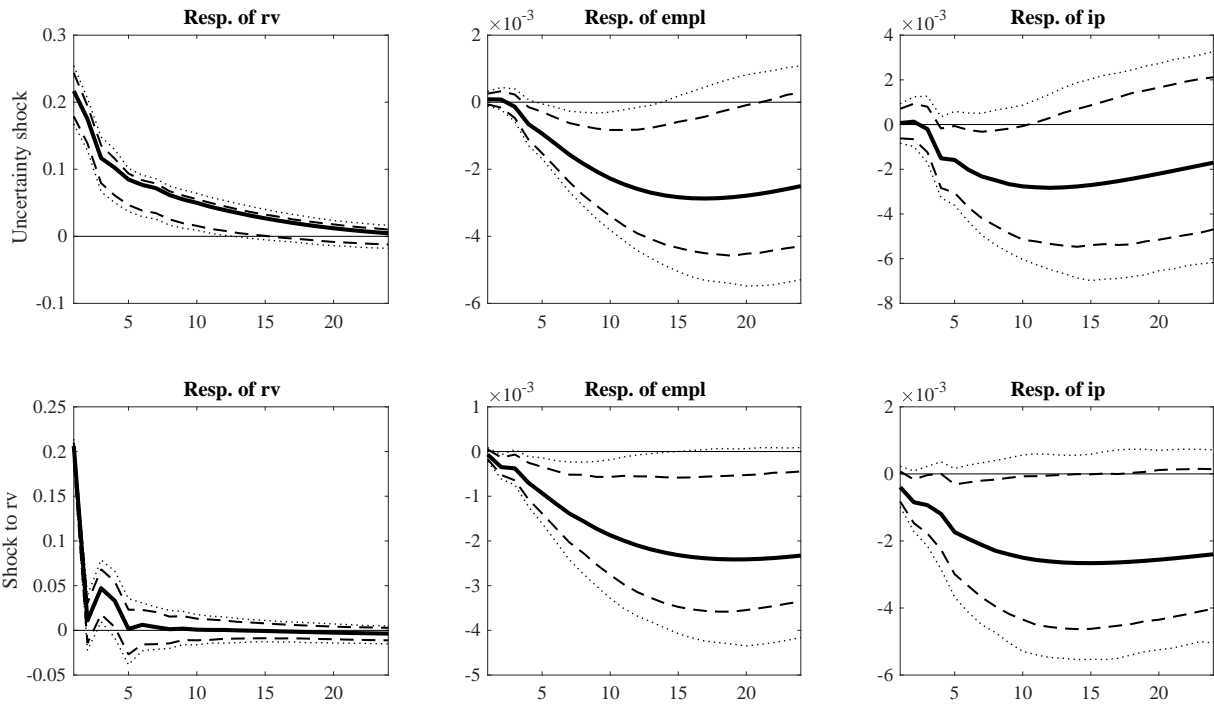
Note: Same as figure 2, but on data simulated from Basu and Bundick (2017).

Figure A.7: Impulse response functions in the Bloom et al. (2017) model




Note: Same as figure 2, but on data simulated from Bloom et al. (2017).

Figure A.8: Impulse response functions from VAR ordering uncertainty first and rv second




Note: See figure 2. Unlike in the baseline identification, the identified uncertainty shock is not orthogonalized with respect to rv . The rv shock in this case is the remaining part of reduced-form innovation to rv that is not spanned by the uncertainty shock.


Figure A.9: Robustness: response of employment to rv and uncertainty shocks across specifications




(a) Using the CBOE VIX instead of v_1



(b) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)




(c) Adding the S&P 500 level as first shock




(d) Using RV and V_1 in levels, not logs

Note: Response of employment to RV shocks (left panels) and uncertainty (middle panels) with the difference in the right panel and different model specifications in each row. Row (a) uses the VIX instead of v_1 . Row (b) estimates the VAR in the subsample 1988-2006, which excludes both RV peaks (1987 crash and financial crisis). Row (c) orthogonalizes both the rv and the uncertainty shocks with respect to the reduced-form innovation in the S&P 500, as in Bloom (2009). Row (d) uses RV and V_1 in levels, not logs.


Figure A.10: Robustness: response of IP to rv and uncertainty shocks across specifications




(a) Using the CBOE VIX instead of v_1



(b) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)




(c) Adding the S&P 500 level as first shock




(d) Using RV and V_1 in levels, not logs

Note: See figure A.9. In this case the responses of IP are reported instead of employment.


Figure A.11: Robustness: VAR with quarterly data from Basu and Bundick (2016), using v_1



(a) Investment




(b) Consumption




(c) GDP

Note: See figure 2. Here we use the quarterly data from Basu and Bundick (2016) as the macro time series.


Figure A.12: Robustness: VAR with quarterly data from Basu and Bundick (2016), using v_6



(a) Investment



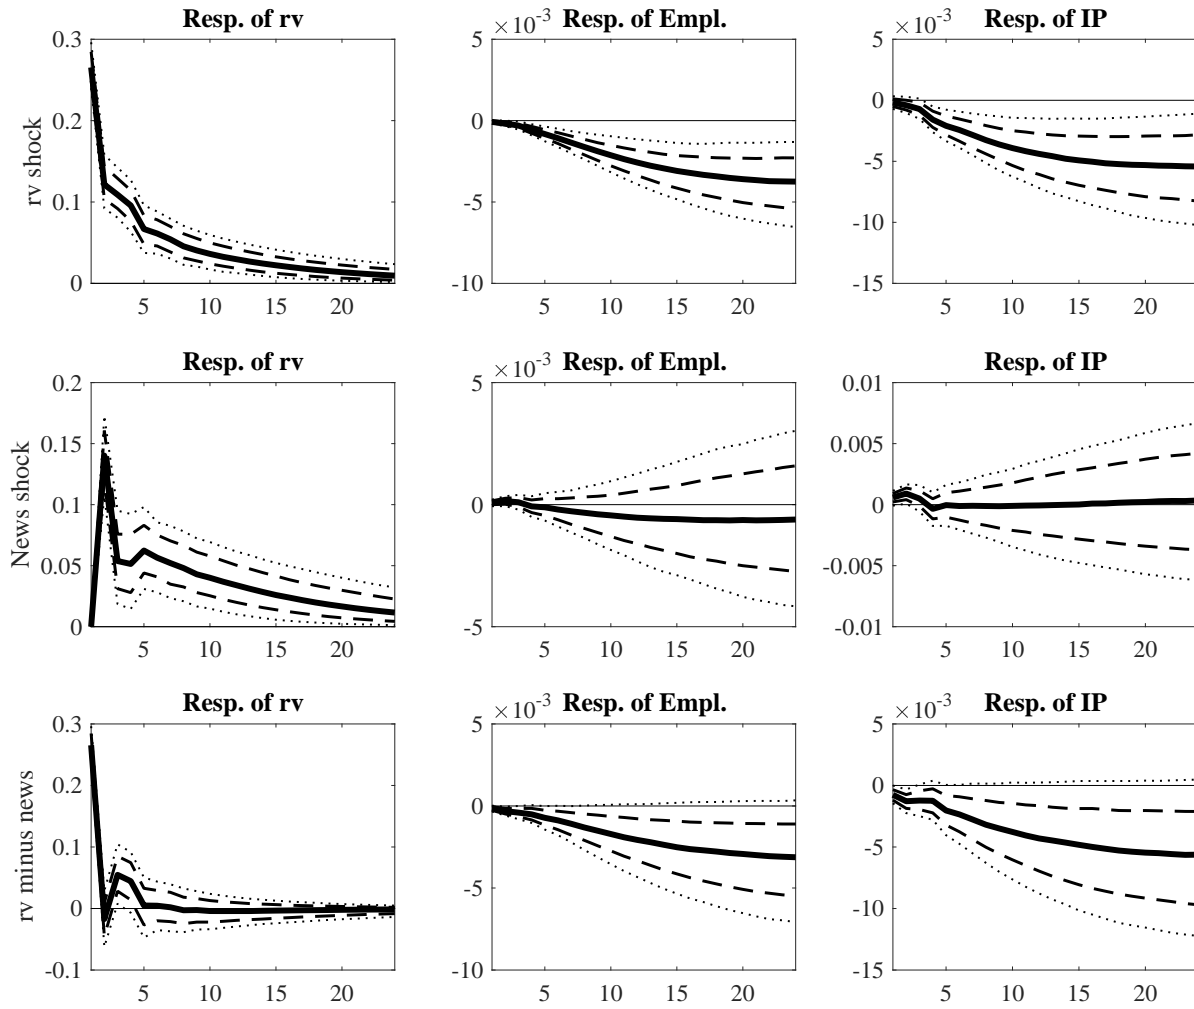
(b) Consumption



(c) GDP

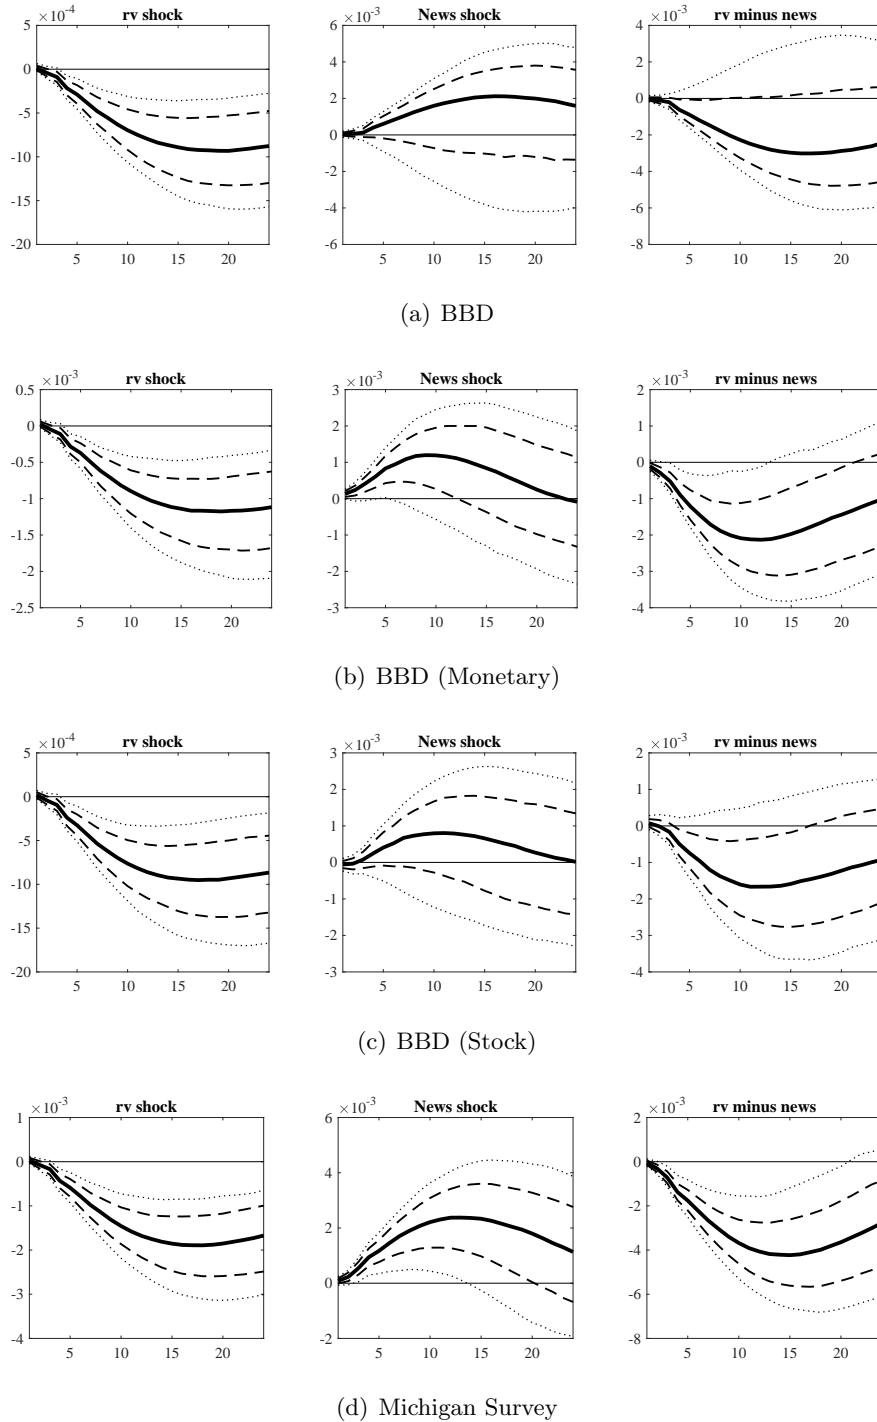
Note: See figure 2. Here we use the quarterly data from Basu and Bundick (2016) as the macro time series. We use v_6 instead of v_1 .

Figure A.13: Impulse response functions from FAVAR




Note: See figure 2. IRFs from a FAVAR in which the state variables are rv , v_1 , federal funds rate, and the first three principal components from the data set of Ludvigson and Ng (2007). Employment and industrial production are then obtained from regressions onto their own lags and those of the state variables.

Figure A.14: Response of employment to rv and alternative measures of uncertainty shocks




Note: Response of employment to RV shocks (left panels) and news shocks (middle panels) with the difference in the right panel and different measures of uncertainty in the VAR instead of v_1 . Row (a) uses the Baker, Bloom and Davis (2015) measure. Rows (b) and (c) use two subcomponents of the BBD measure, capturing monetary and stock market uncertainty respectively. Row (d) uses a measure from the Michigan survey used by Leduc and Liu (2016).


Figure A.15: Response of IP to rv and alternative measures of uncertainty shocks




(a) BBD



(b) BBD (Monetary)




(c) BBD (Stock)




(d) Michigan Survey

Note: See figure A.14.


Figure A.16: Response of rv to rv and alternative measures of uncertainty shocks




(a) BBD



(b) BBD (Monetary)



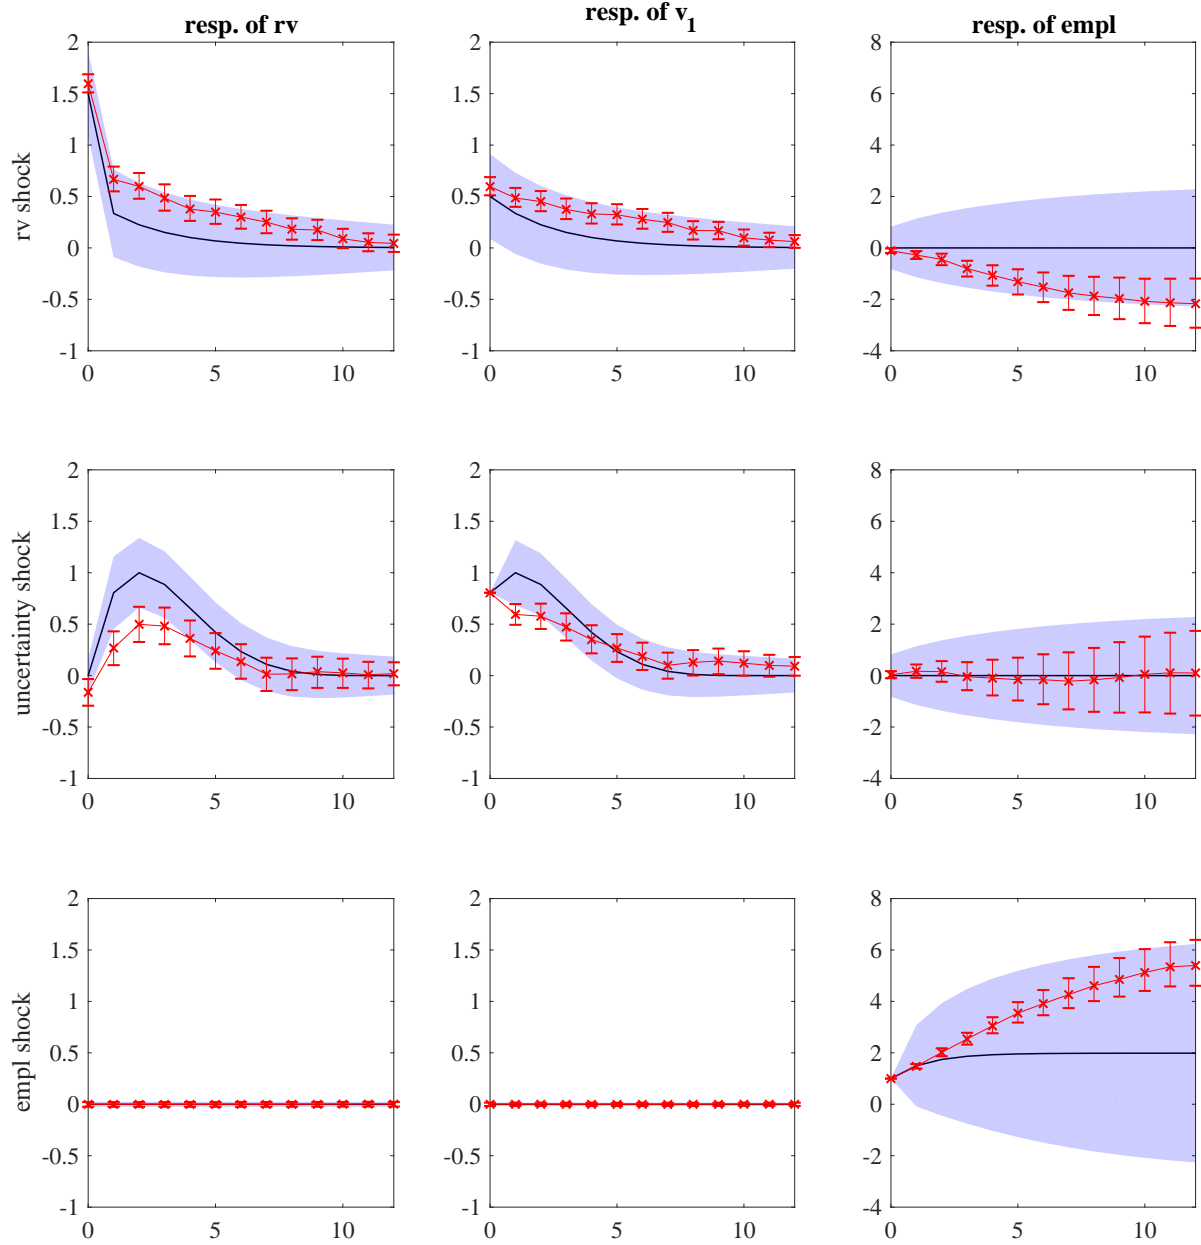
(c) BBD (Stock)



(d) Michigan Survey

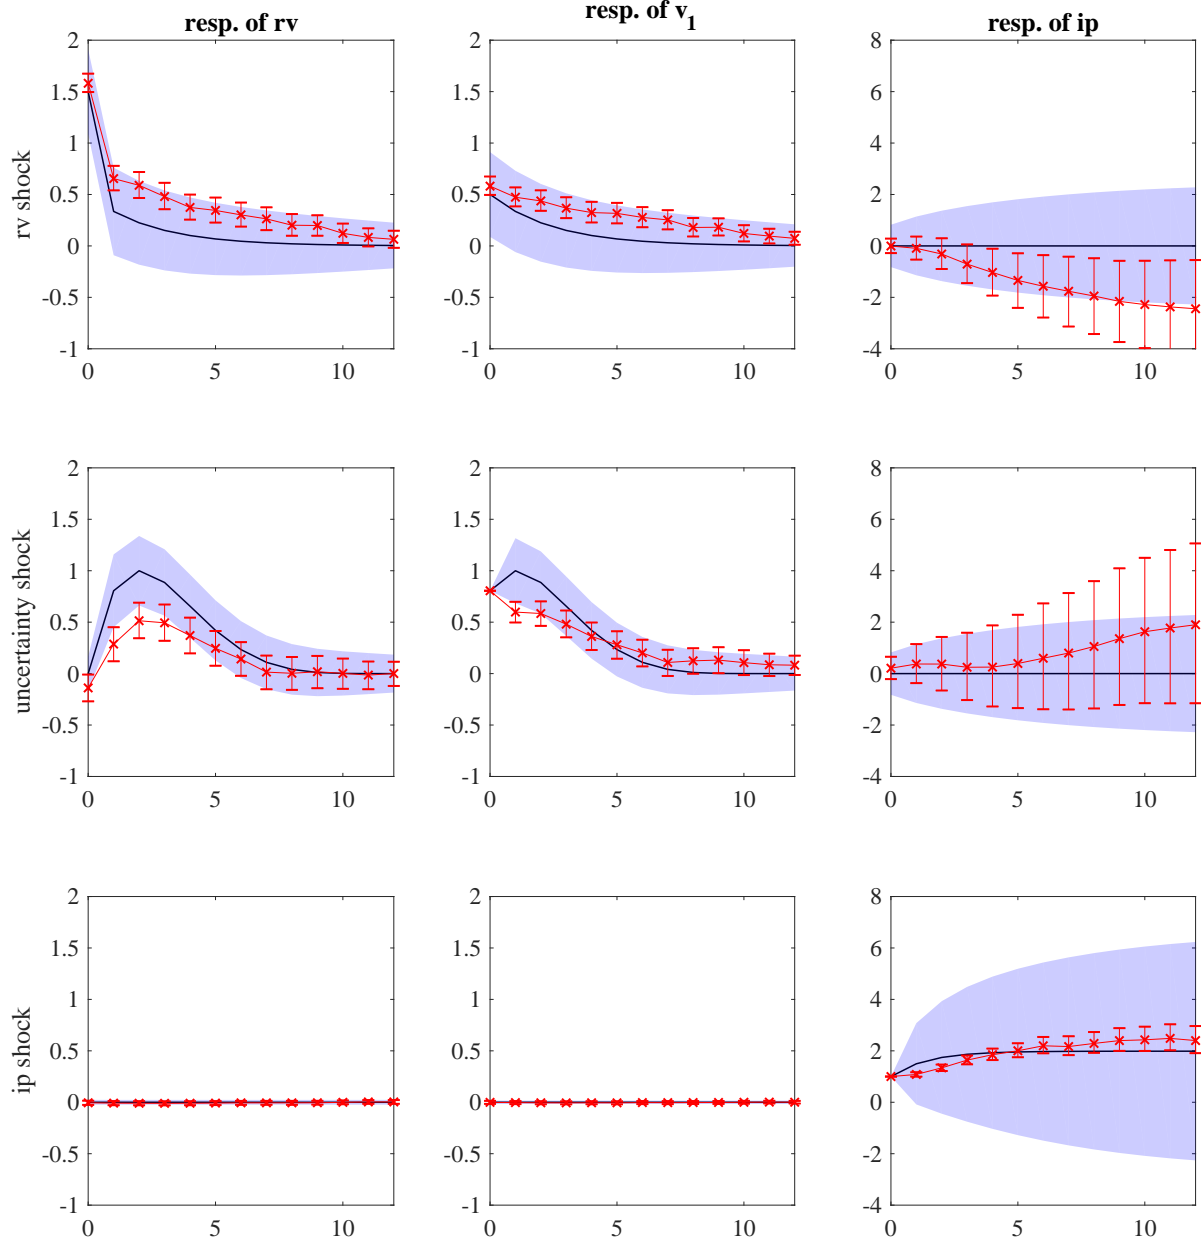
Note: See figure A.14.

Figure A.17: Priors and posteriors from Bayesian estimation with employment



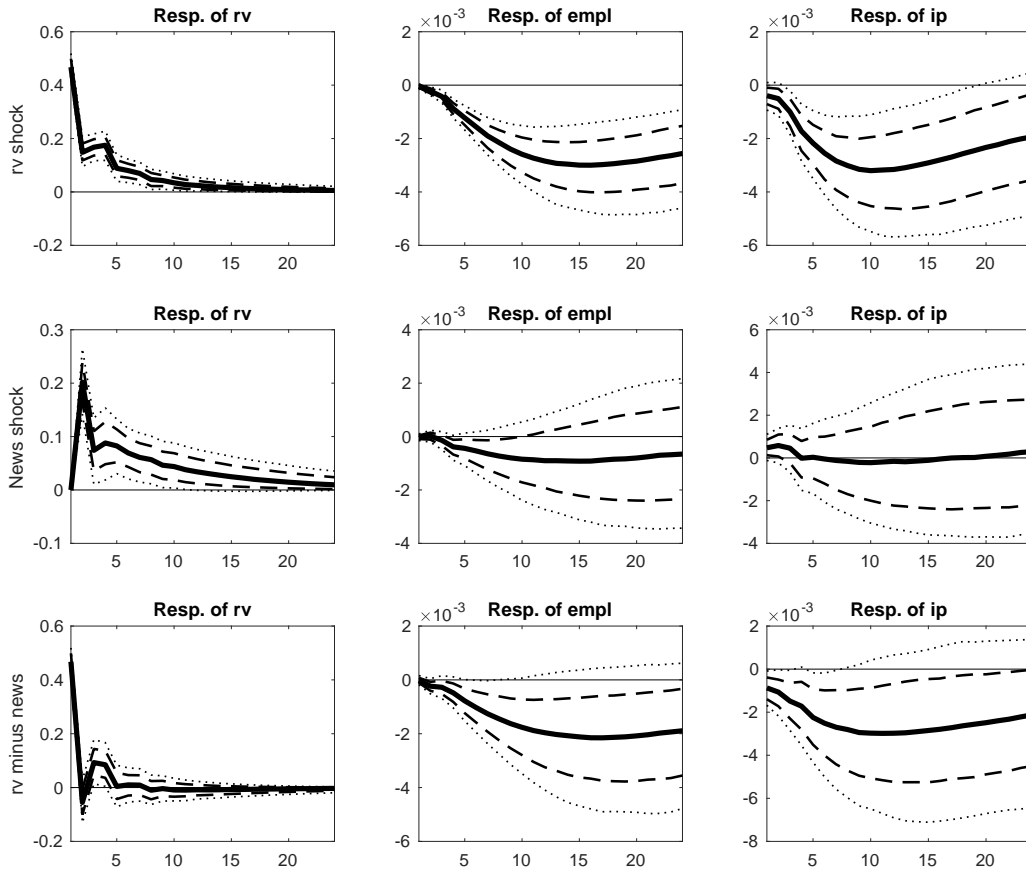
Note: Priors and posteriors from estimates of the Bayesian model with rv , v_1 , and employment. Each panel reports an estimated impulse response to one of the three identified shocks. The black lines and shaded areas represent priors. The red boxes are the posterior mode and the whiskers are 90-percent posterior probability intervals.

Figure A.18: Priors and posteriors from Bayesian estimation with IP




Note: See figure A.17. Here employment is replaced by IP.

Figure A.19: VAR results purged of bias from within-month implied volatility shocks




Note: See figure 2. The time series of RV_t has been replaced with $RV_t^{iv-purged}$, the version purged of the within-month implied volatility shocks (as described in Appendix A.4).


Figure A.20: Response of employment using alternatives to rv



(a) Using downside volatility instead of rv




(b) Using upside volatility instead of rv




(c) Using jumps instead of rv

Note: Response of employment to first-moment shocks (left panels) and news shocks (middle panels) with the difference in the right panel and different model specifications in each row. Row (a) uses downside volatility instead of rv . Row (b) uses upside volatility instead of rv . Row (c) uses a measure of realized jumps instead of rv .


Figure A.21: Response of IP using alternatives to rv



(a) Using downside volatility instead of rv



(b) Using upside volatility instead of rv



(c) Using jumps instead of rv

Note: See figure A.20.

Table A.1: Relationship between employment, industrial production and volatility (using $RV^{iv-purged}$)

	Employment		Industrial Production	
v_1	-0.10*** (0.04)	0.04 (0.05)	-0.09* (0.06)	0.13 (0.12)
rv		-0.15*** (0.05)		-0.25** (0.12)

Note: Same as table 1, but here we use $RV^{iv-purged}$ instead of RV , as described in Appendix A.4.

Table A.2: Predictability of rv with and without v_1 as predictor

Predictors	(1)	(2)	(3)	(4)
rv		0.13 (0.09)		0.09 (0.09)
v_1		0.58*** (0.15)		0.60*** (0.15)
FFR	-0.01 (0.02)	0.01 (0.01)		
Δemp	-34.96** (17.37)	-4.98 (10.87)		
Δip	-3.62 (5.45)	-4.81 (3.57)		
PC 1			-0.029** (0.013)	-0.010 (0.009)
PC 2			-0.014 (0.018)	-0.017 (0.012)
PC 3			-0.014 (0.010)	-0.011 (0.007)
$R_{S\&P}$			-1.56*** (0.43)	0.12 (0.41)
Def			0.21 (0.14)	0.07 (0.09)
N	377	377	377	377
Adj. R^2	0.12	0.45	0.20	0.46

Note: Regressions of 6-month realized volatility on lagged rv , option-implied volatility, and various macroeconomic variables, with Hansen-Hodrick standard errors using a 6-month lag window. PC 1 – PC3 are the first three principal components from a large set of macroeconomic time series. $R_{S\&P}$ is the return on the S&P 500. Def is the default spread, the gap between yields on Aaa and Baa bonds.

Table A.3: Predictability of 6-month-ahead v_1

Predictors	(1)	(2)	(3)	(5)	(6)
rv_t	0.02 (0.09)	0.01 (0.09)	0.03 (0.09)	0.01 (0.09)	0.01 (0.08)
$v_{1,t}$	0.57*** (0.13)	0.49*** (0.15)	0.51** (0.26)	0.54*** (0.13)	0.60*** (0.13)
$v_{6,t}$			0.06 (0.27)		
rv_{t-1}		0.10* (0.06)			
FFR				0.01 (0.01)	
$\Delta empl$				-10.60 (9.90)	
Δip				0.40 (3.31)	
PC1					-0.010 (0.008)
PC2					-0.010 (0.010)
PC3					-0.008 (0.006)
$R_{S\&P}$					0.78** (0.38)
Default spread					-0.05 (0.08)
Adj. R^2	0.35	0.36	0.35	0.36	0.37

Note: Regressions of 6-month-ahead v_1 on lagged rv , option-implied volatility, and various macroeconomic variables, with Hansen–Hodrick (1980) standard errors using a 6-month lag window. PC1–3 are principal components from the data set used in Ludvigson and Ng (2007). The default spread is the difference in yields on Baa and Aaa bonds. The sample is 1983–2014.