BANK OF ENGLAND

Staff Working Paper No. 677

A time varying parameter structural model of the UK economy

Katerina Petrova, George Kapetanios, Riccardo M Masolo and Matthew Waldron

September 2017

BANK OF ENGLAND

Staff Working Paper No. 677

A time varying parameter structural model of the UK economy

Katerina Petrova,[1] George Kapetanios,[2] Riccardo M Masolo[3] and Matthew Waldron[4]

Abstract

We estimate a time varying parameter structural macroeconomic model of the UK economy, using a Bayesian local likelihood methodology. This enables us to estimate a large open-economy DSGE model over a sample that comprises several different regimes and an incomplete set of data. Our estimation identifies a gradual shift to a monetary policy regime characterised by a marked increase in the responsiveness of monetary policy to inflation alongside a decrease in the level of trend inflation down to the 2% target level. The time varying model also performs remarkably well in forecasting and delivers statistically significant accuracy improvements for most variables and horizons in both point and density forecast performance compared to the standard fixed parameter version.

Key words: DSGE models, Bayesian methods, local likelihood, time varying parameters, forecasting.

JEL classification: C11, C53, E27.

(1) University of St Andrews. Email: katerina.petrova@st-andrews.ac.uk
(2) School of Management and Business, King's College, London.
(3) Bank of England and Centre for Macroeconomics. Email: riccardo.masolo@bankofengland.co.uk
(4) Bank of England. Email: matthew.waldron@bankofengland.co.uk

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Dynamic stochastic general equilibrium (DSGE) models have become popular tools for macroeconomic analysis and forecasting. Their success is the result of their capacity to combine macroeconomic theory with a reasonable fit to business cycle fluctuations in the data and a relatively good forecasting performance. Developments in Bayesian methods coupled with innovations in computing have made it possible for medium- to large-sized DSGE models to be estimated with ease.

An assumption underpinning standard Bayesian estimation of DSGE models, such as the one presented in Smets and Wouters (2007), is that the parameters of the model are constant over time. While over relatively short samples or monotonous periods, this is a reasonable assumption, over longer samples and periods characterised by structural change, it is unlikely to hold. For the United Kingdom, there is a-priori reason to believe that the structure of the economy is likely to have changed substantially over recent decades and, as a result, an assumption that the parameters of a DSGE model have remained constant is unlikely to be valid. For example, recent decades have seen a period of substantial changes to the labor market (e.g. declining unionisation), a large-scale shift in production from manufacturing towards services, substantive changes to the financial sector following e.g. the 'Big Bang' and the recent financial crisis, and a rapid expansion in world trade. They have also included several substantive changes to the conduct of monetary policy, beginning with intermediate monetary aggregate targetting in the 1970s, followed by exchange rate targetting, which was formalised in 1989 when the UK entered the Exchange Rate Mechanism (ERM), and ending with inflation targetting – first by the UK government and then by the Monetary Policy Committee of the Bank of England. In this context, it is difficult to justify the assumption that the structural parameters of a model describing the UK economy have remained constant over the past several decades.

In this paper, we investigate structural change in the UK economy through the lens of a DSGE model using a methodology that allows for the model's parameters to vary over time. The model we use is the small open economy DSGE model developed by Burgess, Fernandez-Corugedo, Groth, Harrison, Monti, Theodoridis and Waldron (2013).[1] Reflecting on the brief discussion of recent UK monetary history above, Burgess et al. (2013) restrict the estimation sample to the 1993Q1-

[1] The model was developed for policy analysis and forecasting at the Bank of England and bears similarities to other open-economy models used in policy institutions, such as Adolfson, Andersson, Linde, Villani and Vredin (2007).

2007Q4 period that begins after the adoption of inflation targetting and ends before the Great Recession. This has the drawback of being a relatively short sample (e.g. compared to similar studies on US data) and may not be representative (with implications for forecasting), given that it only incorporates data from the Great Moderation.[2] We address this shortcoming by extending the sample backwards to 1975 and forwards to 2014.

The literature on estimating reduced form models such as vector autoregressions with time variation in the parameters has become popular with papers such as Cogley and Sargent (2002), Primiceri (2005), Cogley and Sbordone (2008), Benati and Surico (2009), Gali and Gambetti (2009), Canova and Gambetti (2009) and Mumtaz and Surico (2009). One example of a paper that considers a similar research question to ours is Ellis, Mumtaz and Zabczyk (2014) who use a time varying factor augmented VAR to study structural changes in the transmission of monetary policy shocks in the UK. The literature on DSGE models with drifts in the parameters is less developed, possibly due to: (i) the additional complexity that arises from the algorithms used for the solution and estimation of these models, and (ii) the additional assumptions required about the way the agents in the model form expectations about the future parameter values. One way in which time variation in the parameters of a DSGE model has been modelled is by specifying stochastic processes known to agents in the model for a subset of the parameters (e.g. Justiniano and Primiceri (2008), Fernandez-Villaverde and Rubio-Ramirez (2008)). For instance, Fernandez-Villaverde and Rubio-Ramirez (2008) assume that agents in the model take into account future parameter variation when forming their expectations. Similar assumptions are made by Schorfheide (2005), Bianchi (2013), Foerster, Rubio-Ramirez, Waggoner and Zha (2014), but the parameters are modelled as Markov-switching processes. There are two drawbacks from this approach. First, for every time varying parameter, the state vector is augmented and an additional shock is introduced, which increases the complexity of the DSGE model and is subject to the 'curse of dimensionality' so that only a small of subset of the model's parameters can be modelled in this way. Second, it imposes additional structure by relying on the assumption that the law of motion for the parameters' time variation is correctly specified[3].

In contrast, Canova (2006), Canova and Sala (2009) and Castelnuovo (2012) allow for parameter

[2]Due to the challenges discussed, there are only a handful of papers similar to ours in scope. Harrison and Oomen (2010) can be considered a predecessor of Burgess et al. (2013), while DiCecio and Nelson (2007) estimates a closed-economy model.

[3]Petrova (2017) shows in a Monte Carlo exercise that treating parameters as state variables when the law of motion is misspecified may result in invalid estimates of the parameters' time variation, even asymptotically.

variation by estimating DSGE models over rolling samples. In recent work, Galvão, Giraitis, Kapetanios and Petrova (2017) have proposed a new methodology that allows the time-varying Bayesian estimation of large structural models, as demonstrated by the estimation of a Smets and Wouters (2007) DSGE model in Galvão et al. (2016) and (2017). Their approach is an extension and formalisation of rolling window estimation, generalised by combining kernel-generated local likelihoods with appropriately chosen priors to generate a sequence of posterior distributions for the objects of interest over time, following the methodology developed in Giraitis, Kapetanios and Yates (2014), Giraitis, Kapetanios, Wetherilt and Zikes (2016) and Petrova (2017). The advantages of the kernel method are that: (i) it does not require parametric assumptions about the parameters' law of motion and it performs well for many different deterministic and stochastic processes, and (ii) it allows for estimation of drifts in all DSGE parameters. Both the kernel and the rolling window approaches, when applied to structural models, assume that, instead of being endowed with perfect knowledge about the economy's data generating process, agents take parameter variation as exogenous when forming their expectations about the future. This assumption facilitates estimation and can be rationalised from the perspective of models featuring learning problems[4]. For example, Cogley and Sargent (2009) employ Kreps (1998)'s anticipated utility approach, where at each period agents employ their current beliefs as the true (time invariant) parameters. They show that in the presence of parameter uncertainty, the anticipated utility approach outperforms the rational expectation approximation. A recent application of the anticipated utility approach is Johannes, Lochstoer and Mou (2016), where assets are priced at each period, using current posterior means for the parameters and assuming that current values will last indefinitely in the future; thereafter agents learn the new parameter values and adjust their expectations.

In this paper, we employ the Galvão et al. (2017) approach and apply it to COMPASS to investigate the structural nature of the parameters of the model. The flexibility of the approach in the face of structural change permits the estimation of COMPASS over a longer period, alleviating the need to restrict the sample to post-1992 and pre-crisis data. Given that this approach is based on the Kalman filter, it also allows us to deal with missing observations, which is required due to unavailability of some series used in the estimation of COMPASS prior to 1987 (the world output

[4] Note that the estimation technique we employ lends itself to an *anticipated utility* interpretation as it induces smooth (and slow-moving) variation of parameter estimates over time so that, from the perspective of the economic agent, assuming constant parameters to compute expectations is more sensible than if estimation techniques generating more abrupt parameter changes were used.

and export price inflation).

Our results are noteworthy for two main reasons. First, our estimates clearly show evidence of time variation in parameters which translates into changes in the transmission of shocks as well as in the evolution of the relative importance of structural shocks. Second, we demonstrate that our time varying model outperforms its constant parameter counterpart when it comes to forecasting, which is important since forecasting is one of the main uses for COMPASS in the Bank of England (Burgess et al. (2013)).

The remainder of the paper is organised as follows. Section 2 describes the Bayesian Local Likelihood approach of Galvão et al. (2017), Section 3 presents COMPASS, Section 4 contains the empirical results and forecasting comparison and Section 5 concludes.

2 Modelling time variation in DSGE parameters

This section outlines the estimation strategy based on the local quasi-Bayesian Local Likelihood (QBLL) method. The methodology is fully developed in Galvão et al. (2016) and (2017) and Petrova (2017), and we provide a brief description below for ease of reference.

The linearised rational expectation model can be written in the form

$$A(\theta_t)x_{t+1} = B(\theta_t)x_t + C(\theta_t)v_t + D(\theta_t)\eta_{t+1}, \qquad v_t \sim N(0, Q(\theta_t))$$

where x_t is a $n \times 1$ the vector containing the model's endogenous and exogenous variables, v_t is a $k \times 1$ vector of structural shocks, η_{t+1} is an $l \times 1$ vector of expectation errors, θ_t is a vector of parameters, including parameters governing preferences and the shocks' stochastic processes, A, B, C and D are matrix functions of θ_t, and $Q(\theta_t)$ is a diagonal covariance matrix. Observe that the parameter vector θ_t is indexed by time t.

A numerical solution of the rational expectations model can be obtained by one of the available methods (for instance, Blanchard and Kahn (1980) or Sims (2002)). The resulting state equation is given by

$$x_t = F(\theta_t)x_{t-1} + G(\theta_t)v_t \tag{1}$$

where the $n \times n$ matrix F and the $n \times k$ matrix G can be computed numerically for a given parameter vector θ_t. The system is augmented with a measurement equation:

$$Y_t = K(\theta_t) + Z(\theta_t)x_t + \vartheta_t \qquad \vartheta_t \sim N(0, R(\theta_t)) \tag{2}$$

where Y_t is a an $m \times 1$ vector of observables, normally of a smaller dimension than x_t (i.e. $m < n$), Z is an $m \times n$ matrix that links those observables to the latent variables in the model x_t, $K(\theta_t)$ is a vector of time varying intercepts and ϑ_t is a vector of measurement errors with variance $R(\theta_t)$.

Equations (1) and (2) define the state space representation of the model, which is linear and Gaussian. Therefore, the Kalman filter can be employed to recursively build the likelihood of the sample of observables $\{Y_j\}_{j=1}^T$. The likelihood of the sample - the product of the likelihood functions of each observation - is given by

$$L_t(Y|\theta_t) = \prod_{j=1}^{T} L(Y_j|Y^{j-1}, \theta_t)^{w_{tj}} \text{ for } t = 1, ..., T$$

where w_{tj} is an element of the $T \times T$ weighting matrix $W = [w_{tj}]$, computed using a kernel function

$$\tilde{w}_{tj} = K\left(\frac{t-j}{H}\right) \quad \text{for} \quad j, t = 1, ..., T \tag{3}$$

with a bandwidth parameter H. Petrova (2017) shows that in the Bayesian setup the resulting quasi-posterior distributions are asymptotically Normal and valid for confidence interval construction as long as the weights are then normalised to sum to $\varkappa_t := \left(\sum_{j=1}^T w_{tj}^2\right)^{-1}$ for each t, i.e,

$$w_{tj} = \left(\sum\nolimits_{j=1}^T w_{tj}^2\right)^{-1} \left(\tilde{w}_{tj} / \sum_{j=1}^{T} \tilde{w}_{tj}\right) \quad \text{for} \quad j, t = 1, ..., T.$$

The normalisation employed to maintain the relative balance between the likelihood and the prior and to obtain the same rate of convergence as in the frequentist work of Giraitis et al. (2014). In this paper, the normal kernel function

$$K(x) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}x^2\right\} \tag{4}$$

is used to generate the weights w_{tj}.

The local likelihood of the DSGE model at time t, denoted by $L_t(Y|\theta_t)$, is augmented with the prior distribution of the structural parameters, $p(\theta_t)$, to get the posterior at time t, $p(\theta_t|Y)$:

$$p(\theta_t|Y) = \frac{L_t(Y|\theta_t)p(\theta_t)}{p(Y)} \propto \prod_{j=1}^{T} L(Y_j|Y^{j-1}, \theta_t)^{w_{tj}} p(\theta_t).$$

It should be noted that for our DSGE investigation, we assume the prior $p(\theta_t)$ to be fixed over time[5], i.e., $p(\theta_t) = p(\theta)$ for all t.

[5] It is possible to allow the prior to be time-varying. However, since we focus only on the possibility of parameter change driven by the data, we assume that the prior is constant over time.

We use the following algorithm to numerically approximate the time-varying posterior distribution, $p(\theta_t|Y)$:

Step 1 The posterior is log-linearised and passed to a numerical optimisation routine. Optimisation with respect to θ_t is performed to obtain the posterior mode,

$$\widehat{\theta}_t = \arg\min_{\theta} \left(-\sum_{j=1}^{T} w_{tj} \log L(Y_j|Y^{j-1}, \theta_t) - \log p(\theta_t) \right)$$

Step 2 Numerically compute $\widehat{\Sigma}_t$, the inverse of the (negative) Hessian, evaluated at the posterior mode, $\widehat{\theta}_t$.

Step 3 Draw an initial value θ_t^0 from $N(\widehat{\theta}_t, c_0^2\widehat{\Sigma}_t)$.

Step 4 For $k = 1, ..., n_{sim}$, draw ζ_t from the proposal distribution $N(\theta_t^{(k-1)}, c^2\widehat{\Sigma}_t)$.

Compute

$$r(\theta_t^{k-1}, \zeta_t|Y_{1:T}) = p(\zeta_t|Y)/p(\theta_t^{k-1}|Y) = \prod_{j=1}^{T} L(Y_j|Y^{j-1}, \zeta_t)^{w_{tj}} p(\zeta_t) / \prod_{j=1}^{T} L(Y_j|Y^{j-1}, \theta_t^{k-1})^{w_{tj}} p(\theta_t^{k-1}),$$

which is the ratio between the *weighted* posterior at the proposal draw ζ_t and the previous draw θ_t^{k-1}.

The step from $\theta_t^{(k-1)}$ is accepted (setting $\theta_t^k = \zeta_t$) with probability $\tau = min\{1, r(\theta_t^{(k-1)}, \zeta_t|Y_{1:T})\}$ and rejected ($\theta_t^{k-1} = \theta_t^k$) with probability $1 - \tau$.

Once the posterior distribution of the parameters is obtained, out-of-sample forecasts can be generated. For each forecast, we only need the posterior distribution at the end of the corresponding in-sample period. Therefore, for generating DSGE based forecasts, our method is no more computationally intensive than a standard fixed parameter DSGE forecasting: it requires the computation of the posterior only once. The predictive distribution of the sample $p(Y_{T+1:T+h}|Y_{1:T})$, 1 to h horizons ahead, is given by the conditional probability of the forecasts, averaged over all possible values of the parameters, the variables in the state vector at the end of the sample x_T, and all possible future paths of the variables in the state vector $x_{T+1:T+h}$: $p(Y_{T+1:T+h}|Y_{1:T}) =$

$$\int_{(x_T,\theta_T)} \left[\int_{s=x_{T+1:T+h}} p(Y_{T+1:T+h}|s) p(s|x_T, \theta_T, Y_{1:T}) ds \right] p(x_T|\theta_T, Y_{1:T}) p(\theta_T|Y_{1:T}) d(x_T, \theta_T)$$

where $p(\theta_T|Y_{1:T})$ is the posterior of the parameters at the end point T of the sample used. We employ a slightly modified version of the algorithm for generating draws from the predictive distribution outlined in Del Negro and Schorfheide (2013). The algorithm is as follows.

Step 1 Using the saved draws from the posterior at the end of the sample $p(\theta_T|Y_{1:T})$, for every draw $k = 1, .., n_{sim}$ (or for every i^{th} draw), use the Kalman filter to compute the moments of the unobserved variables at T: $p(x_T|\theta_T^k, Y_{1:T})$.

Step 2 Draw a sequence of shocks $v_{T+1:T+h}^k$ from a $N(0, Q(\theta_T^k))$, where $Q(\theta_T^k)$ is a draw from the estimated posterior distribution of the diagonal variance-covariance matrix of the shocks at T. For each draw k from $p(\theta_T|Y_{1:T})$ and from $p(x_T|\theta_T^k, Y_{1:T})$, use the state equation to obtain forecasts for the state variables

$$\widehat{x}_{T+1:T+h}^k = F(\theta_T^k)x_{T:T+h-1}^k + G(\theta_T^k)v_{T+1:T+h}^k.$$

Step 3 Draw a sequence of shocks $\vartheta_{T+1:T+h}^k$ from a $N(0, R(\theta_T^k))$, where $R(\theta_T^k)$ is a draw from the estimated posterior distribution of the measurement error variance-covariance matrix of the shocks at T.

Step 4 Use the forecast simulations for $\widehat{x}_{T+1:T+h}^k$ and the shocks $\vartheta_{T+1:T+h}^k$ in the measurement equation

$$\widehat{Y}_{T+1:T+h}^k = K(\theta_T^k) + Z(\theta_T^k)\widehat{x}_{T+1:T+h}^k + \vartheta_{T+1:T+h}^k.$$

Using the above algorithm, we obtain a predictive density of $m \times n_{sim}$ draws of $\widehat{Y}_{T+1:T+H}^k$, which can be used to obtain numerical approximations of moments, quantiles and densities of the out-of-sample forecasts. Finally, point forecasts can be computed using the mean of the predictive density for each forecasting horizon.

3 Model and Data

We apply our methodology to an operational medium size DSGE model of the UK economy known as COMPASS, which is short for Central Organizing Model for Projection Analysis and Scenario Simulation. The model is presented in detail in Burgess et al. (2013).

For our purposes, a brief description of the model setup and an overview of the key mechanisms is sufficient. In particular, for the most part we will report log-linearised conditions, referring to Burgess et al. (2013) for the derivation from first principles. The economy is made up of five main economic actors: households, firms, the monetary policy maker, the government and the rest of the world. We will briefly describe each of them.

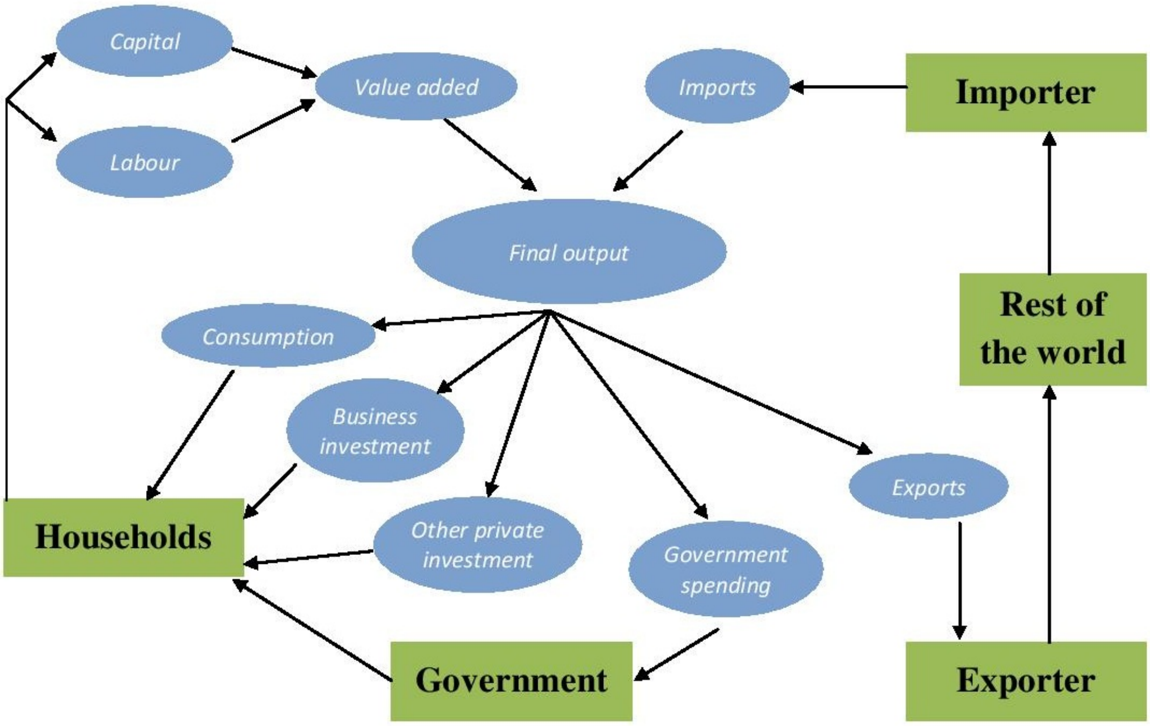

Figure 1. Flow of Goods and Services in the COMPASS

3.1 Households

Households are of two types: optimizing and hand-to-mouth. Optimizing households make the following key economic decisions:

1. *Intertemporal Consumption Decision*

$$c_t = \frac{1}{1+\psi_C+\epsilon_\beta(1-\psi_C)\epsilon_C^{-1}}\left[\mathbb{E}_t c_{t+1}+\psi_C c_{t-1}\right] - \frac{\omega_o\left(1-\psi_C\right)}{\left(1+\psi_C\right)\epsilon_C+\epsilon_\beta(1-\psi_C)}$$
$$\times\left[r_t-\mathbb{E}_t\pi^Z_{t+1}+\hat{\varepsilon}^B_t-\mathbb{E}_t\gamma^Z_{t+1}\right]+(1-\omega_o)\frac{wL}{C}\left[w_t+l_t-\frac{\mathbb{E}_t w_{t+1}+\mathbb{E}_t l_{t+1}+\psi_C\left(w_{t-1}+l_{t-1}\right)}{1+\psi_C+\epsilon_\beta(1-\psi_C)\epsilon_C^{-1}}\right] \tag{5}$$

Equation (5) is a consumption Euler equation. The first term on the RHS contains past consumption because agents' utility is subject to habit formation while the second term contains a risk-premium term $\hat{\varepsilon}^B_t$, which measures a wedge between the rate set by policy-makers and that faced by consumers. The third term on the RHS depends on the fact that a share $(1-\omega_0)$ of the agents are hand-to-mouth so their consumption depends on their labor income.

2. *Investment Decision*

   Optimizing Households can smooth consumption over time by either investing in physical capital or in financial assets. Investment is subject to adjustment costs, which result in the following standard-looking Euler equation:

$$i_t = \frac{\beta\Gamma^H}{1+\beta\Gamma^H}\left(i_{t+1}+\gamma^Z_{t+1}\right) + \frac{1}{1+\beta\Gamma^H}\left(i_{t-1}-\gamma^Z_t\right) + \frac{1}{\left(1+\beta\Gamma^H\right)\left(\Gamma^H\Gamma^Z\Gamma^I\right)^2}\left(\frac{tq_t}{\psi_I}+\hat{\varepsilon}^I_t\right) \tag{6}$$

   where tq_t is Tobin's q value of one unit of capital, which depends on the difference between the future expected returns on capital r^K_t and the real interest rate, adjusted for the risk-premium shock:

$$tq_t = \frac{1-\delta_K}{r^K+(1-\delta_K)}\mathbb{E}_t tq_{t+1} - \left(r_t - \mathbb{E}_t\pi^Z_{t+1} + \hat{\varepsilon}^B_t\right) + \frac{r^K}{r^K+(1-\delta_K)}\mathbb{E}_t r^K_{t+1}$$

3. *Portfolio Decision*

   Households delegate their portfolio decision to risk-neutral portfolio packagers who collect deposits from households and buy domestic and foreign bonds. The end-result is the following UIP condition:

$$q_t = \mathbb{E}_t q_{t+1} + \left(r_t - \mathbb{E}_t\pi^Z_{t+1}\right) - \hat{\varepsilon}^{B^F}_t \tag{7}$$

   which is an arbitrage condition between returns on domestic and foreign assets.

4. *Wage Setting*

   Households supply differentiated labor services in a monopolistically competitive setting. As a result, they have a degree of wage-setting power, i.e. they set their nominal wage at a markup over their marginal-rate of substitution between consumption and leisure (see Erceg, Henderson and Levin (2000)). The wage setting process is also subject to an adjustment cost (Rotemberg (1982)) which, when allowing for indexation to the previous periods' wage rate governed by ξ^W, results in the following wage Phillips Curve:

$$\pi^W_t = \hat{\mu}^W_t + \frac{\hat{\varepsilon}^L_t + \epsilon_L l_t + \frac{\epsilon_C\left(c^o_t - \psi_C c^o_{t-1}\right)}{1-\psi_C} - w_t}{\phi_W\left(1+\beta\Gamma^H\xi^W\right)} + \frac{\xi^W}{1+\beta\Gamma^H\xi^W}\pi^W_{t-1} + \frac{\beta\Gamma^H}{1+\beta\Gamma^H\xi^W}\mathbb{E}_t\pi^W_{t+1}$$

   where the first term on the RHS is the markup, which is allowed to vary over time, and the second represents the marginal rate of substitution.

Hand-to-mouth households do not have access to financial markets by assumption. They will consume their labor income in every period and will receive government transfers to ensure their income grows in line with that of optimizing households along the balanced growth path.

3.2 Firms

The production sector in COMPASS is more complicated than in most medium-size DSGE models (e.g. Smets and Wouters (2007)) because of interactions with the rest of the world and because the model is required to provide a detailed breakdown of various components of GDP.

1. *Value Added Producers*

   This is the most standard of sectors. Firms hire capital and labor, which are used in a Cobb-Douglas production function:

   $$v_t = (1 - \alpha_L)\, k_{t-1} + \alpha_L l_t + \hat{\varepsilon}_t^{TFP} \tag{8}$$

   Firms face monopolistic competition and price adjustment costs which result in the following value-added inflation Phillips Curve:

   $$\pi_t^V = \hat{\mu}_t^V + \frac{1}{\phi_V\left(1 + \beta\Gamma^H\xi_V\right)} mc_t^V + \frac{\xi_V}{1 + \beta\Gamma^H\xi_V}\pi_{t-1}^V + \frac{\beta\Gamma^H}{1 + \beta\Gamma^H\xi_V}\mathbb{E}_t\pi_{t+1}^V \tag{9}$$

2. *Importers*

   Importers buy goods and services from the rest of the world and sell them domestically. They set prices in domestic currency at a markup over the marginal cost and are also subject to Rotemberg-style price adjustment costs. Their pricing decision can thus be summarised by a Phillips curve analogous to that in equation (9).

3. *Final Output Producers*

   They combine value-added output and imports using a Cobb-Douglas technology that mimics that in the Value Added sector. They also face monopolistic competition and price-adjustment costs so that a standard New-Keynesian Phillips curve (similar to that in equation (9)) can be derived for this sector too.

4. *Retailers*

   Retailers operate in a competitive market and transform final output into consumption, business and other investment, government spending and exports, as Figure 1 illustrates. In this way, they operate linear technologies which differ in their productivities. This is a technical expedient to accommodate different trend growth rates in the corresponding observable variables.

5. *Exporters*

   They buy export goods from the corresponding retail sector and sell it to the rest of the world by setting the price for their differentiated goods in the foreign currency. They operate in a monopolistically-competitive market subject to price-adjustment frictions, which results in a Phillips Curve along the lines of equation (9).

3.3 Monetary Policy

In COMPASS policy rates are set according to a simple linear reaction function:

$$r_t = \theta_R r_{t-1} + (1-\theta_R)\left[\theta_\Pi\left(\frac{1}{4}\sum_{j=0}^{3}\pi^Z_{t-j}\right) + \theta_Y \hat{y}_t\right] + \hat{\varepsilon}^R_t \tag{10}$$

which features a response to annual inflation in deviation from its target, the output gap and a degree of interest-rate smoothing governed by θ_R.

3.4 Government Spending

Real-government spending, in deviations from trend, is assumed to follow an autoregressive process:

$$g_t - g_{t-1} + \gamma^Z_t = (\rho_G - 1)\, g_{t-1} + \hat{\varepsilon}^G_t$$

where γ^Z_t measures labor-augmenting productivity and spending is financed via lump-sum taxes on optimising households.

3.5 Rest of the World

We model the UK economy as a small open economy. This implies that world output and prices are independent of domestic shocks, with one important exception, which is necessary for balanced growth: namely that the world economy inherits the domestic permanent labor productivity shock according to a term ω^F_t which ensures the catching up of the world to the domestic productivity shock does not happen instantaneously.

As a result, the world economy is described by three simple equations:

$$z_t^F = \omega_t^F + \rho_{Z^F} z_{t-1}^F + \hat{\varepsilon}_t^{Z^F} \quad (11)$$

$$p_t^{X^F} = \rho_{PX^F} p_{t-1}^{X^F} + \hat{\varepsilon}_t^{PX^F} \quad (12)$$

$$x_t = z_t^F + \hat{\varepsilon}_t^{\kappa^F} - \epsilon_F \left(p_t^{EXP} - p_t^{X^F} \right) \quad (13)$$

which describe world output (which includes the ω_t^F term described above), world prices and the equation governing the demand for UK exports. This is an increasing function of world output and a decreasing function of the prices of UK exports (p_t^{EXP}) relative to world prices, the $\hat{\varepsilon}_t^{\kappa^F}$ terms representing exogenous disturbances.

4 Data

The model is estimated using fifteen macroeconomic quarterly time series[6] for the period from 1975Q1 to 2014Q4, which is considerably longer than the dataset in Burgess et al. (2013). One challenge this presents is that two of the series required for the estimation (world output and the world export price deflator) are unavailable prior to 1987. To circumvent this, we resort to a Kalman filter algorithm that can handle missing observations (see for example Chapter 6 of Harvey (2008)). The variables, data transformations and measurement equations are described in the Appendix. All variables except for the policy rate are log-differenced. As in Burgess et al. (2013), we also remove variable-specific trends from some of the variables, (e.g. exports), to deal with the fact that, while the model permits growth rates to differ across sectors, it implies a set of restrictions that do not hold in the data over the estimation sample[7]. These additional, non- modelled trends are assumed to be constant over the sample[8].

In addition, Burgess et al. (2013) subtract a time-varying trend from inflation as a means to correct for the fact that prior to 1993 there was no explicit inflation target and hence inflation can deviate from its steady state which in their paper is a model's parameter calibrated at 2%[9]. Since our approach can explicitly accommodate structural change and our sample covers a number of

[6] Notice that COMPASS features 18 structural shocks and 7 measurement errors, so the number of shocks is greater than the number of observables.

[7] The economic rationale for these trends is the rapid expansion in world trade over the sample period, which is not captured in the model's supply-side structure – see Section 4.3.1 of Burgess et al. (2013) for further discussion.

[8] These trends are subtracted before the data is taken to the model and so do not depend on parameter estimates.

[9] In Burgess et al. (2013) that has only a marginal effect since it primarily affects the training sample from 1987 to 1992.

different regimes as we outlined in Section 1, we can estimate the time-variation in the inflation steady state coefficient, which we interpret as a measure of trend inflation (Ascari and Sbordone (2014)).

5 A time varying COMPASS Model

5.1 Estimation Results

In this section we present our estimation results of the model estimated with the QBLL method presented above and contrast them with a standard time-invariant full-sample estimation[10]. We use the Random Walk Metropolis algorithm to draw four chains of 220,000 MCMC draws (dropping the first 20,000 and applying a thinning rate of 50%)[11]. We set the MH scaling parameters so that acceptance rates are around 25%[12]. For our time-varying estimation, we apply the QBLL method using the Normal kernel function presented in equation (4). We set the bandwidth $H = \sqrt{T}$, in line with the optimal bandwidth parameter choice used for inference of time-varying random-coefficient models in Giraitis et al. (2014).Figures 2-5 report the posterior mean and 90% confidence intervals of the parameters of our time-varying model as well as the fixed-coefficient specification.

The first point worth highlighting relates to monetary policy. Over time, we can clearly see an increase in the estimated responsiveness of interest rates to inflation, a reduction in the inflation trend which has stabilised around its target level and a decline in the volatility of monetary policy disturbances. All three are normally associated with more effective monetary policy as they are synonymous with an economic environment characterised by low and stable inflation, well anchored around its target[13] (see DiCecio and Nelson (2009) for a detailed comparison of the US and UK experiences). Moreover, our estimate of UK trend inflation is broadly in line with estimates for the US economy (surveyed in Ascari and Sbordone (2014)) with two differences: (i) the peak we estimate in the 1970s is higher than in the US (about 8% in annual terms compared to 5% for the baseline estimate in Ascari and Sbordone (2014)) which is in line with evidence that the Great Inflation

[10] See the Appendix for details on the prior distributions used for both specifications.

[11] This implies an effective number of 400,000 draws after thinning and burning.

[12] This is motivated by Roberts, Gelman and Gilks (1997), who show that the optimal asymptotic acceptance rate is 0.234; their results serve as a rough benchmark in the literature.

[13] The period of ultra low and constant rates, at 50bps between 2009 and the end of our sample is reflected in a marked increase in the interest rate smoothing coefficient as well as in moderate increase in the variability of monetary policy shocks, which increases from the value attained for the estimates centered around the year 2000 but is still well below its constant parameter counterpart.

was more marked in the UK, and (ii) the decline towards the 2% target takes longer to achieve[14]. Moreover, as demonstrated in Section 5.4, allowing the trend inflation coefficient to vary has a significant effect on the both point and density forecasts for CPI inflation, as well as import and export inflation, since the coefficient appears in the intercept of the corresponding measurement equations.

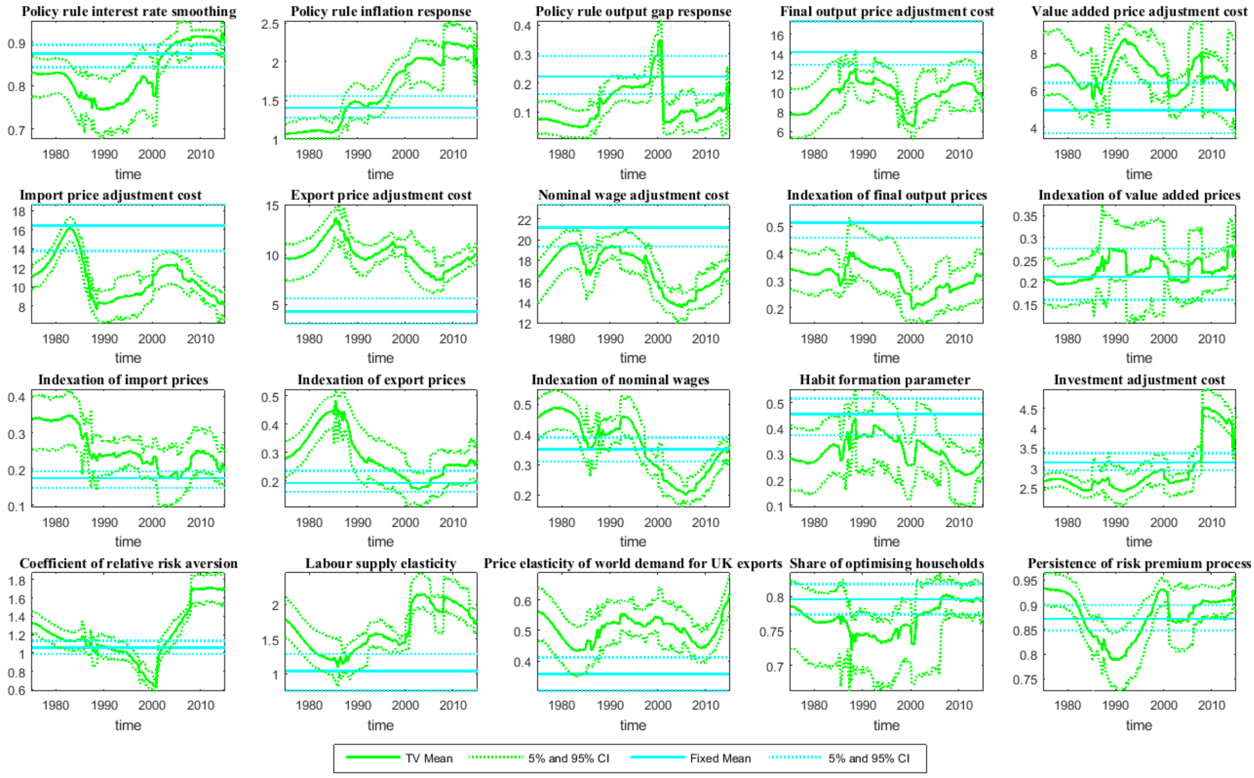

Figure 2. Posterior Estimates. The solid green line is the posterior mean obtained with QBLL, the green dotted lines are the 5% and 95% posterior bands, the solid light blue line is the posterior mean obtained with standard fixed parameter Bayesian estimation and the dotted light blue lines are the corresponding 5% and 95% posterior bands around it.

Through the mid 1980s the policy responsiveness to inflation is close to unity while the estimated annual inflation trend is as high as 8% percent. Over time monetary policy becomes more responsive to inflation variations, the coefficient crossing the 1.5 value popularised by Taylor (1993) around the time of the adoption of the inflation target (1992), while the inflation trend gradually falls from about 3.5% down to its target.

The interest rate smoothing parameter, which enters the model as a coefficient on the lagged policy rate in the Taylor rule, increases during the financial crisis with values close to unity. As a result, during this period, the Taylor rule resembles a difference rule, which is a consequence of

[14] The interpretation of this is not clear-cut because, until 2003, the Bank of England's inflation target was 2.5% on the RPI-X index.

interest rates being close to the Zero Lower Bound.

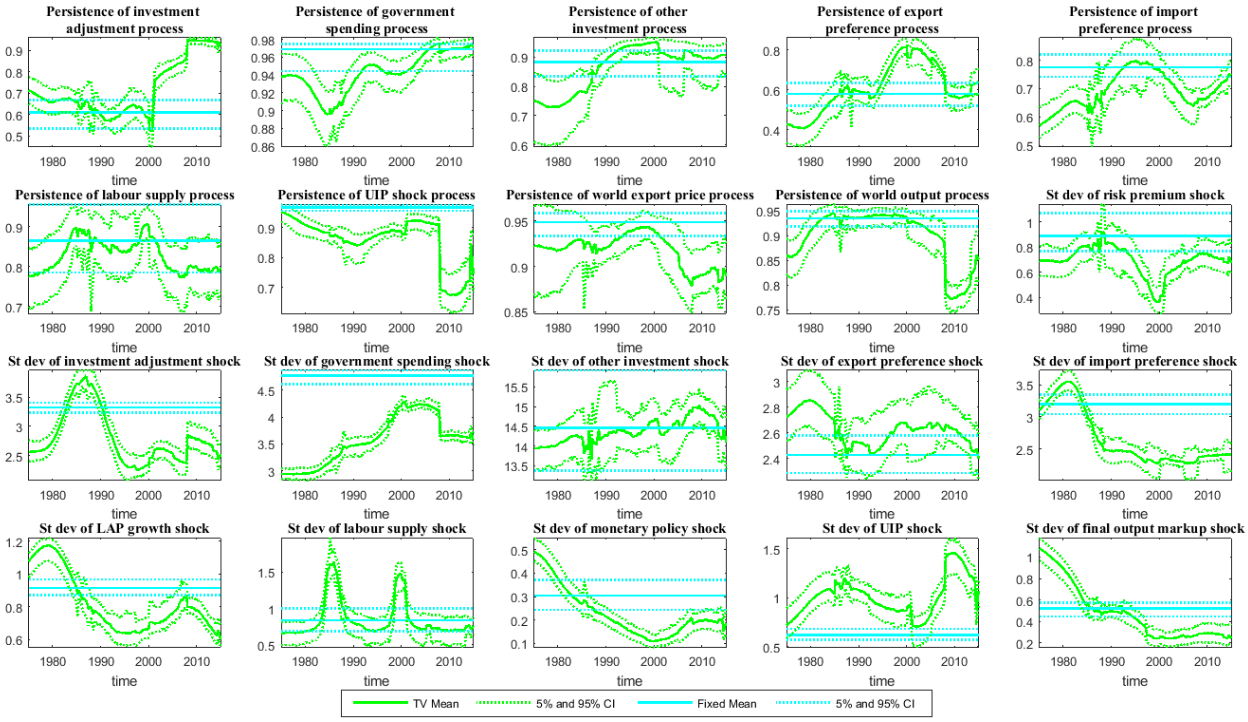

Figure 3. Posterior Estimates. The solid green line is the posterior mean obtained with QBLL, the green dotted lines are the 5% and 95% posterior bands, the solid light blue line is the posterior mean obtained with standard fixed parameter Bayesian estimation and the dotted light blue lines are the corresponding 5% and 95% posterior bands around it.

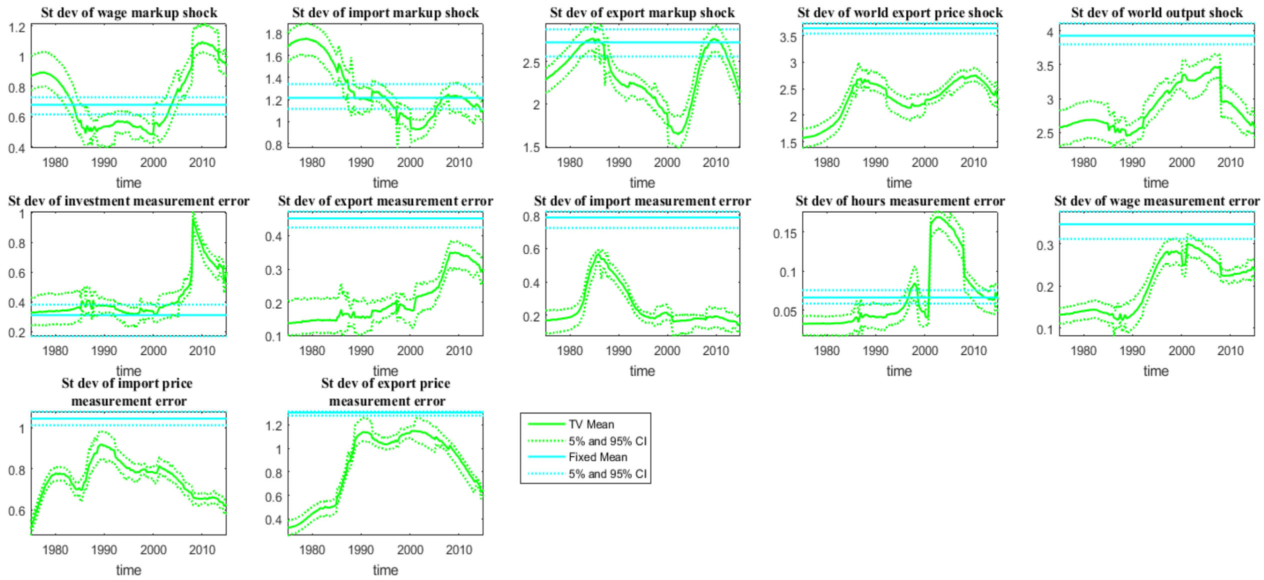

Figure 4. Posterior Estimates. The solid green line is the posterior mean obtained with QBLL, the green dotted lines are the 5% and 95% posterior bands, the solid light blue line is the posterior mean obtained with standard fixed parameter Bayesian estimation and the dotted light blue lines are the corresponding 5% and 95% posterior bands around it.

Another interesting fact is that the reduction in the volatility of shocks is not limited to the monetary policy disturbance. By the mid-1990s, standard deviations for most of the exogenous processes in COMPASS are estimated to be well below their full-sample estimate. Indeed, this pattern emerges gradually and we can thus identify the 1990s as the Great Moderation period in the UK economy. Interestingly, the Great Moderation seems to start later in the UK than in the US.

Two exceptions to this pattern are the wage markup shock, whose spike in volatility over the latest part of our sample captures the peculiarly weak wage growth profile that characterised the UK economy in the aftermath of the Great Recession, and the exchange rate and risk premia shocks. Unsurprisingly, we also observe an increase in the volatility of most structural shocks during the 2008 financial crisis.

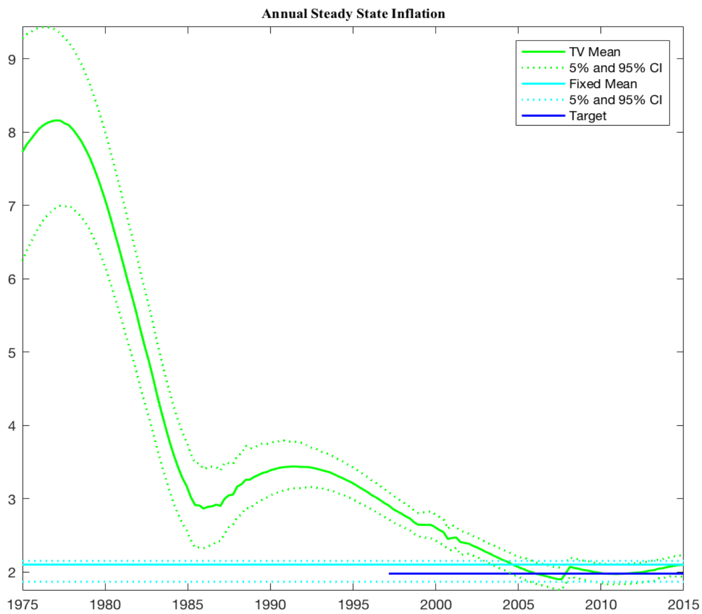

Figure 5. Annual Steady State Inflation Coefficient. The solid green line is the posterior mean obtained with QBLL, the green dotted lines are the 5% and 95% posterior bands, the solid light blue line is the posterior mean obtained with standard fixed parameter Bayesian estimation, the dotted light blue lines are the corresponding 5% and 95% posterior bands around it and the solid dark blue line is the 2% target.

Moreover, both the investment-adjustment cost and the coefficient of relative risk aversion display a significant increase around the financial crisis, which implies that both consumption and investment became less responsive to changes in the policy rate, consistent with a weakening in the transmission mechanism after the crisis.

Finally it is worth noting how the added flexibility built into our estimation procedure has a marked effect on the estimated standard deviations of the measurement errors. With the exception of the investment and hours measurement error series, time-varying estimates for the standard deviations of the measurement error components are lower than their fixed parameter counterparts throughout the sample. This suggests that in the standard time invariant estimation they might not only be picking up noise in the data, but also underlying changes in the structure of the economy.

5.2 Time variation in the monetary transmission mechanism

We can use the time-varying parameters from the estimation to study changes in the monetary transmission mechanism over time. Figures 6 and 7 display the impulse responses for output, prices, the nominal interest rate and the exchange rate to a monetary policy shock in each quarter of the estimation sample.

Figure 6. Impulse response functions of variables to 1 st. dev. of monetary policy shock over time

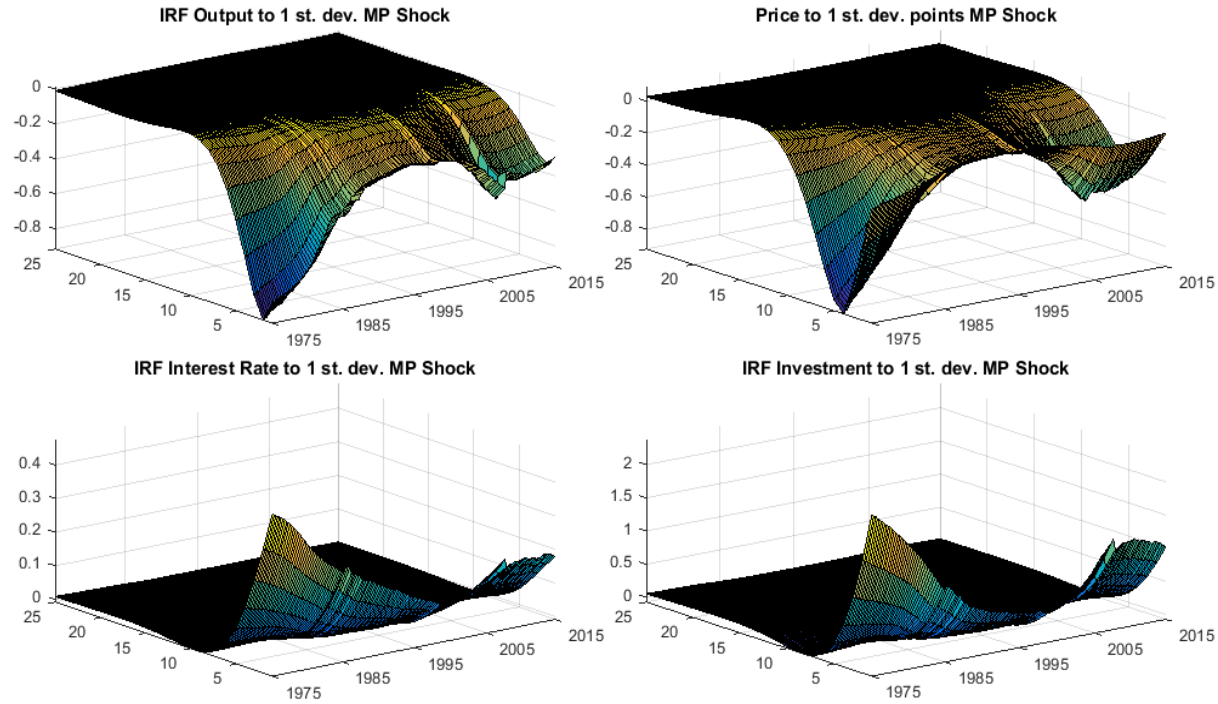

Figure 6 displays responses to a one standard deviation shock and captures the effect of the policy shock on the variables of interest, while also taking into account the changing size of the shock. Responses of output and inflation to monetary surprises are estimated to have been as much as four times as large in the 1970s than around the turn of century. Indeed, these results are consistent with evidence presented in Boivin and Giannoni (2006) for the U.S., who interpret the decreased responsiveness of inflation and output to a monetary policy shock after the 1980s as a result of the monetary authority becoming more effective and systematically more responsive in managing economic fluctuations. This reflects both the change in the systematic component of monetary policy - an increase in the inflation coefficient in the policy response function - and the change in the size of policy surprises - a decline in the standard deviation of the monetary policy shock described above.

Figure 7. Impulse response functions of variables to 25 basis points of monetary policy shock over time

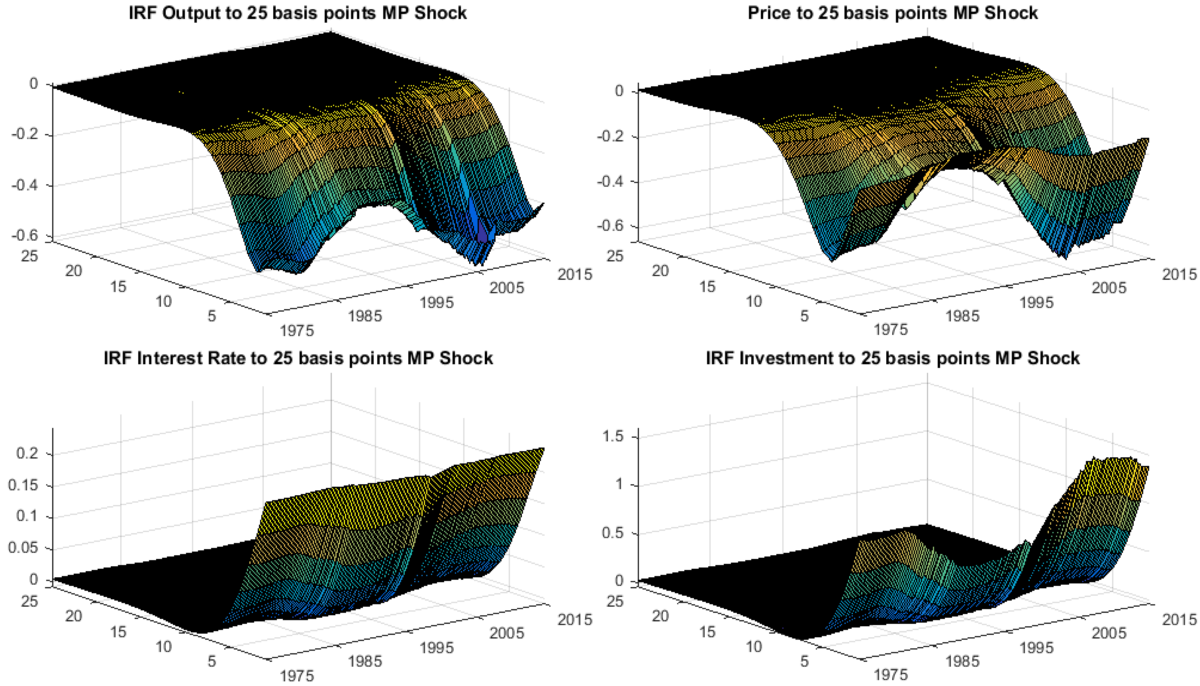

To try and separate the two effects we can contrast the *equally likely*, one standard deviation, shocks in Figure 6 with *equally sized*, 25 basis point, shocks in Figure 7. A noticeable reduction in the effects of a surprise 25 basis-point increase in the policy rate on output, investment and inflation between the 1970s and the 1990s still emerges. Yet, the responses of output, investment and, most notably, inflation show an increase over the most recent period despite the fact that the estimation results suggest that consumption and investment have become less responsive to interest rates. This reflects the marked increase in the policy rule smoothing coefficient (from a value of about 0.75 in the 90s to above 0.9 in the latter part of our sample), which directly increases the persistence of the interest rate. This is in turn partly a consequence of the policy rate having been at or close to its effective lower bound since 2009.

5.3 Time variation in variance decompositions

In this section, we investigate the changing variance decompositions of key model observables over time. Figures 8-10 display the proportion of the variance of GDP growth, inflation and the policy rate explained by the various exogenous shocks over time. At all horizons, the variance of GDP growth is explained primarily by demand shocks. Relative to later in the sample, we see that the risk premium shock plays a prominent role at longer horizons at the beginning of the sample, consistent with their high estimated persistence and with the relatively weaker systematic response of monetary policy to inflation.

Figure 8. Variance decomposition of output growth over time

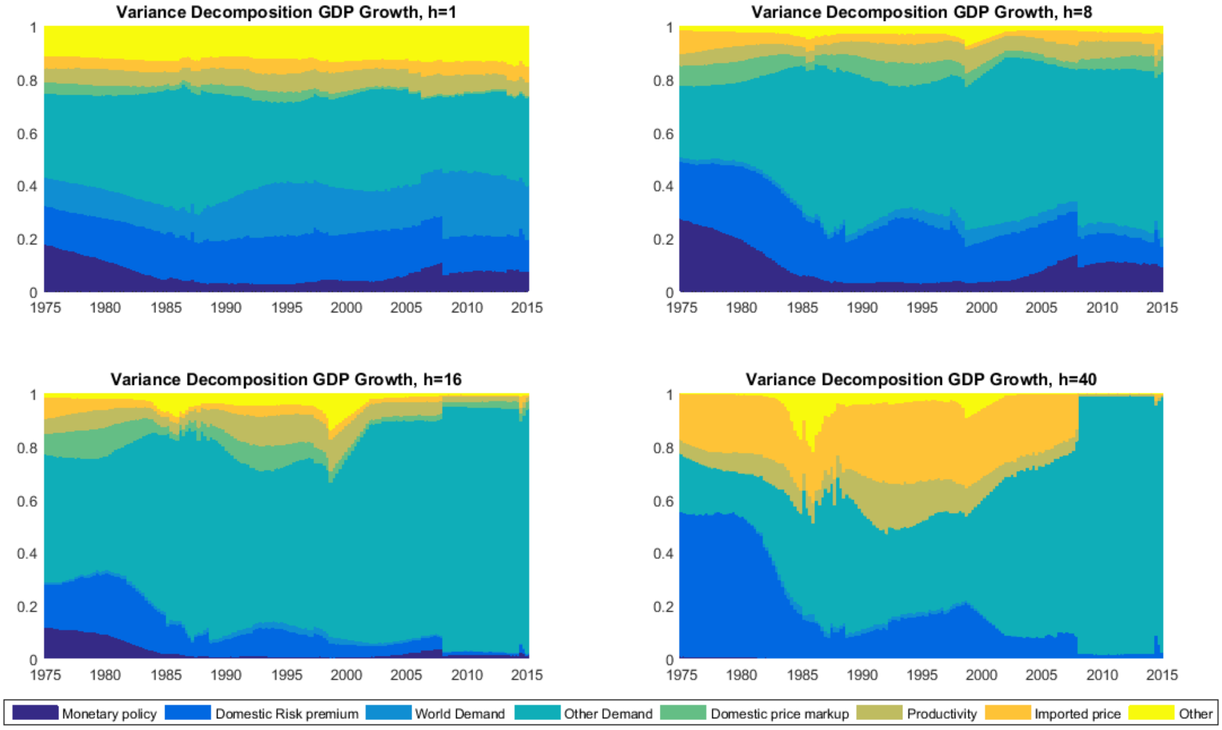

Inflation's variation is absorbed almost entirely by the variance of domestic mark up shocks at one quarter ahead, while at two and four years, we observe that the monetary policy shock also has an effect, especially during the 1970s, early 1980s and the recent financial crisis, while the contribution of risk premia is roughly constant and particularly marked at business cycle frequencies. The pattern of the contribution of 'imported inflation' to the overall variation in headline inflation is particularly interesting. At long horizons, up to three quarters of inflation was explained by foreign factors in the 1970s, consistent with the widely-documented effects that oil prices had on UK inflation at that time.

Figure 9. Variance decomposition of inflation over time

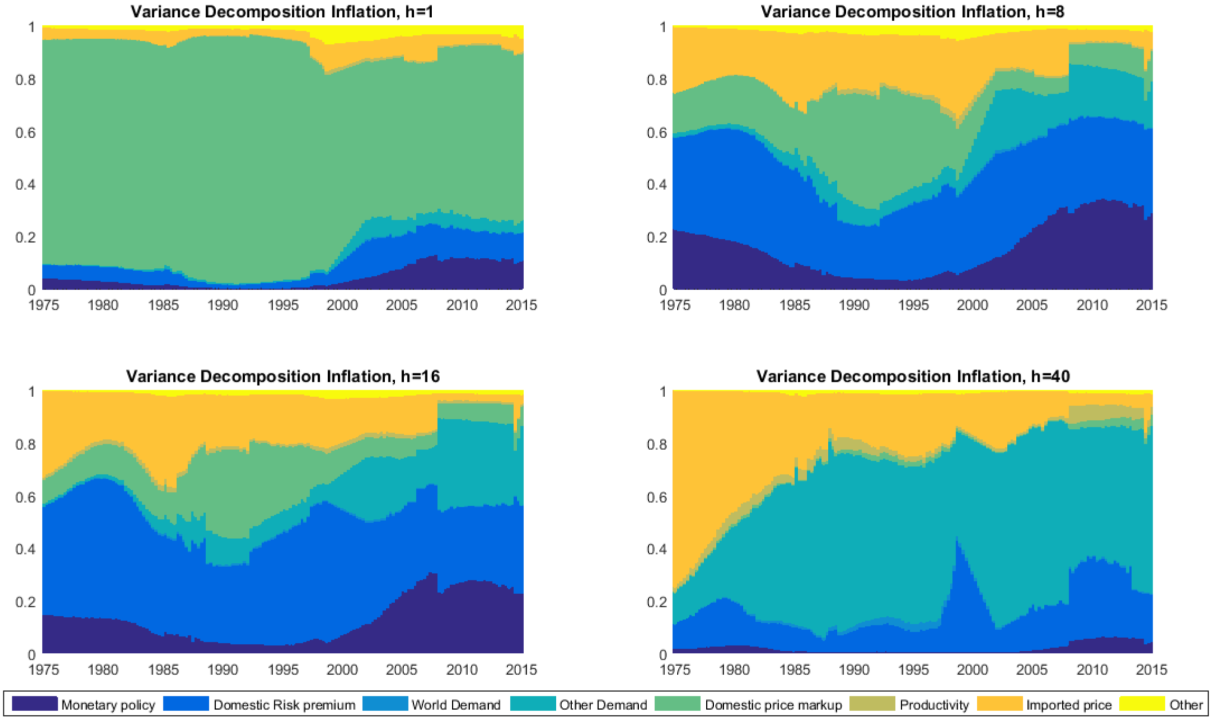

Figure 10. Variance decomposition of policy rate over time

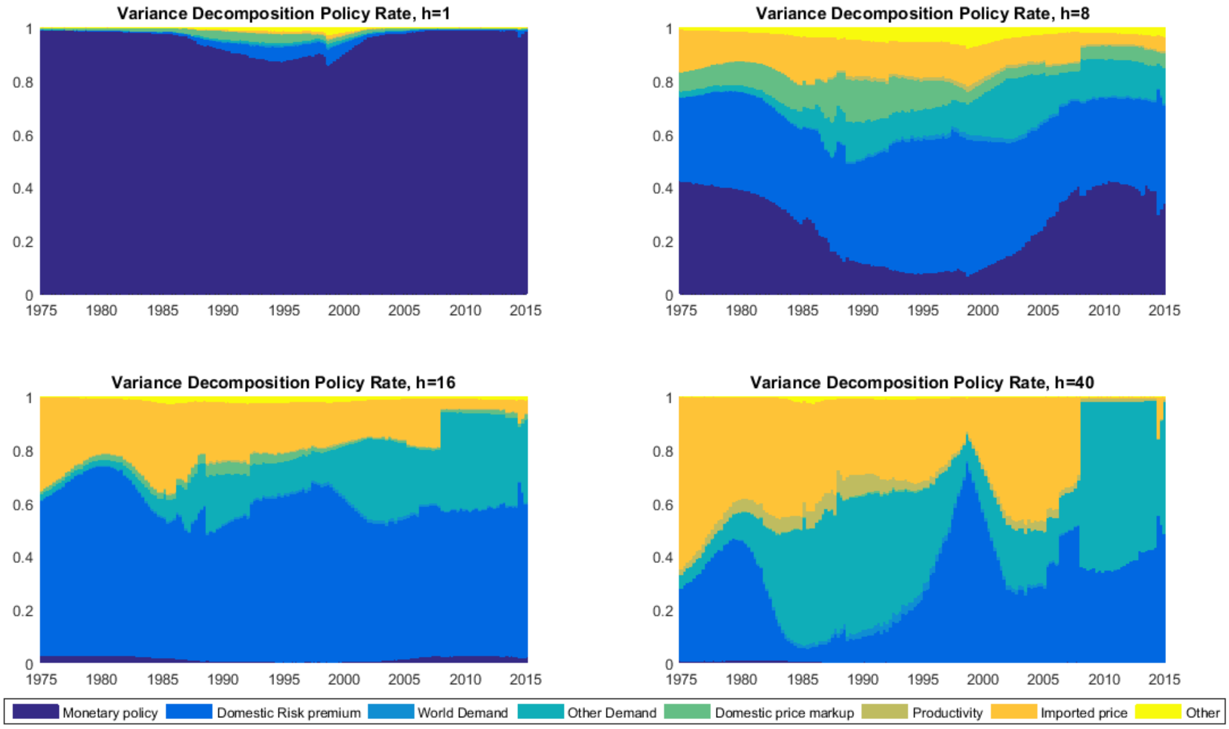

Regarding the variance of the policy rate, it is interesting to note how the overwhelming influence of monetary policy shocks in explaining its variations diminishes at longer horizons while risk premia and other demand shocks take on a much more prominent role, just as expected[15].

5.4 Forecasting

In this section, we evaluate the relative forecasting performance of our time varying parameter COMPASS (TVP-COMPASS) model. In addition, we compare the forecasting record of COMPASS against the fixed-parameter COMPASS (F-COMPASS) specification[16]. We measure accuracy of point forecasts using the root mean squared forecast error (RMSFE). The accuracy of density forecasts are measured by log predictive scores. We compute the logscore with the help of a nonparametric estimator to smooth the draws from the predictive density obtained for each forecast and horizon. We test whether the TVP-COMPASS model is statistically more accurate than the benchmark F-COMPASS with the Diebold and Mariano (1995) statistic computed with the Newey-West estimator to obtain standard errors. We provide the results of the Diebold-Mariano two-sided test for the RMSFEs and logscores.

RMSFEs Forecast Origins: 1985Q1-2012Q4

	horizon	Y	C	I	INFL	EXCH	INT	IM INFL	EX INFL	W INFL	H
F-COMPASS	1	0.90	0.82	5.94	0.60	3.69	0.21	2.45	3.05	1.00	0.99
	2	0.70	0.95	5.25	0.67	3.69	0.38	2.70	2.91	1.07	0.68
	3	0.64	0.99	5.31	0.59	3.63	0.66	2.32	2.45	0.95	0.70
	4	0.64	0.81	5.09	0.47	3.47	0.97	2.00	1.95	0.85	0.66
TVP-COMPASS	1	0.95	0.95	1.04*	0.61*	0.97	0.97	0.80*	0.97	0.94*	0.95
	2	1.00	0.80*	1.04	0.60*	0.99	0.92*	0.81*	0.89*	0.92*	0.98
	3	1.02	0.83*	1.02	0.74*	1.01	0.85*	0.86*	0.88	0.96	0.86
	4	1.09	0.97	1.12	1.02	1.00	0.86*	0.96	0.96	1.10*	0.98

Table 1. RMSFEs. The figures under F-COMPASS are absolute RMSFEs, computed as the mean of the predictive density, the numbers under TVP-COMPASS are ratios over the benchmark fixed parameter COMPASS model. '*', '**' and '***' indicate rejection of the null of equal performance against the two-sided alternative at 10%, 5% and 1% significance level respectively, using a Diebold - Mariano test.

In addition, we also informally assess the density forecast performance of the two models by

[15] Shocks to import prices also play an important role in the early part of our sample via the variation they induce inflation.

[16] See Fawcett, Koerber, Masolo and Waldron (2015) for an evaluation of the forecast performance of a fixed parameter version of COMPASS against statistical and judgemental benchmarks.

looking at the probability integral transformation (PIT) computed as the cumulative density function of the nonparametric estimator for the predictive density at the ex-post realised value of the target variable obtained for each forecast and horizon (Figure 11).

Table 1 presents the absolute performance of the fixed parameter F-COMPASS model (in RMSFEs) and the relative performance of our TVP-COMPASS approach over different horizons (numbers smaller than one imply superior performance of the TVP-COMPASS relative to the F-COMPASS). One, two and three stars indicate that we reject the null of equal accuracy in favour of the better performing model at significance levels of 10%, 5% and 1% respectively.

From Table 1, it is clear that the time varying specification of COMPASS can deliver better point forecasts for most variables and horizons. The gains for inflation forecast accuracy are nearly 40% at short horizons. This better inflation forecast performance can in part be attributed to the considerable time variation uncovered in the inflation trend.

Log Predictive Scores Forecast Origins: 1985Q1-2012Q4

	horizon	Y	C	I	INFL	EXCH	INT	IM INFL	EX INFL	W INFL	H
F-COMPASS	1	-1.60	-1.60	-3.08	-0.97	-2.81	-0.01	-2.35	-2.66	-1.47	-1.78
	2	-1.59	-1.66	-3.00	-1.06	-2.78	-0.53	-2.47	-2.65	-1.53	-1.77
	3	-1.60	-1.70	-3.04	-1.08	-2.78	-1.31	-2.25	-2.62	-1.51	-1.81
	4	-1.63	-1.69	-3.06	-1.07	-2.80	-1.85	-2.13	-2.59	-1.48	-1.83
TVP-COMPASS	1	0.29***	0.32***	-0.07	0.53***	0.14**	0.05	0.07	0.18***	0.11**	0.28***
	2	0.32***	0.40***	0.07	0.53***	0.00	-0.22	0.03	0.23***	0.13**	0.33***
	3	0.32***	0.36***	0.11**	0.38**	0.02	-0.29	0.02	0.28***	0.20***	0.35***
	4	0.33***	0.39***	0.02	0.28	0.08	-0.71	-0.06	0.27***	0.18*	0.35***

Table 2: Log Predictive Scores. The figures under F-COMPASS are absolute log predictive scores, computed as the log of the predictive density evaluated at the ex-post realised observation, the figures under TVP-COMPASS are differences of log scores of the TVP-COMPASS over the benchmark fixed parameter model. '*', '**' and '***' indicate rejection of the null of equal performance against the two-sided alternative at 10%, 5% and 1% significance level respectively, using a Diebold-Mariano test.

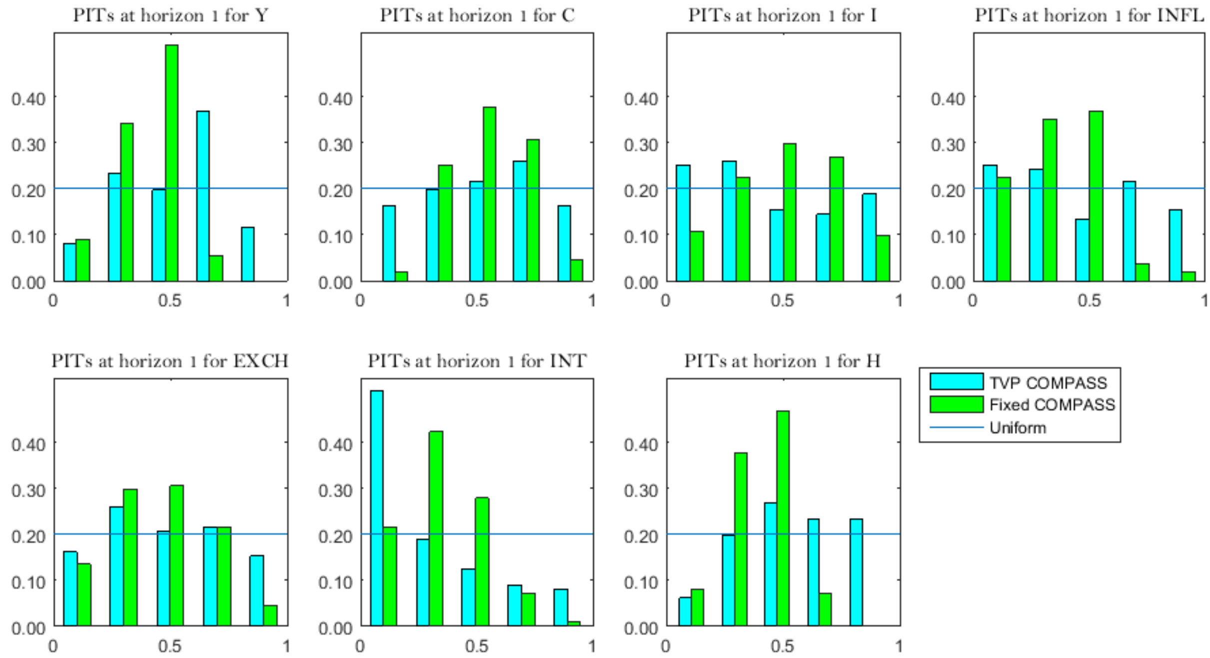

Figure 11. Probability Integral Transformations. The histogram with green bars displays the PITs for the TVP-COMPASS model, computed as the cdf of the predictive density evaluated at the ex-post realised observation, the histogram with blue bars displays the PITs for the F-COMPASS model; the blue dotted line is the cdf of a uniform distribution.

Table 2 accesses the quality of the density forecasts measured by logscores of the predictive density. The table displays absolute log predictive score for the benchmark F-COMPASS model and differences in logscores over the alternative TVP-COMPASS model, so numbers greater than zero imply superior performance of the time-varying parameter model. It is evident from Table 2 that allowing for time variation in the parameters of COMPASS delivers large and statistically significant improvements in the density forecasts for almost all variables and all horizons. This is likely to be a consequence of the ability of the TVP model to capture changes in the volatility of the shocks.

Another way of assessing density forecast performance of the two models is by looking at the probability integral transformation (PITs) in Figure 11, computed as the CDF of the predictive density evaluated at the ex-post realised observation. For a well-calibrated density forecast and a long enough sample, the outturns in all parts of the distribution at all frequencies should match the relevant probabilities, implying uniform PITs. The one-step ahead[17] PITs in Figure 11 for selected variables reveal that neither model is very close to delivering a uniform CDF. However,

[17] For the sake of breivity, we report the one step ahead PITs. The results for other horizons reveal similar pattern and are available upon request.

the TVP-COMPASS model is closer to uniform than the F-COMPASS variant, suggesting that allowing for time variation improves forecast density accuracy at the one-step ahead horizon at least to some degree.

6 Conclusion

Standard Bayesian estimation of DSGE models assume that the parameters are time invariant. Given that the UK economy has undergone substantial structural changes over recent decades, not least those associated with changes in monetary regime, the constant parameter assumption is likely to be invalid except in short sub-samples.

To address this shortcoming, in this paper we applied a nonparametric procedure developed by Galvão et al. (2017) to an open economy DSGE model of the UK. We also modify the Kalman filter in the procedure in order to handle missing observations, which allows us to deal with the data unavailability of some series prior to 1987 and to extend the estimation sample back to 1975.

Our estimation detects the transition to a monetary policy regime characterised by long-term inflation expectations anchored at the target, an increased responsiveness of policy rates to inflation and a reduction in the importance of the non-systematic component of monetary policy.

In our forecasting exercise we demonstrate that allowing for time-variation improves both point and density forecast performance in a statistically significant way for most variables and horizons.

References

Adolfson, M., Andersson, M., Linde, J., Villani, M. and Vredin, A. (2007). Modern forecasting models in action: Improving macroeconomic analyses at Central Banks, *International Journal of Central Banking* **3**(4): 111–144.

Ascari, G. and Sbordone, A. M. (2014). The macroeconomics of trend inflation., *Journal of Economic Literature* **52**(3): 679 – 739.

Benati, L. and Surico, P. (2009). VAR analysis and the Great Moderation, *American Economic Review* **99**(4): 1636–1652.

Bianchi, F. (2013). Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics, *Review of Economic Studies* **80**: 463–490.

Blanchard, O. J. and Kahn, C. M. (1980). The solution of linear difference models under rational expectations, *Econometica* **48**(5): 1305–1312.

Boivin, J. and Giannoni, M. (2006). Has monetary policy become more effective?, *Review of Economics and Statistics* **88**(3): 445–462.

Burgess, S., Fernandez-Corugedo, E., Groth, C., Harrison, R., Monti, F., Theodoridis, K. and Waldron, M. (2013). The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models, *Bank of England Working Paper Series, No471* .

Canova, F. (2006). Monetary policy and the evolution of the US economy, *CEPR Discussion Papers 5467* .

Canova, F. and Gambetti, L. (2009). Structural changes in the US economy: Is there a role for monetary policy, *Journal of Economic Dynamics and Control* **33**(2): 477–490.

Canova, F. and Sala, L. (2009). Back to square one: Identifcation issues in DSGE models, *Journal of Monetary Economics* **56**(4): 431–449.

Castelnuovo, E. (2012). Estimating the evolution of money's role in the U.S. monetary business cycle, *Journal of Money, Credit and Banking* **44**(1): 23–52.

Cogley, T. and Sargent, T. J. (2002). Evolving post-World War II U.S. inflation dynamics, *in* B. S. Bernanke and K. Rogoff (eds), *NBER Macroeconomics Annual*, MIT Press: Cambridge, pp. 331–88.

Cogley, T. and Sargent, T. J. (2009). Anticipated utility and rational expectations as approximations of Bayesian decision making, *International Economic Review* **49**: 185–221.

Cogley, T. and Sbordone, A. M. (2008). Trend inflation, indexation, and inflation persistence in the new Keynesian Phillips curve, *American Economic Review* **95**(5): 2101–2126.

Del Negro, M. and Schorfheide, F. (2013). DSGE model-based forecasting, *in* G. Elliott and A. Timmermann (eds), *Handbook of Economic Forecasting*, Vol. 2, Elsevier, pp. 57–140.

DiCecio, R. and Nelson, E. (2007). An estimated DSGE model for the United Kingdom, *Working Papers 2007-006*, Federal Reserve Bank of St. Louis.

DiCecio, R. and Nelson, E. (2009). The Great Inflation in the United States and the United Kingdom: Reconciling Policy Decisions and Data Outcomes, *NBER Working Papers 14895*, National Bureau of Economic Research, Inc.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy, *Journal of Business and Economic Statistics* **13**: 253–263.

Ellis, C., Mumtaz, H. and Zabczyk, P. (2014). What lies beneath? a time-varying favar model for the uk transmission mechanism, *The Economic Journal* **124**(576): 668–699.

Erceg, C. J., Henderson, D. W. and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts, *Journal of Monetary Economics* **46**(2): 281–313.

Fawcett, N., Koerber, L., Masolo, R. and Waldron, W. (2015). Evaluating UK point and density forecasts from an estimated DSGE model: the role of off-model information over the financial crisis, *Bank of England Working Paper Series, No538* .

Fernandez-Villaverde, J. and Rubio-Ramirez, J. F. (2008). How structural are the structural parameters, *in* K. R. D. Acemoglu and M. Woodford (eds), *NBER Macroeconomics Annual 2007*, Vol. 22, Chicago: University of Chicago Press.

Foerster, A. T., Rubio-Ramirez, J. F., Waggoner, D. F. and Zha, T. A. (2014). Perturbation methods for Markov-switching DSGE models, *NBER Working Paper 258* .

Gali, J. and Gambetti, L. (2009). On the sources of the Great Moderation, *American Economic Journal* **1**(1): 26–57.

Galvão, A. B., Giraitis, L., Kapetanios, G. and Petrova, K. (2016). A time varying DSGE model with financial frictions, *Journal of Empirical Finance* **38**: 690–716.

Galvão, A. B., Giraitis, L., Kapetanios, G. and Petrova, K. (2017). A quasi-Bayesian local likelihood method for medelling parameter time variation in DSGE models, *Working Paper* .

Giraitis, L., Kapetanios, G., Wetherilt, A. and Zikes, F. (2016). Estimating the dynamics and persistence of financial networks, with an application to the Sterling money market, *Journal of Applied Econometrics* **31**(1): 58–84.

Giraitis, L., Kapetanios, G. and Yates, T. (2014). Inference on stochastic time-varying coefficient models, *Journal of Econometrics* **179**(1): 46–65.

Harrison, R. and Oomen, Ã. (2010). Evaluating and estimating a DSGE model for the United Kingdom, *Bank of England working papers 380*, Bank of England.

Harvey, C. A. (2008). *Forecasting, Structural Time Series*, University of Cambridge Press.

Johannes, M., Lochstoer, L. and Mou, Y. (2016). Learning about consumption dynamics, *Journal of Finance* **71**(2): 551–600.

Justiniano, A. and Primiceri, G. E. (2008). The time-varying volatility of macroeconomic fluctuations, *American Economic Review* **98**(3): 604–641.

Kreps, D. (1998). Anticipated utility and dynamic choice, *Frontiers of Research in Economic Theory, Cambridge University Press* pp. 242–274.

Mumtaz, H. and Surico, P. (2009). Time-varying yield curve dynamics and monetary policy, *Journal of Applied Econometrics* **24**(6): 895–913.

Petrova, K. (2017). A quasi-bayesian local likelihood approach to time varying parameter VAR models, *Working Paper* .

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

Roberts, G., Gelman, A. and Gilks, W. (1997). Weak convergence and optimal scaling of random walk metropolis algorithms, *The Annals of Applied Probability* **7**(1): 110–120.

Rotemberg, J. (1982). Monopolistic price adjustment and aggregate output, *Review of Economic Studies* **49**(4): 517–531.

Schorfheide, F. (2005). Learning and monetary policy shifts, *Review of Economic Dynamics* **8**(2): 392–419.

Sims, C. (2002). Solving linear rational expectations models, *Computational Economics* **20**(1-2): 1–20.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach, *American Economic Review* **97**(3): 586–606.

7 Appendix

Parameter	Description	Value
ω_{CZ}	Steady state share of consumption in final output	0.5031
ω_{IZ}	Steady state share of business investment in final output	0.0845
ω_{IOZ}	Steady state share of ‘other’ investment in final output	0.0370
ω_{GZ}	Steady state share of goverment spending in final output	0.1662
ω_{XZ}	Steady state share of exports in final output	0.2092
ω_{MZ}	Steady state share of imports in final output	0.2197
Γ^H	Trend population growth	1.0020
Γ^Z	Trend productivity growth	1.0070
Γ^I	Trend investment growth relative to final output growth	1.0036
Γ^X	Trend export growth relative to final output growth	1.0025
Γ^G	Trend government spending growth relative to final output	0.9950
β	Household discount factor	0.9986
ω_{LV}	Steady state labour share	0.6774
ω_{VZ}	Steady state value added share	0.7599
μ^Z	Steady state final output price mark-up	1.0050
δ_K	Capital depreciation rate	0.0077
ζ_{ω_F}	Speed at which rest of the world inherits LAP shocks	0.9000
β_{factor}	‘Over-discounting’ factor	0.0100
σ_{μ^V}	Standard deviation of value added price mark-up shock	0.0500
σ_{TFP}	Standard deviation of TFP shock	0.0500
ρ_{TFP}	Persistence of TFP forcing process	0.9000
ρ_{LAP}	Persistence of LAP forcing process	0.0000

Table 1: Calibrated Parameters

Parameter	Description	Distribution	Prior Mean	Std
Π^*	Inflation Target	Normal	1.005	0.25
θ_R	Policy rule interest rate smoothing	Beta	0.800	0.100
θ_Π	Policy rule inflation response	Normal	1.500	0.250
θ_Y	Policy rule output gap response	Beta	0.125	0.075
ϕ_Z	Final output price adjustment cost	Gamma	7.000	2.000
ϕ_V	Value added price adjustment cost	Gamma	7.000	2.000
ϕ_M	Import price adjustment cost	Gamma	10.00	2.000
ϕ_X	Export price adjustment cost	Gamma	10.00	2.000
ϕ_W	Nominal wage adjustment cost	Gamma	14.00	2.000
ξ_Z	Indexation of final output prices	Beta	0.25	0.075
ξ_V	Indexation of value added prices	Beta	0.25	0.075
ξ_M	Indexation of import prices	Beta	0.25	0.075
ξ_X	Indexation of export prices	Beta	0.25	0.075
ξ_W	Indexation of nominal wages	Beta	0.25	0.075
ψ_C	Habit formation parameter	Beta	0.70	0.150
ψ_I	Investment adjustment cost	Gamma	2.00	0.400
ϵ_C	Coefficient of relative risk aversion	Gamma	1.50	0.200
ϵ_L	Labour supply elasticity	Gamma	2.00	0.300
ϵ_F	Price elasticity world demand, UK exports	Gamma	0.75	0.100
ω_o	Share of optimising households	Beta	0.70	0.050
ρ_B	Persistence of risk premium forcing process	Beta	0.75	0.100
ρ_I	Persistence of investment adjustment shock	Beta	0.75	0.100
ρ_G	Persistence of government spending shock	Beta	0.90	0.050
ρ_{IO}	Persistence of other investment shock	Beta	0.75	0.100
ρ_{κ^F}	Persistence of export preference shock	Beta	0.75	0.100
ρ_M	Persistence of import preference shock	Beta	0.75	0.100
ρ_L	Persistence of labour supply shock	Beta	0.75	0.100

Table 2: Priors and posteriors for estimated parameters

Parameter	Description	Distribution	Prior Mean	Std
ρ_{B^F}	Persistence of UIP shock	Beta	0.75	0.10
ρ_{PX^F}	Persistence of world export price shock	Beta	0.90	0.05
ρ_{Z^F}	Persistence of world output shock	Beta	0.90	0.05
σ_B	St dev of risk premium shock	Gamma	0.50	0.20
σ_I	St dev of investment adjustment shock	Gamma	1.90	0.20
σ_G	St dev of government spending shock	Gamma	3.00	0.20
σ_{IO}	St dev of other investment shock	Gamma	14.0	1.00
σ_{κ^F}	St dev of export preference shock	Gamma	2.20	0.20
σ_M	St dev of import preference shock	Gamma	2.20	0.20
σ_{LAP}	St dev of LAP growth shock	Gamma	0.35	0.10
σ_L	St dev of labour supply shock	Gamma	0.75	0.20
σ_R	St dev of monetary policy shock	Gamma	0.10	0.10
σ_{B^F}	St dev of UIP shock	Gamma	0.65	0.20
σ_{μ^Z}	St dev of final output markup shock	Gamma	0.10	0.10
σ_{μ^W}	St dev of wage markup shock	Gamma	0.30	0.10
σ_{μ^M}	St dev of import markup shock	Gamma	1.30	0.20
σ_{μ^X}	St dev of export markup shock	Gamma	1.30	0.20
σ_{PX^F}	St dev of world export price shock	Gamma	1.60	0.20
σ_{Z^F}	St dev of world output shock	Gamma	2.50	0.20
σ_i^{me}	St dev of investment measurement error	Gamma	0.35	0.10
σ_X^{me}	St dev of export measurement error	Gamma	0.18	0.055
σ_M^{me}	St dev of import measurement error	Gamma	0.18	0.055
σ_L^{me}	St dev of hours measurement error	Gamma	0.045	0.013
σ_W^{me}	St dev of wage measurement error	Gamma	0.125	0.0275
σ_{PM}^{me}	St dev of import price measurement error	Gamma	0.34	0.075
σ_{PX}^{me}	St dev of export price measurement error	Gamma	0.34	0.075

Table 3: Priors and posteriors for estimated parameters

Variable	Description	Data transformation equation	Measurement equation
gdpkp	Real GDP	$\text{dlngdpkp}_t \equiv 100\Delta \ln \text{gdpkp}_t$	$\Delta v_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \left(\Gamma^X \right)^{-\frac{1-\alpha_V}{\alpha_V}} \right)$
ckp	Real cons.	$\text{dlnckp}_t \equiv 100\Delta \ln \text{ckp}_t$	$\Delta c_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \right)$
ikkp	Real inv.	$\text{dlnikkp}_t \equiv 100\Delta \ln \text{ikkp}_t$	$\Delta i_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \Gamma^I \right) + \sigma_I^{me} me_t^I$
gonskp	Real spending	$\text{dlngonskp}_t \equiv 100\Delta \ln \text{gonskp}_t$	$\Delta g_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \Gamma^G \right)$
xkp	Real exports	$\text{dlnxkp}_t \equiv 100\Delta \ln \text{xkp}_t - \text{dlnxkp}_t^{\text{tt}}$	$\Delta x_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \Gamma^X \right) + \sigma_X^{me} me_t^X$
mkp	Real imports	$\text{dlnmkp}_t \equiv 100\Delta \ln \text{mkp}_t - \text{dlnmkp}_t^{\text{tt}}$	$\Delta m_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \Gamma^X \right) + \sigma_M^{me} me_t^M$
pxdef	Export deflator	$\text{dlnpxdef}_t \equiv 100\Delta \ln \text{pxdef}_t - \Pi_t^{*,\text{tt}} - \Pi_t^{x,\text{tt}}$	$\Delta p_t^{EX} - \Delta q_t + \pi_t^Z + 100 \ln \frac{\Pi^*}{\Gamma^X} + \sigma_{PX}^{me} me_t^{PX}$
pmdef	Import deflator	$\text{dlnpmdef}_t \equiv 100\Delta \ln \text{pmdef}_t - \Pi_t^{*,\text{tt}} - \Pi_t^{m,\text{tt}}$	$\pi_t^M + 100 \ln \frac{\Pi^*}{\Gamma^X} + \sigma_{PM}^{me} me_t^{PM}$
awe	Nom. wage	$\text{dlnawe}_t \equiv \Delta \ln \text{awe}_t - \Pi_t^{*,\text{tt}}$	$\Delta w_t + \gamma_t^Z + \pi_t^Z + 100 \ln \left(\Gamma^Z \Pi^* \right) + \sigma_W^{me} me_t^W$
cpisa	SA CPI	$\text{dlncpisa}_t \equiv 100\Delta \ln \text{cpisa}_t - \Pi_t^{*,\text{tt}}$	$\pi_t^C + 100 \ln \Pi^*$
rga	Bank Rate	$\text{robs}_t \equiv 100 \ln \left(1 + \frac{\text{rga}_t}{100} \right)^{\frac{1}{4}} - \Pi_t^{*,\text{tt}}$	$r_t + 100 \ln R$
eer	Sterling ERI	$\text{dlneer}_t \equiv 100\Delta \ln \text{eer}_t$	$\Delta q_t - \pi_t^Z$
hrs	Hours worked	$\text{dlnhrs}_t \equiv 100\Delta \ln \text{hrs}_t$	$\Delta l_t + 100 \ln \Gamma^H + \sigma_L^{me} me_t^L$
yf	World output	$\text{dlnyf}_t \equiv 100\Delta \ln \text{yf}_t - \text{dlnyf}_t^{\text{tt}}$	$\Delta z_t^F + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H \right)$
pxfdef	World exp. def.	$\text{dlnpxfdef}_t \equiv 100\Delta \ln \text{pxfdef}_t - \Pi_t^{\text{xf,tt}}$	$\Delta p_t^{X^F} + 100 \ln \frac{\Pi^*}{\Gamma^X}$

Table 4: Observables, data transformation and measurement equations