Uncertainty shocks as second-moment news shocks
Online Appendix

David Berger, Ian Dew-Becker, and Stefano Giglio

February 6, 2019

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,

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

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

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. $Var_t^Q [s_{t+n}]$ 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 $Var_t^Q [s_{t+n}]$ 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

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

[2] See 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)

[3] In 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-Mï¿œller (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

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

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-Mï¿œller'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 constuct 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). Note that this is a nonlinear model with a skewed unconditional distribution for returns, as in the main text.

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

$$r_t^2 - E_{t-1}r_t^2 = \sigma_{t-1}^2\left(\varepsilon_t^2 - 1\right) \tag{A.8}$$
$$\sigma_t^2 - E_{t-1}\sigma_{t-1}^2 = k\sigma_{t-1}^2\left(\varepsilon_t^2 - 1\right) + v_t \tag{A.9}$$

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,

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

where

$$C = \begin{bmatrix} 0 \\ (1-\rho)\,\bar{\sigma}^2 \end{bmatrix} \tag{A.11}$$
$$\Gamma = \begin{bmatrix} 0 & 1 \\ 0 & \rho+k \end{bmatrix} \tag{A.12}$$
$$\Sigma = \begin{bmatrix} var\left(\sigma_{t-1}^2\left(\varepsilon_t^2-1\right)\right) & 0 \\ k^2 var\left(\sigma_{t-1}^2\left(\varepsilon_t^2-1\right)\right) & 1 \end{bmatrix} \tag{A.13}$$

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. This result relies on the daily VAR being correctly specified. That is, it must actually recover the true rv and iv shocks. As in the main analysis, what is required is a symmetry assumption in that the shock v_t is uncorrelated with the shock to r_t^2. That is, the basic assumptions from the monthly model are still applied here, but now they must hold at the daily level, and can in fact be violated at the monthly level.

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

$$RV_s \equiv \sum_{t \in s} r_t^2 \tag{A.14}$$

Similarly, one can constuct 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]

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

$RV_s^{iv-purged}$ is then, by construction, purged of influence of the within-month uncertainty shocks. That fact can trivially be confirmed in a numerical simulation of the above model for volatility dynamics. $RV_s^{iv-purged}$ 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 identifes 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.

Empirically, for the daily VAR we use a heterogeneous autoregressive (Corsi (2004)) specification. Specifically, the daily VAR includes iv and rv. For the two regressions, the included lags are the previous day's observation, the average over the previous week, and the average over the previous month of rv and iv. The purged series are then constructed as defined as above and inserted

[4]More 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 $t \in$ 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.

into the baseline regressions and VAR, as discussed in the main text. The HAR specification is a parsimonious method of allowing a large number of lags to be included in the daily VAR while only requiring a few coefficients to be estimated. It can accommodate a wide range of dynamic processes for volatility.

References

Andersen, Torben G. and Oleg Bondarenko, "Construction and Interpretation of Model-Free Implied Volatility," Working Paper, National Bureau of Economic Research September 2007.

Bakshi, Gurdip, Nikunj Kapadia, 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.

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.

Carr, Peter and Liuren Wu, "Variance Risk Premiums," *Review of Financial Studies*, 2009, *22(3)*, 1311–1341.

Corsi, Fulvio, "A simple long memory model of realized volatility," 2004. Workinf 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.

Homescu, Cristian, "Implied Volatility Surface: Construction Methodologies and Characteristics," 2011. Working paper.

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.

Lee, Roger W., "The Moment Formula for Implied Volatility at Extreme Strikes," *Mathematical Finance*, 2004, *14* (3), 469–480.

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.

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

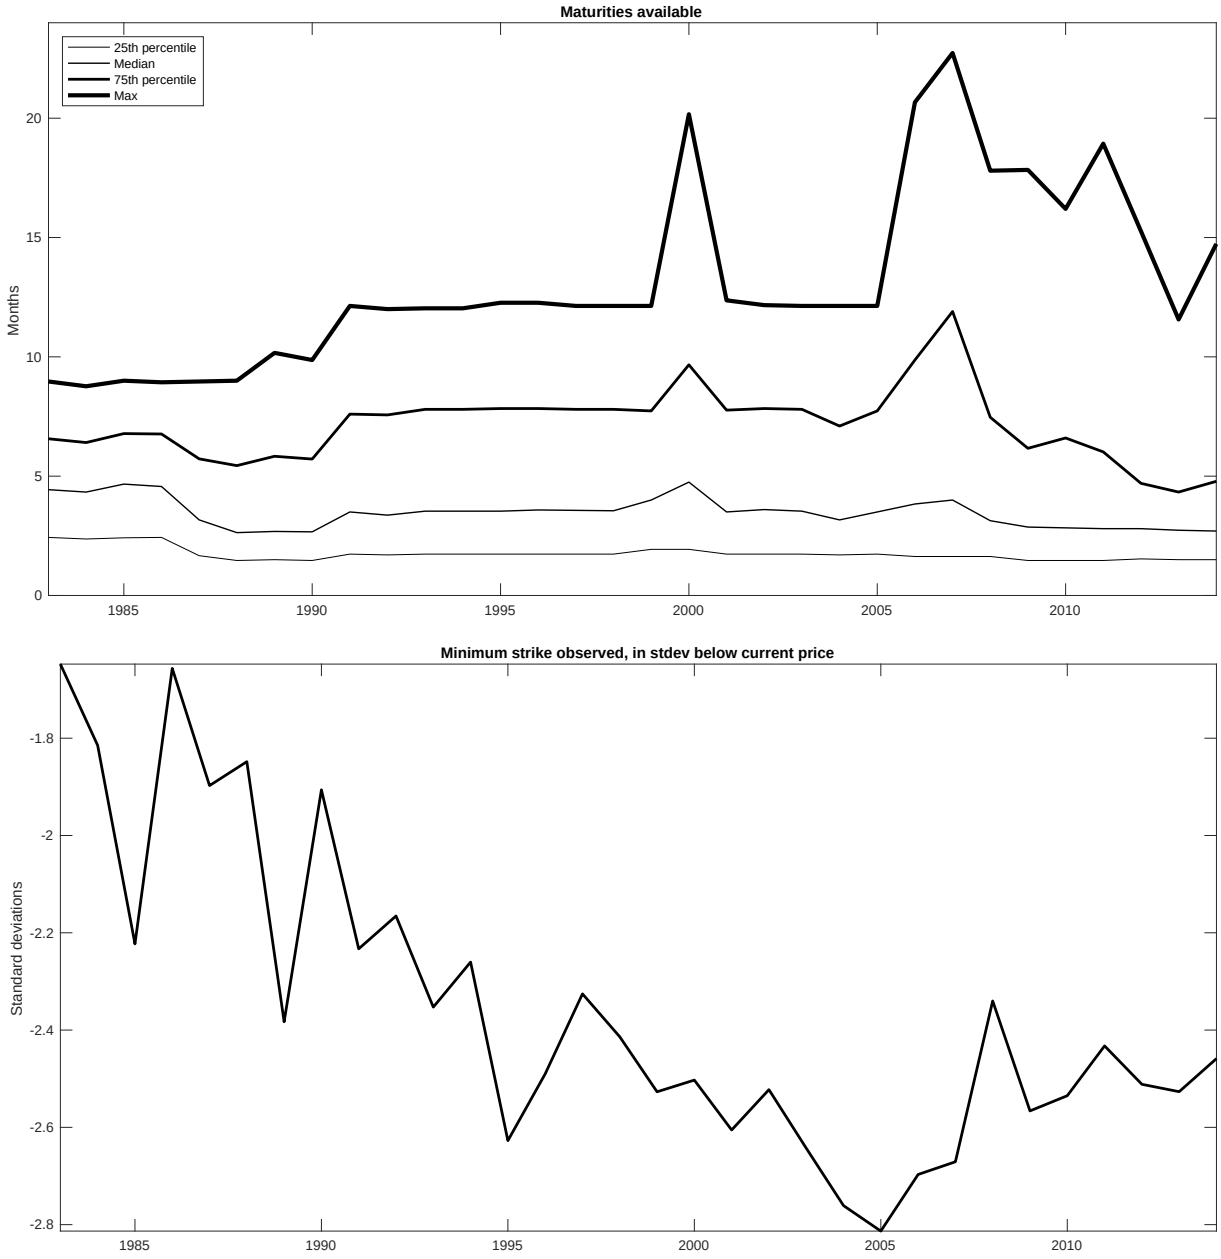

Note: The top panel reports the distribution of maturities of options used to compute implied volatility in each year, in months. The bottom panel reports the average minimum strike in each year, in standard deviations below the forward price. The number is obtained by computing the minimum observed 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

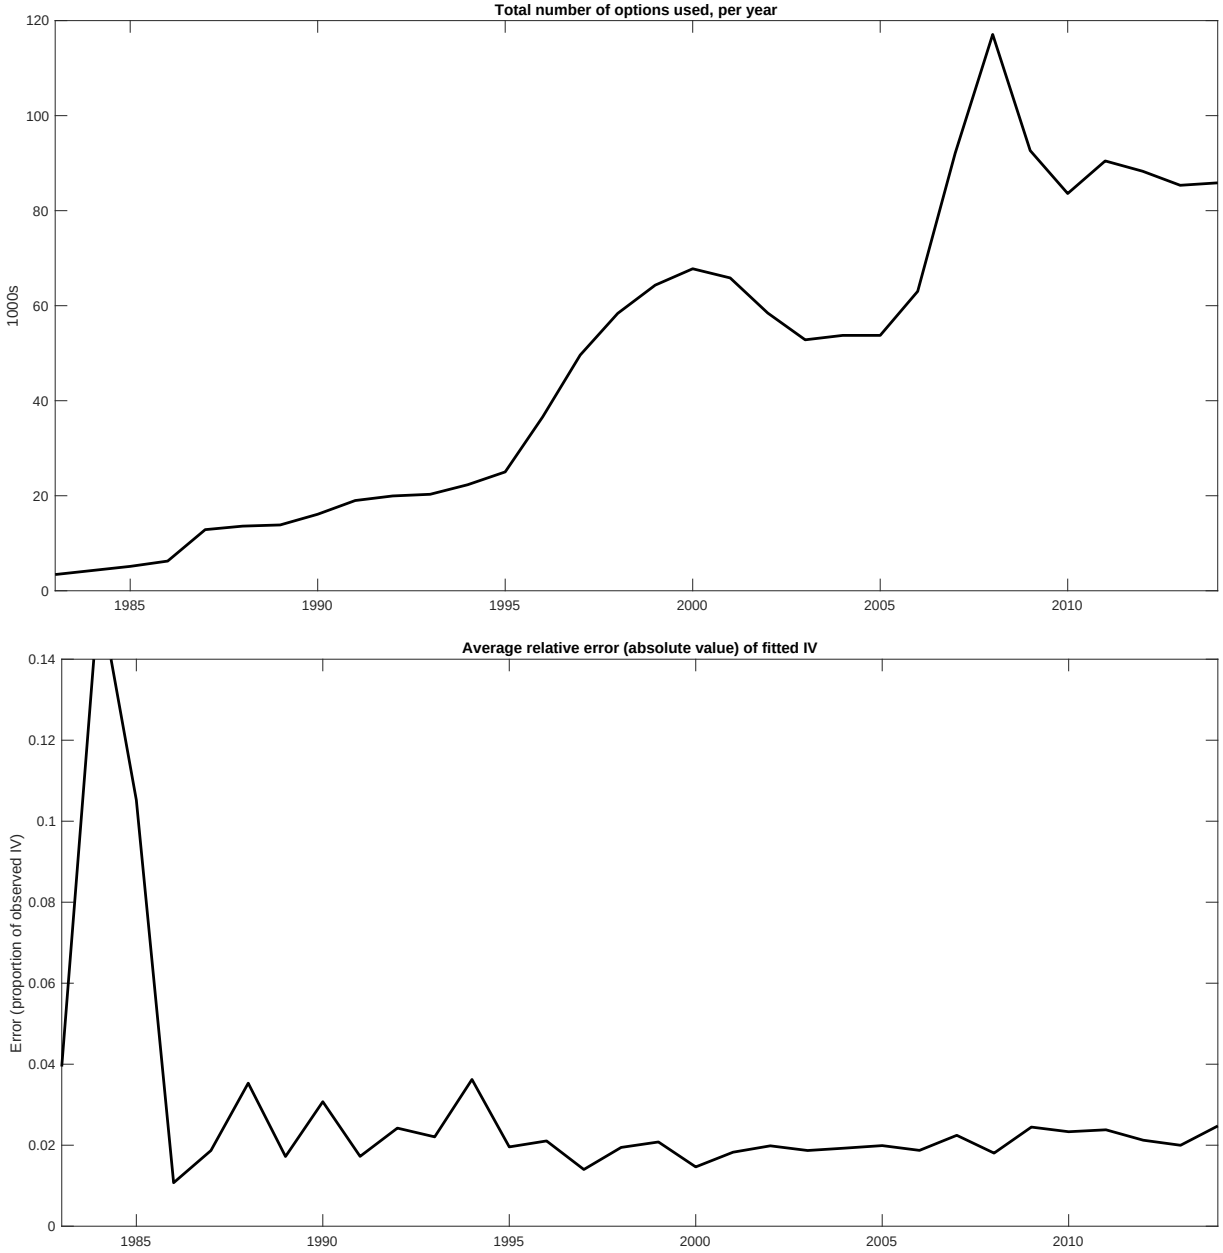

Note: The top panel reports the number of options used to compute implied volatility in each year, in thousands. The bottom panel reports the average absolute value of the pricing error of the SVI fitted line relative to the observed implied variances, in proportional terms (i.e. 0.02 means absolute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

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

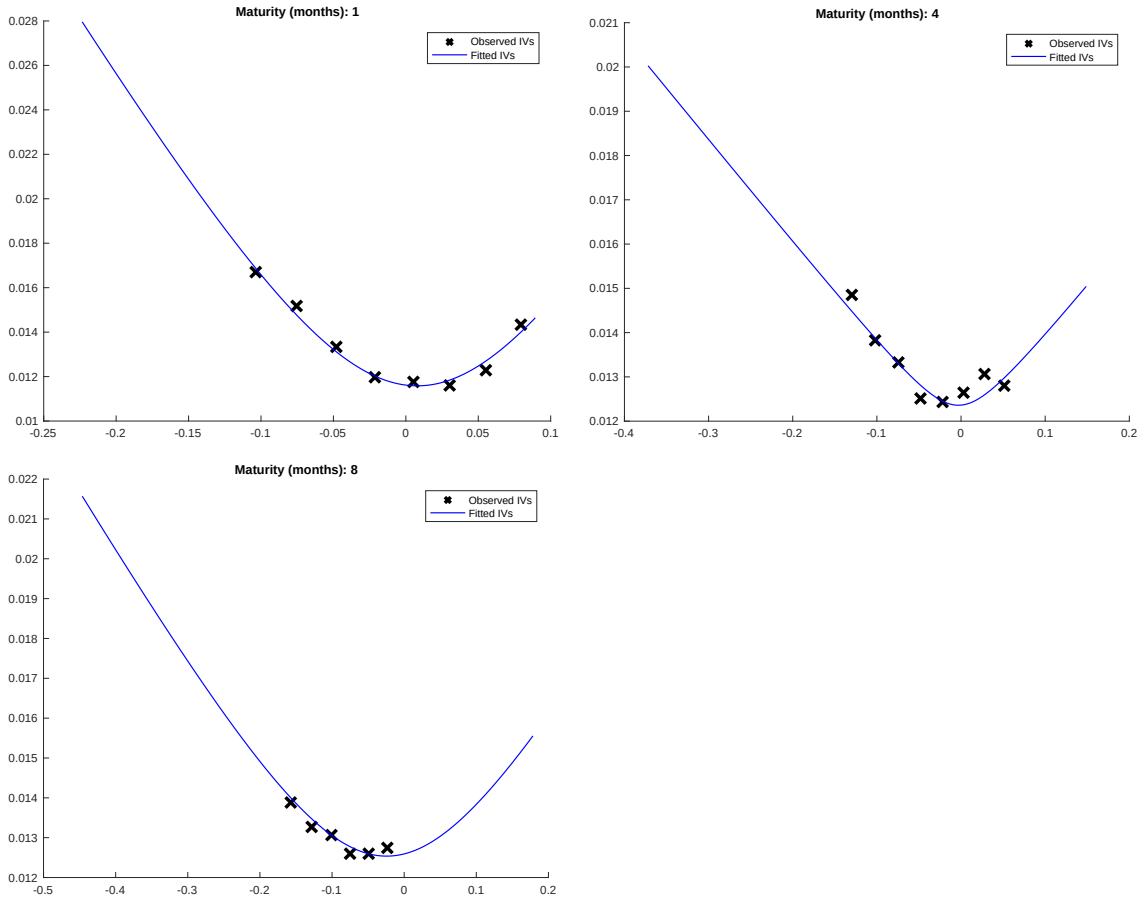

Note: Fitted implied variance curve on 11/7/1987, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve.

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

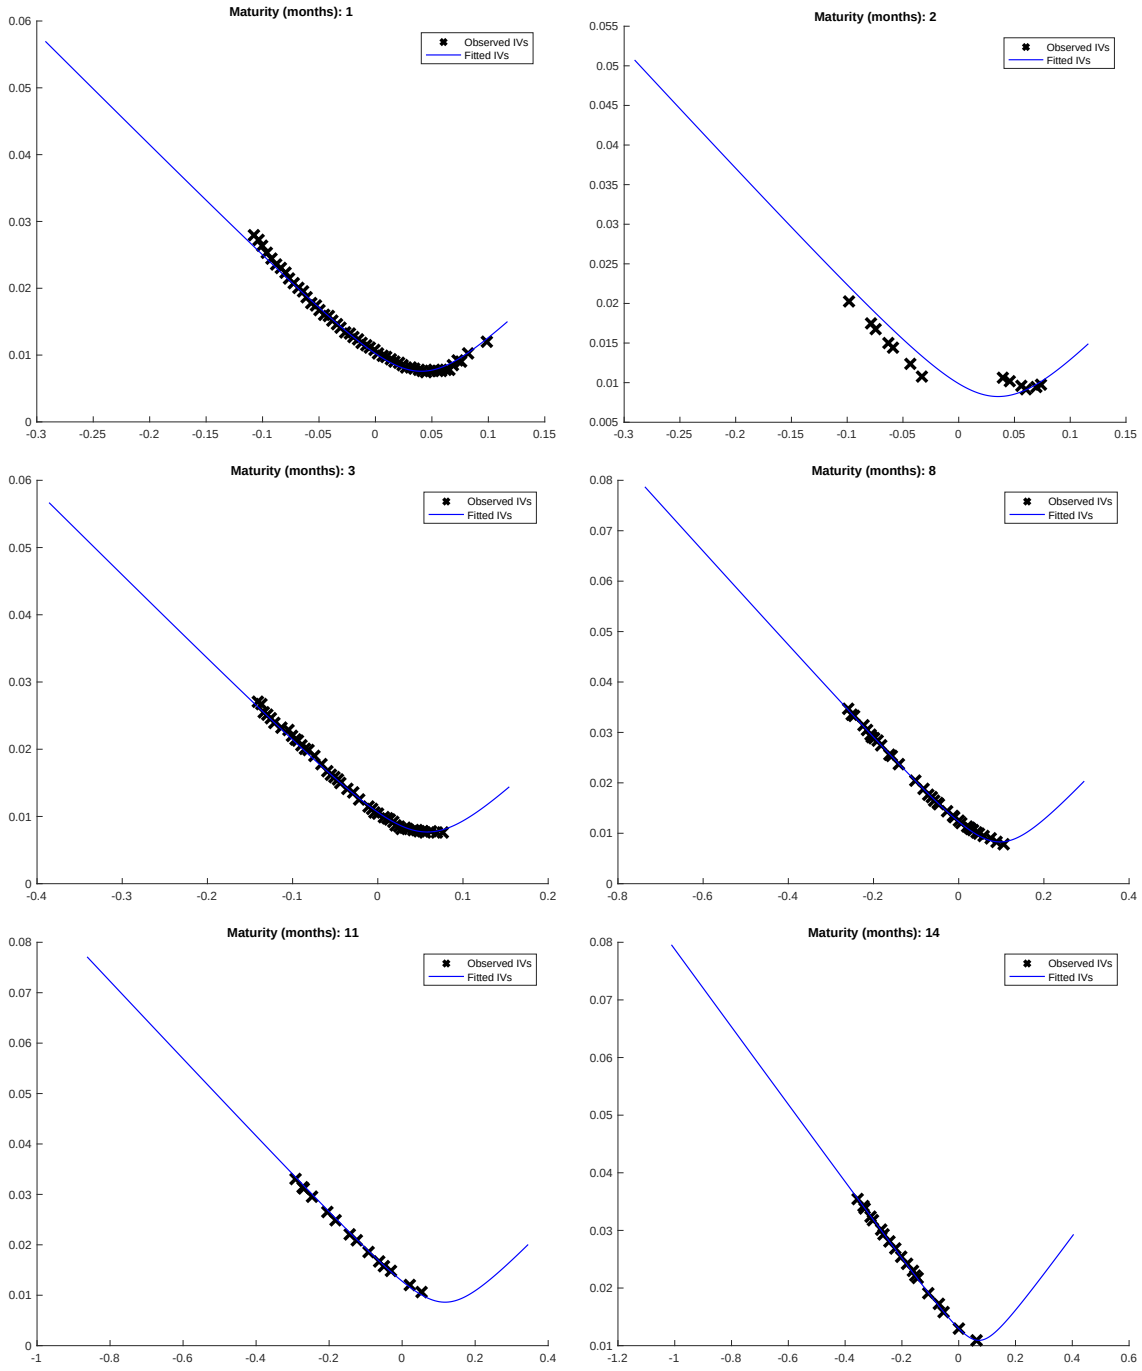

Note: Fitted implied variance curve on 11/1/2006, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve. On 11/1/2006 also a maturity of 5 months was available (not plotted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

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

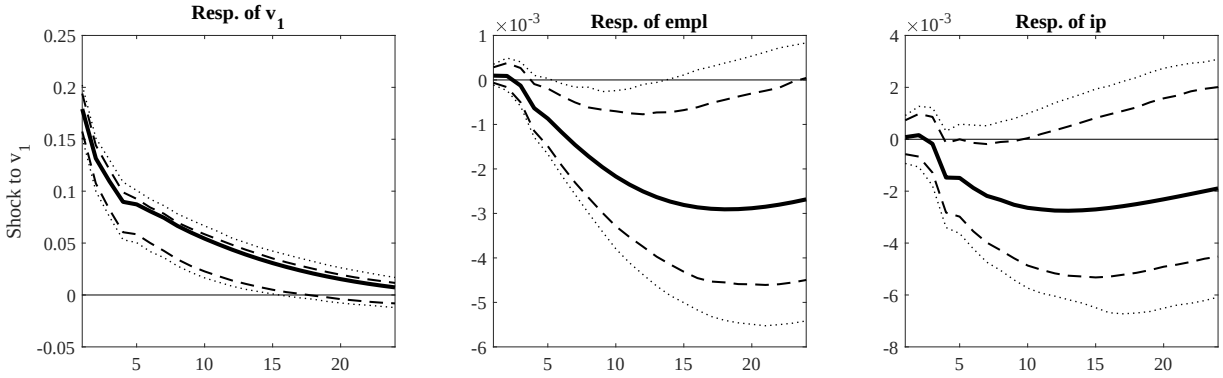

Note: The figure shows responses of volatility (measured by v_1), log employment, and log industrial production to a reduced-form shock to v_1 in a VAR with v_1, the Fed funds rate, log employment, and log industrial production with 68% and 90% confidence intervals. Sample period 1986-2014.

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

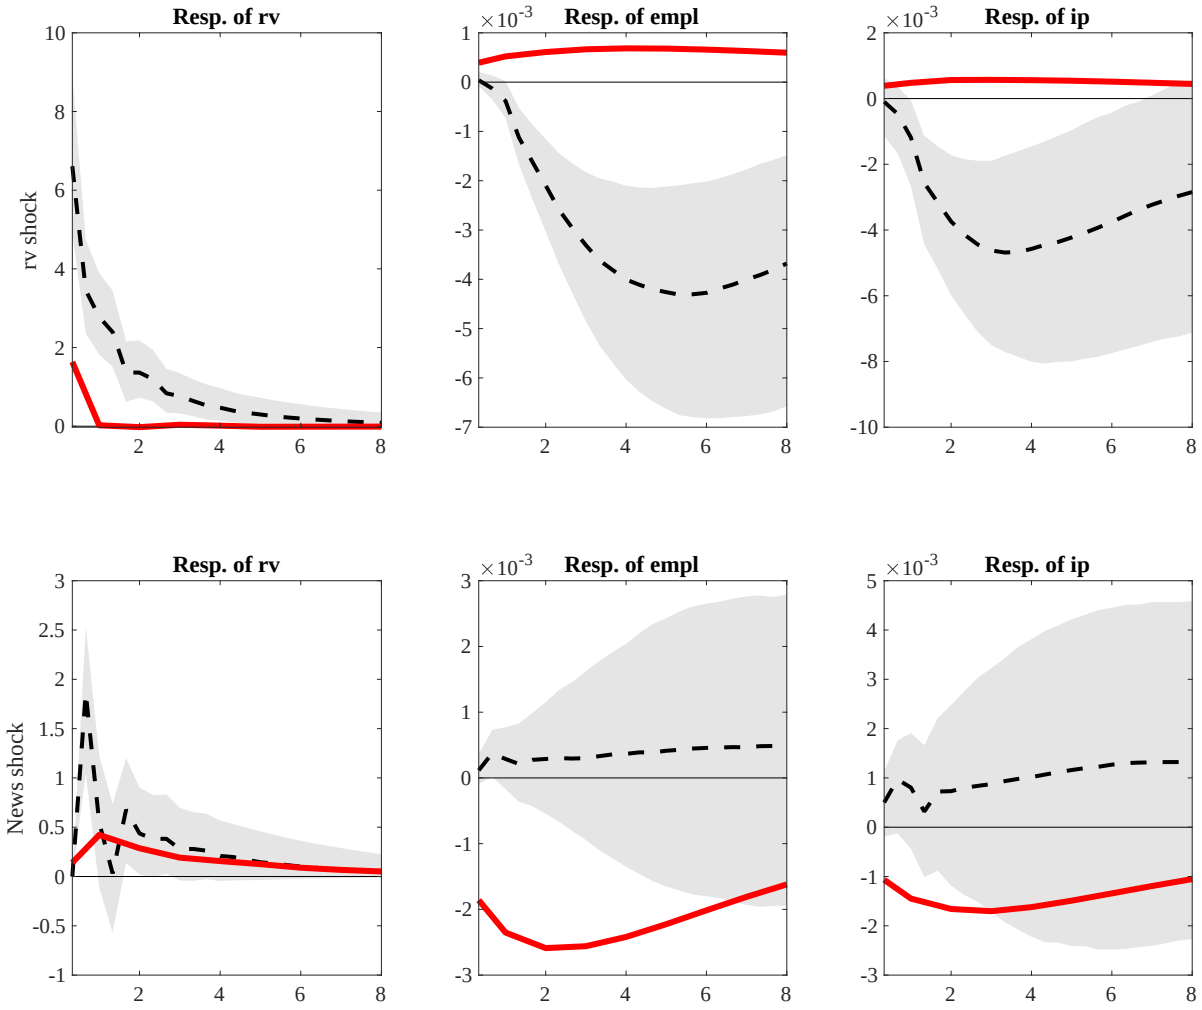

Note: Same as figure 2, but on data simulated from Basu and Bundick (2017).

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

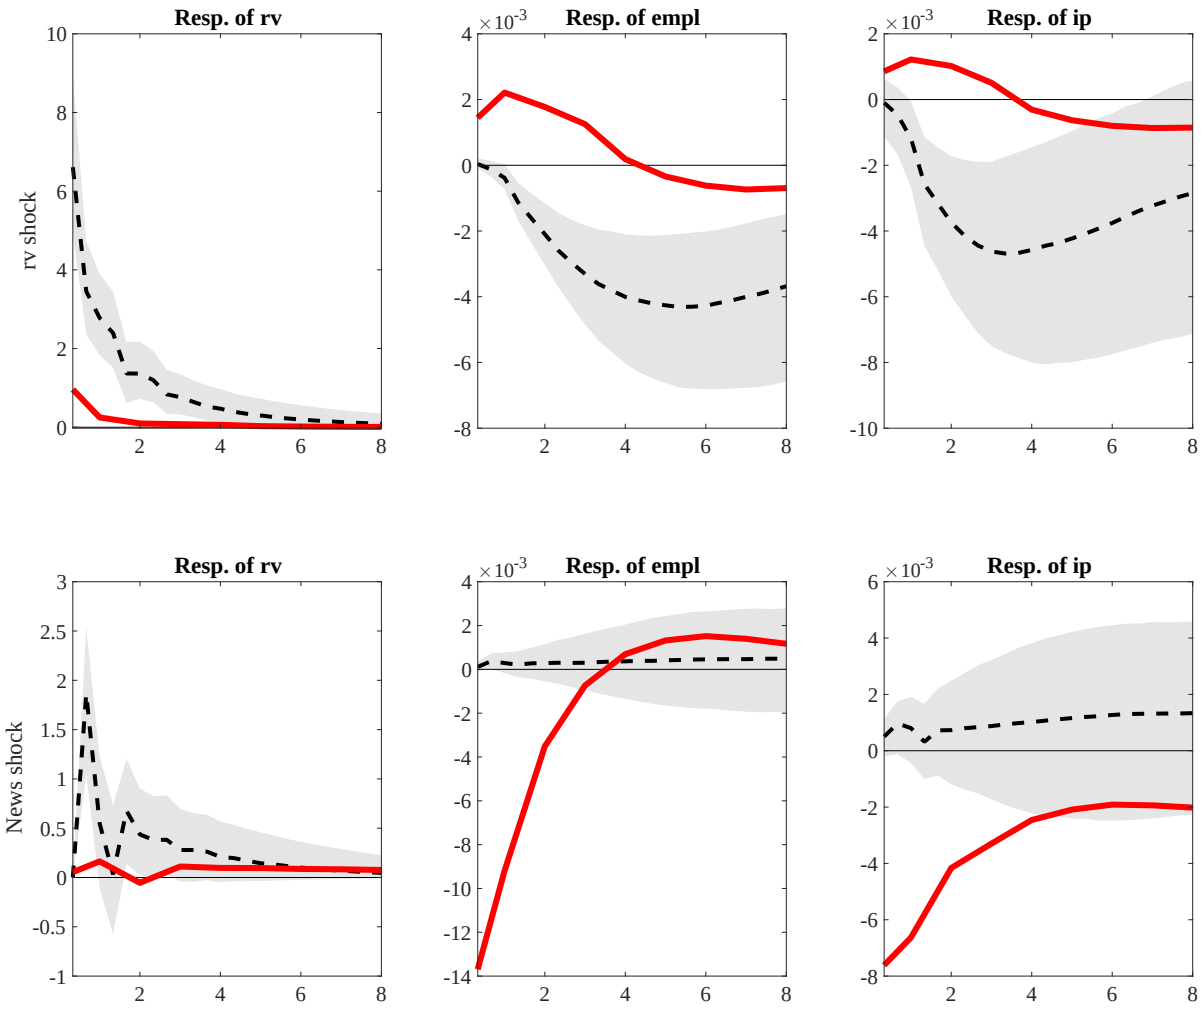

Note: Same as figure 2, but on data simulated from Bloom et al. (2017).

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

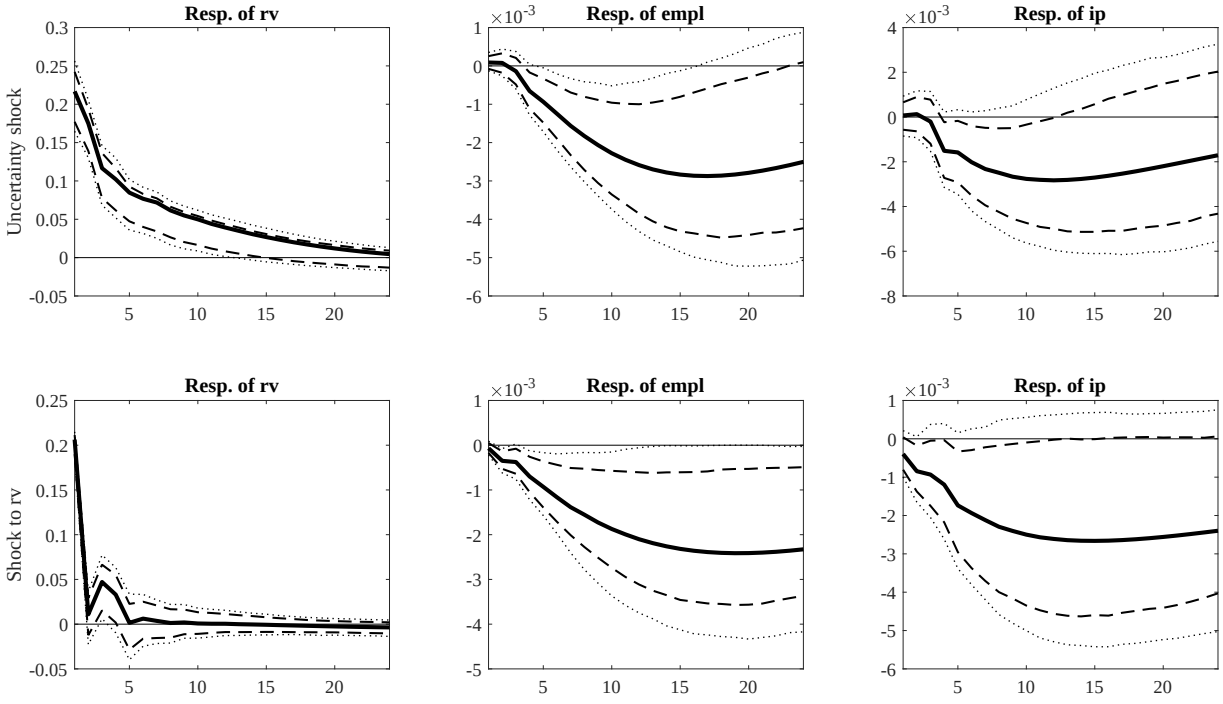

Note: See figure 2. Unlike in the baseline identification, the identified uncertainty shock is not orthogonalized with respect to rv. The rv shock in this case is the remaining part of reduced-form innovation to rv that is not spanned by the uncertainty shock.

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

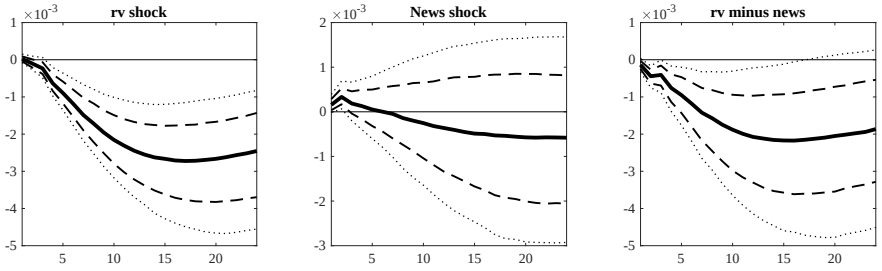

(a) Using the CBOE VIX instead of v_1

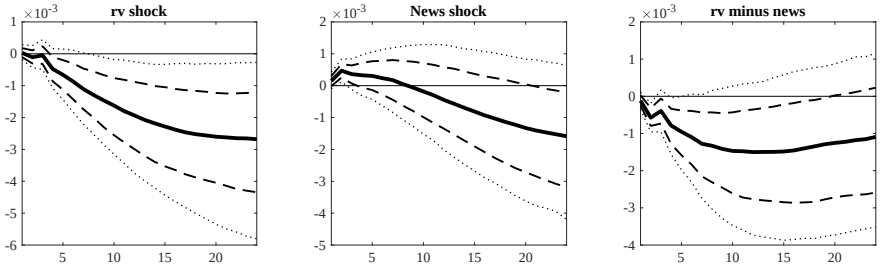

(b) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)

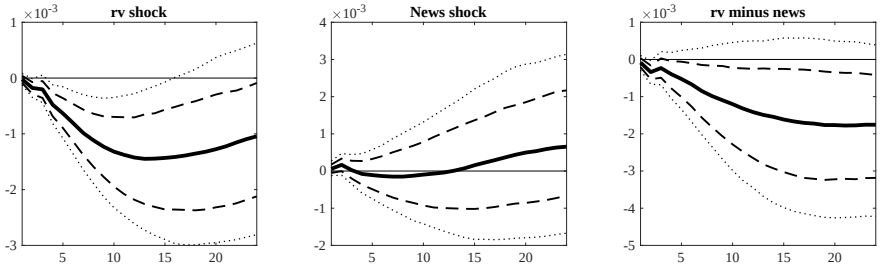

(c) Adding the S&P 500 level as first shock

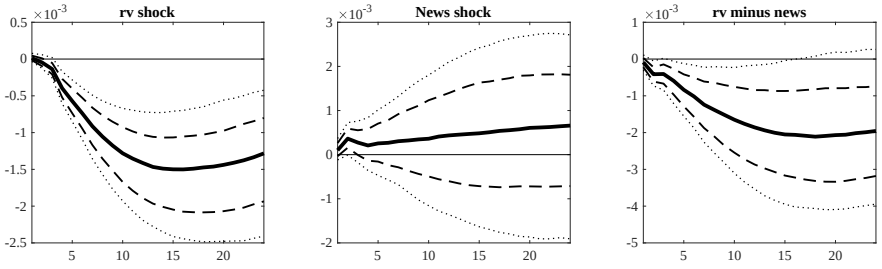

(d) Using RV and V_1 in levels, not logs

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

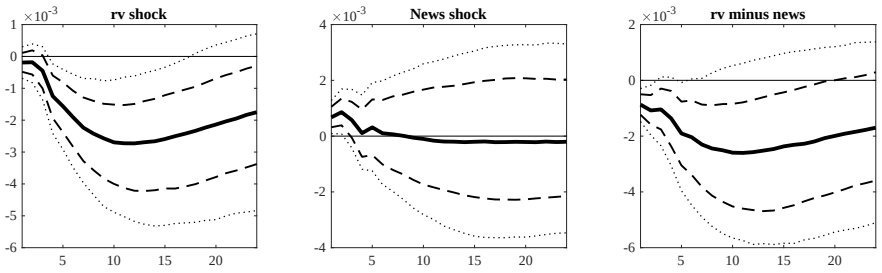

(a) Using the CBOE VIX instead of v_1

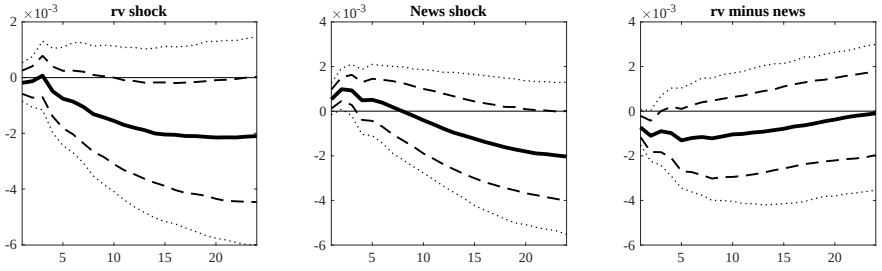

(b) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)

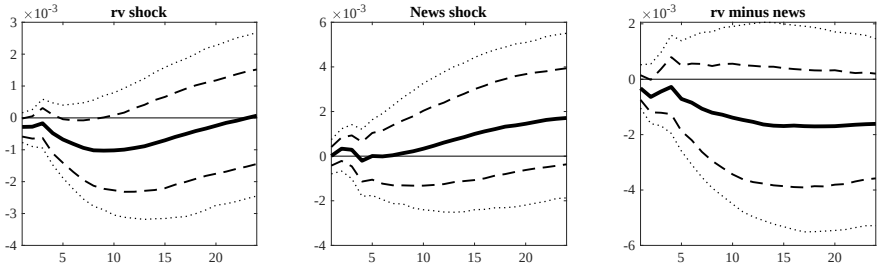

(c) Adding the S&P 500 level as first shock

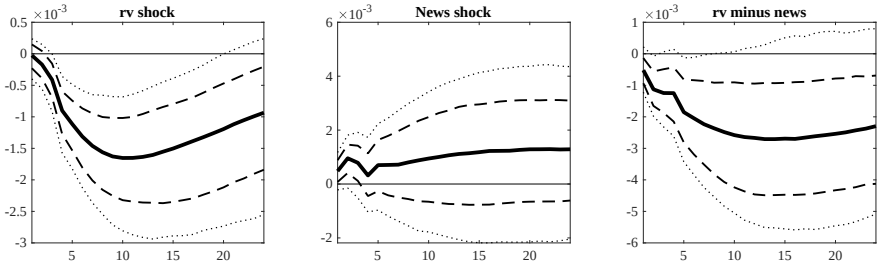

(d) Using RV and V_1 in levels, not logs

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

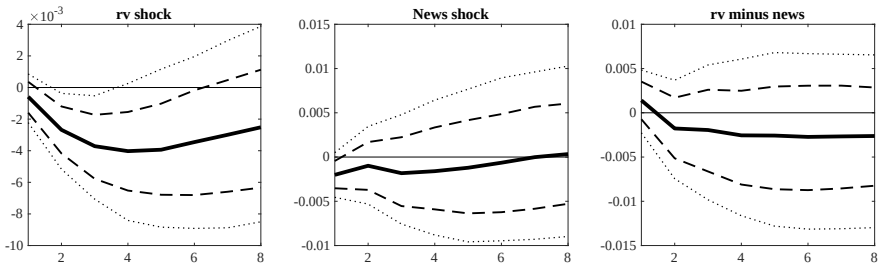

(a) Investment

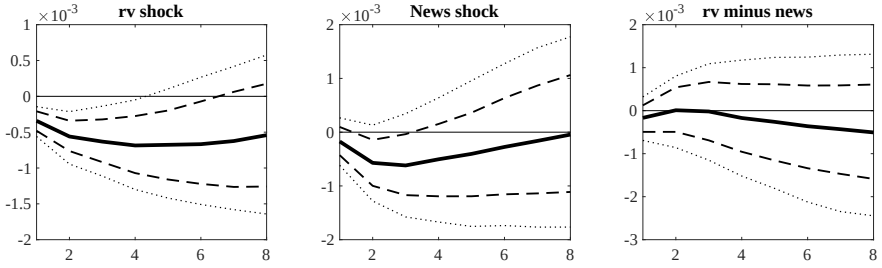

(b) Consumption

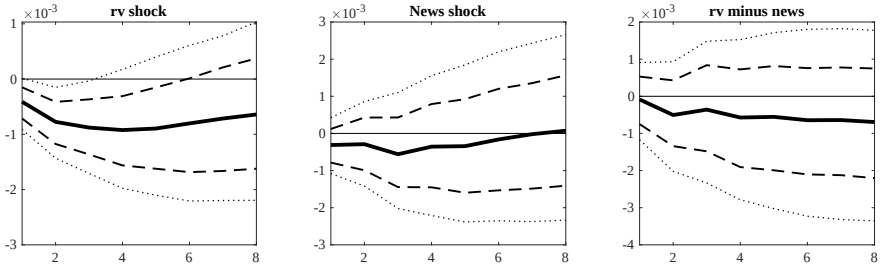

(c) GDP

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

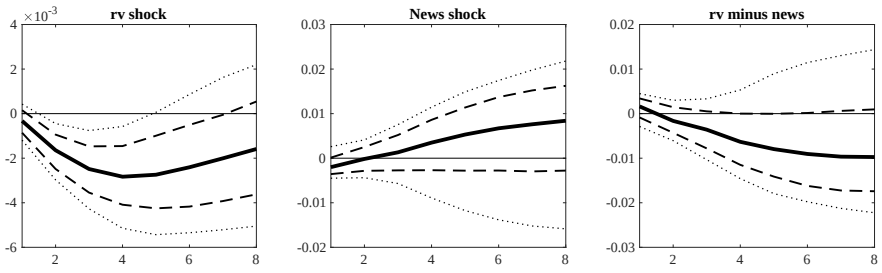

(a) Investment

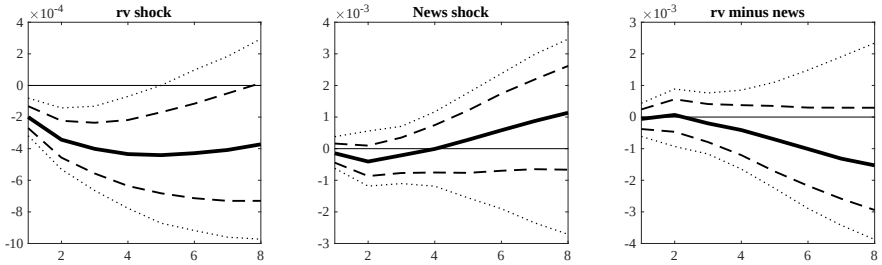

(b) Consumption

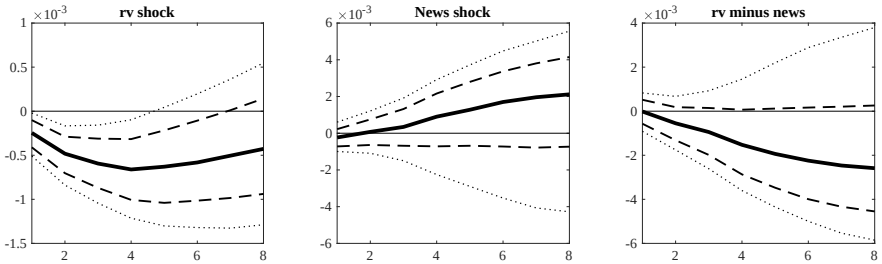

(c) GDP

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

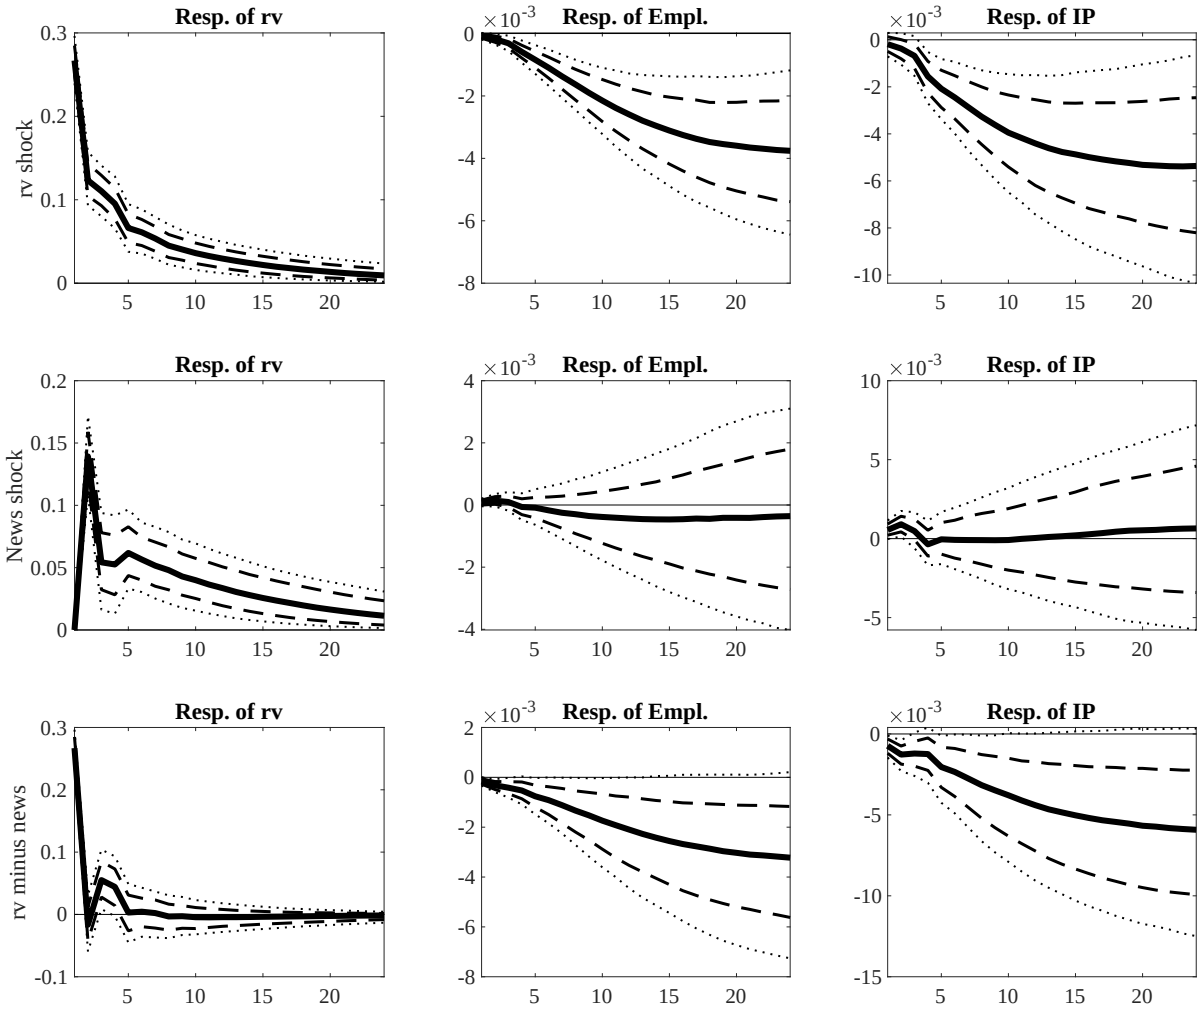

Note: See figure 2. IRFs from a FAVAR in which the state variables are rv, v_1, federal funds rate, and the first three principal components from the data set of Ludvigson and Ng (2007). Employment and industrial production are then obtained from regressions onto their own lags and those of the state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

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

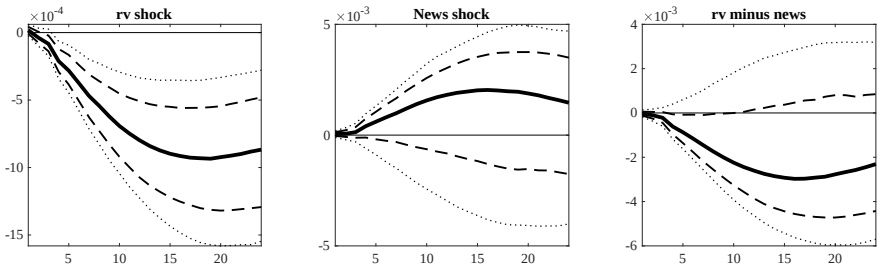

(a) BBD

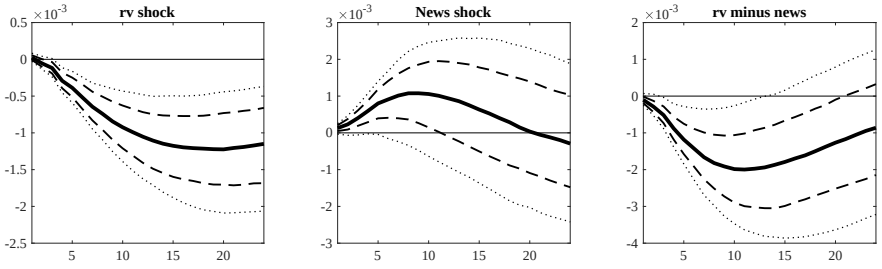

(b) BBD (Monetary)

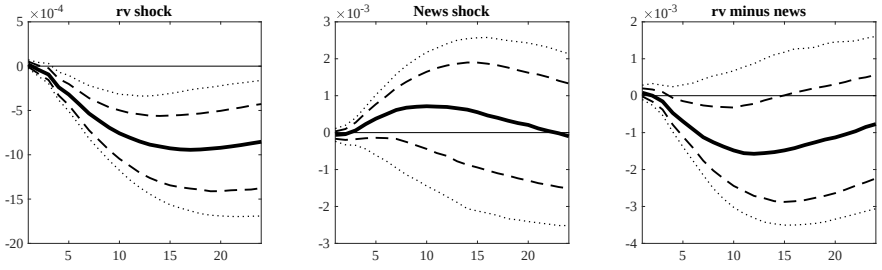

(c) BBD (Stock)

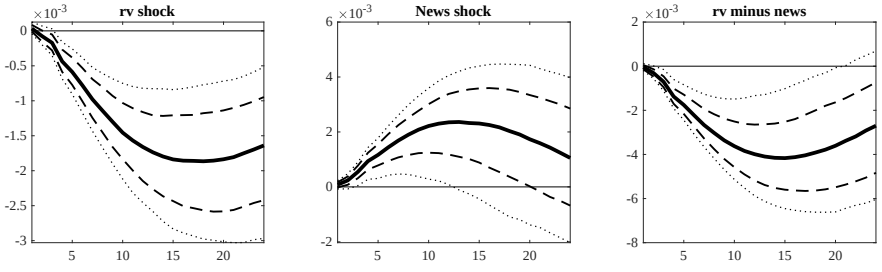

(d) Michigan Survey

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

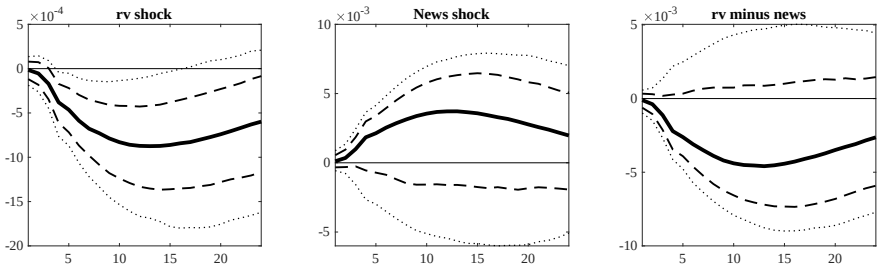

(a) BBD

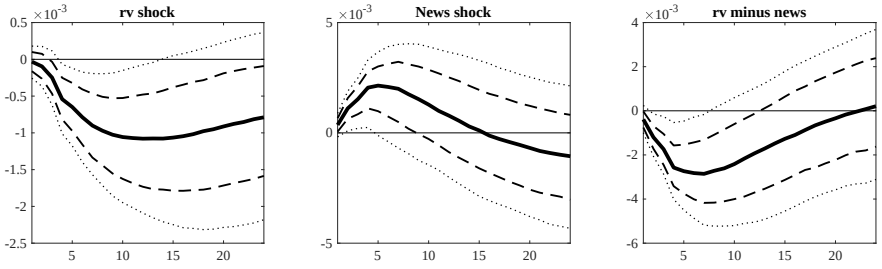

(b) BBD (Monetary)

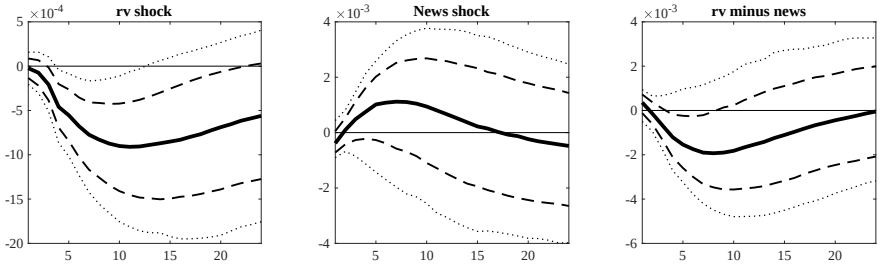

(c) BBD (Stock)

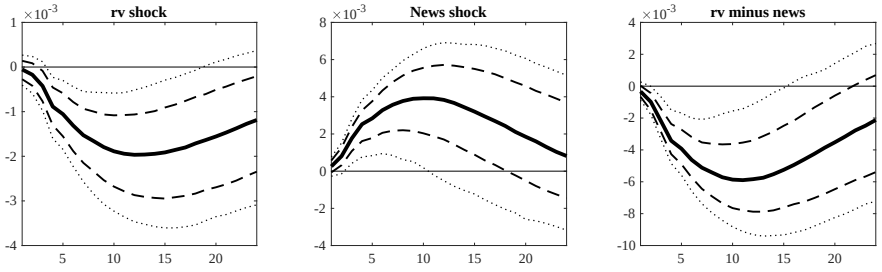

(d) Michigan Survey

Note: See figure A.14.

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

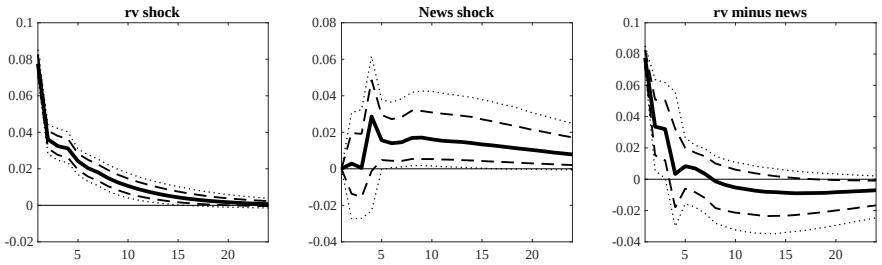

(a) BBD

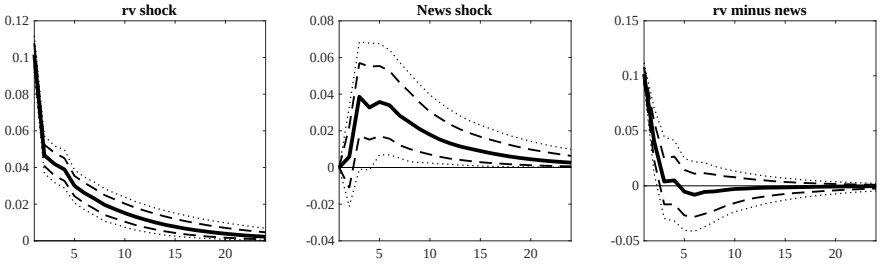

(b) BBD (Monetary)

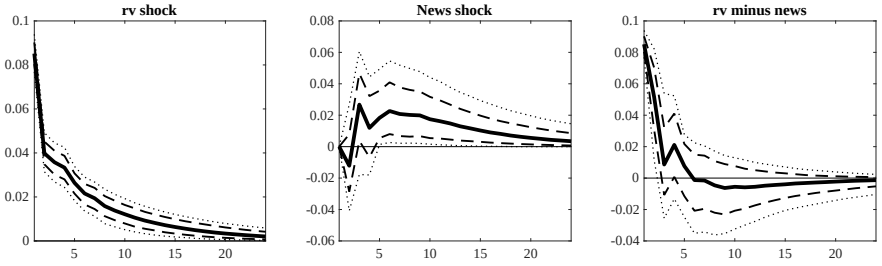

(c) BBD (Stock)

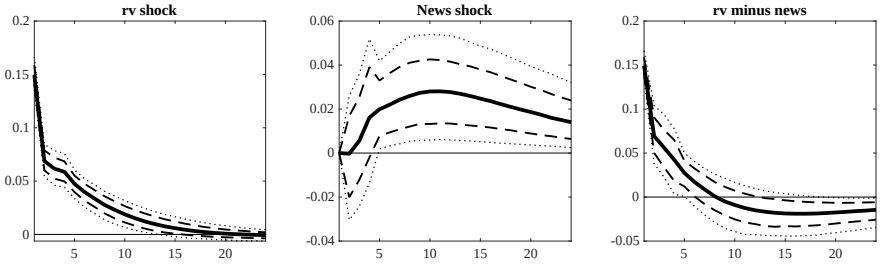

(d) Michigan Survey

Note: See figure A.14.

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

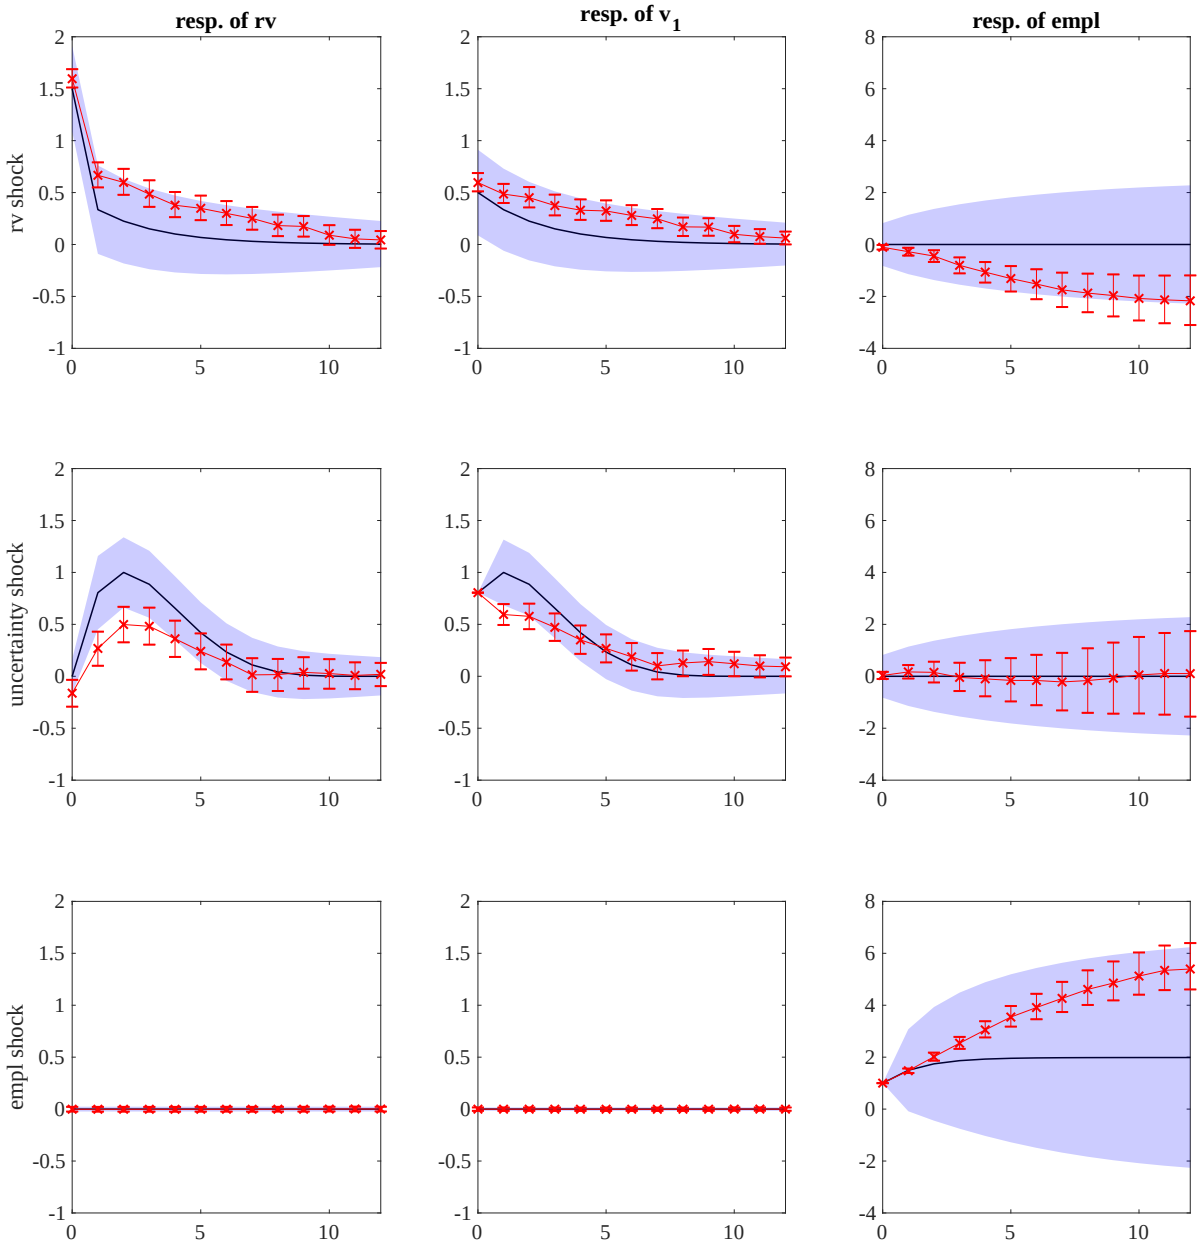

Note: Priors and posteriors from estimates of the Bayesian model with rv, v_1, and employment. Each panel reports an estimated impulse response to one of the three identified shocks. The black lines and shaded areas represent priors. The red boxes are the posterior mode and the whiskers are 90-percent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

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

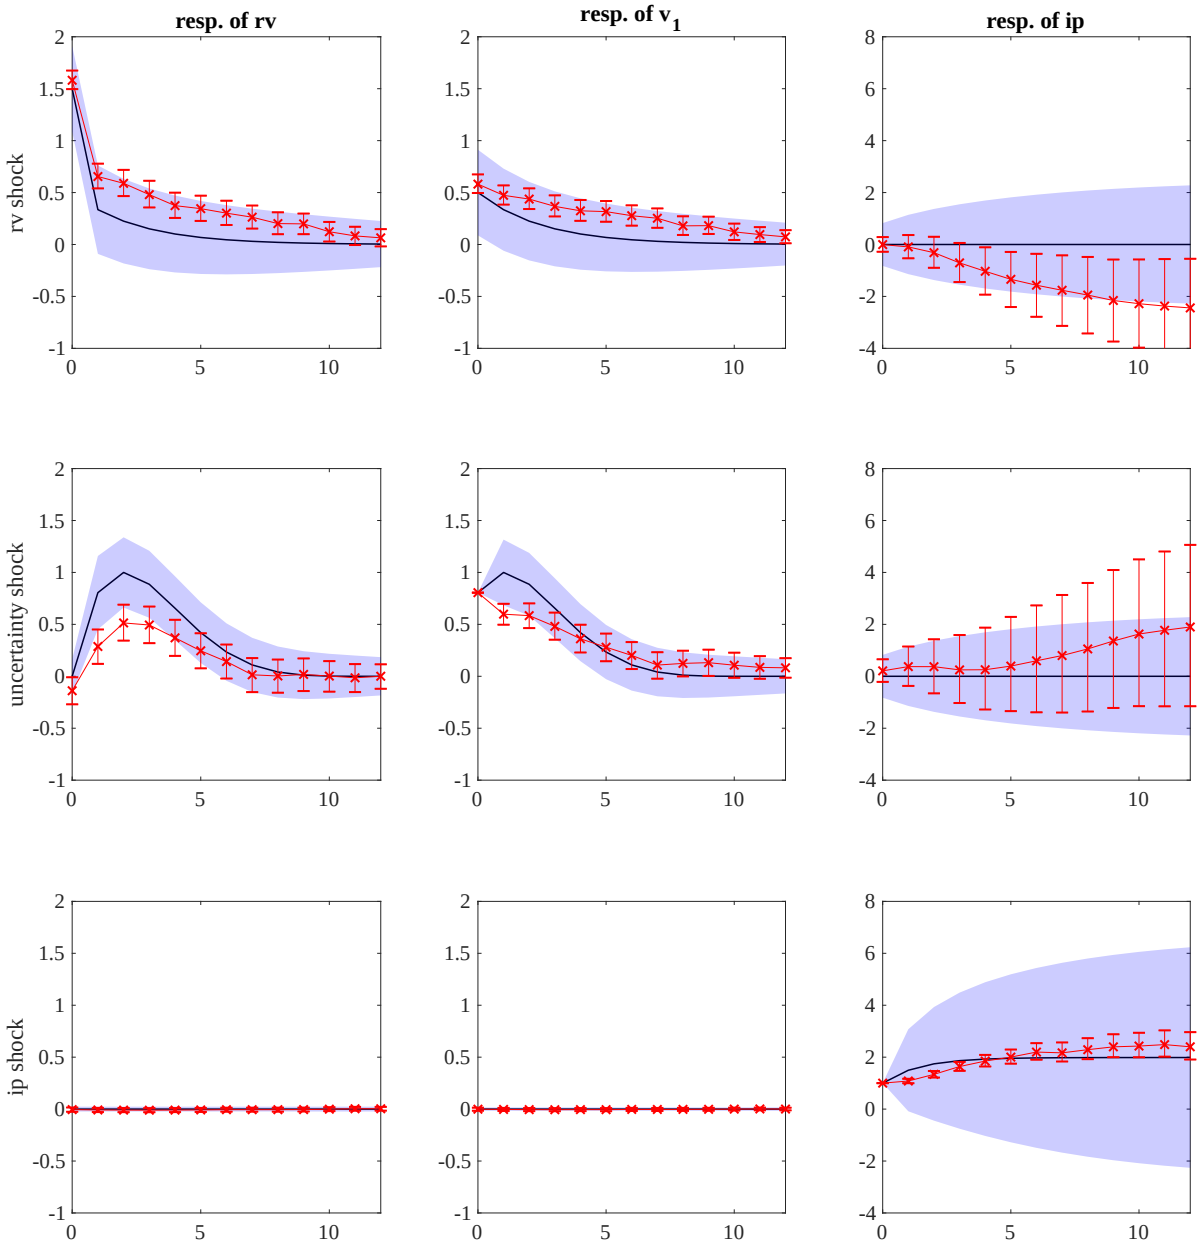

Note: See figure A.17. Here employment is replaced by IP.

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

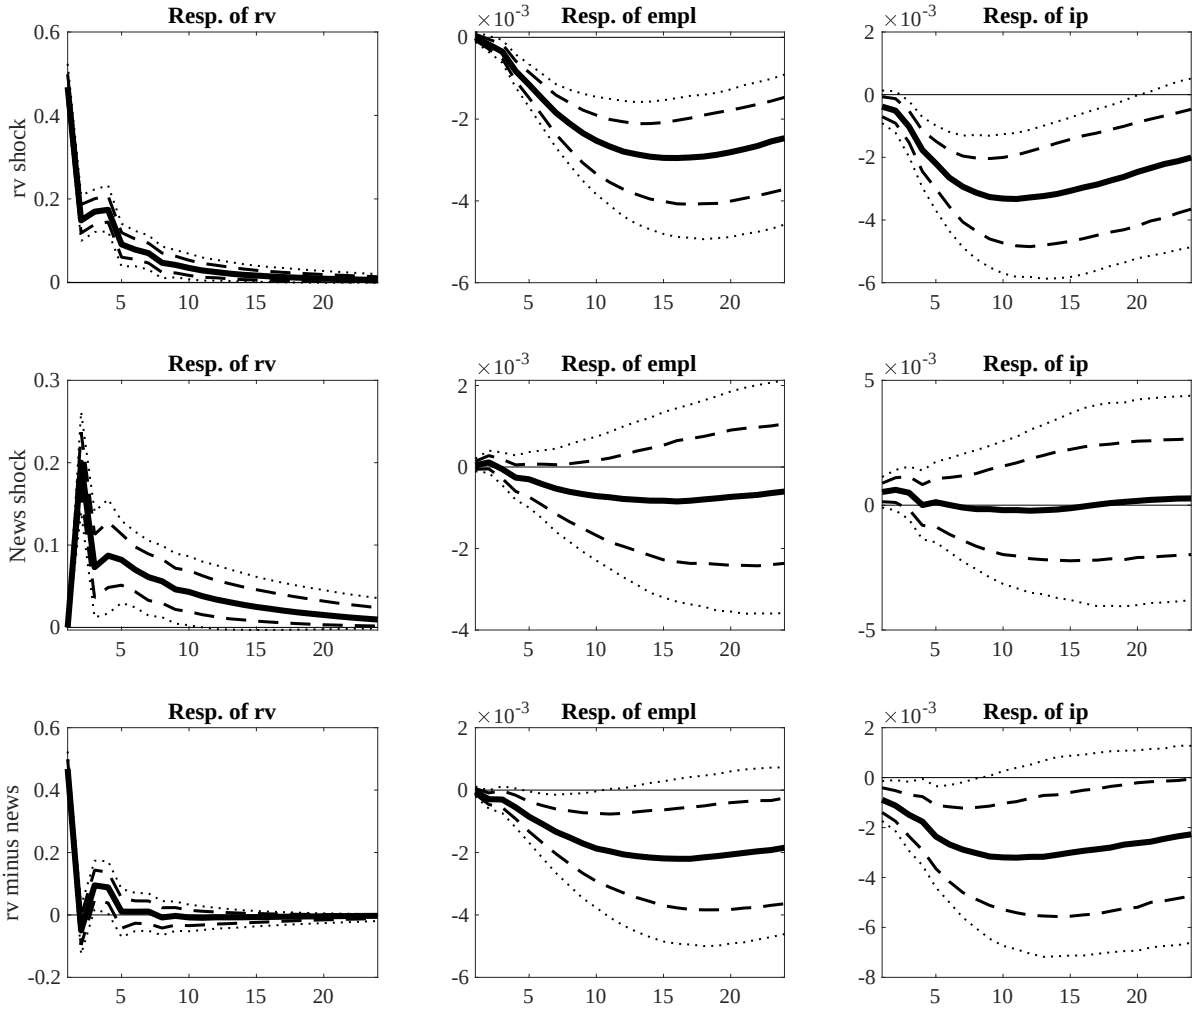

Note: See figure 2. The time series of RV_t has been replaced with $RV_t^{iv-purged}$, the version purged of the within-month implied volatility shocks (as described in Appendix A.4).

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

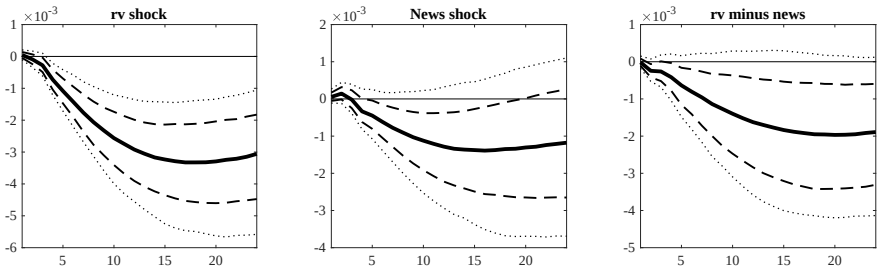

(a) Using downside volatility instead of rv

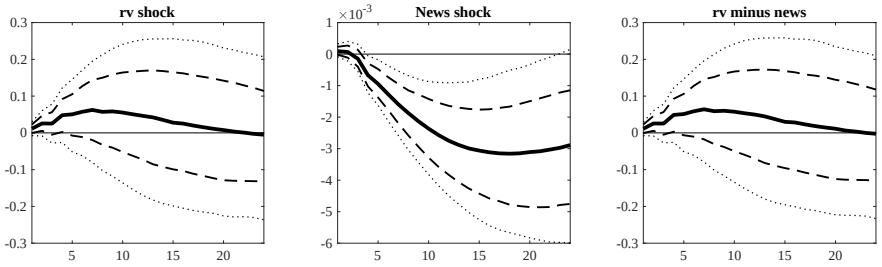

(b) Using upside volatility instead of rv

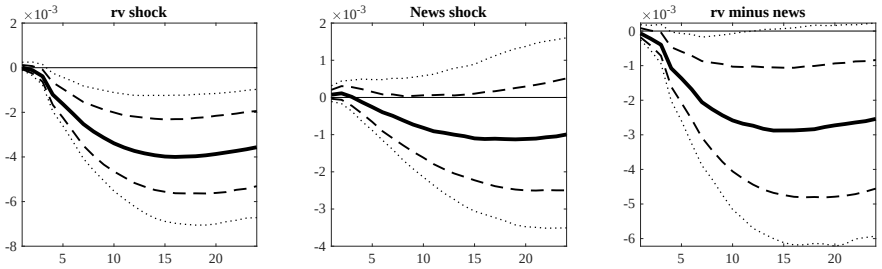

(c) Using jumps instead of rv

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 (c) 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

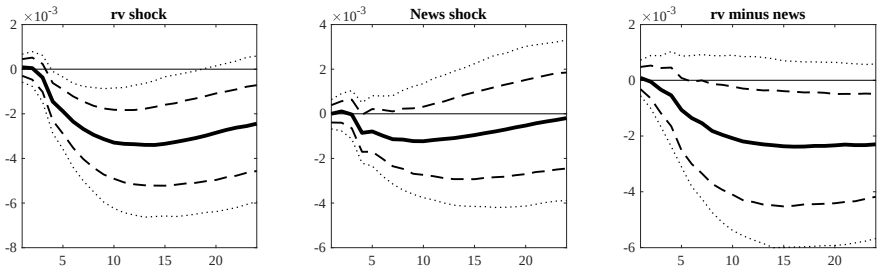

(a) Using downside volatility instead of rv

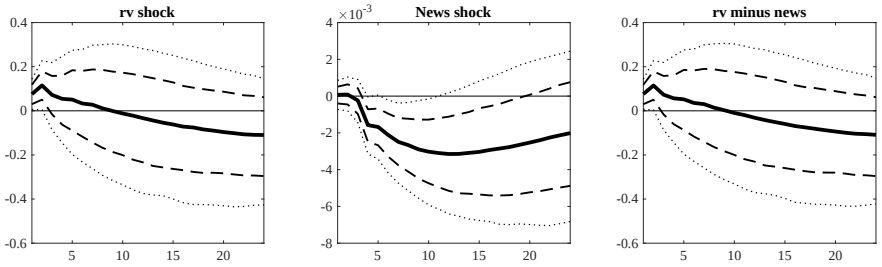

(b) Using upside volatility instead of rv

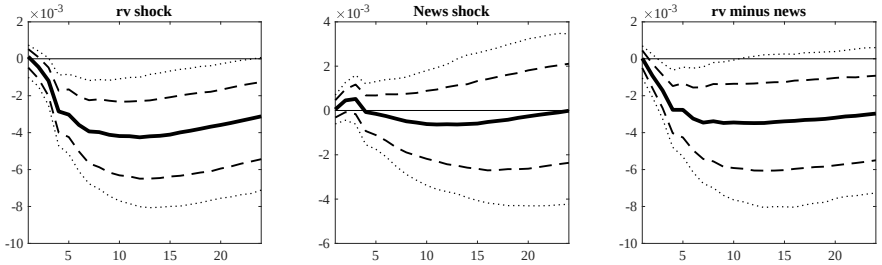

(c) Using jumps instead of rv

Note: See figure A.20.

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

	Employment		Industrial Production	
v_1	-0.09**	0.05	-0.07	0.12
	(0.04)	(0.04)	(0.05)	(0.08)
rv		-0.14***		-0.21**
		(0.04)		(0.10)

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***		0.60***
		(0.15)		(0.15)
FFR	-0.001	0.01		
	(0.02)	(0.01)		
Δemp	-34.96**	-4.98		
	(17.37)	(10.87)		
Δip	-3.62	-4.81		
	(5.45)	(3.57)		
PC 1			-0.029**	-0.010
			(0.013)	(0.009)
PC 2			-0.014	-0.017
			(0.018)	(0.012)
PC 3			-0.014	-0.011
			(0.010)	(0.007)
$R_{S\&P}$			-1.56***	0.12
			(0.43)	(0.41)
Def			0.21	0.07
			(0.14)	(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	0.01	0.03	0.01	0.01
	(0.09)	(0.09)	(0.09)	(0.09)	(0.08)
$v_{1,t}$	0.57***	0.49***	0.51**	0.54***	0.60***
	(0.13)	(0.15)	(0.26)	(0.13)	(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.