Time-Varying Local Projections

Germano Ruisi

Working paper No. 891 July 2019 ISSN1473-0278

School of Economics and Finance

Queen Mary
University of London

Time-Varying Local Projections

Germano Ruisi

Queen Mary University of London
School of Economics and Finance

Abstract

In recent years local projections have become a more and more popular methodology for the estimation of impulse responses. Besides being relatively easy to implement, the main strength of this approach relative to the traditional VAR one is that there is no need to impose any specific assumption on the dynamics of the data. This paper models local projections in a time-varying framework and provides a Gibbs sampler routine to estimate them. A simulation study shows how the performance of the algorithm is satisfactory while the usefulness of the model developed here is shown through an application to fiscal policy shocks.

Keywords: Time-Varying Coefficients, Local Projections
JEL Classification Codes: C11, C32, C36, E32
Author's Address: Queen Mary University of London, School of Economics and Finance, Mile End Road, London E1 4NS
Author's E-mail: germano.ruisi@qmul.ac.uk

1 Introduction

Impulse response functions are a typical tool used in Macroeconometrics to study the dynamic responses of economic variables to structural shocks of interest. In the empirical macroeconomic literature, the most used way to obtain them is through the estimation of vector autoregressions that allow the researcher to both identify the shock of interest and to trace the subsequent dynamic responses. In recent years, besides the latter approach, another one has increased its popularity. In fact, the local projections developed in Jorda (2005) have become a more and more widely used tool to trace the responses to structural shocks of interest. Two are the main reasons. First, as a sequence of predictive regressions, local projections are relatively easy to implement and, second, they do not impose specific dynamics on the variables in the system analysed. As such, they have proven to be relatively robust to model misspecifications.

Despite the increase in popularity, time variation in such a framework has been almost totally neglected. To the best of my knowledge, the only two attempts in the literature are represented by Auerbach & Gorodnichenko (2012) and Ramey & Zubairy (2018). In these works local projections accommodate nonlinearities with the aim of disentangling how the dynamic responses of macroeconomic variables of interest vary from an expansionary to a recessionary regime.

In this work my contribution is the introduction of time variation in the local projections framework. More precisely, I provide a Bayesian setting to estimate time-varying local projections where the shock of interest is identified through the use of a proxy used as an instrumental variable. This results in a framework that allows the researcher not only to trace the time-varying dynamic response of structural shocks of interest but also to assess how strong the narrative measure used to proxy such a shock is and at what response horizon. Moreover, and in a similar way to the insights taken from the "proxy SVAR" literature initiated with Stock & Watson (2008) and Mertens & Ravn (2013), the usage of a proxy to be used within this instrumental variable procedure alleviates the estimation bias due to the likely measurement error made when creating the narrative measure.

In order to assess the performance of the model developed in this work, I conduct a simulation study through two small Monte Carlo experiments where data are artificially generated starting from arbitrarily chosen parameters. The two experiments differs in the way they treat the errors of the predictive regressions as the projection horizon grows large. Moreover, with the aim of providing evidence of the usefulness of this setting, I carry out an empirical application to illustrate how this methodology works when studying the time-varying effects of fiscal policy shocks. The latter is proxied by a narrative measure that has already been developed and is available in the literature.

This paper is related to the growing literature that relies on local projections to estimate impulse responses to structural shocks of interest. Since Jorda's (2005) seminal paper, other works have tried to improve upon this methodology. As already mentioned, Auerbach & Gorodnichenko (2012) and Ramey & Zubairy (2018) accommodate regime-shifts. Some other works, e.g., Barnichon & Matthes (2017) and Miranda-Agrippino & Ricco (2018), propose ways to regularise impulse responses estimated from local projections while some others, e.g., Plagborg-Moller (2017), Barnichon & Brownlees (2018), El-Shagi (2018) and Tanaka (2018), provide strategies to select methods to estimate smoother responses. Finally, Ganics et al. (2018) propose a method to construct confidence intervals to assess the strength of the instruments proxying a shock should the local projections be estimated through a two-stage

least squares procedure within a frequentist framework.

In addition, this paper is related to the strand of literature, initiated with Cogley & Sargent (2005) and Primiceri (2005), that uses time-varying parameter models to provide support to the evidence of how the dynamic response to macroeconomic shocks is likely to change over time. With regard to this strand of literature, two recent contributions deserve a particular mention, i.e., Paul (2017) and Mumtaz & Petrova (2018), as they develop models that nest the proxy SVAR approach into a time-varying parameters framework.

The rest of the paper is organised as follows. Section 2 represents the main block of this work. More precisely, it presents the model developed here, i.e., the time-varying local projections, as well as the settings and routine to estimate it. Section 3 presents a simulation study with artificially generated data while section 4 presents an empirical application. It will show the usefulness of this model and how it provides the researcher with superior information with respect to that obtainable from local projections that do not account for time variation. Finally, section 5 concludes.

2 Econometric Methodology

This section lays out the econometric methodology. To be more precise, it describes the model and how time variation is included in the local projections framework. In addition, it describes the priors chosen and provides a Gibbs sampler routine to estimate the model as well as the settings that are necessary to initialise it.

2.1 The Model

In Jorda's (2005) seminal paper, local projections allow the researcher to estimate impulse responses directly from a set of linear regressions. In the time-varying version developed in this work I begin from the same starting point. As such, it follows that the model for an individual observation looks as follows:

$$Y_{t+h} = \alpha_{t,h} + \beta_{t,h} X_{t-1} + \gamma_{t,h} W_t + v_{t+h} \tag{1}$$

where $h = 0, 1, ..., H$ represents the projection horizon, Y_{t+h} is the variable of interest, i.e., the response variable, while X_{t-1} is a vector of control variables, e.g., lags of dependent and independent variables. The contemporaneous value W_t is endogenous and described by:

$$W_t = \theta + \delta Z_t + e_t \tag{2}$$

where Z_t is a set of instruments which are assumed to be correlated with W_t and uncorrelated with v_{t+h}. In equation 1, $B_{t,h} = [\alpha_{t,h}, \beta_{t,h}, \gamma_{t,h}]$ represents all the time-varying coefficients stacked into one vector, i.e., intercept and slopes. The latter are assumed to evolve according to a random walk with no drift:

$$B_{t,h} = B_{t-1,h} + u_{t,h} \tag{3}$$

where $u_{t,h}$ is a vector of innovations which are assumed to have variance $Var(u_{t,h}) = Q_h$ and to be uncorrelated both with v_{t+h} and e_t, i.e., $Cov(v_{t+h}, u_t) = 0$ and $Cov(e_t, u_t) = 0$, at each horizon h. In equation 2, θ and δ respectively represent intercept and slope coefficients and summarize the relation between the instrument(s) Z_t and the endogenous variable W_t. Note how both θ and δ are not assumed to be time-varying and the reason is that I am only looking for the relation between instrument and endogenous variable in its purest form. With the aim of keeping things simple, this allows me to assess the strength of the instrument(s) by just looking at the highest posterior density distribution of the coefficients in 2[1]. Note how v_{t+h}, at impulse response horizons bigger than 0, i.e., $h > 0$, is serially correlated and heteroscedastic as it represents the sum of the errors made when forecasting the response variable at horizons from 0 to h. Finally, the structural form errors are assumed to be correlated and to have a variance-covariance matrix of the form:

$$Cov(e_t, v_{t+h}) = \Omega_h = \begin{pmatrix} \sigma_{11,h} & \sigma_{12,h} \\ \sigma_{21,h} & \sigma_{22,h} \end{pmatrix} \tag{4}$$

Note how equation 4 is what distinguishes this Bayesian setting from a pure frequentist instrumental variable approach. It is so, as the latter would assume a diagonal Ω_h matrix as no correlation between errors is allowed.

The main goal of the model is the identification of the parameter of interest γ_{t+h} which represents the $h-th$ horizon time-varying response of Y_{t+h} to a shock in W_t which, in turn, is instrumented by Z_t.

2.2 Priors

As it is commonly done in time-varying modelling, following Cogley & Sargent (2005) and Primiceri (2005), I set the priors by using a training sample without assuming time variation in the coefficients. After selecting a proper training sample, for each horizon $h = 0, 1, ..., H$, the starting point is the estimation of

$$Y_{t+h} = \alpha_h + \beta_h X_{t-1} + \gamma_h W_t + v_{t+h} \tag{5}$$

where

$$W_t = \theta + \delta Z_t + e_t \tag{6}$$

is used as a first stage regression. It follows that, by substituting 6 into 5 we get:

$$\begin{cases} W_t = \theta + \delta Z_t + e_t \\ Y_{t+h} = \alpha_h + \gamma_h \theta + \beta_h X_{t-1} + \gamma_h \delta Z_t + \gamma_h e_t + v_{t+h} \end{cases}$$

[1]This assumption may obviously be lifted and time variation in θ and δ can be allowed. However, this would add a layer of complexity as, in this case, the strength of the instrument set Z_t needs to be assessed at each time t of the sample chosen.

which implies that, by denoting the reduced form residuals as $\begin{pmatrix} \omega_t^1 \\ \omega_{t+h}^2 \end{pmatrix}$, the relation between the latter and the structural ones is

$$\begin{pmatrix} \omega_t^1 \\ \omega_{t+h}^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \gamma_h & 1 \end{pmatrix} \begin{pmatrix} e_t \\ v_{t+h} \end{pmatrix} \quad or, equivalently, \quad \begin{pmatrix} 1 & 0 \\ -\gamma_h & 1 \end{pmatrix} \begin{pmatrix} \omega_t^1 \\ \omega_{t+h}^2 \end{pmatrix} = \begin{pmatrix} e_t \\ v_{t+h} \end{pmatrix}$$

As discussed in Rossi et al. (2005), this represents the basis for setting a prior for the covariance matrix of the structural errors which is explicitly dependent on γ_h. This prior is assumed to follow an Inverse Wishart distribution:

$$\Omega_h \sim IW(\Omega_{h0}, T_{\Omega_h 0}) \tag{7}$$

By indicating with A the matrix capturing the relation between reduced and structural form residuals $\begin{pmatrix} 1 & 0 \\ -\gamma_h & 1 \end{pmatrix}$, the prior scale is $\Omega_{h0} = A \begin{pmatrix} \sigma^2_{\omega_t^1} & 0 \\ 0 & \sigma^2_{\omega_{t+h}^2} \end{pmatrix} A'$ where $\sigma^2_{\omega_t^1}$ and $\sigma^2_{\omega_{t+h}^2}$ represent the variances of the reduced form residuals. As discussed in the previous subsection, v_{t+h} is serially correlated and heteroscedastic at impulse response horizons bigger than 0 and so does ω^2_{t+h}. Neglecting to account for this characteristic results in misspecifying the model as this would lead to assume a wrong variance of the error. With this aim in mind, it is necessary to HAC-correct (i.e., to correct for heteroscedasticity and serial correlation) $\sigma^2_{\omega^2_{t+h}}$ for $h > 0$ and I try to achieve this goal in the following way. Denote $\widehat{\sigma_j}$ as an estimate of the $j-th$ order autocovariance of ω^2_{t+h}, then

$$\sigma^2_{\omega^2_{t+h}} = \widehat{\sigma_0} + 2 \sum_{j=1}^{T-1} k_{j,h} \widehat{\sigma_j} \tag{8}$$

where $k_{j,h}$ is the Bartlett kernel where the bandwidth parameter is set to be equal to h, i.e., $k_{j,h} = \begin{cases} (1 - \frac{j}{h+1}) & for\, 0 \leq j \leq h \\ 0 & for\, j > h \end{cases}$. This particular choice of the bandwidth parameter will ensure that the exact order of the serial correlation in the error is taken into account while the decaying value of $k_{j,h}$ assigns a lower and lower weight as such order j increases. This way of accounting for heteroscedasticity and serial correlation is clearly an approximation but, in its simplicity, it avoids to add further layers of complexity by explicitly modelling the behaviour of v_{t+k}, and consequently ω^2_{t+h}, at each horizon h. The prior degrees of freedom are $T_{\Omega_h 0} = 2 + 1$ where 2 is the number of equations in the system.

I set a normal prior for the intercept and the slope coefficients in equation 2 where mean and variance are obtained by simply estimating 6 through ordinary least squares, i.e.,

$$\underset{\Delta}{\begin{pmatrix} \theta \\ \delta \end{pmatrix}} \sim N\left(\underset{\Delta_0}{\begin{pmatrix} \theta_0 \\ \delta_0 \end{pmatrix}}, \underset{\Sigma_{\Delta 0}}{\begin{pmatrix} \Sigma_\theta & 0 \\ 0 & \Sigma_\delta \end{pmatrix}} \right) \tag{9}$$

The initial values of the time-varying intercept and slope coefficients relative to equation 1 are normally distributed

$$B_{0|0} \sim N(B_{0,h}, V_h) \tag{10}$$

where the mean $B_{0,h}$ is obtained by performing a 2-stage least squares of equation 5 by using Z_t as an instrument for W_t. The reason why I do not follow equation 6 as a first stage regression in this case is the necessity to find, within a frequentist environment, the fitted values of W_t which can be considered to be uncorrelated with v_{t+h} in order to properly identify γ_h. This is of crucial importance as the latter will eventually play the role of starting value. V_h is the variance-covariance matrix of the coefficients $B_{0,h} = [\alpha_h, \beta_h, \gamma_h]$ obtained through this 2-stage least squares procedure and is corrected for heteroscedasticity and serial correlation in the error v_{t+h} for $h > 0$. Such a correction is implemented through the Newey & West (1987) estimator where the bandwidth parameter is set to be equal to the projection horizon h in order to account for the exact order of serial correlation of the error.

Finally, the prior for the variance of the innovations $u_{t,h}$, i.e., the innovations relative to the coefficients' law of motion in 3, is assumed to be an Inverse Wishart distribution

$$Q_h \sim IW(Q_{h0}, T_{h0}) \tag{11}$$

where the prior scale is $Q_{h0} = V_h * Train * SF$. V_h is the variance-covariance matrix of the coefficients in $B_{0,h} = [\alpha_h, \beta_h, \gamma_h]$ as previously discussed, $Train$ is the training sample length while SF is a scaling factor which is meant to govern the a priori degree of time variation of the coefficients in 1. The more SF tends to zero the lower the time variation shown by the local projections. Finally, the prior degrees of freedom T_{h0} is simply equal to $Train + 1$.

2.3 Gibbs Sampler

The Gibbs sampling algorithm presented here, which implies drawing from the posterior distributions, can be considered as an extension of that developed in Rossi et al. (2005). As such, for each projection horizon h, it implies the following steps:

1. Given a draw of Ω_h and the residuals e_t the model can be written conditional on e_t. In particular, as $v_{t+h}|e_t \sim N(\frac{\sigma_{12,h}}{\sigma_{11,h}} e_t; \sigma_{22,h} - \frac{\sigma^2_{12,h}}{\sigma_{11,h}})$ it is possible to transform the data in the following way:

$$Y^*_{t+h} = \frac{Y_{t+h} - \frac{\sigma_{12,h}}{\sigma_{11,h}} e_t}{(\sigma_{22} - \frac{\sigma^2_{12,h}}{\sigma_{11,h}})^{\frac{1}{2}}} \quad and \quad \Gamma^*_t = \frac{\Gamma_t - \frac{\sigma_{12,h}}{\sigma_{11,h}} e_t}{(\sigma_{22} - \frac{\sigma^2_{12,h}}{\sigma_{11,h}})^{\frac{1}{2}}}$$

where $\Gamma_t = [1, X_{t-1}, W_t]$. Equation 1 then becomes:

$$Y^*_{t+h} = \Gamma^*_t B_{t,h} + v^*_{t+h}$$

where $B_{t,h} = B_{t-1,h} + u_{t,h}$ and $v^*_{t+h} \sim N(0; R^*_h)$. Note that, as $v^*_{t+h} = \frac{v_{t+h} - \frac{\sigma_{12,h}}{\sigma_{11,h}} e_t}{(\sigma_{22} - \frac{\sigma^2_{12,h}}{\sigma_{11,h}})^{\frac{1}{2}}}$, it follows that for $h = 0$ $var(v^*_{t+0}) = R^*_0 = 1$ while for $h > 0$ R^*_h must be HAC-corrected as v^*_{t+h} comprises a combination of past forecast errors[2].

2. Use the Carter & Kohn (1994) algorithm to sample $\widetilde{B}_{t,h} = [\widetilde{\alpha_{t,h}}, \widetilde{\beta_{t,h}}, \widetilde{\gamma_{t,h}}]$, i.e., the time series of the time-varying coefficients.

3. Sample $Q_h = Var(u_{t,h})$ from its posterior distribution $Q_h \tilde{} IW(Q_{h,Posterior}, T_{Q_h,Posterior})$ where $Q_{h,Posterior} = (\widetilde{B}_{t,h} - \widetilde{B}_{t-1,h})'(\widetilde{B}_{t,h} - \widetilde{B}_{t-1,h})' + Q_{h0}$ and $T_{Q_h,Posterior} = Train + T_{h0}$.

4. Transform the data in the following way:

$$\widetilde{Y_{t+h}} = \frac{Y_{t+h} - (\widetilde{\alpha_{t,h}} + \widetilde{\beta_{t,h}} X_{t-1})}{\widetilde{\gamma_{t,h}}}$$

and then rewrite the model in 1 and 2 as

$$\begin{cases} W_t = \theta + \delta Z_t + e_t \\ \widetilde{Y_{t+h}} = \theta + \delta Z_t + e_t + \frac{1}{\gamma_{t,h}} v_{t+h} \end{cases}$$

where $Cov(e_t, e_t + \frac{1}{\gamma_{t,h}} v_{t+h}) = A_{t,h} \Omega_h A'_{t,h}$ and $A_{t,h} = \begin{pmatrix} 1 & 0 \\ 1 & \frac{1}{\gamma_{t,h}} \end{pmatrix}$. This transformed model can be seen as a Seemingly Unrelated Regressions (SUR) System. Pre-multiplying both sides of the system by $\widetilde{A_{t,h}} = \left(\left(A_{t,h} \Omega_h A'_{t,h} \right)^{\frac{1}{2}} \right)^{-1}$ transforms the model into a stacked regression with a unit variance and the draw for $\Delta = \begin{pmatrix} \theta \\ \delta \end{pmatrix}$ is standard, i.e., $\Delta \sim N(M_\Delta, V_\Delta)$ where

$$M_\Delta = (\Sigma^{-1}_{\Delta 0} + Z'_{SUR,t} Z_{SUR,t})^{-1} (\Sigma^{-1}_{\Delta 0} \Delta_0 + Z'_{SUR,t} Y_{SUR,t})$$

$$V_\Delta = (\Sigma^{-1}_{\Delta 0} + Z'_{SUR,t} Z_{SUR,t})^{-1}$$

where $Y_{SUR,t} = vec(\widetilde{A_{t,h}}[W_t, \widetilde{Y_{t+h}}])$ and $Z_{SUR,t} = [vec(\widetilde{A_{t,h}}[1,1]), vec(\widetilde{A_{t,h}}[Z_t, Z_t])]$.

[2]This HAC-correction is implemented in a similar way to equation 8 shown in the previous subsection.

5. Sample Ω_h from $IW(\Omega_{h,Posterior}, T_{\Omega_h,Posterior})$ with $T_{\Omega_h,Posterior} = T + T_{\Omega_h 0}$, where T is the sample length once the training sample has been removed, and $\Omega_{h,Posterior}$ is the HAC-corrected posterior scale for $h > 0$.

In a similar way to what discussed in the previous subsection, neglecting to explicitly model the autocorrelation and heteroscedasticity in v_{t+h} implies that the time-varying local projections end up being misspecified. As described in Muller (2013) and Miranda-Agrippino & Ricco (2018), in presence of heteroscedasticity and serial correlation the true likelihood is still Gaussian and centred at the Maximum Likelihood Estimator but has a larger variance. In order to soften such a misspecification problem I adjust the posterior scale of Ω_h in the 5^{th} step of the Gibbs sampler in the following way: denote with ε_{t+h} the structural form residuals $\begin{pmatrix} e_t \\ v_{t+h} \end{pmatrix}$ and with S_j the $j-th$ order scale parameter given by $S_j = \varepsilon'_{t+h}\varepsilon_{t+h-j}$, then

$$\Omega_{h,Posterior} = S_0 + 2\sum_{j=1}^{T-1} k_{j,h} S_j \tag{12}$$

where $k_{j,h}$ is the Bartlett kernel with bandwidth parameter equal to h. In a similar way to what seen in the previous subsection, this easy and more frequentist way of accounting for heteroscedasticity and serial correlation avoids the extra burden to add further steps in the Gibbs sampler in order to explicitly model the behaviour of the error. As such, this correction allows to remain agnostic with respect to the nature of the heteroscedasticity and serial correlation of v_{t+h}.

Finally, notice that the estimation of the variance of e_t does not formally require any HAC-correction as the latter is assumed to be homoscedastic and serially uncorrelated independently of the projection horizon h. In order to show how the posterior scale correction applied in the algorithm does not provide biased estimates of the covariance matrix between the structural errors, in the next section I will perform a simulation study showing, inter alia, how the the elements of $Cov(e_t, v_{t+h})$ are all well estimated independently of the error being well behaved or not.

2.4 Estimation Setup

For each response horizon $h = 0, 1, ..., H$, the model is estimated by using the following setup. The number of Gibbs sampler iterations is set to $20,000$ and the first $10,000$ are discarded in order to minimise the effect of the initial draws. Moreover, in order to minimise the serial correlation across draws, a thinning factor of 20 is used. As a result, once past the burn-in stage, inference is conducted only on the retained 500 draws while the others are discarded. Finally, the scaling factor that governs the amount of time variation in the coefficients in equation 1 is set to $3.5e-04$, i.e., a small number to reflect the fact that the training sample is generally short and the estimates might be relatively imprecise.

3 Simulation Study

The aim of this section is to provide evidence that the model developed here works correctly. To this end, I conduct two Monte Carlo experiments with artificially created data. On the one hand, the first experiment is intended to evaluate how the algorithm performs in presence of a normally distributed error v_t which is not heteroscedastic nor serially correlated. To this end, it is meant to show how the algorithm performs when $h = 0$. On the other hand, the second experiment aims at assessing how the model works when the HAC-corrections shown previously are in order, which is the case when $h > 0$.

In both cases data are generated from a system with one observation equation that comprises one constant term, two exogenous regressors and one endogenous one. In order to show how the model is flexible enough to accommodate any number of instruments, the endogenous regressor is in turn explained by two instruments. Only the coefficients of the observation equation are assumed to be time-varying and to be following a driftless random walk where the starting values are set to be equal to zero.

For the sake of brevity, both experiments are conducted by only considering responses on impact. Note that considering also longer horizon projections would just involve the estimation of a set of equations where the only difference is that the response variable is shifted onwards as many times as the projection horizon. As the true responses cannot be retrieved for horizons different from the impact one, this exercise would prove of little utility with artificially generated data as it is not possible to assess how the estimated responses are close to the true ones. The experiments are conducted by generating 600 observations from the data generating processes and by discarding the first 100 in order to remove the effect of initial conditions. In addition, equipped with the remaining 500 observations, I use the first 100 as a training sample that will be necessary for setting the priors. Both experiments are repeated 100 times and, for each of them, I run the algorithm and use the estimation setup outlined in section 2. This means that, for each Monte Carlo iteration, inference is conducted by using only the remaining 500 draws.

The data for the first Monte Carlo experiment are created from the following data generating process:

$$\begin{cases} W_t = 0.2 + 0.9Z_{1,t} - 0.1Z_{2,t} + e_t \\ Y_t = \alpha_t + \beta_{1,t}X_{1,t-1} + \beta_{2,t}X_{2,t-1} + \gamma_t W_t + v_t \end{cases}$$

where

$$\begin{pmatrix} e_t \\ v_t \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.015 & 0.030 \\ 0.030 & 0.160 \end{pmatrix}\right]$$

$$\begin{pmatrix} Z_{1,t} \\ Z_{2,t} \\ X_{1,t-1} \\ X_{2,t-1} \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}\right]$$

$$\begin{pmatrix} \alpha_t \\ \beta_{1,t} \\ \beta_{2,t} \\ \gamma_t \end{pmatrix} = \begin{pmatrix} \alpha_{t-1} \\ \beta_{1,t-1} \\ \beta_{2,t-1} \\ \gamma_{t-1} \end{pmatrix} + \begin{pmatrix} u_t^{\alpha} \\ u_t^{\beta_1} \\ u_t^{\beta_2} \\ u_t^{\gamma} \end{pmatrix} \quad where \quad \begin{pmatrix} \alpha_0 \\ \beta_{1,0} \\ \beta_{2,0} \\ \gamma_0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} u_t^{\alpha} \\ u_t^{\beta_1} \\ u_t^{\beta_2} \\ u_t^{\gamma} \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.001 & 0 & 0 & 0 \\ 0 & 0.001 & 0 & 0 \\ 0 & 0 & 0.001 & 0 \\ 0 & 0 & 0 & 0.001 \end{pmatrix} \right]$$

Figure 1 shows the median estimated time-varying intercept and slope coefficients $\widetilde{B_t} = [\widetilde{\alpha_t}, \widetilde{\beta_{1t}}, \widetilde{\beta_{2t}}, \widetilde{\gamma_t}]$ together with the 90% credible set. It is evident how the median estimates are able to closely match the true ones. Whenever the estimated parameters less closely trace the true ones, the latter appear to be well within the credible set. For the sake of clarity, here, the picture is related to only the first Monte Carlo iteration[3].

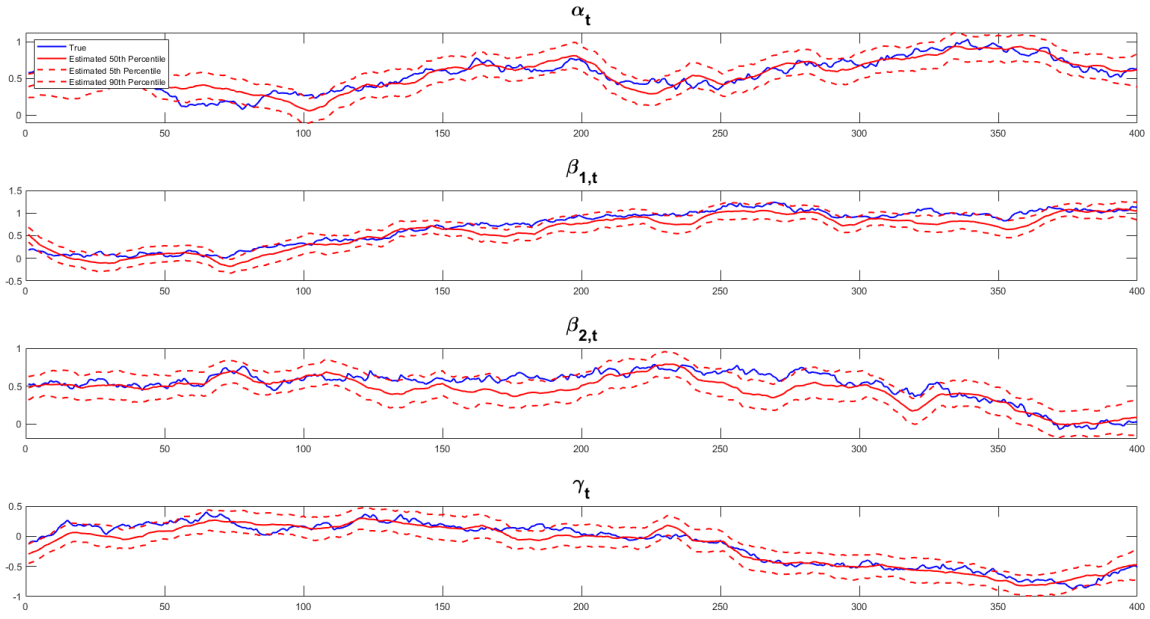

Figure 1: Time-Varying True (Blue) and Estimated Median Coefficients and 90% Credible Set (Red) in Presence of Homoscedastic and Serially Uncorrelated Error

In order to further assess how well the algorithm performs I follow the suggestions in Primiceri (2005) and look at the autocorrelation of the retained draws. Such convergence diagnostics are shown

[3] The inclusion of all the remaining 99 iterations would imply to have 100 different sets of time-varying coefficients in the same graph. This would make hard any attempt to draw conclusions regarding the performance of the algorithm in estimating the true coefficients.

in figure 2. More precisely, the picture shows the $20th$ order autocorrelation of the retained draws of the time-varying coefficients $\widetilde{B_t}$. Rather than showing all of them in one graph, or in as many graphs as the number of time-varying coefficients, I decide to split them into two graphs. To be more precise, the top panel shows those relative to the retained draws of the time-varying $[\widetilde{\alpha_t},\widetilde{\beta_{1t}},\widetilde{\beta_{2t}}]$ stuck into one vector while those relative to $\widetilde{\gamma_t}$ are shown in the bottom one. By doing so, I can better highlight how well the algorithm works with respect to the parameter of interest, i.e., the time-varying impulse response γ_t. Low correlations are evidence of high efficiency of the algorithm as the draws are almost independent. Similarly to the previous picture, for the sake of clarity, the convergence diagnostics shown here are related to the first Monte Carlo iteration. For this exercise, all the autocorrelations are well within the $[-0.2, 0.2]$ interval and therefore they indicate a satisfactory performance of the algorithm.

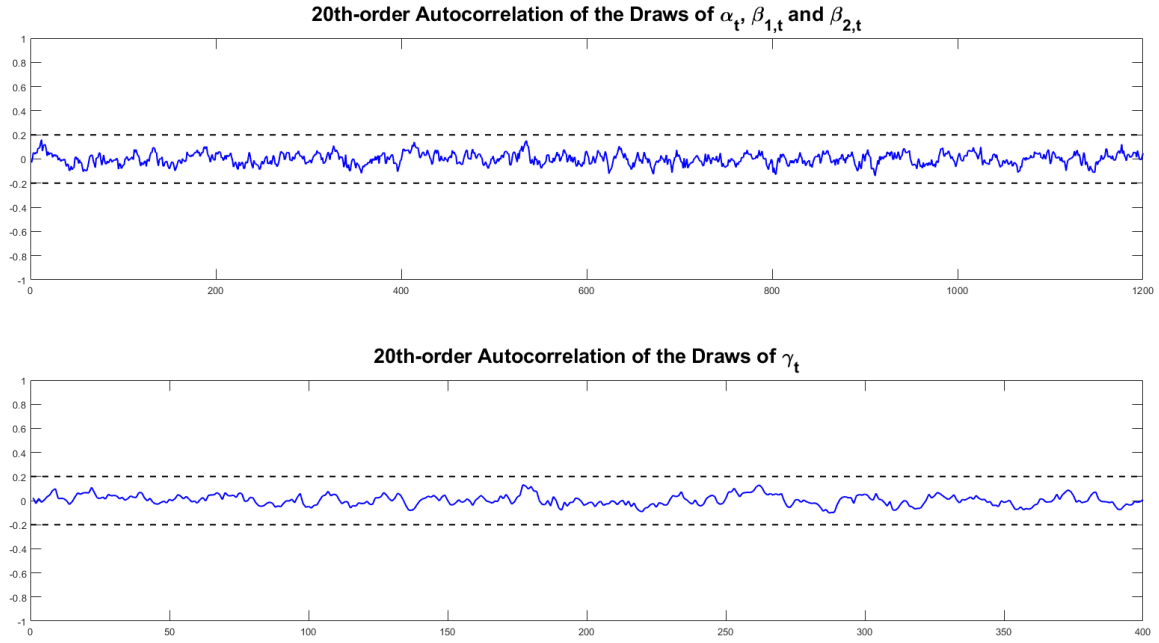

Figure 2: $20th$-order Autocorrelation of the Retained Draws of $(\alpha_t, \beta_{1,t}, \beta_{2,t})$ (Upper Graph) and of γ_t (Bottom Graph)

Finally, the next three pictures are relative to the full Monte Carlo experiment as they depict all the draws that lead to the same posterior distributions independently of the artificially generated data used. Figure 3 shows the posterior distributions of the intercept and slope coefficients relative to equation 2, i.e., $\Delta = (\theta, \delta_1, \delta_2)$. Figure 4 shows the posterior distributions of the disturbances of the law of motion of the time-varying coefficients, i.e., $Q = var(u_t)$, while figure 5 shows the posterior distribution of the covariance matrix of the structural residuals, i.e., $\Omega = Cov(e_t, v_t)$. All the posterior distributions look well identified and relatively well centred around the values arbitrarily chosen above, the latter being indicated by vertical red lines.

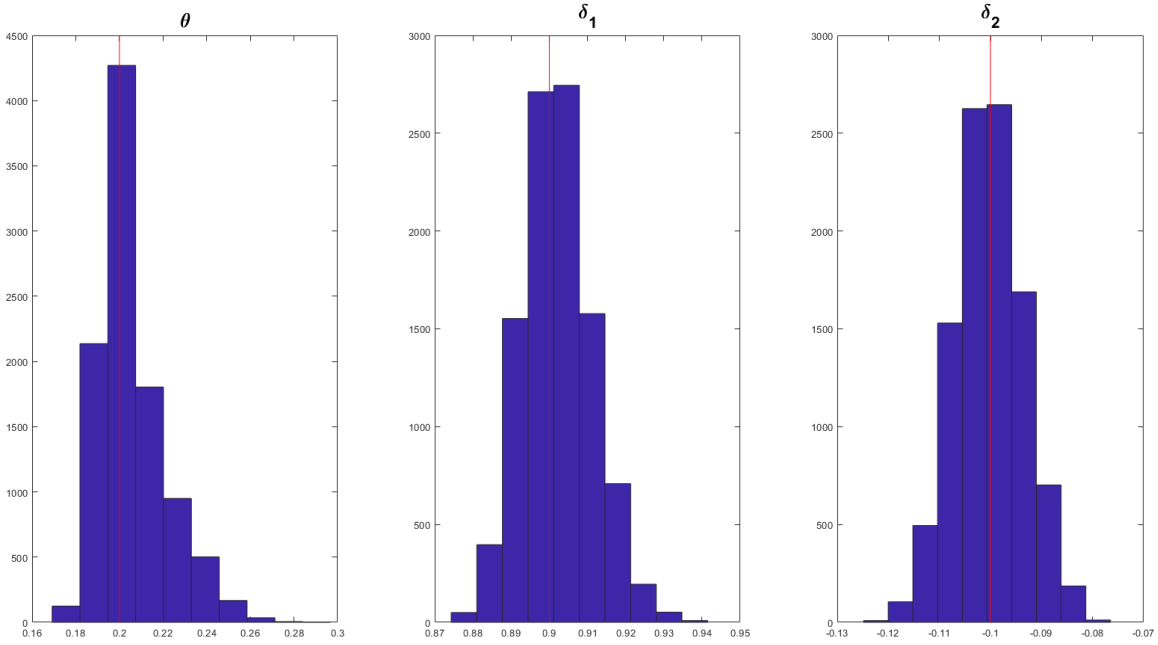

Figure 3: Posterior Distributions of the $\Delta = (\theta, \delta_1, \delta_2)$ Coefficients

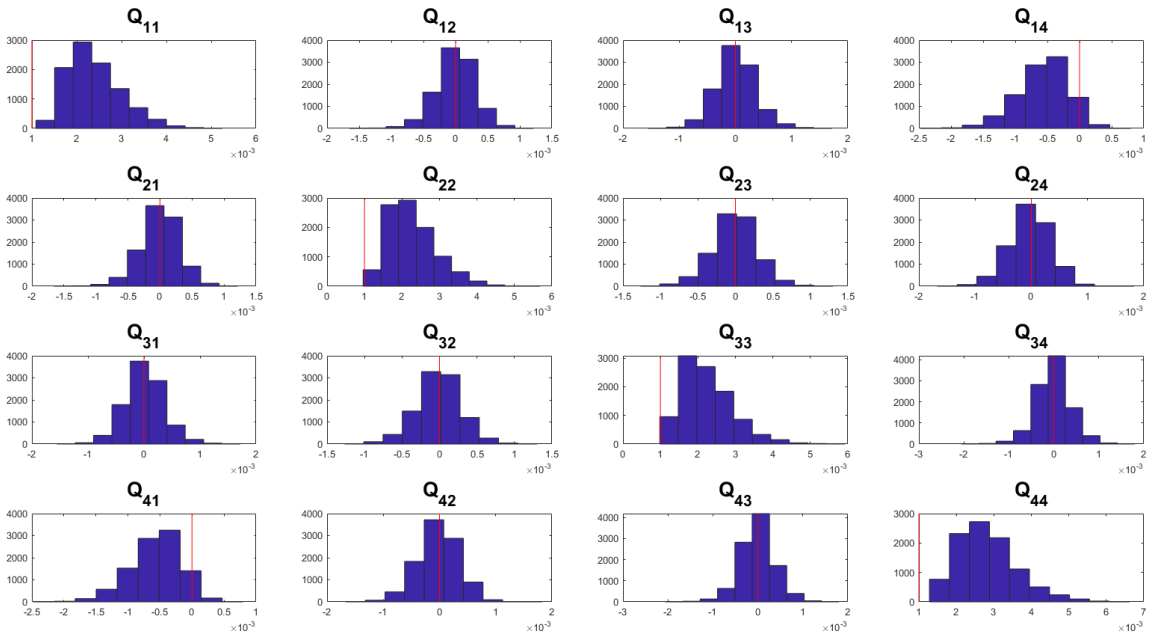

Figure 4: Posterior Distributions of the Elements of $Q = var(u_t)$

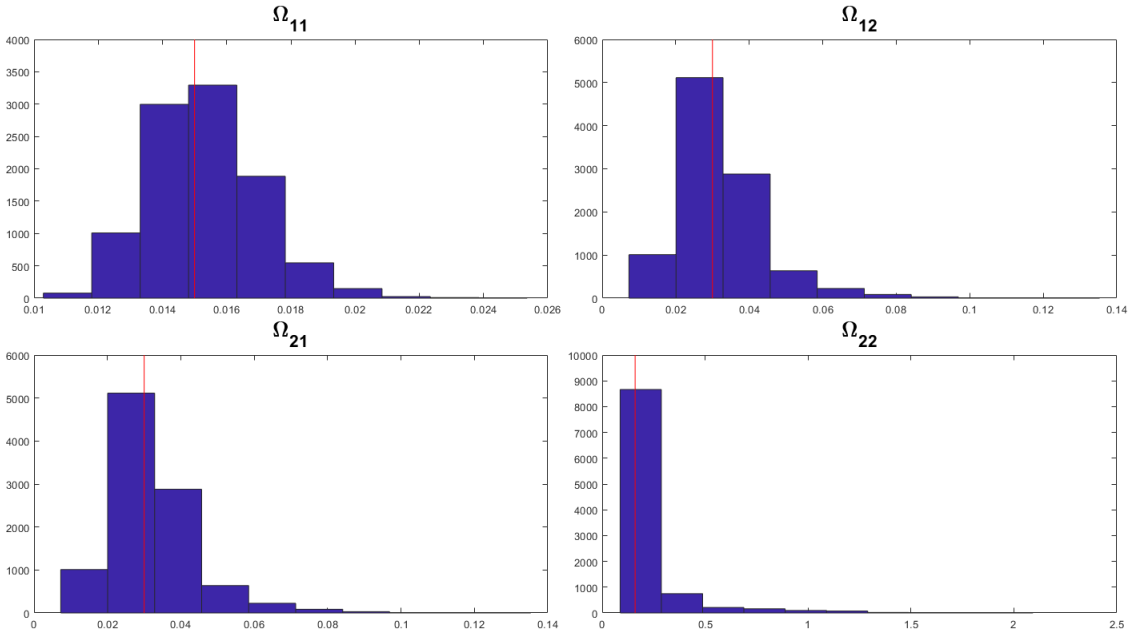

Figure 5: Posterior Distributions of the Elements of $\Omega = Cov(e_t, v_t)$ in Presence of Homoscedastic and Serially Uncorrelated Error

As previously stated, the second Monte Carlo experiment aims at showing how the algorithm performs in presence of heteroscedastic and serially correlated error of the observation equation, which is the case when $h > 0$. With this aim in mind, data are generated from the following data generating process:

$$\begin{cases} W_t = 0.2 + 0.9Z_{1,t} - 0.1Z_{2,t} + e_t \\ Y_t^{HAC} = \alpha_t + \beta_{1,t}X_{1,t-1} + \beta_{2,t}X_{2,t-1} + \gamma_t W_t + v_t^{HAC} \end{cases}$$

where, without loss of generality, I arbitrarily induce heteroscedasticity and serial correlation in v_t^{HAC} in the following way:

$$v_t^{HAC} = 0.6v_{t-1} - 0.2v_{t-2} + 0.1v_{t-3} + \xi_t$$

where

$$\begin{pmatrix} e_t \\ v_t \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.015 & 0.030 \\ 0.030 & 0.160 \end{pmatrix}\right]$$

$$\xi_t \sim N(0, \sigma_t^2) \quad where \quad \sigma_t^2 \sim IG(1,1)$$

All the rest is set as in the previous Monte Carlo experiment. Note how the superscript on the response variable only aims at distinguishing it from that of the previous exercise. To be clearer, in this second experiment, endogenous and exogenous regressors, instruments and coefficients of the two equations are all generated in the same way as before. The only difference between Y_t and Y_t^{HAC} is that the latter is now including a heteroscedastic and serially correlated error, i.e., v_t^{HAC}.

As previously done, the first picture is meant to show how the estimated time-varying coefficients perform in matching the true ones and, for the sake of clarity, the picture is related to only the first Monte Carlo iteration. In order to compare these results with the case of homoscedastic and uncorrelated error, I use the same set of artificially created data with the only difference that I replace v_t with v_t^{HAC} and, as a consequence, Y_t with Y_t^{HAC}. This means that the true time-varying betas that need to be estimated are exactly the same ones as those in figure 1.

Figure 6 shows the median estimated time-varying intercept and slope coefficients $\widetilde{B_t} = [\widetilde{\alpha_t}, \widetilde{\beta_{1t}}, \widetilde{\beta_{2t}}, \widetilde{\gamma_t}]$ together with the 90% credible set. Notice how, differently from the previous case, i.e., as a result of v_t^{HAC} being heteroscedastic and serially correlated, the estimated time-varying coefficients appear to match less closely the true ones. However, the 90% credible set appears wider and this helps keep the true coefficients within the bands.

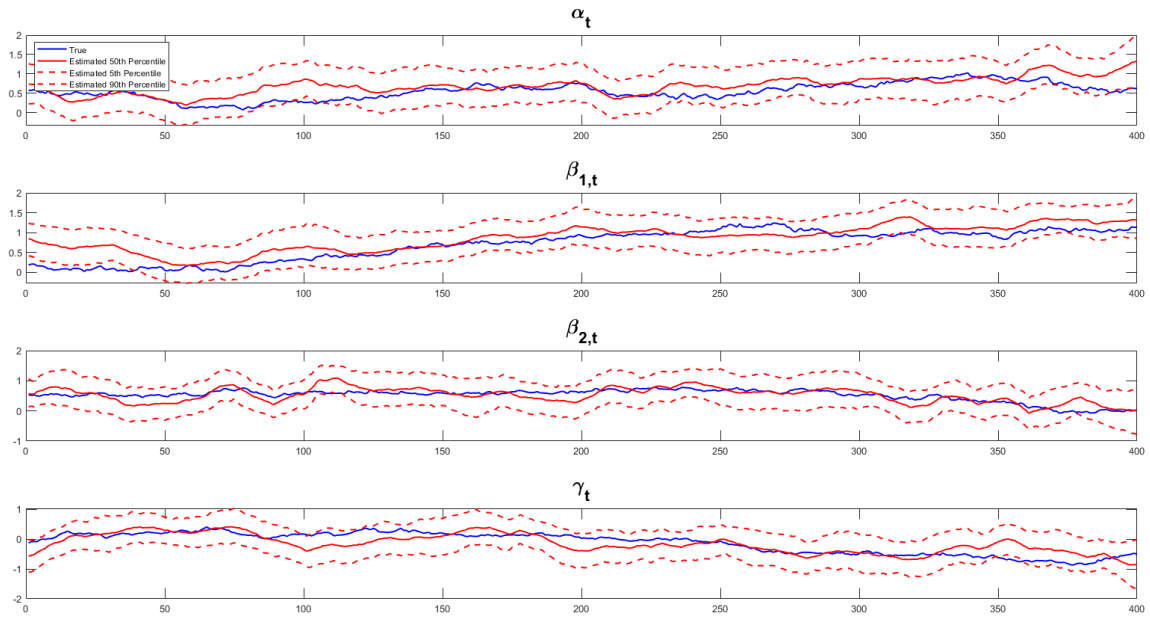

Figure 6: Time-Varying True (Blue) and Estimated Median Coefficients and 90% Credible Set (Red) in Presence of Heteroscedastic and Serially Correlated Error

As an additional check of the good performance of the algorithm in presence of heteroscedastic and serially correlated error, I check the posterior distributions of the elements of the covariance matrix relative to the structural errors. Under this setting the covariance between artificially created structural errors is[4]:

[4]Differently from the previous Monte Carlo exercise where $Cov(e_t, v_t)$ was arbitrarily chosen, here $Cov(e_t, v_t^{HAC})$

$$Cov(e_t, v_t^{HAC}) = \begin{pmatrix} 0.015 & 0.003 \\ 0.003 & 2.069 \end{pmatrix}$$

Figure 7 shows the posterior distribution of the covariance matrix of the structural residuals, i.e., $\Omega = Cov(e_t, v_t^{HAC})$. As in the previous case, this picture is related to the full Monte Carlo experiment and the vertical red lines represent the true values. The picture shows how, even in this case, the algorithm is able to estimate the values of Ω in a good way as the posterior distributions appear to be well centred around the true values. Moreover, it is important to emphasise how the HAC-correction works well. To be clearer, the variance of e_t does not show any difference with respect to the previous case while the variance of v_t^{HAC} instead does. More precisely, while the variance of e_t has remained the same as the one of the previous exercise, that of v_t^{HAC} is bigger than that of v_t and the algorithm is able to capture it well.

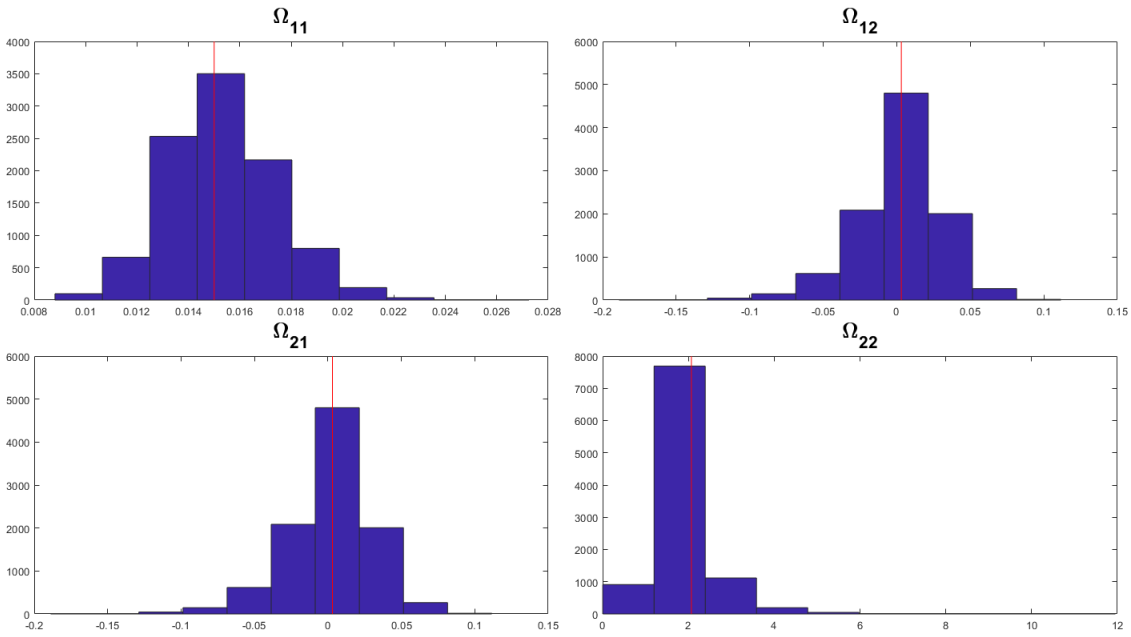

Figure 7: Posterior Distributions of the Elements in $\Omega = Cov(e_t, v_t^{HAC})$ in Presence of Heteroscedastic and Serially Correlated Error

For the sake of space, for this second experiment, the convergence diagnostics and the posterior distributions of $\Delta = (\theta, \delta_1, \delta_2)$ and $Q = var(u_t)$ are not shown as there are no relevant differences with respect to the previous case[5].

All in all, given the results of the two Monte Carlo experiments, there is evidence of how the algorithm provides a good performance even in the case when the error is heteroscedastic and serially

needs to be estimated. To this end, I generate a sample of 10,000,000 observations and then I estimate the covariance between the artificially created e_t and v_t^{HAC}. Such long time series aim at improving the precision of the estimate.

[5]Nevertheless, they are available upon request.

correlated. It is so as, independently of the characteristics of the error, the algorithm is able to provide relatively good estimates of the time-varying coefficients $\widetilde{B_t} = [\widetilde{\alpha_t}, \widetilde{\beta_{1t}}, \widetilde{\beta_{2t}}, \widetilde{\gamma_t}]$ with almost independent draws from their respective posterior distributions and, in addition, it is able to well identify the posterior distributions of $\Delta = (\theta, \delta_1, \delta_2)$, $Q = var(u_t)$ and $\Omega = Cov(e_t, v_t)$ or, alternatively, $\Omega = Cov(e_t, v_t^{HAC})$ when the error is heteroscedastic and serially correlated.

4 Empirical Application

In this section I perform an empirical application in order to demonstrate the usefulness of the model developed in this work. The application aims at understanding the time-varying impact of military news shocks for the United States of America. The exercise is conducted by replicating the main body of the analysis, within the time-varying parameters framework developed here, of an empirical study that is well known in the literature, i.e., Ramey & Zubairy (2018). For the sake of space, only the main results are shown while the convergence diagnostics are presented in appendix 2.

In the macroeconometric literature many works develop narrative measures with the aim of proxying fiscal shocks for subsequently estimating the effect of contractionary or expansionary policies. Ramey & Shapiro (1998) create a dummy variable to capture major exogenous military buildups while a number of follow-up papers, Edelberg et al. (1999) ad Burnside et al. (2004), included dummy variables to capture war dates. The intuition behind these works is that major military events lead, among other things, to an increase in the economic activity. More recent contributions have also looked at the effect of tax changes. In Romer & Romer (2010) the authors develop a narrative measure of exogenous tax changes which are those changes that are not related to or justified by current and/or prospective economic conditions. In two follow-up works, Mertens & Ravn (2012) and Mertens & Ravn (2013), the authors take the Romer & Romer (2010) series and first distinguish between anticipated and unanticipated tax changes and then, in their second work, they further divide the unanticipated tax changes series into personal and income tax changes.

The reason why, in order to show the usefulness of this model, I decide to replicate the study in Ramey & Zubairy (2018) is one. Such a work, besides representing an example of renewed attention on the effect of government spending shocks, develops a rich narrative measure that aims at capturing the "news" component of the government spending shocks. This measure, being an extension of that developed in Ramey (2011), provides an extraordinary base to assess whether, and to what extent, there have been changes in the effectiveness of fiscal policy shocks. It is so because it covers a particularly long time span, i.e., more than one hundred years of observations[6].

By using the same notation outlined in section 2, the matrix of exogenous regressors X_{t-1} contains four lags of the Ramey and Zubairy's (2018) series, of the gross domestic product and of the government spending where all of them are expressed in real terms and scaled by potential real GDP (Gordon & Krenn (2010) transformation)[7]. The inclusion of lags of the narrative measure helps remove any potential serial correlation in the latter. The endogenous regressor W_t is the current change in the

[6] More details regarding data sources and transformations are shown in appendix 1.

[7] Such a transformation allows the researcher to estimate both responses and fiscal multiplies by using the same variables which are directly expressed in dollars. More precisely, in order to estimate the fiscal multipliers it is not necessary to resort to ex-post conversions of the impulse responses. For a more detailed explanation see Ramey and Zubairy (2018).

government spending series while Z_t is the current value of the military news shock measure. The data is at quarterly frequency and the sample runs from 1889Q1 to 2015Q4. The first 60 observations are used as a training sample while the maximum response horizon is set to 20 quarters. The use of such a training sample together with four lags of the exogenous regressors make the effective sample run from 1905Q1 to 2015Q4. Figure 8 shows the median responses of government spending and output to a unit shock in the former where the shock is proxied by the military news shock series[8].

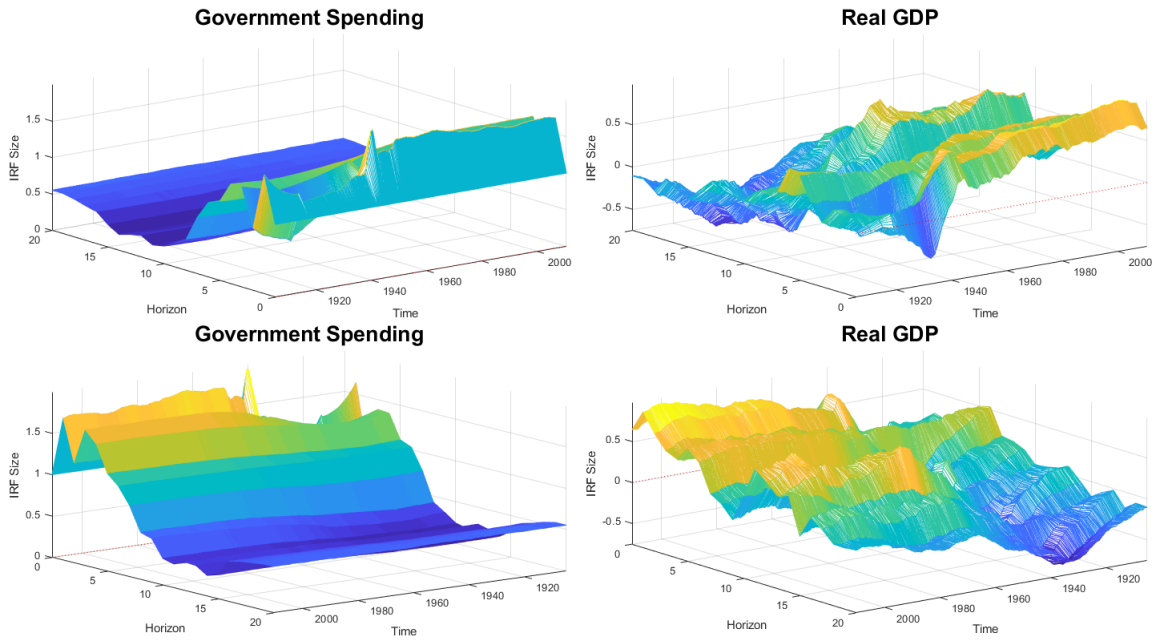

Figure 8: Time-Varying Responses to Military News Shocks - Median Estimates - Effective Sample: 1905Q1-2015Q4

The graphs show the impulse response horizons on the Y-axis, the period of time the sample refers to on the X-axis while the magnitude of the responses is on the Z-axis. The responses of government spending are depicted on the left column while those of the real GDP on the right one. The two graphs at the bottom represent the same local projections as those in the upper part of the figure but they have been rotated by 180 degrees in order to better understand what the responses are at longer horizons. When focusing on government spending, the left side of figure 8 shows how there is evidence (even though little) of time variation at short impulse horizons. More precisely, over the four quarters after the shock hits, the time-varying responses are smaller for the period prior to the second World War but then they go up to around 1.5. When considering longer horizons, instead, things appear to be slightly different as the time variation seems to dampen. Independently of the moment a military news shock occurs, the response of the government spending variable decreases until reaching a value of about 0.2 after 14 quarters. On the right hand side of figure 8, instead, the response of the real GDP presents a clearer evidence of time variation throughout both the entire time span and the projection

[8]For the sake of clarity, credible sets are not shown.

horizons considered. To be more precise, regardless of the projection horizon, output responses are clearly the lowest during the Great Recession. If we focus on the responses on impact, they even reach negative values of -0.7. In addition, figure 8 makes it evident how output responses are bigger at all projection horizons after the Great Recession.

In order to better understand how time-variation produces a richer story about the responses to a military news shock over time, as a comparison, figure 9 shows the same responses but estimated in a linear framework. The responses are obtained from a simplified version of the Gibbs sampler routine presented in section 2 that does not account for time variation in the coefficients of the observation equation 1[9].

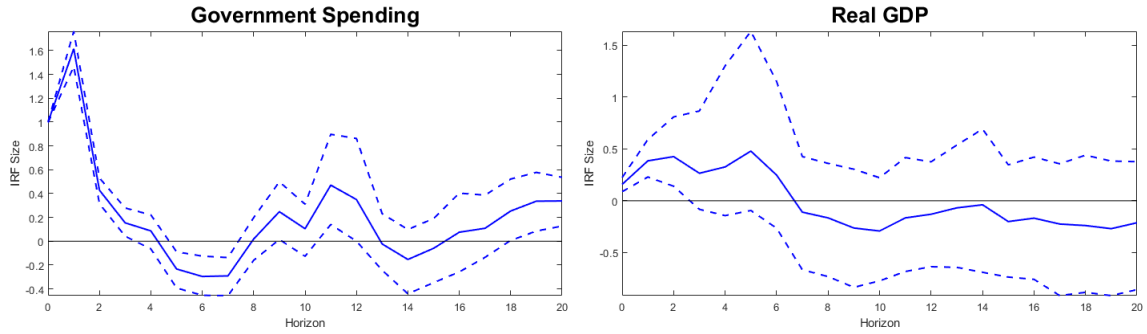

Figure 9: Linear Responses to Military News Shocks - Median Estimates and 68% Bands - Effective Sample: 1905Q1-2015Q4

Figure 9 clearly shows how the linear responses resemble the time-varying ones shown in figure 8. However, as they can be considered as an average response over the entire sample, they are not able to capture how the transmission of a military news shock has changed over time. If, on the one hand, and for this specific exercise, this does not seem to be particularly problematic for the response of the government spending variable, on the other hand, it does for that of the real GDP. More precisely, these linear projections are not able to capture how the responses of the latter become of bigger magnitude after the 40s at any horizon. Notice how this is reflected in the width of the 68% bands which are wider for the real GDP and narrower for the government spending which, as shown in figure 8, presents smaller time variation in the responses.

[9] More precisely, in this linear version of the Gibbs sampler, step 2 does not imply the use of the Carter & Kohn (1994) algorithm as any time variation is no longer required. For this reason, the draw for B_h is obtained from $B_h \sim N(M_{B_h}, V_{B_h})$ where posterior mean and variance are given by

$$M_{B_h} = (\Sigma_{B_{h0}}^{-1} + \tfrac{1}{R_h^*}\Gamma_t^{*\prime}\Gamma_t^*)^{-1}(\Sigma_{B_{h0}}^{-1}B_{h0} + \tfrac{1}{R_h^*}\Gamma_t^{*\prime}Y_{t+h}^*)$$

$$V_{B_h} = (\Sigma_{B_{h0}}^{-1} + \tfrac{1}{R_h^*}\Gamma_t^{*\prime}\Gamma_t^*)^{-1}$$

Moreover, R_h^* is defined as in step 2 while the prior $B_h \sim N(B_{h0}, \Sigma_{B_{h0}})$ is obtained by means of a training sample in a similar way to what done with other priors. Finally, step 3 has been suppressed while in step 4 the covariance matrix between the SUR system errors is no longer time-varying as now $A = \begin{pmatrix} 1 & 0 \\ 1 & \frac{1}{\gamma_h} \end{pmatrix}$.

The results shown in figure 8, however, do not help understand whether the narrative measure to proxy the shock of interest can be considered strong. To this end, in order to assess the strength of the instrument chosen, table 1 presents the median draws and the 90% highest posterior density interval of the parameter δ relative to the local projection where the real GDP is the response variable[10]. For the sake of space, the results are shown only for selected horizons.

Horizon	Median Draw	HPDI 5^{th} Percentile	HPDI 95^{th} Percentile	Median/Std Error
0	0.09	0.07	0.11	7.54
4	0.09	0.06	0.12	4.96
8	0.09	0.06	0.12	5.36
12	0.09	0.06	0.12	5.24
16	0.09	0.06	0.12	5.74
20	0.09	0.06	0.12	5.28

Table 1: Strength of the Narrative Measure at Selected Response Horizons - Effective Sample: 1905Q1-2015Q4

Table 1 clearly evidences how the highest posterior density interval suggests that the hypothesis that $\delta = 0$ is rejected for all the horizons considered. It is so as all the highest posterior density intervals do not include zero. In other words, this means that the instrument maintains its strength even when the response horizon gets bigger.

As a further step to illustrate the usefulness of this framework, and in a similar way to in Ramey & Zubairy (2018), I proceed with the estimation of the cumulated fiscal multipliers. In a similar way to the cited work, a fiscal multiplier is defined as the integral of the output response divided by the integral of the government spending one or, alternatively, as the response of the cumulated output given a shock in the cumulated government spending up to horizon h. I follow the second definition and, in order to estimate horizon h fiscal multipliers, I transform the model in the following way:

$$\begin{cases} \sum_{i=0}^{h} W_{t+i} = \theta + \delta Z_t + e_t \\ \sum_{i=0}^{h} Y_{t+i} = \alpha_{t,h} + \beta_{t,h} X_{t-1} + \gamma_{t,h} \sum_{i=0}^{h} W_{t+i} + v_{t+h} \end{cases} \tag{13}$$

where $\sum_{i=0}^{h} Y_{t+i}$ and $\sum_{i=0}^{h} W_{t+i}$ respectively represent the cumulated real GDP and government spending measures up to response horizon h while X_{t-1} and Z_t contain the same variables as before. This allows to directly estimate the time-varying fiscal multipliers $\gamma_{t,h}$. Figure 10 shows the results:

[10] As equation 2 does not change with the response variable of the local projection, the statistics for the government spending projection are not shown as they would deliver the same insights.

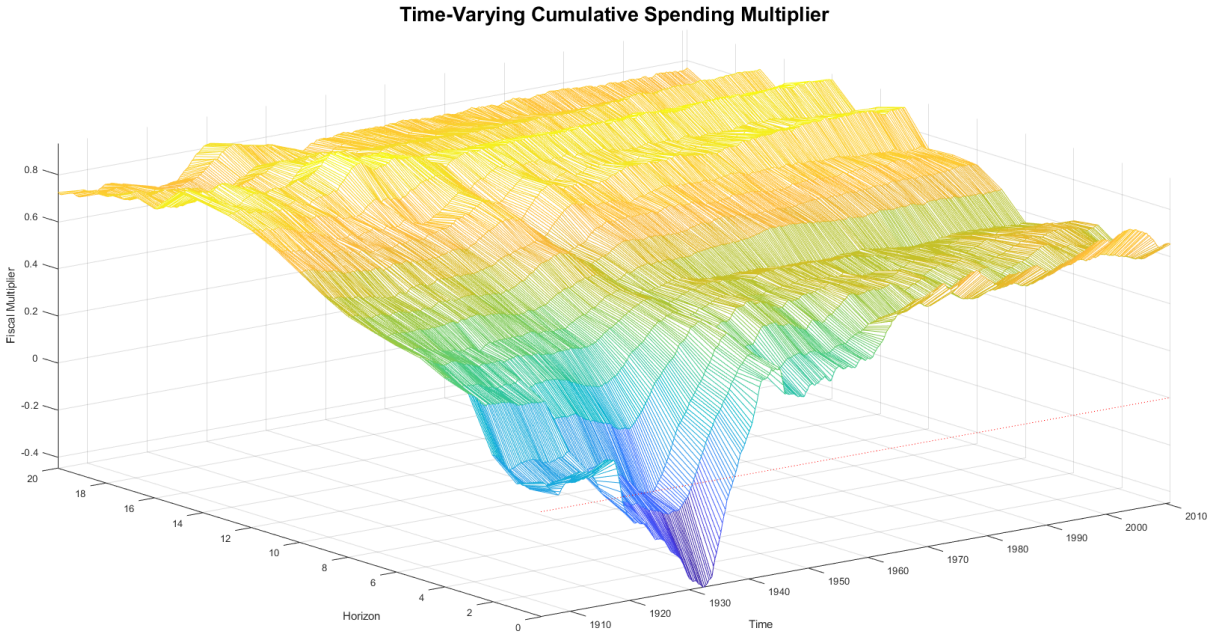

Figure 10: Time-Varying Cumulative Spending Multipliers - Median Estimates - Effective Sample: 1905Q1-2015Q4

If we focus on short horizons, figure 10 clearly shows evidence of how the impact multipliers varied according to the time a fiscal policy shock occurred. More precisely, in line with the insights drawn from figure 8, impact multipliers are lower prior to the WWII and reach the lowest values during the Great Recession. After the WWII they steadily increase. On the other hand, when focusing on longer horizons, the time variation of multipliers seems to be more limited. It is evident from figure 10 how fiscal multipliers tend to flatten as the horizon grows large.

As a comparison, figure 11 shows the cumulative spending multipliers estimated by means of the linear projections model previously mentioned.

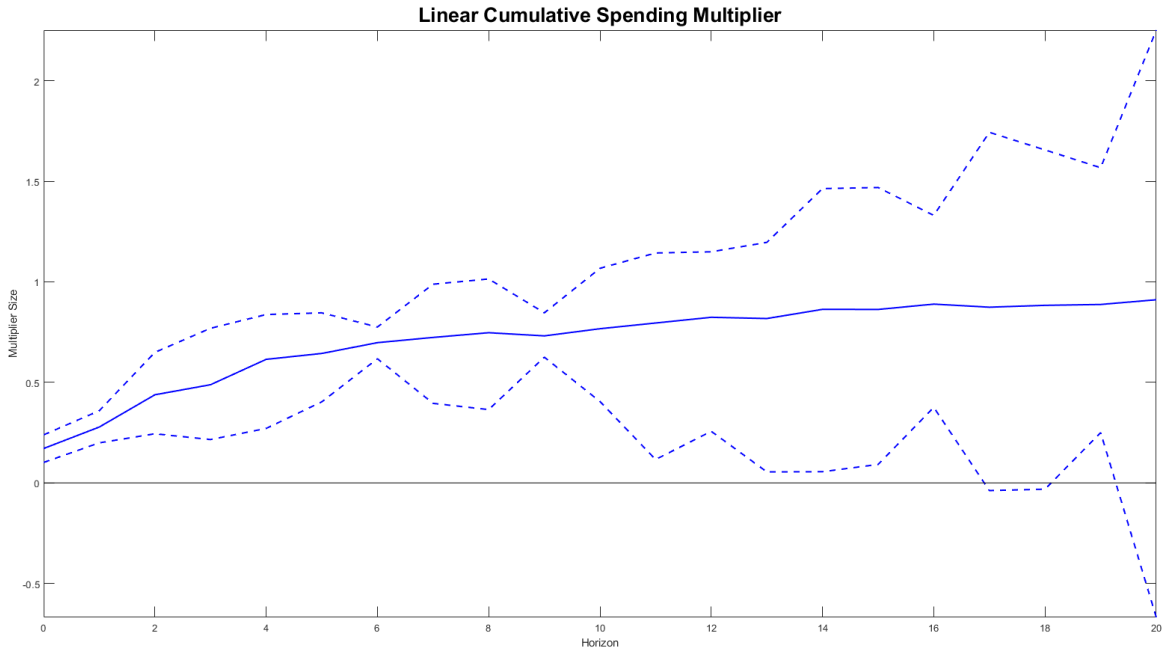

Figure 11: Linear Cumulative Spending Multipliers - Median Estimates and 68% Bands - Effective Sample: 1905Q1-2015Q4

Similarly to what seen with the impulse responses, the estimated linear multipliers resemble the time-varying ones but are not able to capture the time variation that, as shown in figure 10, is evident especially at short horizons.

In order to better evaluate how the spending multipliers have changed over time, figure 12 shows them at selected horizons, i.e., 0, 4, 8, 12, 16, 20 quarters, together with the 68% credible sets. The shaded areas represent the NBER recession dates while, as a comparison, the red horizontal lines the estimated median linear multipliers.

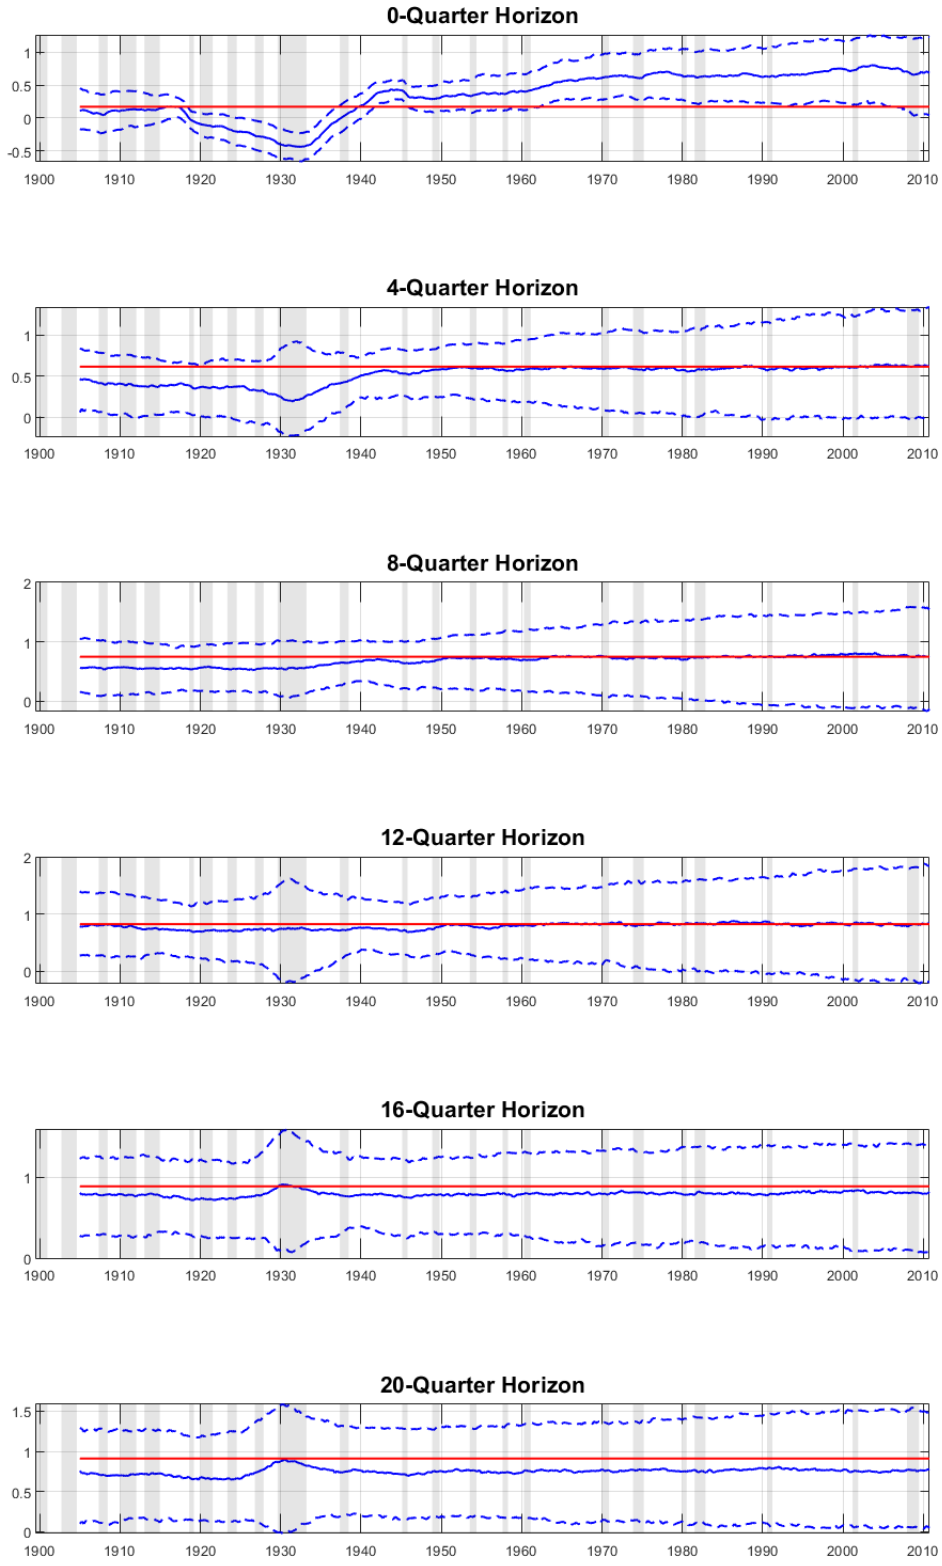

Figure 12: Time-Varying Spending Multipliers at Selected Horizons - Median Estimates and 68% Credible Sets - Effective Sample: 1905Q1-2015Q4

Figure 12 confirms what seen in figures 10 and 11, i.e., shorter horizon multipliers tend to be lower in the pre WWII period while they somehow flatten, throughout the entire sample, at longer horizons. The linear multipliers, being average multipliers over the entire period, closely match the time-varying ones as the horizon grows large. On the contrary, they do not capture the remarkable time variation displayed at shorter horizons and that is especially evident at impact. When focusing on the first graph, impact multipliers are not significantly different from zero until the mid 20s and then reach a value of about -0.4 during the Great Recession. Afterwards, they steadily increase and remain significantly different from zero. The highest values of the impact multipliers are found in the 2000s when they hover 0.75. Similar behaviour is found for the 4-quarter multipliers. More precisely, they are lower during the pre WWII period but then they steadily increase. It is interesting to notice how they are positive and hovering a value of 0.25 during the Great Recession but such estimates are not significantly different from zero. From the 8-quarter horizon on the multipliers seem to flatten throughout the entire sample period and then they seem to stabilise at around the a value of nearly 0.8 as the horizon gets bigger (12, 16 and 20-quarter horizon). Finally, the 16 and 20-quarter multipliers show how they reach a peak of nearly 0.9 during the early 30s.

5 Conclusions

In recent years local projections have become a more and more popular methodology for the estimation of impulse responses. Since Jorda's (2005) seminal paper, a number of works have tried to improve upon such a methodology. Some studies have focused on the regularisation of local projections, some others on techniques aiming at making the estimated responses smoother while some others on the assessment of the strength of the instruments proxying a shock should the local projections be estimated through a two-stage least squares procedure. Despite the increase of this strand of literature, time variation has been almost entirely neglected. This paper tries to fill this gap by modelling local projections in a time-varying framework. To this end, this work has provided a Bayesian approach to estimate time-varying local projections where the shock of interest is identified through the use of a proxy used as an instrumental variable.

The model's performance is assessed through two small Monte Carlo exercises where the data are artificially generated. The two exercises differ in the way the error of the observation equation is treated, i.e., whether a correction for heteroscedasticity and serial correlation is necessary or not. In both cases the algorithm shows a satisfactory performance. Subsequently, the usefulness of the model is illustrated through an empirical application. The latter replicates the main body of a relevant recent study relative to the transmission of military news shocks where the latter are identified through a narrative measure which is used as a proxy, i.e., Ramey & Zubairy (2018). This empirical application shows how the model developed in this work can provide a richer information regarding the effects of fiscal policies with respect to that drawn from a model that does not account for time variation.

In conclusion, this study provides an additional step for further developing Bayesian local projections. Further improvements might include the possibility to make this framework more flexible. A possibility, that is left for further research, could be to make this setting account for other shock identification strategies that are commonly used in the literature besides the narrative one employed here.

References

Auerbach, A. J. & Gorodnichenko, Y. (2012), 'Measuring the output responses to fiscal policy', *American Economic Journal: Economic Policy* **4**(2), 1–27.

Barnichon, R. & Brownlees, C. T. (2018), 'Impulse response estimation by smooth local projections'.

Barnichon, R. & Matthes, C. (2017), 'Functional approximation of impulse responses: Insights into the effects of monetary shocks', *Techical Report* .

Burnside, C., Eichenbaum, M. & Fisher, J. D. (2004), 'Fiscal shocks and their consequences', *Journal of Economic Theory* **115**(1), 89 – 117.

Carter, C. K. & Kohn, R. (1994), 'On gibbs sampling for state space models', *Biometrika* **81**(3), 541–553.

Cogley, T. & Sargent, T. J. (2005), 'Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S', *Review of Economic Dynamics* **8**(2), 262–302.

Edelberg, W., Eichenbaum, M. & Fisher, J. D. (1999), 'Understanding the effects of a shock to government purchases', *Review of Economic Dynamics* **2**(1), 166 – 206.

El-Shagi, M. (2018), 'A simple estimator for smooth local projections', *Applied Economics Letters* .

Ganics, G., Inoue, A. & Rossi, B. (2018), 'Confidence intervals for bias and size distortion in iv and local projections - iv models', *Banco de Espana Working Paper* (1841).

Gordon, R. J. & Krenn, R. (2010), The end of the great depression 1939-41: Policy contributions and fiscal multipliers, Working Paper 16380, National Bureau of Economic Research.

Jorda, O. (2005), 'Estimation and inference of impulse responses by local projections', *American Economic Review* **95**(1), 161–182.

Mertens, K. & Ravn, M. O. (2012), 'Empirical evidence on the aggregate effects of anticipated and unanticipated us tax policy shocks', *American Economic Journal: Economic Policy* **4**(2), 145–81.

Mertens, K. & Ravn, M. O. (2013), 'The dynamic effects of personal and corporate income tax changes in the united states', *American Economic Review* **103**(4), 1212–47.

Miranda-Agrippino, S. & Ricco, G. (2018), 'The transmission of monetary policy shocks', *Bank of England, Staff Working Paper* (657).

Muller, U. K. (2013), 'Risk of bayesian inference in misspecified models, and the sandwich covariance matrix', *Econometrica* **81**(5), 1805–1849.

Mumtaz, H. & Petrova, K. (2018), 'Changing impact of shocks: A time-varying proxy svar approach', *Queen Mary University of London - Working Paper Series* (875).

Newey, W. K. & West, K. D. (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**(3), 703–708.

Paul, P. (2017), The Time-Varying Effect of Monetary Policy on Asset Prices, Working Paper Series 2017-9, Federal Reserve Bank of San Francisco.

Plagborg-Moller, M. (2017), 'Essays in macroeconometrics, chapter 3. phd dissertation, harvard university', *Techical Report* .

Primiceri, G. E. (2005), 'Time varying structural vector autoregressions and monetary policy', *The Review of Economic Studies* **72**(3), 821–852.

Ramey, V. A. (2011), 'Identifying government spending shocks: It's all in the timing', *The Quarterly Journal of Economics* **126**(1), 1–50.

Ramey, V. A. & Shapiro, M. D. (1998), 'Costly capital reallocation and the effects of government spending', *Carnegie-Rochester Conference Series on Public Policy* **48**, 145 – 194.

Ramey, V. A. & Zubairy, S. (2018), 'Government spending multipliers in good times and in bad: Evidence from us historical data', *Journal of Political Economy* **126**(2), 850–901.

Romer, C. D. & Romer, D. H. (2010), 'The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks', *American Economic Review* **100**(3), 763–801.

Rossi, P. E., Allenby, G. M. & McCulloch, R. (2005), 'Bayesian statistics and marketing', *Wiley Series in Probability and Statistics* (13).

Stock, J. H. & Watson, M. W. (2008), 'Lecture 7 - recent developments in structural var modeling', *What's New in Econometrics - Time Series* .

Tanaka, M. (2018), 'Bayesian inference of local projections with roughness penalty priors'.

Appendix 1 - Empirical Application - Data

The data used for the empirical application in section 4 are taken from the Valery Ramey's website and are part of the replication files and data relative to Ramey & Zubairy (2018). Figure 13 plots the series used in the estimation, i.e., government spending, gross domestic product and the Ramey's military news shock series. As in Gordon & Krenn (2010), all of the three series are expressed in real terms (by dividing each of them by a $1-$quarter lag of the GDP deflator) and scaled by a measure of potential real gross domestic product where the latter is obtained by using a 6-degree polynomial for the logarithm of the real GDP.

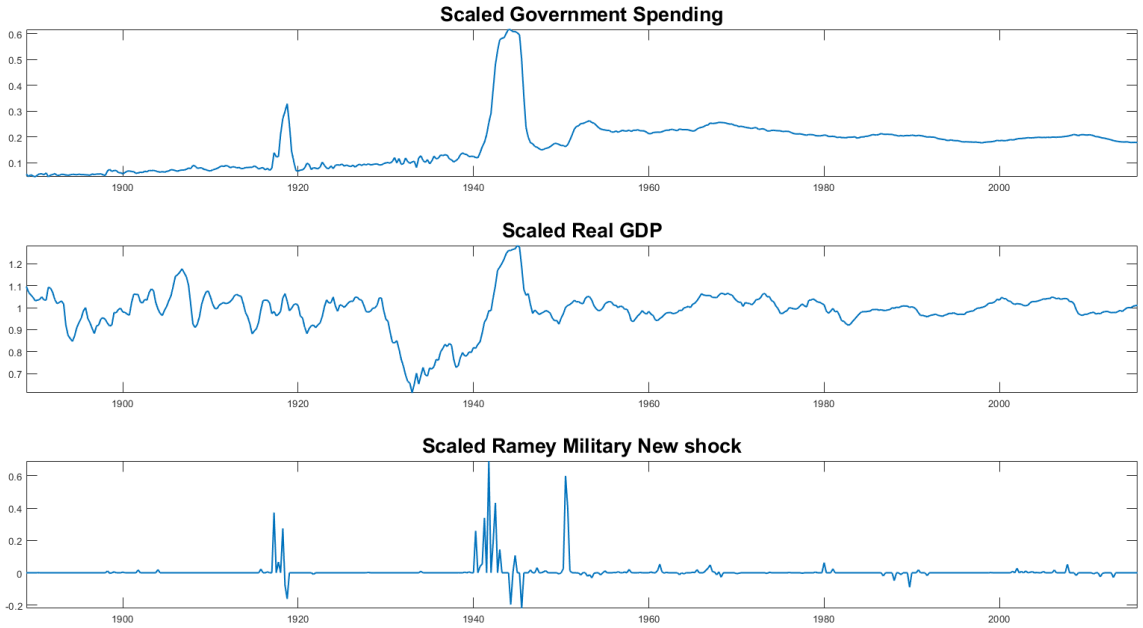

Figure 13: Empirical Application - Plot of Variables - Sample: 1889Q1-2015Q4

Appendix 2 - Empirical Application - Convergence Diagnostics

This appendix shows the convergence diagnostics relative to the empirical application presented in section 4. As previously done for the simulation study, in order to assess the satisfactory performance of the algorithm, the $20th$-order autocorrelation of the retained draws is employed. For the sake of space, figure 14 shows such a measure only relative to the $\gamma_{t,h}$ parameter, i.e., the time-varying projection coefficient, for each of the three objects of interest, i.e., the responses of government spending and gross domestic product as well as the spending multipliers. The convergence diagnostics relative to the other parameters, as well as those relative to the linear model, are available upon request. The red vertical lines separate the autocorrelations by response horizon h in order to facilitate the comprehension of the algorithm's performance as the projection horizon lengthens.

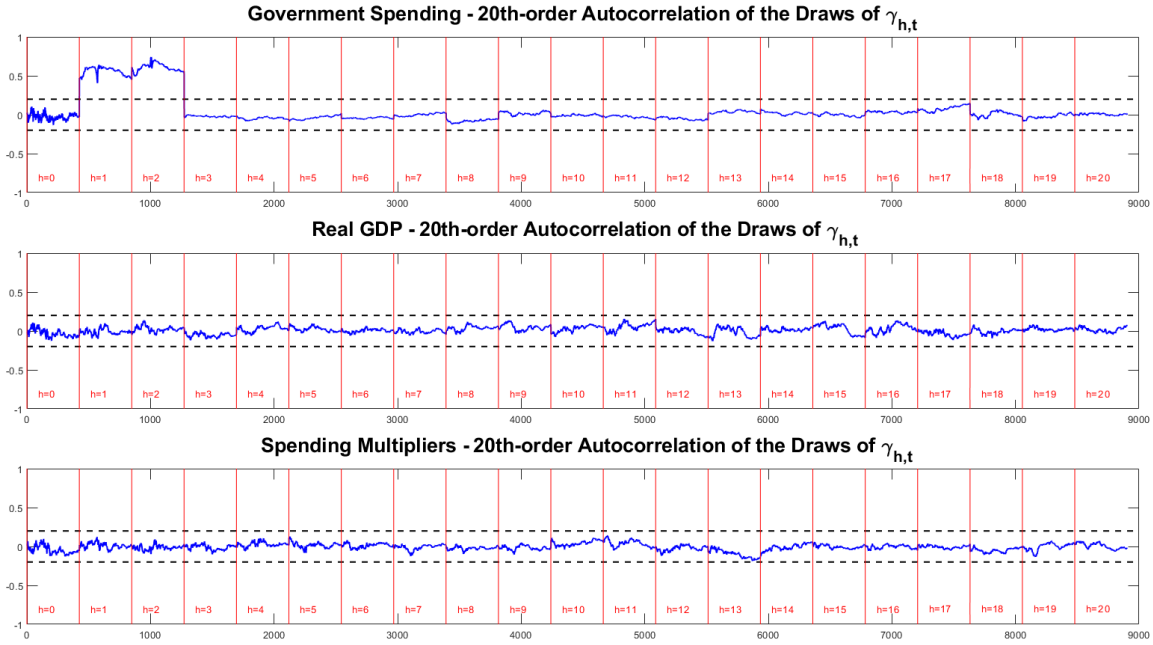

Figure 14: Empirical Application - $20th$-order Autocorrelation of the Retained Draws of $\gamma_{h,t}$ at each Projection Horizon

Figure 14 shows how for both real gross domestic product and for the multipliers, all the autocorrelations of the retained draws lie within the $[-0.2, 0.2]$ interval. However, the same thing cannot be said for the government spending variable. To be more precise, the autocorrelations related to the one and two quarter horizon do not lie within the mentioned interval but show higher values hovering 0.5. To conclude, except these two cases, all the draws appear to be nearly independent. On the basis of these results, and in light of the high dimensionality of the model, the algorithm seems to perform in a satisfactory way.

School of Economics and Finance

Queen Mary
University of London

This working paper has been produced by
the School of Economics and Finance at
Queen Mary University of London

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/