Sergey Egiev

ON PERSISTENCE OF UNCERTAINTY SHOCKS

BASIC RESEARCH PROGRAM
WORKING PAPERS

On Persistence of Uncertainty Shocks

Sergey Egiev *

Preliminary and Incomplete

Abstract

I study real effects of uncertainty shocks. Using time-varying volatility of the forecast error, I construct a two-part uncertainty metric that consists of persistent and volatile, burst-like components. These indices are used to study empirically several predictions of uncertainty models: that uncertainty shocks have real effects, that these effects realize in a downturn/overshoot pattern and that persistence of uncertainty shocks decreases this pattern's frequency and increases its amplitude. Using the constructed metric in a simple VAR framework I show that real effects are there, that shock to the volatile uncertainty causes significant downturn/overshoot pattern, and that shock to the persistent component causes severe and prolonged damage.

JEL Classification: C53, E32, G12, G35, L25
Keywords: uncertainty, business-cycles, real business cycles

* Department of Economics, Higher School of Economics, Moscow. Email: s.egiev@hse.ru

1 Introduction

In this paper I study real effects of uncertainty shocks. I use time-varying volatility of the forecast error as an uncertainty proxy (Jurado, Ludvigson & Ng, 2015) and construct a two-part uncertainty index that consists of persistent and volatile components. Persistent component is slowly-moving, and its increase is a signal of a prolonged period of elevated uncertainty. In contrast, the volatile, burst-like metric, proxies uncertainty that tends to get resolved quickly.

These indices allow me to study empirically three predictions of rational expectations equilibrium (REE) uncertainty models such as Bloom (2009) or Bloom et. al. (2014). First, that uncertainty shocks have significant real effects. Second, that these effects realize in a particular pattern: downturn followed by overshoot. Third, that more persistent shocks result in this pattern having lower frequency and larger amplitude[1]. I use the constructed persistent and volatile indices in a VAR setting with essentially the same ordering as in JLN (2015) and Bloom (2009). I show that a) shocks to both proxies negatively affect production and employment on impact; b) shock to the volatile burst-like component creates a significant slowdown/overshoot pattern; c) shock to the persistent component results in a more prolonged and severe downturn when compared to that of the original JLN metric.

A growing body of theoretical literature studies uncertainty shocks or 'volatility risk' (Liu, Miao, 2015). These shocks have been shown to have sizable real effects of business-cycle frequency both in partial (Bloom, 2009) and general equilibrium settings (Bloom et. al., 2014). Various empirical studies that attempted to verify these predictions in the REE setting used implied volatilities (VIX/VXO) as a proxy for uncertainty (see, for example, the original contribution of Bloom (2009) or recent Caggiano et. al. (2015) study with ST-VARs). They use variation in such indices to estimate effects of uncertainty shocks on real variables. While good as a first-order approximation, VIX/VXO may be too crude a proxy given other options available. Nodari (2013) and Colombo (2013) use multi-layered Bloom et. al. (2013) index that tracks variation in uncertainty by combining newspaper-based data with other measures such as forecast disperison. Bachmann et. al. (2013) construct their own, micro-based index of uncertainty using German business climate survey (IFO).

Jurado, Ludvigson & Ng (2015) provide a state-of-the art uncertainty metric. First, they

---

[1]Kozlowski, Veldkamp & Venkateswaran (2015) show, that in a model with constant belief updating even *non-persistent* shock in uncertainty can create *persistent* change in beliefs that will have long-run consequences.

Figure 1: IRFs to the 1 s.d. uncertainty shock, 60 month-period
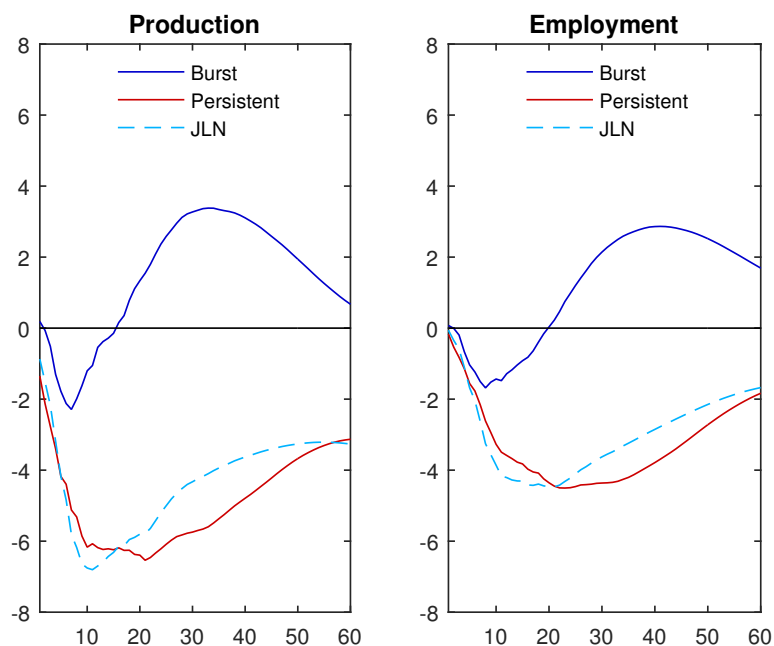
Notes: Metrics are based on the 3-month forecast error variation.

recursively forecast future values of several hundreds of macroeconomic and financial series and isolate the forecast error for each of them. In order to accurately estimate the forecastable part of a series, informational set of agents should be spanned as closely as possible. Hence, approximate dynamic factor model (DFM) that is suitable for data-rich environment is used as a forecasting device. JLN explicitly allow for heteroskedasticity in the forecasting equation in order to catch time-variation in the volatility of forecast error. Then, given the forecast error data, they estimate volatility of the unforecastable component of each series with a stochastic volatility model (Kim et. al., 1998). Finally, they aggregate across series and use the resulting index as a proxy for macroeconomic uncertainty.

I use the JLN framework to study three major properties of theoretical uncertainty shocks. First, that they have significant effect on real variables. Second, that this effect takes form of a particular patter. The original Bloom's model (2009) shows in a partial equilibrium setting that short-run uncertainty shock will result in a swift slowdown and, then, an overshoot. This happens because elevated uncertainty causes wait-and-see behaviour, so when uncertainty gets finally resolved, firms react with overshooting. Third, that more persistent uncertainty shock results in this pattern having lower frequency and higher amplitude. This happens because, wait-

and-see behavior tends to accumulate on various production units while uncertainty remains elevated. This accumulation amplifies downturn/overshoot fluctuation.

To study these implications, I decompose the original JLN metric into two parts. First, I construct a persistent uncertainty index by allowing for heavy-tailed innovations in the forecasting equations of JLN. The volatility of the forecast error is then estimated with a stochastic volatility model with heavy-tails (Chib et. al. (2002)). While in a normally distributed case tails decay exponentially, so almost each large realization of forecast error is treated as a signal of higher uncertainty, in heavy-tailed case there should be several forecast errors of considerable magnitude in a row for the model to decide that uncertainty has indeed increased. This results in a more persistent uncertainty metric than the original JLN was.

Second, I use the difference between original JLN metric that used SV-normal model and persistent metric that uses heavy-tailed SV (this difference taken directly as in Chib, Nardari & Shephard, 2002 or estimated as a scaling component in the gaussian mixture approximation to the measurement equation error) as a proxy to the volatile, burst-like component. Broadly speaking, this index accounts for uncertainty that results from very short-run, sudden forecast errors that do not persist. Although such bursts can have large magnitude (e.g. 2008 crisis), they are not likely to stay there for long and are ex-ante expected to recede quickly.

I use persistent and burst uncertainty indices to study in more detail effects of uncertainty shocks in an otherwise standard VAR setting. First, both components significantly affect real variables, such as employment and production, on impact. Second, the burst component results exactly in what one might expect given Bloom (2009) and Bloom et. al. (2014) contributions: swift slowdown followed by an overshoot. I find this downturn/overshoot pattern to be statistically significant both in production and employment given 1 standard error confidence bands. Third, shock to the persistent component of the original JLN metric results in a more severe and prolonged downturn, than the JLN itself which suggests effect of persistence on frequency and amplitude as described in Bloom (2009).

The rest of the paper is organized as follows. Section 2 describes the forecasting procedure. Section 3 describes construction of uncertainty indices from the forecast error data using stochastic volatility models. Section 4 describes VAR simulations. Section 5 concludes.

2   Estimating the forecast error

To proxy uncertainty, Jurado, Ludvigson & Ng (2015) construct a measure how 'unpredictable' the economy is based on the forecasting errors. Hence, the notion of 'uncertainty horizon' inevitably arises; in particular, the h-period ahead forecast uncertainty is defined as conditional volatility of the unforecastable component of each series $j \leq N$:

$$U_{jt}^y(h) = \sqrt{E\left(\left(y_{j,t+h} - E[y_{j,t+h}|I_t]\right)^2 |I_t\right)} \tag{1}$$

The aggregate measure is then constructed from individual uncertainty series as an (equally) weighted average. To accurately estimate individual uncertainty, the forecastable component of each data series should be isolated as closely as possible. Hence, the informational set of agents (I_t) should be spanned well with a standard technique in a data-rich environment being dynamic factor model (DFM). The model is approximate (Chamberlain & Rotschild, 1983) in the sense of allowing for limited correlation between errors in the DFM equation which is suitable given large size of the data set.

$$X_{it} = \Lambda_t F_t + e_{it} \tag{2}$$

X consists of a large number of series, whereas F_t are the small number of factors with splashed information in them. F_t have autoregressive structure of their own. The standard practice in DFM forecasting is to estimate factors with PCA and then forecast with OLS treating factors as observed (Bai, Ng, 2006). In particular, mean-squared optimal forecast $y_{T+h} = \alpha' F_t + \beta' W_t$, where F_t are factors and W_t possibly additional regressors is replaced with a feasible prediction $\tilde{y}_{T+h} = \hat{\alpha}' \tilde{F}_t + \hat{\beta}' W_T$. Stock & Watson (2002a) showed that such forecast $\tilde{y}_{T+h}$ is a consistent estimate of the y_{T+h}. The number of factors is 12 due to Bai, Ng (2002) criterion. In practice, the forecasting equation takes form of:

$$y_{j,t+t} = \gamma(L)y_{j,t} + \gamma^F(L)F_t + \gamma^W(L)W_t + \varepsilon_t \tag{3}$$

I explicitly allow for heavy-tailed time-variation in the innovations y_{t+1}. Moreover, to forecast more than one period ahead ($h > 1$), the F_t itself is moved forward with a simple AR(4) model. I also allow for heavy-tailed heteroskedasticity in the F_t-forecasting AR(4) model. Heavy-tailed

heteroskedasticity results in possible inefficiency on the OLS' side, but does not affect consistency of estimates.

3   Estimating uncertainty

There is a large literature on stochastic volatility models (see Bos, 2011 for a good review). The key difference between the SV and GARCH approaches to modeling time-variation of the second moment is that SV model explicitly allows for a separate shock to variance. GARCH does not, hence, the suitability of SV model for the study of uncertainty shocks. It can be neatly expressed in the state space form as:

$$y_t = \sigma_t \varepsilon_t \tag{4}$$

$$\log \sigma_t^2 = h_t = \gamma + \phi h_{t-1} + \eta_t \tag{5}$$

Table 1: Sample statistics of the persistent and JLN indices for 1, 3 and 12 month ahead forecast error series.

Series	Mean	Standard Dev.	Kurtosis	Persistence
h = 1				
persistent	66.9	9.0	6.8	31%
JLN index	68.8	9.4	7.1	27%
h = 3				
persistent	81.6	9.6	6.3	45%
JLN index	84.6	10.0	6.6	36%
h = 12				
persistent	91.5	6.4	4.9	60%
JLN index	95.4	6.5	5.0	53%

Notes: Persistence is measured with the slope coefficient of the $d(x) = const + d(x(-2))$ equation.

The key obstacle to its estimation is that it is either non-Gaussian, or non-linear, so the standard EM strategy is not feasible. See that $\log y_t^2 = h_t + \log \varepsilon_t^2$, where $\log \varepsilon_t^2$ is clearly not normal. Kim, Chib & Shephard (1998) proposed to approximate it with a mixture of Gaussians and used a multi-step MCMC sampler to obtain the joint posterior of the parameters and latent volatilities. With data augmentation (Tanner & Wong, 1987), they trade the state space size for conditional normality of the model. While at each cycle another state variable should be drawn (indicator of which of the gaussians acts now), after that the whole vector $\mathbf{h}_t$ can be drawn at

Figure 2: Persistent and JLN uncertainty indices for various forecast horizons
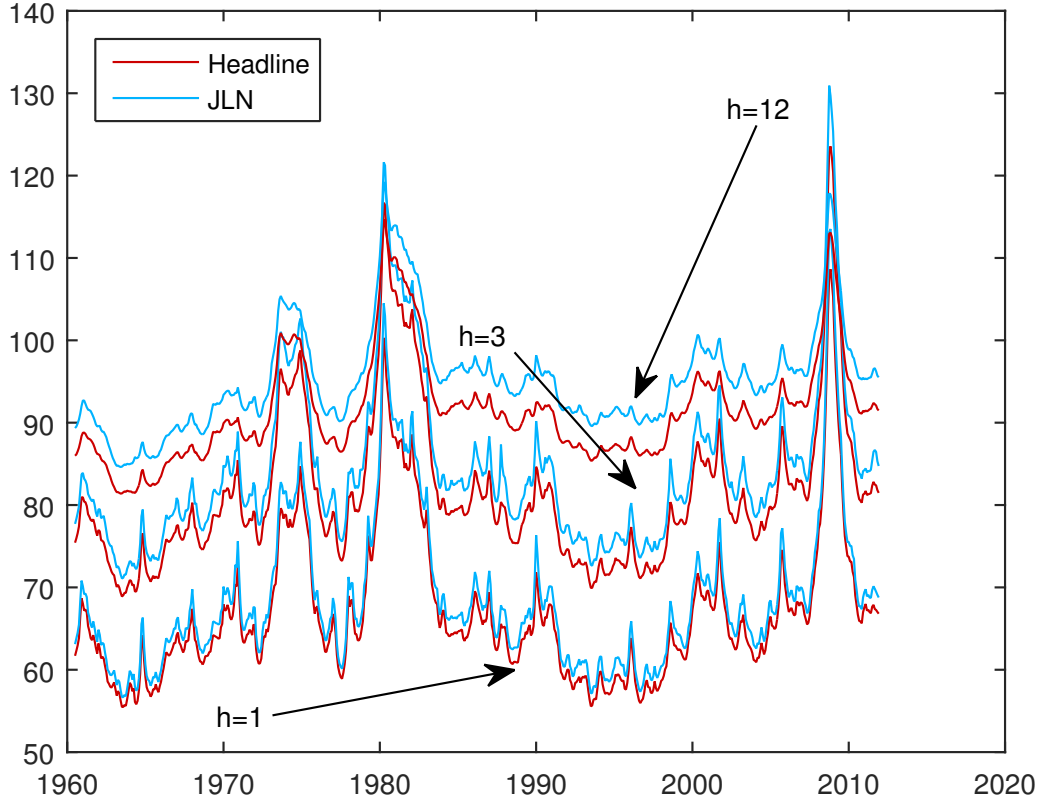
once. This multi-step sampler turns out to have lower autocorrelation and converge much faster that a single-step sampler of Jacquier, Polson & Rossi, 1994. The small approximation error that emerges can be corrected with importance weights. Kastner & Fruhwirth-Schnatter (2014) further enhance sampling efficiency by implementing the ancillarity-sufficiency (ASIS) strategy (Yu, Meng, 2010) of mixing centered ($\sim N(0, exp(h))$) and non-centered ($\sim N(\gamma + \phi(h - \gamma), \sigma_\eta)$) for the SV model (see Appendix II for the MCMC procedure).

As the error in the measurement equation has to be approximated with a mixture distribution, the model can be readily extended for the heavy-tail innovations in measurement equation (Chib et. al., 2002). This is due to the standard result of representing t-distribution with a mixture of normal and inverse-gamma (IG) distributions. y_t becomes $y_t = \sqrt{\lambda_t}\sigma_t\varepsilon_t$ where $\sqrt{\lambda_t}$ is an IG random variable with $(\nu/2, \nu/2)$ parameters. Hence, the mixture becomes a t with ν degrees of freedom (a standard practice is to assume ν to be uniformly distributed on the [2,

128] interval). Kastner (2015) re-implements the ASIS strategy for the heavy-tailed model.

Figure 2 provides estimates of persistent and Jurado, Ludvigson & Ng (2015) uncertainties. It is clear that (a) time-varying volatility of the forecast error is on average smaller in case of heavy-tail uncertainty, (b) heavy-tail uncertainty index is more persistent than the JLN one, (c) the longer is the forecasting horizon, the more pronounced these differences are. Table 1 reports the relevant sample statistics with standard deviation and kurtosis lower, persistence higher in case of the heavy-tail metric.

After the persistent uncertainty is estimated, I construct the burst component. In the simplest case, I proxy it with a simple difference of the normal and $\sim t$ uncertainty indices given that data and forecasting specification are identical expect for the assumption on the distribution of innovations in the forecasting equation. In substance, the burst component is the part of the normally distributed index that results from it being sensitive to non-persistent shocks of

Figure 3: Burst uncertainty index for various forecast horizons
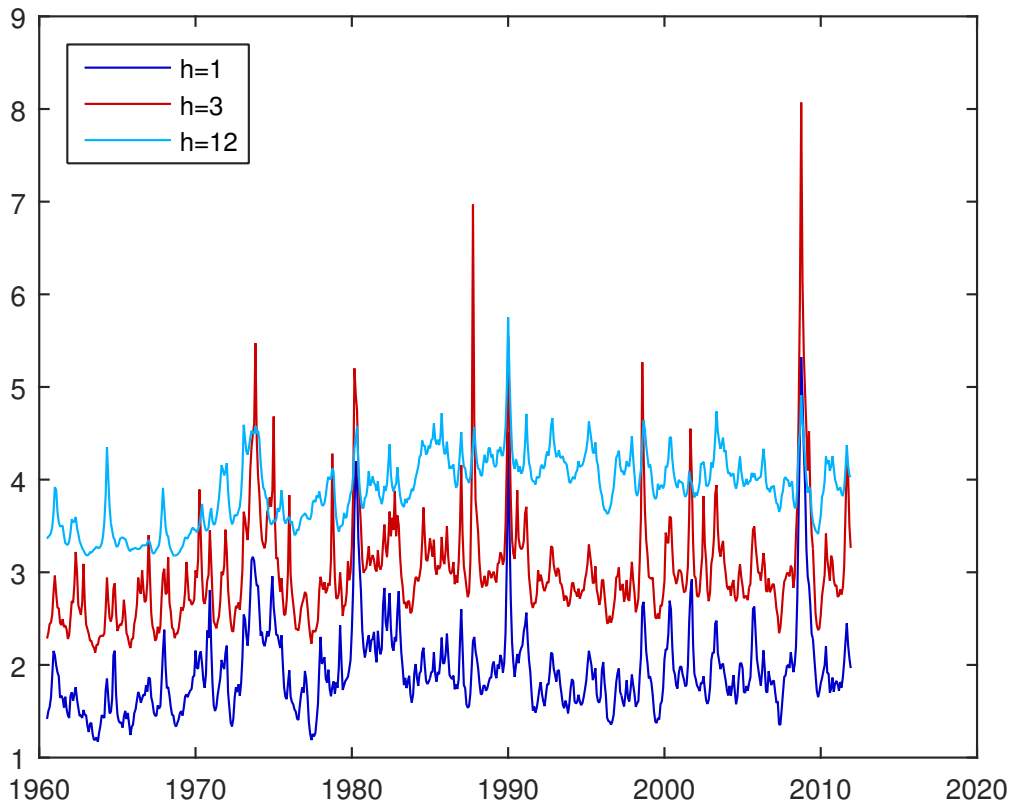
the forecast error. For robustness check I re-estimate the burst component with averaged draws of the time-varying λ_t parameter that 'downscales' sigma in the t-distributed case vs. N-case.

Note, that contrary to the original JLN and persistent metrics, this component is highly volatile: it jumps (with the most prominent spike in 2008) and recedes almost instantaneously. I use the burst uncertainty as a proxy for ex-ante very short-run, soon-resolved uncertainty spikes. In case of Bloom (2009) model, such shocks would've induced a) wait-and-see behavior that causes downturn but b) as it recedes very quickly, c) the downturn is followed by overshoot and decay of the shock's effect.
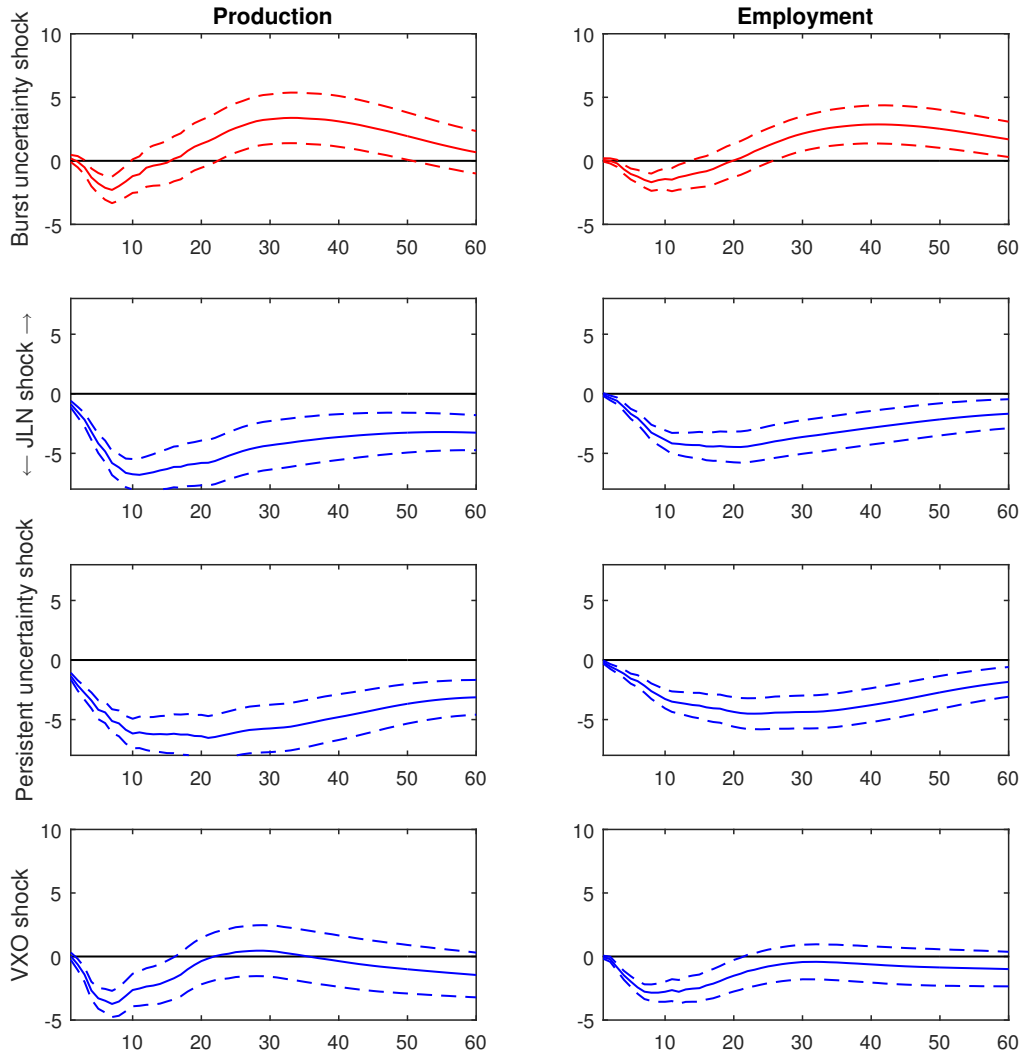
4   Effects of uncertainty

I estimate effects of the persistent and burst components to study three theoretical effects of Bloom (2009) or Bloom et. al. (2014) models. First, that uncertainty shocks have significant real effects. Second, that these effects realize in a particular pattern: downturn followed by overshoot. Third, that more persistent shocks result in this pattern having lower frequency and larger amplitude.

I use the constructed persistent and volatile components in a VAR setting with essentially the same ordering as in JLN and Bloom. I also compare effects of shocks to the persistent and volatile proxies with that of the original JLN metric and to the VXO index. To make results comparable, simple VARs are used with essentially the same ordering and lag structure (12) as in these studies. IRFs show response to the 1 s.d. shock and during the 60-month (5-year) period. As I decompose uncertainty index into two components, the ordering is updated:

$$
\begin{bmatrix}
\log(\text{stock}) \\
\text{single uncertainty proxy} \\
\text{FFR} \\
\log(\text{wages}) \\
\log(\text{CPI}) \\
\text{hours} \\
\log(\text{employment}) \\
\log(\text{industrial production})
\end{bmatrix}
\implies
\begin{bmatrix}
\log(\text{stock}) \\
\text{uncertainty (burst)} \\
\text{uncertainty (persistent)} \\
\text{FFR} \\
\log(\text{wages}) \\
\log(\text{CPI}) \\
\text{hours} \\
\log(\text{employment}) \\
\log(\text{industrial production})
\end{bmatrix}
$$

Figure 4: IRFs to uncertainty shocks: forecast error based uncertainty indices (3-month horizon) and VXO index

Notes: Confidence intervals for 1 standard error.

Figure 4 presents IRFs of industrial production and employment to the 1 standard deviation shocks of uncertainty indices. First thing to notice is that both persistent and volatile uncertainty shocks negatively affect production and unemployment on impact. This effect is significant for all (1-12 months) forecasting horizons in case of persistent uncertainty shock, but significance of shock to the volatile component vanishes with increasing horizons.

Second, shock to the volatile burst-like uncertainty creates a downturn followed by an over-

shoot in real activity. JLN note, that a) they don't find overshooting behavior with their own metric and b) they find overshooting behavior using VXO metric that is not statistically significant. Using a separate, non-persistent metric I find short-term shock being able to create the overshoot that corresponds to the Bloom's one[2].

Third, persistent uncertainty shock results in a more prolonged and severe drop in both industrial production and employment, than shocks to other metrics including original JLN. This result fits into theoretical picture well: shock to persistent uncertainty index means, on the agents' side, that ex-ante uncertainty is likely to stay elevated in future. In response, wait-and-see behavior accumulates on production units which results in a more severe and prolonged downturn.

As is mentioned above, statistical significance of the overshoot depends on the forecast error horizon. In particular (see Figures 5 and 6 in Appendix I), the effect is *not* insignificant with very short and long horizons (say, 1 and 12 months). Yet in between, sudden non-persistent jumps in the 3-5-month forecast error volatility create both slowdown and overshoot. This result may suggest, that the horizon of the 'fog of the future' affects agents' decision; in particular, if agent suddenly looses forecasting power on the 3-5-month horizon he switches to the wait-and-see pattern, while if he experiences a temporary loss of forecasting power on long horizons, this doesn't cause significant wait-and-see today.

5   Conclusions

In this paper I've studied real effects of uncertainty shocks. Using time-varying volatility of the forecast error, two-part uncertainty proxy is constructed; it consists of the persistent and volatile, burst-like components. I've used these two metrics to empirically study three predictions of uncertainty models (Bloom, 2009, Bloom et. al. 2014): first, that uncertainty shocks can have real effects; second, that uncertainty shocks affect real variables in a downturn/overshoot patter; third, that persistence of the shock decreases frequency and increases amplitude of this pattern. In a simple VAR setting these predictions appear to be present and significant: both metrics negatively affect employment and production on impact, short-term bursts create the expected overshoot/downturn, while persistent shock instantiates a prolonged, severe downturn.

---

[2]This may also suggest that VXO/VIX-type of indices act like a noisy signal of burst, rather than slowly-moving uncertainty.
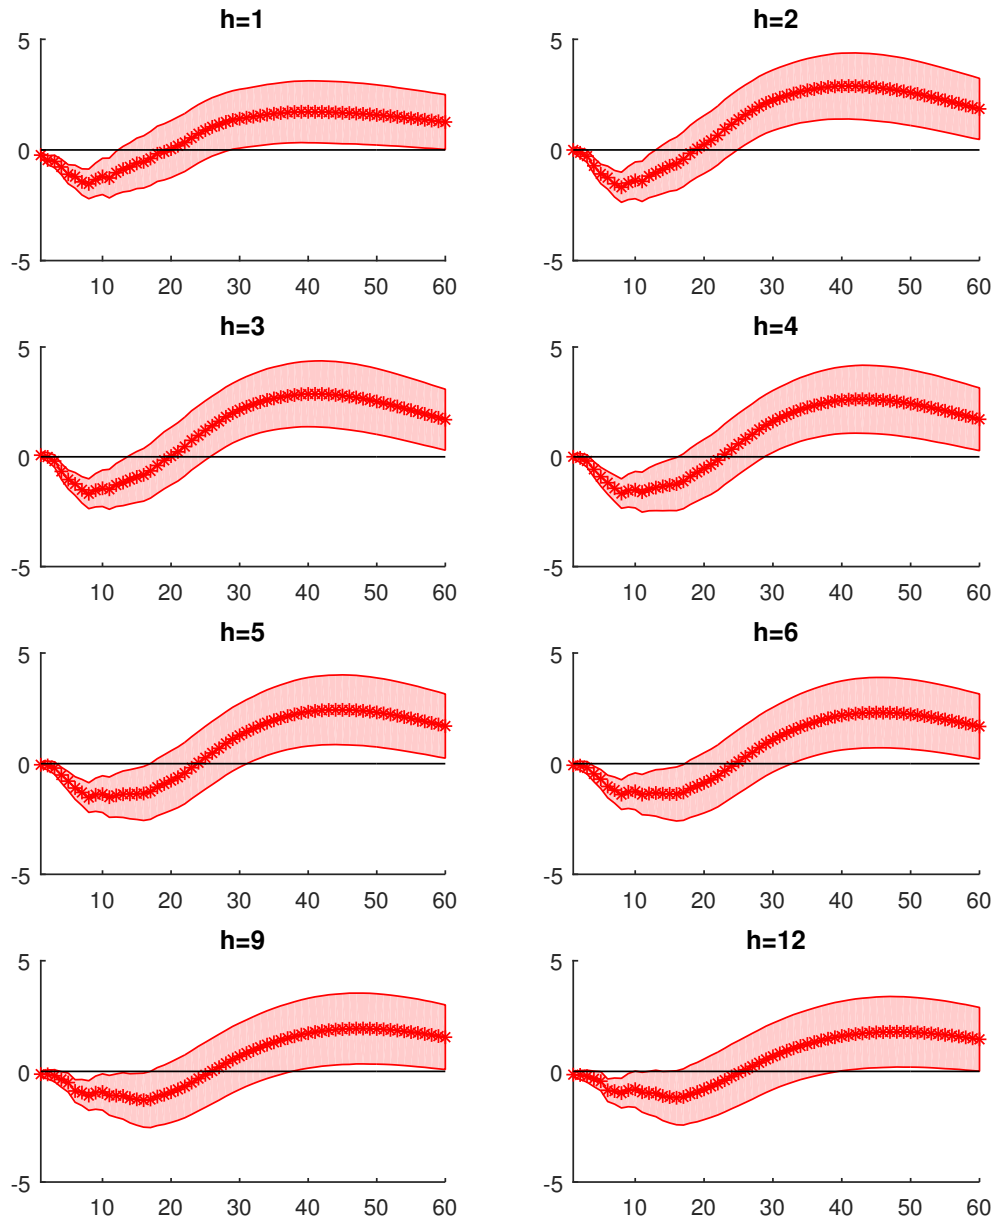
6   Appendix I

Figure 5: Response of **industrial production** to burst uncertainty shock for various forecast error horizons (h months)
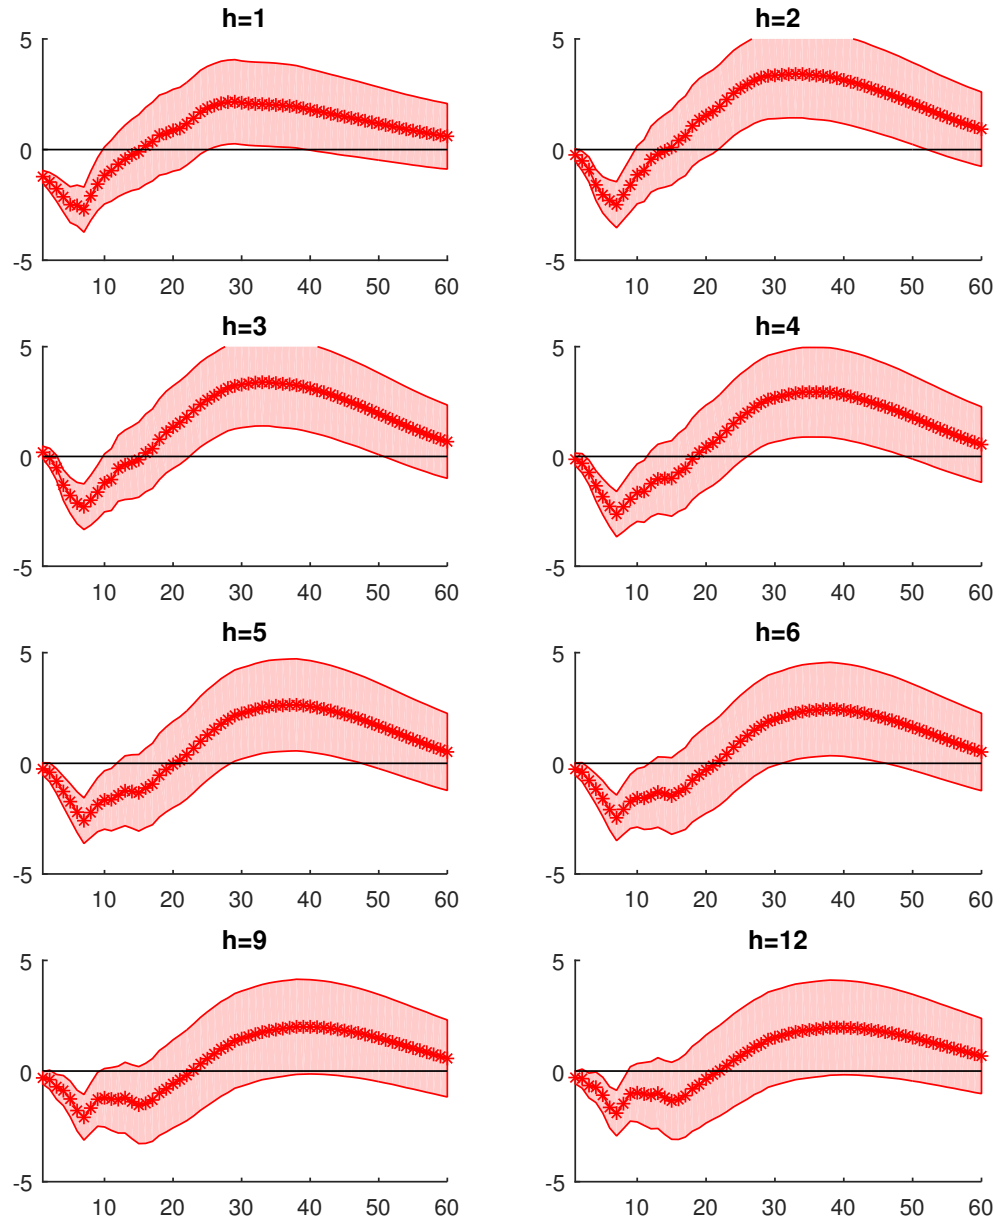
Notes: Confidence intervals for 1 standard error.

Figure 6: Response of **employment** to burst uncertainty shock for various forecast error horizons (h months)

Notes: Confidence intervals for 1 standard error.

7   Appendix II

MCMC procedure generally follows that of Chib et. al. (2002) but is updated in the fashion of Kastner & Fruhwirth-Schnatter (2014) and Kastner (2015). The idea is to extend the Kim et. al. (1998) MCMC algorithm by adding one more Gibbs-sample block that draws ν and λ conditional on latent volatilities and data. Drawing ν is not an easy task, at even conditional is not available in closed form, so Metropolis-Hasting is used at this step. In particaular, the algorithm takes form of:

- Draw vector of volatilities h s.t. parameters θ, data y and indicator r that sets which of 10 Gaussian (Omori et. al., 2007) from the mixture acts in this run of Gibbs sampler in the Non-Centered model parametrization. Note that s.t. r the model becomes linear-Gaussian so Kalman filter is used to recover latent states - volatilities.

- Draw parameters θ given data y, indicator r and volatilities h in the Non-Centred parametrization with a Bayesian regression. Posterior is not available in the closed form so Metropolis-Hastings is needed in this step within a general Gibbs sampler.

- Switch to the Centered parametrization.

- Draw parameters θ given data y, indicator r and volatilities h in the Centred parametrization. ϕ is drawn with Metropolis-Hastings while μ and σ can be drawn by running a Gibbs between conditionals.

- Switch to the Non-Centered parametrization.

- Draw the r indicator of the currently active Gaussian.

8 References

[1] Rudiger Bachmann, Steffen Elstner, and Eric R. Sims. Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics*, 5(2):217–49, April 2013.

[2] Jushan Bai. Inferential Theory for Factor Models of Large Dimensions. *Econometrica*, 71(1):135–171, January 2003.

[3] Jushan Bai and Serena Ng. Determining the Number of Factors in Approximate Factor Models. *Econometrica*, 70(1):191–221, January 2002.

[4] Scott R. Baker, Nicholas Bloom, and Steven J Davis. Measuring Economic Policy Uncertainty. CEPR Discussion Papers 10900, C.E.P.R. Discussion Papers, October 2015.

[5] Nicholas Bloom. The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685, 05 2009.

[6] Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry. Really Uncertain Business Cycles. CEP Discussion Papers dp1195, Centre for Economic Performance, LSE, March 2014.

[7] Nicholas Bloom and John Van Reenen. Why Do Management Practices Differ across Firms and Countries? *Journal of Economic Perspectives*, 24(1):203–24, Winter 2010.

[8] Charles S. Bos. Relating Stochastic Volatility Estimation Methods. Tinbergen Institute Discussion Papers 11-049/4, Tinbergen Institute, March 2011.

[9] Siddhartha Chib, Federico Nardari, and Neil Shephard. Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2):281–316, June 2002.

[10] Valentina Colombo. Economic policy uncertainty in the US: Does it matter for the Euro area? *Economics Letters*, 121(1):39–42, 2013.

[11] Jesus Fernandez-Villaverde, Pablo Guerron-Quintana, Juan F. Rubio-Ramirez, and Martin Uribe. Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review*, 101(6):2530–61, October 2011.

[12] Mario Forni, Marc Hallin, Marco Lippi, and Lucrezia Reichlin. The Generalized Dynamic-Factor Model: Identification And Estimation. *The Review of Economics and Statistics*, 82(4):540–554, November 2000.

[13] Luigi Guiso and Giuseppe Parigi. Investment And Demand Uncertainty. *The Quarterly Journal of Economics*, 114(1):185–227, February 1999.

[14] Andrew Harvey, Esther Ruiz, and Neil Shephard. Multivariate Stochastic Variance Models. *Review of Economic Studies*, 61(2):247–64, April 1994.

[15] Eric Jacquier, Nicholas G. Polson, and P.E.Peter E. Rossi. Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics*, 122(1):185–212, September 2004.

[16] Kyle Jurado, Sydney C. Ludvigson, and Serena Ng. Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216, March 2015.

[17] Gregor Kastner. Heavy-tailed innovations in the R package stochvol. R package vignette, CRAN, 2015.

[18] Gregor Kastner and Sylvia Fruhwirth-Schnatter. Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis*, 76(C):408–423, 2014.

[19] Sangjoon Kim, Neil Shephard, and Siddhartha Chib. Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–93, July 1998.

[20] Hening Liu and Jianjun Miao. Growth uncertainty, generalized disappointment aversion and production-based asset pricing. *Journal of Monetary Economics*, 69(C):70–89, 2015.

[21] Gabriela Nodari. Financial regulation policy uncertainty and credit spreads in the US. *Journal of Macroeconomics*, 41(C):122–132, 2014.

[22] Yasuhiro Omori, Siddhartha Chib, Neil Shephard, and Jouchi Nakajima. Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140(2):425–449, October 2007.

[23] James H Stock and Mark W Watson. Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business & Economic Statistics*, 20(2):147–62, April 2002.

Sergey Egiev

Department of Economics, Higher School of Economics, Moscow. Email: s.egiev@hse.ru