Detecting time variation in the price puzzle: an improved prior choice for time varying parameter VAR models

Reusens P, Croux C.

Detecting time variation in the price puzzle: An improved prior choice for time varying parameter VAR models

Peter Reusens^{a,*}, Christophe Croux^a

^a*Faculty of Economics and Business, KU Leuven, Belgium.*

Abstract

This paper compares Bayesian estimators with different prior choices for the time variation of the coefficients of Time Varying Parameter Vector Autoregression models using Monte Carlo simulations. Since the commonly used prior choice only allows for a tiny amount of time variation, less informative priors are proposed. Additional empirical evidence on the time varying response of inflation to an interest rate shock is provided for USA. While a major and statistically significant ‘price puzzle’ is detected for the period 1972-1979, the estimated response of inflation to an interest rate shock is negative for most other time periods.

Keywords: Inverse Wishart prior, Monte Carlo simulation, Price puzzle, Time varying parameter, Vector autoregression

1. Introduction

In order to account for changing macroeconomic relationships, Time Varying Parameter Vector Autoregression (TVP VAR) models have been introduced to relax the assumption of time invariance in the vector autoregression coefficients. TVP VAR models typically impose a random walk assumption on the coefficients, where the covariance matrix of the innovations of the coefficients is denoted by Q and controls the amount of time variation in the coefficients. Our paper is the first Monte Carlo exercise that compares Bayesian estimators with different prior choices for this ‘time variation parameter’ Q . We advance the usage of estimators with more uninformative prior choices for Q . In a three variable empirical application, we apply these estimators to study the time varying response of inflation to an interest rate shock in the USA.

*Corresponding author: Faculty of Economics and Business, KU Leuven, Naamsestraat 69 B-3000 Leuven Belgium, Tel: +32 16 326728.

Email addresses: peter.reusens@kuleuven.be (Peter Reusens), christophe.croux@kuleuven.be (Christophe Croux)

For typical macroeconomic data, we show that the prior choice for the time variation parameter Q proposed by Primiceri (2005) and frequently used in empirical research, only allows for a tiny amount of time variation in the coefficients. In other words, we argue that the posterior estimate for the time variation parameter Q of these papers is almost fully determined by the prior and that it is nearly unaffected by the true time variation of the data generating process. Therefore, the tiny amount of estimated time variation reported in these papers does not inform on the true time variation present in the coefficients. As a motivation for his strict prior choice against time variation, Primiceri (2005) stated that a less informative prior would lead to overestimation of Q . But, we show that less informative priors provide estimators with lower bias and lower mean squared error. As a second motivation for his prior choice, Primiceri (2005) states that the estimated time variation should be small for the model to have good predictive power, in line with the study performed by Stock and Watson (1996) on many bivariate macroeconomic time series. However, D’Agostino *et al.* (2013) reports that his model with substantial time variation in the coefficients has good forecasting performance for US unemployment, inflation and interest rate. Also, a forecasting exercise performed on our simulated data shows that the TVP VAR models with substantial time variation perform well. Finally, a third motivation for using the Primiceri (2005) prior could be to allow for some small amount of time variation in the coefficients, while not being interested in estimating the time variation parameter Q . We argue that it then would be more transparent to set Q to a prespecified small value rather than giving the false impression that the posterior estimate of the time variation is driven by the data.

We are the first to perform a detailed simulation study on the estimation of the time variation parameter Q in TVP VAR models. We compare the performance of Bayesian estimators with different prior choices on simulated data for both the univariate local level model setting and a three variable TVP VAR model setting. As the true amount of time variation in macroeconomic data differs between macroeconomic applications, we compare the performance of the estimators over the range of values of the time variation that is typically found in macroeconomic data. In this way, our simulation exercise allows us to select the estimator that performs well across this range of data configurations, avoiding to select an estimator that coincidentally does well for one particular setting. A related study is Korobilis (2014), who compares different prior choices for Q using a forecasting exercise on one specific USA macroeconomic dataset. Other related studies are the online appendix of Baumeister and Peersman (2013) and Nakajima (2011), who both perform a simulation study to evaluate whether a structural break data generating process can be fitted well by the time varying parameter regression model.

A second contribution of the paper is that we use the estimator of the TVP VAR model with our improved prior choices to bring additional empirical evidence on the question how the response of inflation to an interest rate shock varies over time in the USA. Starting with Sims (1992) and Eichenbaum (1992), the empirical research has often encountered a ‘price puzzle’ in this relationship, in the sense that a positive interest rate shock is followed by a sustained inflation rate rise (Rusnak *et al.*, 2013). Recently, several researchers have investigated whether the presence of this price puzzle varies over time. First, using an estimated TVP VAR model, Primiceri (2005) reports a small and statistically insignificant price puzzle which almost does not change over time. We argue that his near time invariant response of inflation is an artefact of his strict prior choice against time variation. Second, estimating a standard VAR model on two different subsamples, Boivin and Giannoni (2006), Barth and Ramey (2002), Castelnuovo and Surico (2010) and Hanson (2004) report that the price puzzle only occurs for the subsample before 1979. An advantage of using TVP VAR over constant VAR applied on pre-defined subsamples is that we allow for continuous time variation and that our results are not dependent on the sometimes arbitrary choice of the subsample boundaries. Our main finding is that the presence of the price puzzle is predominantly associated with the period 1972-1979.

In addition to the TVP VAR model with constant volatility used by for instance Cogley and Sargent (2001) and Koop and Korobilis (2010), a stochastic volatility extension was introduced to allow for a time varying conditional covariance matrix of the series. Estimators for these TVP VAR models with stochastic volatility have been developed by Primiceri (2005) and Cogley and Sargent (2005) and they have been applied in macroeconomic applications by Clark and Ravazzolo (2014), D’Agostino *et al.* (2013), Mumtaz and Sunder-Plassmann (2013) and Sargent and Surico (2011) among others. As the focus of the paper is on the estimation of the time variation parameter Q , this paper mostly abstracts from the stochastic volatility and it estimates the TVP VAR model with constant volatility. Our robustness check including stochastic volatility shows that our main conclusions are transferable to the setting with stochastic volatility.

Our paper is organized as follows. Section 2 presents the time varying parameter model and the estimators with the different prior choices. Afterwards, Section 3 discusses the simulation setup both for the univariate local level model and the three variable TVP VAR model. The results of these simulation exercises are presented in Section 4. Next, Section 5 discusses the empirical application in which we estimate the time varying response of inflation to an interest rate shock using a three variable TVP VAR model for the USA. Finally, Section 6 concludes our findings.

2. Methodology

2.1. Time Varying Parameter Vector Autoregression (TVP VAR) model

We use the homoscedastic version of the model of Primiceri (2005) as discussed in Koop and Korobilis (2010):

$$y_t = X_t B_t + u_t \quad u_t \sim N(0, \Sigma) \quad (1)$$

$$B_t = B_{t-1} + \nu_t \quad \nu_t \sim N(0, Q) \quad (2)$$

where $X_t = I_N \otimes [1, y'_{t-1}, \dots, y'_{t-p}]$, y_t and u_t are a $N \times 1$ vectors, N is the number of variables, p is the lag length, B_t and ν_t are a $K \times 1$ vectors of the coefficients with $K = N(Np + 1)$. u_t and ν_t are independently and normally distributed innovations, respectively with covariance matrix Σ and Q . The state equations are modelled as random walks, which involves permanent shifts in the coefficients and which limits the number of parameters to be estimated. Moreover, we assume that the B_t coefficients satisfy the stability criterion for a stable VAR, in line with Cogley and Sargent (2005). Finally, we will use the ‘mean of the diagonal of Q ’ as a measure for the amount of time variation in the coefficients.

2.2. Prior choices for the TVP VAR parameters

The Bayesian procedures use priors of the general form

$$B_0 \sim N(\hat{B}_{OLS}, 4 * Cov(\hat{B}_{OLS})) \quad (3)$$

$$\Sigma \sim IW(1 + N, I_N) \quad (4)$$

$$Q \sim IW(df, scale) \quad (5)$$

where IW is the inverse Wishart distribution, $\hat{B}_{OLS}$ and $Cov(\hat{B}_{OLS})$ are OLS estimates on a training sample of the first 40 observations and df and $scale$ respectively are the degrees of freedom and scale parameter of the inverse Wishart distribution.¹

In this paper, we compare Bayesian estimators, i.e. the mean of the posterior distribution, between prior choices that differ with respect to the degrees of freedom and the scale of the inverse Wishart

¹For the inverse Wishart distribution of the $K \times K$ dimensional Q to be a proper prior, the degrees of freedom should be larger than $K - 1$ (Muirhead, 1982). Note that for the univariate case, the inverse Wishart distribution with scale equal to a and degrees of freedom equal to b corresponds to an Inverse Gamma distribution with a scale parameter equal to $a/2$ and a shape parameter equal to $b/2$.

prior for Q . Table 1 presents these different prior choices. Estimator 1 uses a prior of the form $IW(K-1+\epsilon, \epsilon I_K)$, where ϵ is chosen to be 0.00001. In the univariate setting, this type of prior has been used in the Bayesian literature as a non-informative prior for the variance parameter (Gelman, 2006). Estimators 2, 3 and 4 have scale parameters of the form $k_Q^2 * df * Cov(\hat{B}_{OLS})$, with k_Q equal to 0.01 and the degrees of freedom equal to ‘ $K-1+0.1$ ’, ‘ $K-1+2$ ’ or ‘ $K-1+20$ ’, respectively. Estimator 4 equals the prior choice proposed by Primiceri (2005) for the ‘three variable TVP VAR with two lags’ setting.

Table 1: Overview of the different prior choices for the inverse Wishart distribution for Q (Equation 5). The first, second and third column show the name of the estimator, the degrees of freedom and the scale of the inverse Wishart prior distribution, respectively.

Estimator name	df	scale
1. df= $K-1+0.00001$, scale= 0.00001	$K - 1 + 0.00001$	$0.00001 I_K$
2. df= $K-1+0.1$, $k_Q=0.01$	$K - 1 + 0.1$	$0.01^2 (K - 1 + 0.1) Cov(\hat{B}_{OLS})$
3. df= $K-1+2$, $k_Q=0.01$	$K - 1 + 2$	$0.01^2 (K - 1 + 2) Cov(\hat{B}_{OLS})$
4. df= $K-1+20$, $k_Q=0.01$ (Primiceri, 2005)	$K - 1 + 20$	$0.01^2 (K - 1 + 20) Cov(\hat{B}_{OLS})$

An important feature of the estimators 2, 3 and 4 is that the scale parameter of the prior of Q is dependent on $Cov(\hat{B}_{OLS})$, the variance of the estimated OLS coefficient estimated on the training sample. One motivation of using such a prior is that the prior information has the same scale as the likelihood information, which is a similar motivation to the use of an empirical Bayes prior and g-prior in classical non time varying models (Koop and Potter, 2004). Another motivation is that coefficients of the data generating process who exhibit less time variation can be more precisely estimated by a non-time varying VAR model. Therefore, it can be reasonable to impose a smaller scale parameter for the prior distribution of these coefficients. However, a disadvantage of this prior choice is that it is data dependent and that the prior can be different even if the data generating process is the same.

Next, we compare the different prior choices according to two aspects. First, we assess whether the prior is vague by analyzing the difference between the 99th and 1st percentile of the prior. Second, we assess whether the prior has some support for the very low values of Q by evaluating the 1st percentile of the prior. This is important as the posterior estimate can be substantially influenced by the steep descent of the inverse Wishart prior around zero (Harvey *et al.*, 2007). Table 2 shows the 1st and 99th percentiles of the mean of the diagonal elements of Q for the different prior distributions both for the

univariate local level setting and for the three variable TVP VAR setting with two lags.²

Table 2: The 1st and 99th percentiles of the prior distribution of the mean of the diagonal elements of Q for the different prior choices are shown for the univariate local level setting in the top table and for the three variable TVP VAR setting with 2 lags in the bottom table.

(a) Univariate Local level setting	df	scale	0.01	0.99
1. df=0.00001,scale=0.00001	0.00001	0.00001	>10	>10
2. df=0.1,kQ=0.01	0.1	3e-08	1e-08	>10
3. df=2,kQ=0.01	2	6e-07	7e-08	3e-05
4. df=20,kQ=0.01	20	6e-06	2e-07	7e-07
(b) Three variable TVP VAR setting	df	scale	0.01	0.99
1. df=20.00001,scale=0.00001	20.00001	0.00001	>10	>10
2. df=20.1,kQ=0.01	20.1	4e-05	7e-05	>10
3. df=22,kQ=0.01	22	5e-05	1e-05	2e-03
4. df=40,kQ=0.01	40	9e-05	4e-06	6e-06

A comparison of the priors 2, 3 and 4 shows that a smaller the degrees of freedom corresponds to a larger relative difference between the 99th and 1st percentile, meaning that the prior is less informative. A lower degrees of freedom of an inverse Wishart distribution corresponds to a less informative prior as shown by Rossi and Allenby (2003): while the 1st and 99th percentile are very close to each other for estimator 4, they are more distant for estimator 3 and especially estimator 2. Next, one can choose parameters for the inverse Wishart distribution such that the prior still has support for the very small values. In particular, the 1st percentile of the priors 2, 3 and 4 are small enough such that also very small values of Q can be estimated accurately. Finally, prior 1 is a vague prior and it also has some support for the small values of Q . In sum, while the priors of estimators 1, 2 and to a lesser extent also estimator 3, have mass for both the very low values and the large values of the time variation coefficient, the prior of estimator 4 is very strict around a tiny value of Q .

2.3. Estimation of the TVP VAR model

A Gibbs sampler algorithm is used which sequentially draws from the conditional distributions $p(B^T|\Sigma, Q)$, $p(Q|B^T, \Sigma)$ and $p(\Sigma|B^T, Q)$, where B^T is the $K \times T$ vector of the coefficients B_t for all time periods. Draws from $p(B^T|\Sigma, Q)$ are performed by the algorithm of Carter and Kohn (1994) (details can be found in Primiceri (2005)) and draws from $p(Q|B^T, \Sigma)$ and $p(\Sigma|B^T, Q)$ are performed

²For the univariate local level setting, the mean of the diagonal elements of Q is evidently equal to Q since Q is unidimensional. For the three variable TVP VAR setting, the quantiles of the mean of the diagonal of the prior distribution of Q are obtained by simulating from the inverse Wishart distribution. For the prior choices 2, 3 and 4, we choose the median value of $Cov(\hat{B}_{OLS})$ over the simulation scheme of Section 3 as the value for $Cov(\hat{B}_{OLS})$ in the formula for the scale parameter.

by sampling from the inverse Wishart densities

$$\Sigma|B^T, Q \sim IW\left(1 + N + T, I_N + \sum_{t=1}^T (y_t - X_t B_t)(y_t - X_t B_t)'\right) \quad (6)$$

$$Q|B^T, \Sigma \sim IW\left(df + T, scale + \sum_{t=1}^{T-1} (B_{t+1} - B_t)(B_{t+1} - B_t)'\right). \quad (7)$$

In the simulation study, we use 2000 draws with a burn-in of length 1000 and a thinning factor of 4 for the local level setting. For the three variable TVP VAR model, we use 5000 iterations of the Gibbs sampler where we discard the first 2000 iterations as a burn-in. Then the posterior mean is used as the Bayesian estimator. For the local level model only, we compute the maximum likelihood estimator as a benchmark estimator.

For the estimation of the TVP VAR, we restrict the coefficients B_t to a stable VAR at each time period by using the algorithm developed by Cogley and Sargent (2005) and used by Mumtaz and Sunder-Plassmann (2013) among others. In particular, we only accept a proposed MCMC draw of the coefficients if the coefficients are stable for every time period and we reject the MCMC draw when the stability criterion is violated for at least one time period.³

3. Simulation setup

We perform a Monte Carlo study to compare the estimators with different prior choices for the estimation of the time variation. Section 3.1 and Section 3.2 discuss the data generating process for the univariate local level setting with $K = 1$ and for a three variable TVP VAR setting with $K = 21$, respectively. The time variation of these data generating processes ranges over values that are typically found in macroeconomic data.

3.1. Data generating process: the local level model setting

We simulate data from the univariate local level model

$$y_t = B_t + u_t \quad u_t \sim N(0, \Sigma) \quad (8)$$

$$B_t = B_{t-1} + \nu_t \quad \nu_t \sim N(0, Q) \quad (9)$$

³The dlm package of Petris (2010) in R is used for the estimation of the univariate local level model and an adapted version of the Matlab code of Koop and Korobilis (2010) is used for the estimation of the three variable TVP VAR model. As discussed in Koop and Potter (2011), it can take many iterations to accept a single draw of the Gibbs sampler, implying that the algorithm with stability conditions on the coefficients can be very computer intensive.

where y_t is a univariate observed variable, B_t is a time varying intercept, u_t and v_t are independently distributed innovations, Σ is the variance of the noise innovations and Q is the variance that governs the amount of time variation in B_t . Note that this local level model is the simplest setting of the TVP VAR model of Section 2.1 with dimension of y_t equal to one and lag length p equal to zero. For each simulation design with a certain Q value, we simulate 200 time series.

We want to generate simulated data similar to real macroeconomic data. Table 6 in Appendix surveys the time varying parameter literature and reports the estimated parameters for typical macroeconomic series expressed in annualized percentage growth rates. For each study, the final column shows the estimated λ parameter, defined as $\lambda = \sqrt{(Q/\Sigma)}T$, is a measure of the signal for time variation relative to the noise in the data. The other columns show: the sample size T , the amount of noise in the data Σ and the amount of time variation in the coefficients Q . First, the estimated λ for papers that use the ‘local level model’ (studies above the dotted line) ranges between 0 and 66. Second, for the more complex univariate models (studies below the dotted line), the estimated λ varies even much more and lies between 0 and 412. In our simulation exercise, we therefore consider 14 different simulation designs with λ values ranging between 0 and 150. By setting the variance of the noise Σ equal to 0.1, the initial value of B_t equal to 1 and the number of observations T equal to 159, this corresponds to Q ranging between 0 and 0.16. If we translate the time variation parameter Q to the 95th percentile of the absolute difference between B_{159} and B_1 , then this difference ranges between 0 and 9.8.⁴

3.2. Data generating process: the three variable TVP VAR model setting

We simulate from the TVP VAR model (1) and (2). For each simulation design with a certain Q value, we simulate 200 time series. We choose the number of variables N equal to 3 and the lag length p equal to 2, in line with the model used by Primiceri (2005) and Cogley and Sargent (2005). Followingly, the dimensions of the parameters in the model are 21×1 for B_t and ν_t , 3×3 for Σ and 21×21 for Q . Similar to the univariate local level setting of Section 3.1, we choose the number of observations T to be equal to 159 and the covariance matrix of the error terms Σ to be a scalar matrix with scalar 0.1. The initial value of B_t is chosen to be a zero vector. Q is chosen to be a scalar matrix implying that all random walk modelled coefficients B_t of our data generating process

⁴The 95th percentile of the absolute difference between B_{159} and B_1 can be easily computed as $\sqrt{159Q}z_{0.975}$, where $z_{0.975}$ is the 97.5th percentile of the standard normal distribution.

are independent and have the same amount of time variation. We analyze different simulation designs where the scalar of the Q matrix ranges between 0 and 0.002, which we believe to be representative for macroeconomic time series that are expressed in growth rates. The corresponding 95th percentile of the absolute difference between the coefficient B_{159}^i and B_1^i ranges then between 0 and 0.98. We do not consider larger values for the scalar of Q since they often result in non-stationary processes, which are not realistic for macroeconomic growth rate time series.

4. Simulation results

This section compares the performance of the Bayesian estimators with different prior choices for the estimation of the true time variation parameter Q of the TVP VAR model. Section 4.1 and Section 4.2 respectively presents the univariate local level setting and the three variable TVP VAR setting.

4.1. Results: the local level model setting

Figure 1 shows the median over the simulation runs of the estimated time variation parameter Q for the Bayesian estimators with the different prior choices of Table 2a and for the Maximum Likelihood estimator. The horizontal axis represents the different values of the true time variation parameter Q of the data generating process, which each corresponds to a different simulation design as discussed in Section 3.1. The performance of the different estimators can be assessed by comparing the vertical distance between the estimated Q and the 45 degree line.

The Bayesian estimators 1 (df=0.00001, scale=0.00001) and 2 (df=0.1, kQ=0.01), which have vague priors, perform very well over the entire range of values for the true time variation parameter Q . In contrast, Bayesian estimator 4 (df=20, kQ=0.01) largely underestimates the time variation for all Q values. In other words, the estimator is then not very dependent on the amount of time variation in the coefficients of the data generating process. Next, while estimator 3 (df=2, kQ=0.01) underestimates the time variation for the simulation designs with lower values of the true time variation parameter Q , it performs well for the simulation designs with larger values. Finally, our benchmark estimator 5 (the maximum likelihood estimator) also performs well across the different simulation designs. However, the distribution of the maximum likelihood estimator has a point mass at 0 when the true time variation is small, which is called the ‘pile-up problem’ (Aguiar and Martins, 2005; Primiceri, 2005; Stock and Watson, 1998). Because of this pile-up problem, we prefer the Bayesian estimators 1 and 2 over the maximum likelihood estimator. Next, each subfigure of Figure 8 in Appendix corresponds

Figure 1: The median over the 200 simulation runs of the five different estimators for Q of the local level model, as a function of the true value of Q . The different lines represent the four different estimators of Table 2a and the maximum likelihood estimator. The thick black 45 degree line maps the true Q value on the vertical axis. Note that the axes have a logarithmic scale and that zero values are shown by a split in the axes.

to one simulation design and shows the boxplot of the estimated Q of the local level model over the 200 simulation runs. In addition to the median over the simulation runs, which is already shown in Figure 1, these boxplots represent simplified visual representations of the distribution of the estimated Q over the simulation runs. For most simulation runs of each simulation design, the estimated Q of Bayesian estimators 1 and 2 and to a lesser extend also estimator 3 are relatively close to the true Q value of the data generating process. In contrast, for all simulation runs of the simulation designs with smaller Q values, the posterior estimates of Bayesian estimator 4 are very concentrated around a tiny value. Although for the simulation designs with larger Q values, the quartiles of these estimators become closer to the Q value of the data generating process, Q is still largely underestimated for several simulation runs. For the first nine simulation designs, the maximum likelihood estimator 5 has several estimates equal to zero due to the above discussed pile-up problem.

For the different simulation designs and for the Bayesian estimators, Figure 2 shows the simulated

Mean Squared Error (MSE) of the logarithm of Q , defined as⁵

$$MSE = \frac{1}{200} \sum_{s=1}^{200} (\log \hat{Q}_s - \log Q)^2, \quad (10)$$

where $\hat{Q}_s$ is the estimated Q for the simulation s . This MSE is another metric to evaluate the different estimators of Q . We prefer this MSE defined on the logarithm of Q because an underestimation and an overestimation of Q by a certain factor have the same impact on MSE. We observe that for all simulation designs, estimator 1 has the lowest MSE value, followed by estimator 2 and estimator 3. Estimator 4 has a much higher MSE values across all simulation designs.

Figure 2: Simulated mean squared error of the logarithm of the four different Bayesian estimators of Q , as a function of the true value of Q .

In summary, this simulation exercise for the univariate local level setting has shown that estimator 4 ($df=20, kQ=0.01$) substantially underestimates the time variation parameter Q and that for almost all simulation designs, estimators 1 ($df=0.00001, scale=0.00001$) and 2 ($df=0.1, kQ=0.01$) and estimator 3 ($df=2, kQ=0.01$) have lower bias and lower mean square error.

⁵Note that the simulation design with the true Q equal to zero is not shown as its MSE is plus infinity. Similarly, the MLE estimator is also not shown on the graph as its estimated MSE is often minus infinity due to the occurrence of too many zeros in the MLE estimate of Q .

4.2. Results: the three variable TVP VAR model setting

Figure 3 presents, for each estimator and as a function of the true value of Q , the median over the 200 simulation runs of the estimated amount of time variation, which we have defined in Section 2.1 as the mean of the diagonal of the estimated Q .⁶

Figure 3: The Median over the 200 simulation runs of the mean of the diagonal of the estimated Q of the three variable TVP VAR model, as a function of the true value of Q . The different lines represent the four different estimators of Table 2b. Finally, The thick black 45 degree line maps the true value for the mean of the diagonal of Q on the vertical axis.

The Bayesian estimators 1 (df=20.00001, scale=0.00001), 2 (df=20.1, kQ=0.01) and 3 (df=22, kQ=0.01) perform relatively well over the entire range of values for the true time variation parameter Q . Especially for the simulation designs with larger values for the scalar of Q , they clearly outperform estimator 4 (df=40, kQ=0.01), which largely underestimates the time variation. It is only for the two simulation designs with the lowest time variation that estimator 4 performs well. The reason is that its prior choice is very strict around a value close to the true value of Q in these two cases. Next, each subfigure of Figure 4 corresponds to one simulation design and shows the boxplot of the estimated time variation over the simulation runs, giving a simplified visual representations of the estimators. The estimated Q of Bayesian estimators 1, 2 and 3 are relatively close to the true Q value of the data generating process for most simulation runs. Next, for all simulation runs of all simulation designs, the

⁶Note that we exclude the elements of the diagonal of Q that correspond to the intercept coefficients as these are different in nature from the other VAR coefficients.

posterior estimate of Bayesian estimator 4 is very concentrated around a tiny value for the simulation designs with smaller Q values.

Figure 4: Each subfigure corresponds to one simulation design and shows the boxplot of the mean of the diagonal of the estimated Q of the three variable TVP VAR model over the 200 simulation runs. The title of each subfigure represents the scalar of the true value of Q , which is a scalar matrix. The horizontal axis represents the different estimators which are labelled 1 ($df=20.00001, scale=0.00001$), 2 ($df=20.1, kQ=0.01$), 3 ($df=22, kQ=0.01$) and 4 ($df=40, kQ=0.01$). Finally, the horizontal line is the true value for the scalar of Q of the data generating process.

Figure 5 shows the mean of the simulated Mean Squared Error, as defined in Equation (10) of the diagonal elements of the logarithm of Q for the different simulation designs and for the different estimators. In line with our discussion of Figure 3 and 4, estimators 1, 2 and 3 outperform estimator 4 for all simulation designs, except the one with the lowest value for the scalar of Q . While for the middle simulation designs, estimator 1 and 3 are best, estimator 2 does better for the largest Q values. Estimator 4 only performs well for the simulation design for which the amount of time variation Q is small, corresponding to the value imposed by its prior.

Figure 5: Simulated mean squared error of the diagonal of the logarithm of the four different estimators of Q , as a function of the true value of Q .

Also, we have performed a forecasting comparison. For the different simulation designs and for the different estimators, we simulate the Mean Squared Forecast Error (MSFE) at horizon 1

$$MSFE = \frac{1}{200 \times 3} \sum_{s=1}^{200} \|y_{T+1,s}^o - \hat{y}_{T+1,s}\|^2 \quad (11)$$

$$\hat{y}_{T+1,s} = X_{T+1} \hat{B}_T \quad (12)$$

where $y_{T+1,s}^o$ is value of the series in the s^{th} simulation at time $T + 1$, $\hat{y}_{t+1,s}$ is its forecast and $\hat{B}_T$ is the posterior mean of the coefficients B_T at the last observation of the sample. The top row of Table 3 shows the MSFE of the forecast that uses the true value of the parameters B_T rather than its estimate as a benchmark. The other rows of the table show the ‘Relative Mean Squared Forecast Error’, defined as the ratio of the MSFE and the benchmark MSFE.

Table 3: The top row of the table presents the MSFE at horizon 1 of the forecast using the true value of the parameters B_T . The following rows show the relative MSFE at horizon 1 of the different estimators. The different columns correspond to different simulation designs.

Time variation	0e+00	2e-05	6e-05	1e-04	3e-04	4e-04	8e-04	2e-03
Benchmark MSFE True B_T	0.10	0.11	0.11	0.11	0.10	0.09	0.10	0.14
1. df=20.00001, scaleQ=0.00001	1.09	1.03	1.07	1.16	1.21	1.31	1.63	1.92
2. df=20.1, kQ=0.01	1.09	1.05	1.09	1.18	1.21	1.32	1.59	1.82
3. df=22, kQ=0.01	1.08	1.03	1.07	1.18	1.20	1.28	1.66	1.91
4. df=40, kQ=0.01	1.08	1.03	1.07	1.18	1.22	1.34	1.80	2.30

For most simulation designs, the benchmark MSFE approximates the scalar value 0.1 of the scalar error term covariance matrix Σ , because B_{T+1} is close to B_T . The relative MSFE of the estimators is worse for simulation designs with larger time variation, since it is more difficult to accurately estimate the coefficients B_T . While we observe that the forecast performance of all estimators is very similar for the simulation designs with low amount of time variation, estimators 1, 2 and 3 outperform estimator 4 for the simulation designs with larger amount of time variation.

Finally, Table 4 shows for each estimator and simulation design, the average number of rejected Gibbs draws, being Gibbs samples B^T that do not meet the stability criterion.⁷

Table 4: The average number of iterations to find a stable draw of the Gibbs sampler.

	0e+00	2e-05	6e-05	1e-04	3e-04	4e-04	8e-04	2e-03
1. df=20.00001,scaleQ=0.00001	0.00	0.00	0.00	0.01	0.04	0.30	2.18	10.15
2. df=20.1,kQ=0.01	0.01	0.01	0.01	0.02	0.12	1.32	4.66	13.85
3. df=22,kQ=0.01	0.00	0.00	0.00	0.00	0.08	0.62	3.06	10.11
4. df=40,kQ=0.01	0.00	0.00	0.00	0.00	0.05	0.08	0.94	1.50

The larger the estimated Q , the more the coefficients can vary over time and the more often iterations are proposed that do not meet the stability criterion. The average number of rejected draws is larger for estimators 1, 2 and 3 than for the prior choice 4 used in Primiceri (2005). Because his very strict prior against time variation implies that the posterior estimated time variation of the coefficients is very small, the probability that the estimated coefficients move into the non-stationary region is low, especially for a finite sample with only a few hundred observations.

Overall, this simulation exercise for the three variable TVP VAR setting has shown that for most simulation designs, estimators 1 (df=20.00001,scale=0.00001), 2 (df=20.1,kQ=0.01) and 3 (df=22,kQ=0.01) outperform estimator 4 (df=40,kQ=0.01) with respect to bias, mean squared error and out of sample forecast performance. For most simulation designs, estimator 4, which is proposed by Primiceri (2005), substantially underestimates the time variation parameter Q . This questions Primiceri (2005), who state that his prior choice on the time variation parameter Q is weak and does not have a lot of impact on the posterior estimate.

⁷We have discarded time series in the rare cases for which the stability conditions were not met in less than 30000 iterations.

5. Data application: the price puzzle

In this section, we estimate the time varying effect of an interest rate shock on inflation using a three variable TVP VAR model with two lags for the USA. We compare the results between the different estimators of Table 2b.⁸ We show that the ‘price puzzle’ phenomenon, which is a contractionary interest rate shock leading to a sustained rise in inflation, is predominantly associated with the time period 1972-1979. We use the same variables as Primiceri (2005): the annualized quarterly growth rate of a seasonally adjusted chain weighted GDP price index, the seasonally adjusted civilian unemployment rate and the seasonally unadjusted yield on the three-month Treasury bills. The data is obtained from the ‘Federal Reserve Economic Data’ database for the sample period from 1953Q1 until 2014Q1. Figure 6 shows the time plot of our data.

Figure 6: Time plot of inflation, unemployment rate and nominal interest rate for the USA.

We analyze the time varying response in inflation to a 1 percent shock in the interest rate. In particular, we show how our proposed prior choices change the results reported in Figure 2 of Primiceri (2005).⁹ Similar to Primiceri (2005), we study orthogonalized impulse response function where we make the identification assumption that interest rate shocks affect inflation and unemployment with at least one period of lag. For calculating impulse responses at each time t , we follow Primiceri (2005) and

⁸For the estimation of the TVP VAR model, we use 4000 iterations of the Gibbs sample where we discard the first 2000 iterations as a burn-in.

⁹We actually compare our results with Del Negro and Primiceri (2013), which uses a corrected version of the Gibbs sampler elaborated in Primiceri (2005) and has impulse response functions slightly different from Primiceri (2005).

Koop and Korobilis (2010) by considering the coefficients at time t as fixed for the entire response horizon. Hence, the response at time t informs on the response of the variables to shocks under the assumption that economic relationships remain the same as at time t . For each estimator, Figure 7 shows how the response of inflation to a one unit interest rate shock after one and four quarters, evolves over the sampling period. The full line is the median of the posterior distribution of the response and the dotted lines represent the 10th and 90th percentiles.

The response of inflation to a monetary shock of estimator 4 ('df=40,kQ=0.01'), which is used by Primiceri (2005), is almost time invariant.¹⁰ Primiceri (2005) considers this 'near time invariance' as evidence that the response of the economy to orthogonal interest rate shocks does not vary much over time. We believe that the almost constant responses over time are an artefact of the prior choice 'df=40,kQ=0.01' as this prior does not allow for much time variation in the coefficients B_t . Hence, we argue that this estimated responses does not inform on the time variation in the effect of orthogonal monetary policy shocks on the economy. For the entire sample period, the response of estimator 4 after one and four quarters shows a small and time invariant 'price puzzle' in the sense that a one percentage interest rate shock leads to an estimated rise in inflation of about 0.1 percent. As in Primiceri (2005), this price puzzle is not significantly different from zero at the 80% confidence level.

For estimators 1 ('df=20.00001,scale=0.00001'), 2 ('df=20.1,kQ=0.01') and 3 ('df=22,kQ=0.01'), the response of inflation to an interest rate shock shows much more time variation. We find a large and statistically significant price puzzle that is predominantly present for the time period 1972-1979. During this time period, the estimated inflation response after one quarter to an interest rate shock lies between 0.1% and 0.2% for estimator 1, between 0% and 0.3% for estimator 2 and between 0.2% and 0.4% for estimator 2. In addition, also the estimated response after four quarters during this time period remains positive but is only statistically significant for estimator 2. The price puzzle mostly disappears for the time periods before 1972 and after 1979: the estimated response of inflation to an interest rate shock becomes mostly negative. For instance, for the four quarter horizon, the estimated response mostly ranges between -0.3% and 0%.

Our findings on the presence of the price puzzle in the 1970s and on the absence of it in the period after 1979 are in line with Boivin and Giannoni (2006), Barth and Ramey (2002), Castelnuovo and Surico (2010) and Hanson (2004), who estimate vector autoregression models on both a subsample

¹⁰Similarly, the responses of unemployment and interest rate to an interest rate shock also did not show much time variation (Results are available upon request).

Figure 7: Time varying inflation response after (a) one and (b) four quarters to an interest rate shock with 10th, 50th and 90th percentiles.

for the time period before 1979 and a subsample for the period after 1979. Below, we discuss three possible explanations that can account for the time variation in the presence of the price puzzle.

- (i) Using an estimated regime switching DSGE model with changes in regimes between determinacy and indeterminacy, Belaygorod and Dueker (2009) reports that a price puzzle is present during the 1972-1981 indeterminacy regime. Many researchers have indeed found that the period before the appointment of former Federal Reserve chairman Paul Volcker in October 1979 was characterised by a passive monetary policy that led to indeterminacy (Clarida *et al.*, 2000; Cogley and Sargent, 2005). After 1979, the monetary policy became active in the sense that the nominal interest rate then responded more than proportionally to inflation changes and this consequently led to determinacy. Belaygorod and Dueker (2009) claim that in the indeterminacy regime of the 1970s, there was a self-fulfilling belief that interest rate shocks are cost-push shocks implying that the price puzzle is a genuine consequence of the indeterminacy regime, rather than a false finding. In contrast to the model of Belaygorod and Dueker (2009) however, the estimated DSGE model with indeterminacy of Castelnovo and Surico (2010) and Lubik and Schorfheide (2004) did not produce a price puzzle during the indeterminacy periods.
- (ii) The price puzzle in the period before 1979 can also be the result of a strong cost channel transmission channel of monetary policy during that period: if the inflationary impact of this cost channel is stronger than that of the demand channel, the inflation would rise following a monetary policy shock. Barth and Ramey (2002) and Tillmann (2009) have showed that this cost channel was strong before 1979 and became weaker afterwards. They argue that the financial deregulation, the absence of credit actions of the Federal Reserve, the change from a fixed to floating exchange rate regime and decreased financial frictions have contributed to this reduced impact of the cost channel after 1979. However, Castelnovo (2012) and Rabanal (2003) report that their New Keynesian DSGE models with a cost channel cannot produce the price puzzle.
- (iii) Many researchers have posited that the price puzzle is a false finding that arises due to the misspecification of the monetary policy shock. In particular, the omission of the variables 'expected future inflation' and 'potential output' is said to spuriously produce a price puzzle in the empirical VAR literature (Giordani, 2004; Sims, 1992). In addition, Carlstrom *et al.* (2009) show that the recursive identification assumption can create a price puzzle when this identification assumption is wrong. However, none of the proposed remedies for these misspecifications of the monetary policy shock have been able to fully solve the price puzzle (Demiralp *et al.*, 2014).

As a robustness check, we estimated the TVP VAR model with stochastic volatility.¹¹ Figure 9 in Appendix shows the time varying inflation response to an interest rate shock after one and four quarters. In line with the baseline model, the response for estimators 1, 2 and 3 shows that the price puzzle is predominantly associated with the period 1972-1979. Also, the response of estimator 4 shows again that the prior choice of Primiceri (2005) does not have substantial time variation.

Another observation is that for the recent period after 2008, the effect of an interest rate shock on inflation increases for most estimators. For estimator 2 and 3, the median response even becomes positive, which might suggest the return of the price puzzle. However, we caution with interpreting interest rate shocks after 2008Q4 as the short term interest rate was stuck at the zero lower bound and monetary policy during this period was more focused on influencing the long-term interest rates through unconventional policy measures such as large-scale asset purchases and forward guidance. As a robustness check, we replace the nominal interest rate by a ‘shadow rate’ as introduced by Wu and Xia (2014) for the ‘zero lower bound period’ after 2008Q4. This shadow rate is estimated using the forward interest rates and it is a measure for monetary policy that incorporates both the traditional nominal interest rate measure and the unconventional monetary policy measures. Figure 10 in Appendix shows the response of inflation to an interest rate shock for this new dataset after one and four quarters. The results for the period before 2008Q4 are very similar to our basecase responses of Figure 7 and confirms our finding that the price puzzle is predominantly associated with the time period 1972-1979. However, for the period after 2008Q4, we observe that for most estimates, the response stays negative and does not increase.

Finally, for each estimator, Table 5 shows the mean of the diagonal of the posterior mean of Q (excluding intercept components) (i) for the baseline model, (ii) for the TVP VAR model including stochastic volatility and (iii) for the model with the shadow interest rate. One sees that estimator 4 returns indeed a very tiny time variation: for the baseline model, the estimate is $6.6 * 10^{-6}$, corresponding to a value of only 0.07 for the 95th percentile of the absolute difference of the autoregressive coefficients over the period 1963Q2-2014Q1. For the other estimators 1, 2 and 3, this number ranges between 0.62 and 0.85, showing that time variation of the autoregressive coefficients is

¹¹In particular, the TVP VAR with stochastic volatility model of Primiceri (2005) is used, where we additionally impose the stability criterion on the coefficients, as in Section 2.1. We follow Del Negro and Primiceri (2013) for the estimation the TVP VAR model with stochastic volatility and we take 10000 iterations of the Gibbs sample with a burn-in of 2000 iterations. We again use the prior choices for Q of Table 2 and we follow Primiceri (2005) for the prior choices of the other parameters.

Table 5: Mean of the diagonal of the estimated Q for the baseline model, for a model including stochastic volatility and for a model with the shadow rate as the interest rate variable after 2008Q4.

Estimator	1	2	3	4
Baseline model	4.9e-04	9.2e-04	7.4e-04	6.6e-06
Stochastic volatility	2.4e-04	2.0e-03	1.9e-03	6.8e-06
Shadow interest rate	4.0e-04	1.6e-03	1.2e-03	6.2e-06

substantially larger. A similar finding holds for the model with stochastic volatility and for the model using the shadow interest rate.

6. Conclusion

In the recent macro-econometric literature, Time Varying Parameter Vector Autoregression (TVP VAR) models have often been applied to model time varying relationships between macroeconomic variables. This paper is the first Monte Carlo simulation study that compare Bayesian estimators with different prior choices for the time variation coefficient Q of these models. In particular, we conducted a simulation study both for the univariate local level model setting and a three variable TVP VAR model setting where the time variation of the different data generating processes ranges over values that are typically found in macroeconomic data.

Our main finding is that, both for the univariate local level setting and the three variable TVP VAR setting, estimator 4 ($df=K-1+20, kQ=0.01$) largely underestimates the amount of time variation in the VAR coefficients for most simulation designs because its prior choice is too strict around a very small value of the time variation. Unfortunately, starting with Primiceri (2005) for a three variable TVP VAR setting, this prior choice has been often used in empirical research. We advance the use of less informative priors for the time variation coefficient. In particular, we have shown that estimator 1 ($df=K-1+0.00001, scale=K-1+0.00001$), 2 ($df=K-1+0.1, kQ=0.01$) and 3 ($df=K-1+2, kQ=0.01$) have better performance for most simulation designs. Conveniently, as our proposed prior choices of Q remain inverse Wishart distributions, the MCMC estimation algorithm developed by Primiceri (2005) to estimate the TVP VAR model does not change.

Our improved prior choices are used to estimate the time varying effect of an interest rate shock on inflation using a three variable TVP VAR model for the USA. We detect considerable time variation in the impulse response function. In particular, we find that the ‘price puzzle’ phenomenon is predominantly associated with the period 1972-1979 and that the response of inflation to an interest rate shock is substantially negative for most other time periods. This finding is in line with empirical

evidence using non time varying VAR models applied on different sub-samples and empirical literature using an estimated DSGE model with indeterminacy. Our finding differs from Primiceri (2005), which near time invariance of its estimated inflation response to an interest rate shock is caused by his very strict prior against time variation.

Using a simulation exercise similar to this paper, future research could analyze the performance of different type priors such as the ‘data-based priors’ recently proposed by Korobilis (2014). In addition to the performance measures on simulated data, the forecast performance of different priors on actual macroeconomic datasets could also be analyzed.

This paper studies TVP VAR models with constant volatility. Our proposed inverse Wishart prior choices perform well in a variety of simulation studies. We have shown the relevance of using our improved prior choice for TVP VAR models by their ability to detect substantial time variation in the presence of the price puzzle.

Acknowledgments

We are thankful to Frank Smets, Joris Wauters, Michele Lenza, Gerdie Everaert, Karel Mertens and Gert Peersman for useful discussions and helpful comments. Also, financial support from the Agency for Innovation by Science and Technology in Flanders (IWT) is gratefully acknowledged.

References

- Aguiar, A. and Martins, M. (2005). Testing the significance and the non-linearity of the Phillips trade-off in the Euro Area. *Empirical Economics*, **30**(3), 665–691.
- Barth, M. J. and Ramey, V. A. (2002). The cost channel of monetary transmission. In *NBER Macroeconomics Annual 2001, Volume 16*, pages 199–256.
- Baumeister, C. and Peersman, G. (2013). Time-varying effects of oil supply shocks on the US economy. *American Economic Journal: Macroeconomics*, **5**(4), 1–28.
- Belaygorod, A. and Dueker, M. (2009). Indeterminacy, change points and the price puzzle in an estimated DSGE model. *Journal of Economic Dynamics and Control*, **33**(3), 624–648.
- Boivin, J. and Giannoni, M. P. (2006). Has monetary policy become more effective? *The Review of Economics and Statistics*, **88**(3), 445–462.

- Carlstrom, C. T., Fuerst, T. S., and Paustian, M. (2009). Monetary policy shocks, Choleski identification, and DNK models. *Journal of Monetary Economics*, **56**(7), 1014–1021.
- Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, **81**, 541–553.
- Castelnuovo, E. (2012). Testing the structural interpretation of the price puzzle with a cost-channel model. *Oxford Bulletin of Economics and Statistics*, **74**(3), 425–452.
- Castelnuovo, E. and Surico, P. (2010). Monetary policy, inflation expectations and the price puzzle. *Economic Journal*, **120**(549), 1262–1283.
- Clarida, R., Gali, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly Journal of Economics*, **115**(1), 147–180.
- Clark, T. E. and Ravazzolo, F. (2014). The macroeconomic forecasting performance of autoregressive models with alternative specifications of time-varying volatility. *Journal of Applied Econometrics*.
- Cogley, T. and Sargent, T. J. (2001). Evolving post-world war II U.S. inflation dynamics. In *NBER Macroeconomics Annual*, volume 16, pages 331–388.
- Cogley, T. and Sargent, T. J. (2005). Drift and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, **8**(2), 262–302.
- Commandeur, J. J. and Koopman, S. J. (2007). *An introduction to state space time series analysis*. Oxford university press.
- D’Agostino, A., Gambetti, L., and Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, **28**(1), 82–101.
- Del Negro, M. and Primiceri, G. (2013). Time-varying structural vector autoregressions and monetary policy: a corrigendum. Staff Reports 619, Federal Reserve Bank of New York.
- Demiralp, S., Hoover, K., and Perez, S. (2014). Still puzzling: evaluating the price puzzle in an empirically identified structural vector autoregression. *Empirical Economics*, **46**(2), 701–731.
- Eichenbaum, M. (1992). ‘interpreting the macroeconomic time series facts: the effects of monetary policy’: by Christopher Sims. *European Economic Review*, **36**(5), 1001–1011.

- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, **1**, 1–19.
- Giordani, P. (2004). An alternative explanation of the price puzzle. *Journal of Monetary Economics*, **51**(6), 1271–1296.
- Hanson, M. S. (2004). The ‘price puzzle’ reconsidered. *Journal of Monetary Economics*, **51**(7), 1385–1413.
- Harvey, A. and Koopman, S. (2009). Unobserved components model in economics and finance: the role of the Kalman filter in time series econometrics. *EEE Control System magazine*, **29**.
- Harvey, A. C. and Jaeger, A. (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics*, **8**(3), 231–47.
- Harvey, A. C., Trimbur, T. M., and Van Dijk, H. (2007). Trends and cycles in economic time series: a Bayesian approach. *Journal of Econometrics*, **140**(2), 618 – 649.
- Koop, G. and Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends(R) in Econometrics*, **3**(4), 267–358.
- Koop, G. and Potter, S. (2004). Forecasting in dynamic factor models using Bayesian model averaging. *Econometrics Journal*, **7**(2), 550–565.
- Koop, G. and Potter, S. M. (2011). Time varying vars with inequality restrictions. *Journal of Economic Dynamics and Control*, **35**(7), 1126–1138.
- Korobilis, D. (2014). Data-based priors for vector autoregressions with drifting coefficients. MPRA Paper 53772, University Library of Munich, Germany.
- Lubik, T. A. and Schorfheide, F. (2004). Testing for indeterminacy: an application to U.S. monetary policy. *American Economic Review*, **94**(1), 190–217.
- Muirhead, R. (1982). *Aspects of multivariate statistical theory*. Wiley.
- Mumtaz, H. and Sunder-Plassmann, L. (2013). Time varying dynamics of the real exchange rate: an empirical analysis. *Journal of Applied Econometrics*, **28**(3), 498–525.
- Nakajima, J. (2011). Time-varying parameter VAR model with stochastic volatility: an overview of methodology and empirical applications. *Monetary and Economics Studies*, **29**, 107–142.

- Petris, G. (2010). An R package for dynamic linear models. *Journal of Statistical Software*, **36**(12).
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, **72**(3), 821–852.
- Rabanal, P. (2003). The cost channel of monetary policy. IMF Working Papers 03/149.
- Rossi, P. E. and Allenby, G. M. (2003). Bayesian statistics and marketing. *Marketing Science*, **22**(3), 304–328.
- Rusnak, M., Havranek, T., and Horvath, R. (2013). How to solve the price puzzle? A meta analysis. *Journal of Money, Credit and Banking*, **45**(1), 37–70.
- Sargent, T. and Surico, P. (2011). Two illustrations of the quantity theory of money: breakdowns and revivals. *American Economic Review*, **101**(1), 109–128.
- Sims, C. A. (1992). Interpreting the macroeconomic time series facts: the effects of monetary policy. *European Economic Review*, **36**(5), 975–1000.
- Stock, J. H. and Watson, M. W. (1996). Evidence on structural instability in macroeconomic time series relations. *Journal of Business & Economic Statistics*, **14**(1), 11–30.
- Stock, J. H. and Watson, M. W. (1998). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association*, **93**(441), pp. 349–358.
- Stock, J. H. and Watson, M. W. (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking*, **39**(s1), 3–33.
- Stock, J. H. and Watson, M. W. (2010). Modeling inflation after the crisis. Working Paper 16488, National Bureau of Economic Research.
- Tillmann, P. (2009). The time-varying cost channel of monetary transmission. *Journal of International Money and Finance*, **28**(6), 941–953.
- Watson, M. W. (2014). Inflation persistence, the NAIRU, and the great recession. *American Economic Review*, **104**(5), 31–36.
- Wu, J. C. and Xia, F. D. (2014). Measuring the macroeconomic impact of monetary policy at the zero lower bound. NBER Working Papers 20117, National Bureau of Economic Research, Inc.

Appendix

Table 6: Estimated parameters reported in the macroeconomic literature for the ‘local level model’ (above the dashed line) and for more complex univariate models (below the dashed line). Each row represents a study and each column shows summary information. When necessary, we have translated the values of Q and Σ to correspond to annualized percentage growth rates. NA means that the information is not available in the paper. When a range of values is given, the results are for multiple variables or for multiple estimators.

Paper	Country	Variable	Model	Estimator
Harvey and Koopman (2009)	USA	Inflation	Local level	Maximum likelihood
Stock and Watson (1996)	USA	76 economic series	Local level	Median unbiased
Stock and Watson (1998)	USA	real GDP	Local level	ML + Median unbiased
Watson (2014)	USA	Linear combination	Local level	Nonparametric
Commandeur and Koopman (2007)	UK	Inflation	Local level + seasonal component	Maximum likelihood
Harvey and Jaeger (1993)	USA, AU	GNP, CPI, money	Local level + trend	Frequency domain
Stock and Watson (2007)	USA	Inflation	Local level + stochastic volatility	Nonlinear filtering methods
Stock and Watson (2010)	USA	Inflation	Local level + stochastic volatility	Nonlinear filtering methods

Paper	Sample Period	T	$\hat{Q}$	$\hat{\Sigma}$	$\sqrt{\hat{Q}/\hat{\Sigma}}$	$\hat{\lambda}$
Harvey and Koopman (2009)	1951Q1-1985Q4	140	NA	NA	0.47	66
Stock and Watson (1996)	1959M1-1993M12	408	NA	NA	0-0.02	0-8
Stock and Watson (1998)	1947Q2-1995Q4	195	0-0.01	NA	0-0.021	0-4
Watson (2014)	1959Q2-2013Q3	218	0.0004-0.007	0.2-0.6	0.03-0.19	7-41
Commandeur and Koopman (2007)	1950Q1-2001Q4	208	0.21	0.34	0.79	163
Harvey and Jaeger (1993)	1954Q1-1989Q4	144	0-0.57	0.15-1.1	0-0.75	0-108
Stock and Watson (2007)	1953Q1-2004Q4	208	0.04-1	0.3	0.4-1.8	83-374
Stock and Watson (2010)	1959Q1-2010Q2	206	0.2-1	0.2-0.7	0.53-2	109-412

Figure 8: Each subfigure corresponds to one simulation design and shows the boxplot of the estimated Q of the local level model over the 200 simulation runs. The horizontal axis represents the different estimators which are labelled 1 ($df=0.1, kQ=0.01$), 2 ($df=0.00001, scale=0.00001$), 3 ($df=20, kQ=0.01$), 4 ($df=20, kQ=0.01$) and 5 (Maximum Likelihood estimator). Finally, the horizontal line is the true Q value of the data generating process.

(a)

(b)

Figure 9: TVP VAR with stochastic volatility robustness check: Time varying inflation response after (a) one and (b) four quarters to an interest rate shock with 10^{th} , 50^{th} and 90^{th} percentiles.

Figure 10: Shadow interest rate robustness check: Time varying inflation response after (a) one and (b) four quarters to an interest rate shock with 10^{th} , 50^{th} and 90^{th} percentiles.

FACULTY OF ECONOMICS AND BUSINESS

Naamsestraat 69 bus 3500

3000 LEUVEN, BELGIË

tel. + 32 16 32 66 12

fax + 32 16 32 67 91

info@econ.kuleuven.be

www.econ.kuleuven.be