

Estimating the Output, Inflation and Unemployment Gaps in Ireland using Bayesian Model Averaging

Michael O’Grady*

Central Bank of Ireland

Abstract: We aim to overcome the issue of model selection in output, wage inflation and unemployment gap estimation for Ireland, using Bayesian model averaging. Employing a stochastic model specification search with time-varying parameters approach, we draw from a number of standard model specifications, based on variable selection, trend output identification and distributional assumptions. From the resulting model averaging with Irish data, we find that the unemployment gap is a strong predictor of the output gap, but conditional on the unemployment gap, the output gap has limited influence on the wage inflation gap. Additionally, we observe a decline in potential output growth from the early 2000s, although growth rates have increased strongly since Q1 2012. Finally we find that shocks to output growth and wage inflation are better characterised by Student’s t -distributions, rather than conventional Gaussian distributions, suggesting that extreme events occur with a more relative frequency that is typically assumed.

I INTRODUCTION

The related concepts of potential output, trend wage inflation, the non-accelerating wage rate of unemployment (NAWRU) and the output, unemployment and wage inflation gaps are key metrics for informing policymakers on the cyclical position of the economy, as well as its productive capacity. These measures play an important role in determining the stance of both microeconomic and macroeconomic policy, with consistent estimation across relevant institutions necessary to keep policy co-ordinated. From a European perspective, potential output and the output gap form key components of the fiscal surveillance process

Acknowledgements: The opinions expressed in this paper are purely those of the author and do not necessarily represent the views of the Central Bank of Ireland or the ESCB. We would like to thank Philip Lane, Luca Onorante, Paul Reddan, and members of the CBI Irish Economic Analysis Division, and participants in the IFAC 2018 Path for the Public Finances Conference for helpful comments and suggestions.

* michael.ogrady@centralbank.ie

resulting from the Stability and Growth Pact, in evaluating the effectiveness of structural reforms. Potential output can be seen as a summary indicator of the capacity to generate long-run, wage inflation-neutral growth, while the output gap can be considered a proxy for the degree of economic overheating or underperformance, relative to this growth potential.

A primary concern regarding the use of these measures lies in their estimation. As potential output is an unobserved variable, it must be determined using either a purely empirical approach, or through the specification of a theoretical model. Either approach requires a number of assumptions to be made: purely empirical models are dependent on the selection of observable variables, estimation techniques, constraint specifications and the choice of trend/cycle decomposition tool (Álvarez and Gómez-Loscos, 2018); theoretical models require the specification of the main equations of the model, long-run equilibrium conditions, short-run adjustment mechanisms and the choice of estimation or calibration (Kiley, 2013; Fueki *et al.*, 2016). These considerations are non-trivial, as the optimal policy response is dependent on the underlying conditions behind a perceived output gap. A large output gap with strong underlying trend growth rates requires the use of demand-side policy instruments, while output gaps caused by declining trend growth rates necessitate the implementation of supply-side policy measures (Coibion *et al.*, 2017).

When estimating the output gap for the Irish economy, additional complexities need to be considered. First, there are twin problems of data duration and revisions. As highlighted by Bergin and FitzGerald (2014), revisions to existing calculations and the availability of new data can cause considerable changes to the estimates of potential output for a given year. Additionally, as reliable output data only extends back to the 1960s for the Irish economy (on an annual basis), there are a large class of models for which small-sample problems could affect both the estimated error distributions and parameter values, leading to model misspecification and biased estimates of potential output and the output gap. Further issues may arise in models that are specified and calibrated to allow for comparison across European countries, given Ireland's unique role as a production and financial intermediary (Lane, 2014), combined with asymmetric economic and financial linkages (Lane and Ruane, 2006), and its relative vulnerability to heightened volatility in international capital flows and skilled labour demand as a small open economy (Byrne and O'Brien, 2017).

Consequently, there are a number of aims to this paper. We first attempt to determine the degree to which fundamental macroeconomic relationships (Okun's law, the wage Philips curve) hold for the Irish economy. As the outcome of this is highly conditional on the set of equations being estimated, we employ an approach that allows for both a large number of possible specifications, but also weights the contribution of each specification by its likelihood, given the underlying data. Thus, we allow for the data to drive our specification, so that arbitrary specification

selection does not drive our results. Using the identified relationships, we then aim to estimate trend/cycle decompositions for the variables in our model. This form of multivariate approach has been shown to reduce filter and estimation uncertainty (Basistha and Startz, 2008), while Bayesian techniques allow for the estimation of confidence intervals for the unobserved states (Planas *et al.*, 2008). We also attempt to account for some of the small open economy issues that affect Ireland, particularly its exposure to large, infrequent shocks and the distortions of the multinational sector. Finally, we compare our estimates with those of the established models from international organisations, to see if the trend/cycle decompositions derived from our model better fit the stylised facts of the Irish business cycle, particularly with regard to turning points and the timing of cycle phases.

Given the previously identified sensitivities in estimating the Irish output gap from any one model specification, we adopt a Bayesian model averaging approach that aggregates estimates across a number of specifications used in the current literature. Specifically, we allow for flexibility in three of the core specification choices: variable selection; trend/cycle decomposition; and error distributions.

The first and, arguably, most critical modelling assumption is based on the set of variables to be included in the model and the specification of inter-relationships within the system. Beginning with real output, wage inflation and unemployment, we allow for specifications that consist of univariate models of output, bivariate models of output and either wage inflation or unemployment, and trivariate models incorporating the complete set of variables.

Secondly, we allow for variation in the specification of potential output. Again, this is critical to our modelling outcomes and policy inferences, as the output gap is estimated as the difference between observed output and unobserved potential output (as a percentage of potential output). While there are numerous possible specification options within the literature, we confine our possibility-space to specifications where potential output follows either a stochastic or deterministic trend, and where the growth rate is either constant or time-varying.

Finally, we incorporate flexibility in the distribution of innovations to the set of variables within the model. While allowing for the possibility of stochastic volatility in the modelled innovations, we also include an empirically determined selection between Gaussian and Student's t -distributed innovations, which alters the frequency with which extreme events are experienced in the system. Models with stochastic volatility and Student's t -distributed errors have been shown to fit the data better (Chiu *et al.*, 2017) and forecast better (Cross and Poon, 2016) than conventional specifications.

For our sample of Irish data, results suggest that, while the unemployment gap is beneficial in refining estimates of the output gap, there is limited evidence to suggest that the output gap aids estimation of the wage inflation gap, once we have accounted for the information contained in the unemployment gap. Based on our model estimates, we find that trend output growth rates peaked in Q1 2000, with

the effects of the Global Financial Crisis (GFC) and Sovereign Debt Crisis (SDC) causing a peak-to-trough decline in potential output of 3.1 percentage points, although our model suggests that potential output has been increasing since 2012. Finally, our model provides strong evidence to suggest that innovations to both output and wage inflation follow a Student's t -distribution, suggesting the relatively frequent occurrence of extreme events. This result should not be surprising, given Ireland's position as a small open economy that remains exposed to global economic events.

The paper is structured as follows. Section II gives a brief overview of the output gap literature relevant to our work. Section III discusses the potential problems in estimating the output, wage inflation and unemployment gaps for the Irish economy, and our solutions to some of these issues. Section IV presents the general model space over which the specification search is conducted, outlining the processes for the model dimension, trend output specification, and innovation distribution searches. Section V gives an overview of our prior distribution choices and an outline of the Bayesian estimation technique used to identify our system. Section VI presents the results from our model, including our estimates of the output gap, potential output growth, trend wage inflation and the NAWRU for the Irish economy. Section VII concludes the paper.

II RELATED LITERATURE

While the field of research into output, inflation and unemployment gaps is vast, limited work has been conducted in this area with specific reference to Ireland. The European Commission, as part of its responsibility to assess stability/convergence programmes of EU Member States, estimates potential output and the non-accelerating wage rate of unemployment (NAWRU) through its production function model (with the IMF and OCED producing research with variants of these models). However, the EC production function model has been criticised across a number of fronts: use of common model parameters for all countries under analysis; imposition of symmetric business cycles; instability of potential/trend variables over short-term horizons; considerable revisions to estimates for smaller economies over time; historical trend revision following sustained growth or decline; and use of a two-tailed HP filter (and its associated "end point" issues) to estimate the NAWRU. Specifically, Bergin and FitzGerald (2014) detail a number of issues that make the EC production function model unsuitable for analysis of the Irish economy.

Beyond this framework, there is not much additional work that jointly estimates all three gaps for the Irish economy. However, some further work has been done at the single-equation level. As part of their multi-country analysis, Ball *et al.* (2013) estimate the relationship between output and unemployment (the Okun equation)

for a sample of 20 economies. Using annual data from 1980 to 2011, their results for Ireland suggest that there is a negative and significant relationship (-0.41) between the output and employment gaps. Similarly, Conefrey *et al.* (2014) specify an Okun equation for Ireland using a sample of data from 1960 to 2012, estimating the Okun coefficient to again be negative (-0.31) over the sample period.

Meyler (1999) uses a reduced-form Phillips curve, with a measure of domestically generated inflation, to identify the inflation-unemployment relationship in Ireland. Modelling the Non-Accelerating Inflation Rate of Unemployment (NAIRU) as a random walk with Gaussian errors, they set their Phillips curve and unemployment equations in state-space form, allowing the NAIRU to be extracted from the underlying models using a Kalman filter. Their results suggest a strong relationship between domestically generated inflation and labour market tightness, while the NAIRU is found to be relatively time-invariant, ranging between 10-13 per cent over the 1979-1999 sample period. Gerlach *et al.* (2016) examine the historical relationship between unemployment and inflation in Ireland, estimating a Phillips curve equation for Ireland using data from 1926 to 2012. Assuming the NAIRU follows a random walk, they specify a Phillips curve equation in first-difference form, where the unemployment gap (and hence the NAIRU) is treated as an unobserved component. Their estimates show a significant, negative coefficient (-0.17) for the effect of the unemployment gap on inflation.

Slevin (2001) estimates the output gap for Ireland using both statistical trends and a Cobb-Douglas production function. Modelling the technology component of the function as a linear time trend, they find a significant relationship between the output gap and inflation in Ireland. Using a linear time trend, a split time trend and a HP filter, their estimates of the output gap show strong, positive correlations with the rate of inflation, while dynamic correlation analysis suggests that the output gap leads inflation, consistent with Phillips curve theory.

With respect to our estimation technique, the stochastic model specification search methodology is relatively new, with Frühwirth-Schnatter and Wagner (2010) developing the process. Only two papers have since used this approach to decompose output, inflation and unemployment into their stochastic trend and cyclical components; Berger *et al.* (2016) and Chan and Grant (2017). Both papers employ a multi-equation approach to develop a reduced-form macroeconomic model for the US economy that accounts for structural macroeconomic changes.

Berger *et al.* (2016) focus on the degree of parameter time variation in a multivariate unobserved components model of the US economy from Q2 1959 to Q3 2014. Decomposing output, inflation and unemployment into stochastic trends, common cyclical factors and idiosyncratic components, they allow for time variation in the growth rate of potential output, the slope of the Phillips curve, the Okun's law coefficient and all variance terms within the model. Their results suggest a persistent decline in potential output growth over the sample, with the

Okun coefficient found to be weaker in expansions and stronger in recessions, while the slope of the Phillips curve is estimated by the model to be constant. They also observe that the inflation gap is invariant to changes in the output gap, driven instead by a persistent AR(1) term. Chan and Grant (2017) use a similar methodology to Berger *et al.* (2016), employing the trivariate system to assess a range of specification choices for identifying the output gap. Using data from Q1 1948 to Q3 2014, they find steady declines in trend output growth rates, from almost 4 per cent in the 1960s to under 3 per cent in the 1980-2000 period, with post-GFC rates declining to just over 1 per cent at sample end.

III TREND / GAP ESTIMATION AND THE IRISH ECONOMY

While the modelling of latent variables has developed in the last decade, becoming an integral component of mainstream economic research, there remain a number of challenges in their estimation. These issues are exacerbated when considering output, inflation and unemployment trend/cycle decompositions for the Irish economy. Included in these complications are: Ireland's status as a small open economy; the distortionary effects of contract manufacturing and intellectual property transfers by MNEs; the domestic credit and housing booms of the early 2000s; and the large scale expansion and reversal of external funding in the Irish banking system.

These confounding factors may affect all three endogenous variables in our system, potentially biasing the amplitude and magnitude of cycles, the nature of time-variation in the system, and coefficient estimates. While preferred solutions to these complications would be to either remove their effects from the data, or account for their influence in the dynamics of our model, this may not always be feasible and could potentially introduce further biases into the system. We discuss some previously identified issues in this section, detailing the problems introduced into the trend/cycle estimation process, and how we attempt to account for their effects on the Irish economy.

3.1 Economic and Financial Openness

International competition, market liberalisation and deregulation, increased labour supply and improvements in monetary and fiscal policies are just some of the suggested transmission channels through which increased globalisation and openness have caused inflation to decline across countries since the early 2000s. Hardouvelis (1988; 1992) and Romer (1993) were some of the first papers to suggest that there are channels through which openness can negatively influence the rate of inflation in a country. Lane (1997) showed that the negative relationship between inflation and trade openness also exists in economies that are not large enough to influence relative prices.

The foremost channel through which globalisation is perceived to affect output and inflation is trade integration. With the integration of high-cost economies with low-cost economies, increases in the flow of trade from low-cost economies causes a favourable terms-of-trade shift in the high-cost countries, increasing potential output and reducing the trend rate of inflation (*ceteris paribus*). Enhanced competition indirectly influences output growth and inflation rates, as greater competition encourages technological progress and introduces potential scale economies. Additionally, labour market integration redistributes labour from low-wage to high-wage countries, equalising wages across economic regions, thereby expanding the labour force, increasing output and reducing wage pressures in high-cost economies. These factors enhance growth prospects and alleviate inflationary pressures through reduced production costs in the non-traded sector of the economy.

Financial openness allows for a greater availability of domestic credit through capital market integration, which can affect inflation in a number of ways. Easier access to credit, reductions in the total cost of borrowing, and elimination of credit constraints may shift the aggregate demand curve, leading to changes in the price level. Additionally, increased FDI flows can increase competition for global goods and services, leading to reduced prices in the tradeables sector. Furthermore, wage inflation may be influenced by capital market integration; increased capital stocks may raise wages through enhanced labour productivity, but ease of capital movement and the risk of capital flight to low-wage countries may diminish wage-inflationary pressures.

Various approaches have been suggested in the literature to account for the impact of openness on potential output and output gap estimation. One common approach taken is to include variables that account for domestic economic and financial interactions with foreign markets. Darvas and Simon (2015) employ a system-of-equations model that incorporates a Phillips curve and a current account equation. Konuki (2010) uses a multivariate Kalman-filter approach that includes nominal unit labour cost growth rates and export market output gap estimates.

We avoid such an augmented-variable approach for a number of reasons, key being the use of such variables in a linear system. The argument presented in a number of papers for using a current account-neutral measure of potential output derives from the findings of Lane and Milesi-Ferretti (2012), whereby countries with pre-crisis current account balances in excess of what could be explained by standard economic fundamentals experienced the largest contractions in their external balance. While this holds for the crisis period, where global capital flows were in retrenchment, there are a number of factors that would suggest the relationship between current account imbalances and potential growth is not linear, and are dependent on a number of prevailing domestic and global conditions. The United Kingdom, for example, has run a persistent current account deficit since 1984, and while their cumulative current account deteriorated at a faster rate (in GDP percentage point terms) than Ireland's between 2006 and 2010, the effects of the global financial crisis (GFC) were not as strongly felt in the UK.

As discussed by Lane (2014), understanding the Irish external balance sheet at a disaggregated level is needed to understand the extent of the current account contraction. While Ireland experienced current account deficits over the 2006-2010 period, the scale of the decline in the Net International Investment Position (through stock-flow adjustments and negative net valuation gains) was unmatched by movements in the Balance of Payments data. These valuation losses can be attributed to the long-equity, short-debt international strategy of Irish national investors, which was excessively exposed to the GFC, with contemporaneous declines in global equity, property values and the availability of debt market funding. Without accounting for these factors, the effect of current account imbalances will be exacerbated during the crisis, biasing both the statistical and economic significance of current account imbalances on potential output, and the output gap, in our model. Similar issues apply to foreign output gaps and trade balances, both of which were disproportionately affected by the GFC, and would not be suitable for inclusion in a linear model, given the state dependence of their effects.

Similarly, there are two distinct phases to the Irish version of the capital flow bonanza that occurred before the crisis. In the first stage (1997-2002), accommodative economic conditions with solid economic fundamentals overlapped with rising rates of inward foreign direct investment (FDI) and declining shares of debt and equity liabilities. The second stage of the bonanza (2002-2007) saw inward FDI decline, replaced by debt and equity liability flows. Assuming that the causal effects of corresponding current account movements would be similar across these periods, despite the obvious underlying differences in the nature of the capital flow types, could severely bias our measures of potent output and the output gap.

Instead, we consider how to account for openness in the structure of our model. Ireland's high level of international integration, FDI and trade activity create transmission channels through which external shocks can cause transitory or permanent changes to domestic macroeconomic variables. Given Ireland's relatively small size, these external shocks can have large effects on both cyclical and trend components of the domestic variables. To account for these issues, we make two key modifications to the system. First, when specifying our priors, we allow for greater variability in the movement of the trend components of output, wages and unemployment, relative to a larger, less-open economy like the US. With this alteration, structural shocks to the system that affect output, inflation and employment (like large capital inflows and inward migration) are more likely to be captured as permanent changes in the growth rate of trend components (or as permanent level shocks to trends).

Furthermore, given the effects that international capital and labour movements can have on output and wage inflation, we allow for stochastic volatility in the errors in these equations, and for these errors to be modelled with Student's t -distributions. By allowing for specifications with Student's t -distributions in our

state space, we increase the likelihood of “fat-tailed” events that affect output and wage inflation, a characterisation of the vulnerability of small open economies to external shocks. Stochastic volatility allows for the variance of shocks to our system to change over time. Given the structural changes to the Irish economy over the last 30 years, the nature and transmission channels of exogenous shocks to the economy will have changed over time, as will the co-dependency of the variables in our system. By allowing for time-varying volatility in our series, we account for the changing effects of these shocks to the Irish economy.

The use of wage inflation (instead of consumer price inflation) in our model is predicated on the influence of exogenous global factors on Irish CPI. As discussed above, there are multiple channels through which openness and globalisation can influence consumer price inflation. With so many external factors that can potentially drive consumer price inflation, its association to other macroeconomic fundamentals (i.e. output and unemployment) may be diminished. As far back as 2001, the IMF Article IV consultation highlighted the issue of estimating the NAIRU in Ireland, using conventional CPI measures. Given the size and degree of openness of the Irish economy, the linkage between inflation and unemployment dynamics was weak, with the NAWRU perceived to be the more appropriate measure of the structural unemployment rate.

Finally, we allow for a New Keynesian Wage Phillips curve (NKPWC) specification when modelling our inflation dynamics. As shown by Orlandi *et al.* (2018), NKPWCs signal cyclical fluctuations in unemployment more clearly, and yield less pro-cyclical estimates of the NAWRU, than standard accelerationist Phillips curves. This suggests that the use of the accelerationist Phillips curve to identify the unemployment gap is inappropriate for the Irish economy, as it delivers poor signals in volatile times.

3.2 Multinational Enterprise Distortions

Since 2010, Ireland has recorded a strong flow of imports from cross-border intellectual property (IP) purchases, by multinational enterprises resident in Ireland from their foreign affiliates. In 2015, coupled with this ongoing increase in investment and imports, there was a sizable balance-sheet relocation, involving the large-scale transfer of intangible assets into Ireland, offset by a corresponding increase in external financial liabilities. The effect on the National Accounts was considerable. Primarily driven by IP assets, a capital increase of approximately €300 billion was recorded in the first quarter of 2015.

As a by product of the MNE relocation, Ireland became an economic principal for a considerable degree of production abroad, carried out through the use of contract manufacturing. The resulting sales of goods produced abroad under contract, when incorporated into the Irish National Accounts, increased the balance of trade in goods and services from €35 billion in 2014 to €70 billion in 2015.

Two further channels through which the impacts of the corporate relocation are observed in the National Accounts are royalty payments and depreciation. Prior to the IP capital being booked in Ireland, the effect of contract manufacturing activities on exports was essentially offset by imports of royalty services used in the production process. Once the IP was Irish-owned, there was no corresponding entry in the import component of the trade balance, leading to an increase in Irish GDP due to the value added from contract manufacturing activities.

Unsurprisingly, there are a number of distortionary effects that the above elements will have on the estimation of potential output (FitzGerald, 2018). Primarily, the level-shift in GDP/GNI resulting from the domestic booking of R&D and IP capital would cause a corresponding increase in potential output. As the empirical model cannot distinguish between intangible capital and conventional capital, naïve estimation will observe a relationship between the increase in the productive capacity of the economy (from a higher capital stock) and the corresponding increase in domestic output levels.

Furthermore, as the capital relocation represents a level shift that is unlikely to have a future effect on output growth, this will distort estimates of the composition of contributions to both current and future potential output. Any analysis that attempts to decompose potential output into constituent elements, e.g. the labour, capital and TFP elements in a standard production-function analysis, will likely be biased, unless there is specific accounting for the differing nature of tangible and intangible capital.

To account for the effect of MNE activity in our model, we modify the output data series that is used to estimate the system of equations. Taking GDP data as our baseline series, we splice the historical data with the CSO's Modified Final Domestic Demand (MDD) series, defined as total domestic demand net of trade in aircraft by leasing companies, trade in investment in intellectual property and changes in the physical value of stocks. This series presents a more reflective account of the domestic demand of Irish residents, and should mitigate the dissociative effects of MNEs on potential output and the output gap.

As the CSO's MDD series is only available back to 1995, we splice the series together using Q1 1995 as the linking point. To account for the small level shift between the series (approximately 3.5 per cent of GDP) the data are merged using the de la Fuente (2009) mixed splicing procedure for economic time series. Rather than using interpolation or backcasting, the procedure allows for the correction to be calculated by weighting both the level and growth rate of the older series on the nature of the underlying factors that account for the disparity between the series at their linking point. Given that MNE activity in Ireland was increasing from our sample start point (Q1 1981) to our linking point (Q1 1995), we allow for the disparity between our composite series and GDP to expand over time. Similarly, we allow the spliced series and GDP to differ at their starting point, as there were a large number of pharmaceutical and electronics MNEs active in Ireland at this

time. Thus, the difference between our spliced output series and GDP is reflective of an emerging sector that grows faster than the economic aggregate, which is reflective of the time path of MNE activity in Ireland. Formally, we can represent our spliced series, in log form, as

$$\left. \begin{aligned} \hat{y}_t^s &= x_t + \hat{d}_t^s \\ \hat{d}_t^s &= d_T \rho^{\frac{T-t}{T}} \end{aligned} \right\} \text{for } 0 \leq t \leq T$$

where \hat{y}_t^s is our spliced series, x_t is the older series (GDP), \hat{d}_t^s is the gap between the old series and the spliced series at time t , d_T is the divergence component between the series being spliced at linkage point T , and ρ is a parameter that controls for the speed of adjustment in the spliced series, between the old series at its starting point and the new series at the linkage point. We set $\rho = 0.5$, to allow for some MNE activity in the GDP series at the beginning of our sample, and for this activity to increase consistently from the initial period of the sample to the linkage point of our data.

3.3 Domestic Credit and House Price Booms

While the interrelated domestic credit and house price booms of the late 1990s and early 2000s characterised the Irish economy at the time, the flow of foreign debt to the balance sheets of domestic banks accounted for much of the unprecedented growth in these markets. Although favourable domestic conditions contributed to this inflow of credit, there is little doubt that excessive credit lending for the purpose of property development and speculative investment in foreign assets accounted for much of the credit and liquidity risk that resulted in the Irish banking crisis.

In the aftermath of the global financial crisis, a number of researchers (Alberola *et al.*, 2013; Borio *et al.*, 2017; Berger and Richter 2017) have attempted to augment the estimation of output gaps and trend growth with estimates of both macro-economic and financial imbalances. Using combinations of economic, financial and international cycles, this work aims to identify “sustainable” or “finance-neutral” measures of trend output and the associated output gap. To date, the most commonly used identifiers of the financial cycle include the availability of credit to the domestic economy and national house prices.

However, this strand of research has faced a number of criticisms. Similar to our discussion of capital flows above, a key concern lies in the potentially idiosyncratic nature of credit and property price booms. While the Irish credit boom appeared superficially similar to, and overlapped with, the US credit bubble, there were considerable differences in the nature of each credit boom. The US credit boom was built on the expansion of credit through financial product innovation; the Irish boom was fuelled through euro area interbank borrowing. US banks

distributed loans to the securitisation industry, removing the risk from their balance sheet and introducing moral hazard to credit markets; Irish banks held the majority of loans on their balance sheets, increasing their risk exposure but limiting moral hazard issues with lending standards. US bank balance sheets were relatively well diversified; Irish banks were heavily exposed to the domestic mortgage market and domestic property developers that were reliant on capital gains for investment returns.

Similarly, there are considerable differences in the Irish credit markets pre- and post-crisis: macro-prudential rules constrain mortgage credit borrowing, limiting bank exposure to the property market and mortgage default risk; bank capitalisation rates are higher; and regulatory standards and implementations have improved with European banking legislation.

Additionally, naïve statistical models that decompose trend rates of credit growth using filtering techniques fail to distinguish between the sustainability of credit growth in different periods. Connor *et al.* (2012) argue that the Irish credit growth of 1997-2002 was sustainable due to the nature of capital flows, economic fundamentals and credit conditions in the Irish economy at the time. However the pre-crisis period of growth in 2002-2007 led to unsustainable conditions in the Irish credit markets, to the point where the resulting crisis would likely have occurred in the absence of the GFC (although the magnitude of the effects would have differed). Similarly, post-crisis growth in the credit and housing markets does not exhibit the same risk characteristics; bank credit is more diversified with less maturity and liquidity mismatch, while domestic speculative investment in the property market is at a fraction of previous values. Assuming that above-trend credit growth must be unsustainable, and lacking in appropriate underlying fundamental conditions, could bias estimates of our output, wage inflation and unemployment gaps, amplifying estimates of overheating during periods of credit growth and understating the level of economic performance during periods of financial market retrenchment. Furthermore, declining levels of credit due to a shift towards equity financing, and declining levels of credit due to a stop in the availability of credit should have very different effects on an economy.

Thus, we avoid using a multivariate approach that incorporates financial or international imbalances. As these measures would generate a sustainable future output trend, the approach would introduce uncertainty regarding the effect of current imbalances, the timing and effects of any future adjustments, and the likelihood of adjustments occurring. Instead, we maintain an approach that retains conventional definitions of potential output, trend wage inflation and the NAWRU as outcome variables.

IV STOCHASTIC MODEL SPECIFICATION SEARCH

This section of the paper details our estimation approach and search strategy, which consists of three elements. In Section 4.1, we develop our system of equations that includes the set of variables under analysis and the structures through which they interact. Section 4.2 details the data-driven approach taken in determining trend components of the system. Finally, Section 4.3 considers the distributional characterisation of innovations within the system.

4.1 Unobserved Components Model Identification

In selecting the equations that govern our system, we adopt an approach consistent with classical macroeconomic research, and model our system as a trivariate set of structural equations that include wage inflation, output and unemployment. This approach has a number of benefits, including the simultaneous estimation of variable interactions, the reduced computational requirements of a small-system approach, and the availability of data over a sufficient time horizon. Using the log of output (y_t), the rate of wage inflation (π_t) and the unemployment rate (u_t), the system can be represented as

$$\begin{aligned}
 y_t - y_t^* &= \beta_{\theta u} \theta^u (u_t - u_t^*) + \omega_t \\
 \omega_t &= \tau_1^\omega \omega_{t-1} + \tau_2^\omega \omega_{t-2} + \varepsilon_t^\omega \\
 \pi_t - \pi_t^* &= \tau^\pi (\pi_{t-1} - \pi_{t-1}^*) + \beta_{\gamma u} \gamma^u (u_t - u_t^*) + \beta_{\gamma y} \gamma^y (y_t - y_t^*) + \varepsilon_t^\pi \\
 u_t - u_t^* &= \tau_1^u (u_{t-1} - u_{t-1}^*) + \tau_2^u (u_{t-2} - u_{t-2}^*) + \varepsilon_t^u
 \end{aligned} \tag{1}$$

where y_t^* is potential output, π_t^* is the trend wage inflation rate, and u_t^* is the non-accelerating wage rate of unemployment (NAWRU).

The first two equations presented in the above system drive deviations in potential output and output growth. In estimating our measure of the output gap, we employ an augmented, time-varying form of Okun's law, whereby the output gap is driven by two factors: the standard component that estimates the unemployment gap, plus an additional term that represents the cyclical component of output. This cyclical component allows for the decomposition of output into stationary and non-stationary elements, with the non-stationary component representing our estimate of potential output. To account for the possibility of permanent shocks to the growth rate of output, as well as a time-varying trend growth rate, we model the stationary component as an AR(2) process with a stochastic error term.

Following Galí (2011), Stella and Stock (2012), Kim et al. (2014) and Orlandi et al. (2018), we allow for wage inflation dynamics to be represented as a New-Keynesian Phillips curve, incorporating a time-varying trend rate of wage inflation. This approach allows for stationarity in the rate of wage inflation (which should

hold if the output gap is a stationary process and current wage inflation represents the discounted future value of future output gaps), improving the fit of the NKPC to wage inflation data without requiring an excessive lag structure.

Finally, we account for unemployment dynamics through an unemployment gap equation. Specifically, we let the unemployment gap follow a stationary AR(2) process. This structure is flexible enough to provide dynamics for the NAWRU that do not simply mirror movements in the unemployment rate (an identified issue in the empirical output gap literature), while also removing inflationary pressures from the labour market when the unemployment rate is at its equilibrium level. As our system possesses a triangular structure, we assume all error terms (ε_t^ω , ε_t^π and ε_t^u) to be independent at all leads and lags.

To allow for the Bayesian model average component of our strategy, we include three β_i terms in the set of equations. These coefficients are binary terms, taking values of either 0 or 1 depending on the model specification under consideration. When all three β_i terms are set equal to zero, the system reduces to the unobserved components model of Watson (1986). When $\beta_{\theta^u} = 0$, the Okun coefficient drops out of the output gap equation, and movements are entirely driven by the unobserved cyclical component of output.

4.2 Trend Specification Search

The next decision we face in estimating the output, wage inflation and unemployment gaps relates to how we choose to functionally represent the trend component of each unobserved series. The choice of trend determination process is particularly important for output growth, as several papers have shown. Using US data, the unobserved components model of Morley et al. (2003) estimates an output gap that is low in amplitude and small in magnitude, contrasting Hodrick-Prescott (HP) filter estimates, which generate large and persistent cycles in the output gap. This difference arises due to the way that trend output is modelled; the unobserved components model represents the trend component of output as a random walk with constant drift, while the HP filter assumes all trends to be random walks with stochastic drift.

Given the degree to which the trend-estimation process can drive our results, we aim to approach the problem from as agnostic a position as possible, allowing the data to speak to the time path of the trend. To this end, we adopt the approach of Frühwirth-Schnatter and Wagner (2010), a method that employs Bayesian stochastic model specification search for state-space models using non-centered parameterisations. This technique is designed for dealing with non-regular testing problems in deciding whether a component is fixed or time-varying. By constructing a general framework that nests the complete set of alternative models under analysis, binary stochastic variables can be incorporated into the trend components of output, wage inflation and unemployment. These variables can then

be sampled together with the model's parameters, for the entire model-space under consideration. Let

$$y_0^* = \beta_{\alpha_0} \alpha_0 t + \beta_{y^*} \sigma_{y^*} \tilde{y}_t^* + \beta_{\alpha} \sigma_{\alpha} \sum_{i=1}^t \tilde{\alpha}_i$$

where β_{y^*} , β_{α_0} and β_{α} are the binary stochastic indicators, and

$$\begin{aligned} \tilde{y}_t^* &= \tilde{y}_{t-1}^* + \varepsilon_t^{\tilde{y}^*} & \varepsilon_t^{\tilde{y}^*} &\sim N(0, 1) \\ \tilde{\alpha}_t &= \tilde{\alpha}_{t-1} + \varepsilon_t^{\tilde{\alpha}} & \varepsilon_t^{\tilde{\alpha}} &\sim N(0, 1) \end{aligned}$$

with initialisation of the state equations such that $\tilde{y}_0^* = \tilde{\alpha}_0 = 0$. Letting $\alpha_t = \beta_{\alpha} \sigma_{\alpha} \tilde{\alpha}_t$, $\gamma_t^{\alpha} = \sigma_{\alpha} \varepsilon_t^{\tilde{\alpha}}$ and $\gamma_t^{y^*} = \sigma_{y^*} \varepsilon_t^{\tilde{y}^*}$, the previous equations reduce to

$$\begin{aligned} y_t^* - y_{t-1}^* &= \beta_{\alpha_0} \alpha_0 + \beta_{y^*} \sigma_{y^*} \varepsilon_t^{\tilde{y}^*} + \beta_{\alpha} \sigma_{\alpha} \tilde{\alpha}_t = \beta_{\alpha_0} \alpha_0 + \alpha_t + \beta_{y^*} \gamma_t^{y^*} & \gamma_t^{y^*} &\sim N(0, \sigma_{y^*}^2) \\ \alpha_t - \alpha_{t-1} &= \beta_{\alpha} \sigma_{\alpha} \varepsilon_t^{\tilde{\alpha}} = \beta_{\alpha} \gamma_t^{\alpha} & \gamma_t^{\alpha} &\sim N(0, \sigma_{\alpha}^2) \end{aligned}$$

With β_{α_0} , β_{α} and β_{y^*} being binary, there are a total of eight trend specifications nested in the modelled framework. If $\beta_{\alpha} = 0$ and $\beta_{y^*} = 1$, potential output follows a random walk, with drift coefficient $\beta_{\alpha_0} \alpha_0$. If $\beta_{\alpha_0} = 1$, $\beta_{\alpha} = \beta_{y^*} = 0$, potential output follows a deterministic trend, with growth rate α_0 . If $\beta_{\alpha_0} = \beta_{\alpha} = \beta_{y^*} = 1$, potential output follows a stochastic trend, with a time varying rate of growth ($\alpha_0 + \alpha_t$). If $\beta_{\alpha} = 1$ and $\beta_{y^*} = 0$, then potential output follows the same trend specification as assumed by the use of the HP filter. Thus, we nest a number of commonly observed trend specifications within the above system, over which we average our estimates of the output gap.

While our trend determination system for output is multi-faceted, we assume a single specification for both the NAWRU and trend wage inflation. On the assumption that long-run growth for both wage inflation and unemployment is expected to be zero, we model each process as an independent driftless random walk

$$\begin{aligned} \pi_t^* &= \pi_{t-1}^* + \varepsilon_t^{\pi^*} & \varepsilon_t^{\pi^*} &\sim N(0, \sigma_{\pi^*}^2) \\ u_t^* &= u_{t-1}^* + \varepsilon_t^{u^*} & \varepsilon_t^{u^*} &\sim N(0, \sigma_{u^*}^2) \end{aligned} \quad (2)$$

where the equations are initialized as $\pi_1^* \sim (\pi_0, V_{\pi^*})$ and $u_1^* \sim (u_0, V_{u^*})$

4.3 Distribution Search

The final decision we are required to make regarding our model specification relates to the distribution of the innovations. There is an increasing body of literature that identifies considerable benefits to modelling macroeconomic time series, particularly output and wage inflation, as having time-varying volatility.

The reduction in the volatility of business cycles in the 1980s, the Great Moderation, has shown that real gross domestic product growth, wage inflation, interest rates and a number of other macroeconomic variables exhibit a high degree of time variation in their volatilities. Consequently, it is important to incorporate this feature into macroeconomic time-series models, allowing the variance of innovations to a series or system to fluctuate over time. This is particularly true of a small open economy like Ireland; heavily integrated financial and goods markets, coupled with highly accessible capital and labour markets, expose Ireland to an array of heterogeneous shocks from external sources.

While most fields have been quick to adopt time varying volatility into their models, the majority of the literature assume that innovations to their systems follow Gaussian distributions. This implies that extreme events occur with low probabilities, and that volatility distributions across all series in the system are uniform. These assumptions contrast with what has been observed in the data, with greater reductions in the volatility of wage inflation than in output growth or unemployment rate volatility.

To overcome these issues, a number of researchers have begun to incorporate alternative distributions into macroeconomic models. Cúrdia et al. (2014) show the importance of rare large shocks in driving US business cycles, using a dynamic stochastic general equilibrium (DSGE) model with innovations that possess stochastic volatility and Student's t -distributions. Chiu et al. (2016) use a Bayesian VAR model with non-Gaussian errors that generates large out-of-sample forecast gains relative to standard forecasting models, especially during tranquil periods.

In line with this thinking, we allow stochastic volatility in the innovations in both the wage inflation (ε_t^π) and cyclical output (ε_t^ω) equations of our system, while also including a distribution search to select between Gaussian and Student's t -distributions for each innovation. To incorporate these features into the system, we add two more binary variables to the model; β_{z^ω} and β_{z^π} . For each variable, if the dummy coefficient equals zero, the innovation term follows a Gaussian distribution. If the dummy term equals one for either variable, then the innovation term follows a Student's t -distribution. We can structure this system, with latent variable representations of the Student's t -distribution terms, as

$$\left. \begin{array}{l} \varepsilon_t^i \sim N(0, e^{h_t^i}) \\ \varepsilon_t^i | z_t^i \sim N(0, z_t^i e^{h_t^i}) \\ z_t^i \sim IG\left(\frac{v_i}{2}, \frac{v_i}{2}\right) \end{array} \right\} \begin{array}{l} \text{iff } \beta_{z^i} = 0 \\ \text{iff } \beta_{z^i} = 1 \end{array} \quad \forall i \in \{\omega, \pi\} \quad (3)$$

Finally, log-volatilities are assumed to follow random walk processes

$$\begin{array}{ll} h_t^\omega = h_{t-1}^\omega + \varepsilon_t^{h^\omega} & \varepsilon_t^{h^\omega} \sim N(0, \sigma_{h^\omega}^2) \\ h_t^\pi = h_{t-1}^\pi + \varepsilon_t^{h^\pi} & \varepsilon_t^{h^\pi} \sim N(0, \sigma_{h^\pi}^2) \end{array} \quad (4)$$

V ESTIMATION STRATEGY

In this section, we provide intuition regarding the prior distributions used in our analysis, and present a description of the Markov Chain Monte Carlo (MCMC) algorithm used to estimate the model.

In the above discussion of our modelling procedure, we have incorporated eight binary indicators into the system, to control for the specification of the individual models nested in the model-space. Collecting this set of indicators into a vector

$$M = (\beta_{\theta^u} \ \beta_{\gamma^u} \ \beta_{\gamma^y} \ \beta_{\alpha_0} \ \beta_{y^*} \ \beta_{\alpha} \ \beta_{z^\omega} \ \beta_{z^\pi})$$

each model is identified by a value for M . With eight indicators, there are $2^8 = 256$ potential models to consider. However, before we can implement the algorithm to sample the model space, we must first specify our priors on: the model parameters; the binary indicators; and the error variances in the state and measurement equations of the system.

5.1 Priors

5.1.1 Priors on Model Parameters

Regarding the distributional assumptions on the parameters of the system, we adopt a diffuse prior approach, modelling all structural coefficients as having independent normal distributions

$$\begin{aligned} \theta^u &\sim N(\mu_{\theta^u}, V_{\theta^u}) & \gamma^u &\sim N(\mu_{\gamma^u}, V_{\gamma^u}) & \gamma^y &\sim N(\mu_{\gamma^y}, V_{\gamma^y}) \\ \tau^\pi &\sim N(\mu_{\tau^\pi}, V_{\tau^\pi}) & \alpha_0 &\sim N(\mu_{\alpha_0}, V_{\alpha_0}) \\ \tau^\omega &\sim N(\mu_{\tau^\omega}, V_{\tau^\omega}) | (\tau^\omega \in \mathbf{R}) & \tau^u &\sim N(\mu_{\tau^u}, V_{\tau^u}) | (\tau^u \in \mathbf{R}) \end{aligned}$$

where \mathbf{R} is the stationary region of the parameter space.

5.1.2 Priors on Binary Indicators

With respect to the binary indicators, we assume independent Bernoulli priors, with success probability, such that $p_i, \forall i \in \{\theta^u \ \gamma^u \ \gamma^y \ \alpha_0 \ y^* \ \alpha \ z^\omega \ z^\pi\}$, such that

$$p_i(\beta_i = 1 | \Theta) = \frac{f(\beta_i = 1 | \Theta)}{f(\beta_i = 0 | \Theta) + f(\beta_i = 1 | \Theta)}$$

where $\Theta = [y^*, \pi^*, u^*, \alpha, \alpha_0, \omega]$.

5.1.3 Priors on Error Variances

Following Frühwirth-Schnatter and Wagner (2010), we consider normal priors with zero mean for the standard deviations σ_{y^*} and σ_α

$$\begin{aligned}\sigma_{y^*} &\sim N(0, V_{\sigma_{y^*}}) \\ \sigma_{\alpha} &\sim N(0, V_{\sigma_{\alpha}})\end{aligned}$$

It can be shown that, if the standard deviations follow the above distributions, then the variance terms will follow Gamma distributions, such that

$$\begin{aligned}\sigma_{y^*}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{y^*}}}\right) \\ \sigma_{\alpha}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{\alpha}}}\right)\end{aligned}$$

Under the Gamma prior, the distribution is more concentrated around small values of $\sigma_{y^*}^2$, than under the standard inverse Gamma prior. As the inverse Gamma distribution does not have probability mass around zero, using it as a prior distribution tends to force the posterior density away from zero. This is of particular importance when estimating the variance of the innovations to the time-varying trend growth rates and to the stochastic volatilities in the output and wage inflation equations.

We also assume Gamma priors for the error variances in the state equations

$$\begin{aligned}\sigma_{y^*}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{u^*}}}\right) & \sigma_{\pi^*}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{\pi^*}}}\right) \\ \sigma_{h\omega}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{h\omega}}}\right) & \sigma_{h\pi}^2 &\sim G\left(0.5, \frac{0.5}{V_{\sigma_{h\pi}}}\right)\end{aligned}$$

Finally, we assume that the error variance in the measurement equation for u_t follows a standard inverse-gamma distribution

$$\sigma_u^2 \sim IG(\mu_{u,0}, S_{u,0})$$

5.2 Bayesian Estimation

In the standard linear Gaussian state-space model, the unobserved states can be easily decomposed from the data using the Kalman filter, after which maximum likelihood estimation of the unknown parameters can be implemented via the likelihood function. However, given the inclusion of stochastic volatilities in our state-space model, the stochastic model specification search procedure outlined in the previous section, and the fact that not all innovations to the system are modelled as having Gaussian distributions, the estimation problem is highly non-linear, making use of the Kalman filter and maximum likelihood approach infeasible.

Instead, we employ a Markov Chain Monte Carlo (MCMC) mixed sampler, which consists of 11 ergodic sampling steps from the precision sampler of Chan and Jeliaskov (2009), the auxiliary mixture sampler of Frühwirth-Schnatter and Wagner (2010), and the auxiliary mixture sampler of Kim *et al.* (1998). While a

full derivation of each step of the sampler is presented in the Appendix, we provide an overview of each step in Table 1.

Table 1: MCMC Mixed Sampler Process

1	Sample $\boldsymbol{\pi}^*$ from its full conditional distribution.
2	Sample \boldsymbol{u}^* from its full conditional distribution.
3	Sample $\tilde{\boldsymbol{y}}^*$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distribution, randomly permuting the signs of $\tilde{\boldsymbol{y}}^*$ and σ_{y^*}
4	Sample $\tilde{\boldsymbol{\alpha}}$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distribution, randomly permuting the signs of $\tilde{\boldsymbol{\alpha}}$ and σ_α
5	Sample $(\beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u})$ and $(\gamma^y, \gamma^u, \theta^u, \tau^\pi)$ jointly, by first sampling $(\beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u})$ marginally of $(\gamma^y, \gamma^u, \theta^u, \tau^\pi)$ then sampling $(\gamma^y, \gamma^u, \theta^u, \tau^\pi)$ from their full conditional distribution.
6	Sample $(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha)$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ jointly, by first sampling $(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha)$ marginally of $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ then sampling $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distribution.
7	Sample $(\beta_{z^\omega}, \beta_{z^\pi}, \boldsymbol{z}^\omega)$ and \boldsymbol{z}^π jointly, by first sampling $(\beta_{z^\omega}, \beta_{z^\pi})$ marginally of \boldsymbol{z}^ω and \boldsymbol{z}^π , then drawing \boldsymbol{z}^ω and \boldsymbol{z}^π from their full conditional distributions.
8	Sample $\boldsymbol{\tau}^u$ and $\boldsymbol{\tau}^\omega$ from their full conditional distributions.
9	Sample \boldsymbol{h}^π and \boldsymbol{h}^ω from their full conditional distributions.
10	Sample the variances $\sigma_{u^*}^2, \sigma_{\pi^*}^2, \sigma_u^2, \sigma_{h^\pi}^2$ and $\sigma_{h^\omega}^2$ using an independence-chain Metropolis-Hastings step.
11	Sample ν_ω and ν_π using an independence-chain Metropolis-Hastings step.

Source: Author's analysis.

VI DATA AND EMPIRICAL RESULTS

In this section, we first discuss the data that are employed in our model, and the modifications to overcome some of the known issues with Irish output data. Next, we present the estimates of our model, including a set of the most likely model specifications and parameter estimates for the weighted average set of models. Finally, we show results from the decomposition of our three variables, with a comparison of the trend and cyclical components of output across IMF and DG-ECFIN model alternatives.

6.1 Data

The three data series that we use to estimate our model are a measures of real output, wage inflation and the total population unemployment rate. Real output is a composite series, incorporating real GDP from the IMF International Financial

Statistics (IFS) series, and real modified final domestic demand (MDD) from the CSO. We combine the series together in Q1 1995 (the starting point for the MDD series), accounting for the slight level shift through the use of the mixed splicing procedure for economic time series developed by de la Fuente (2009). As discussed above, this method allows for the correction from the MDD series to the older GDP series to be linked between levels and growth rates, and is designed to account for distortions in the base series being reflective of emerging sectors or activities that grow faster than the aggregate. Given that we are accounting for distortions due to the MNE sector, this approach is well suited to our purposes. The composite series is then transformed by taking logs. The wage inflation series is taken from the Earnings Hours and Employment Costs Survey (EHECS) database of the CSO and transformed into an annualised growth rate. Unemployment data come from the OECD's Main Economic Indicators (MEI) database. Our data are quarterly, with the sample running from Q1 1981 to Q4 2016. Posterior analysis of the data is based on 100,000 posterior draws, following a burn-in period of 10,000 draws.

6.2 Posterior Model Probabilities

Table 2 presents posterior means for the eight binary indicator variables in our system. Given that they are binary, the posterior mean values can be interpreted as the inclusion probabilities of the associated variables.

Table 2: Posterior Inclusion Probability Estimates of the Indicator Variables

β_{θ^u}	β_{γ^u}	β_{γ^v}	β_{α_0}	β_{y^*}	β_{α}	β_{z^π}	β_{z^ω}
1.00	0.699	0.208	0.869	0.611	0.947	0.671	0.969

Source: Author's analysis.

The strongest result coming from the posterior estimates is that the model selection process strongly accepts the use of the unemployment gap in estimating the output gap. Across the entire model space, the data clearly support the inclusion unemployment gap term in the output equation, as the posterior mean of β_{θ^u} is 1. In contrast, there is less evidence to support the use of the output gap in driving wage inflation. Once we have accounted for past wage inflation and unemployment, the output gap only informs estimation of the wage inflation equation in 21 per cent of the specified models. The information content regarding wage inflation present in the unemployment gap is somewhat stronger, supporting the existence of a wage Phillips curve in Ireland, with the term being included in almost 70 per cent of the estimated models.

With respect to trend output specification, the results strongly suggest that the output series possesses a stochastic trend with a time-varying rate of growth. The drift indicator (β_{α_0}) is included in 87 per cent of the sampled models, while the

time varying growth rate indicator (β_α) is included in 95 per cent of all estimated models. There is weaker support for the role of permanent level shocks to potential output, with the posterior mean of β_{y^*} estimated to be 0.61.

In determining whether Gaussian or Student’s t innovations better characterise the output and wage inflation series, the probability estimates suggest that innovations to both series more closely follow Student’s t -distributions. While the probability estimates provide some evidence in favour of wage inflation innovations having a Student’s t -distribution (with $\beta_{z\pi} = 1$ in 67 per cent of all models), there is far greater support for innovations to the output equation being characterised by a Student’s t -distribution ($\beta_{z\omega} = 1$ in 97 per cent of all models).

Thus, even accounting for stochastic volatility in the output and wage inflation equations, there is still strong evidence to suggest that the likelihood of rare innovations to wage inflation and output is larger than would be assumed under a Gaussian distribution. This implies that “tail event” shocks to wage inflation and output play an important role in Irish business cycles. Without allowing for Student’s t -distributions, large realisations of shocks to output and wage inflation may be incorrectly identified as permanent changes in the level of macroeconomic volatility in models that only account for time-varying variances, or may be incorrectly captured in the effect of explanatory variables in models with constant variance terms. This should be of particular importance in models that attempt to identify the permanent and transitory effects of recessions and expansions in aggregate time-series data.

A final noteworthy feature of the posterior model probabilities relates to the likelihood of observing any one specific model. Table 3 presents the M vector and posterior probabilities for the ten most likely models in the model space.

Table 3: Posterior Model Probabilities

$\beta_{\theta u}$	$\beta_{\gamma u}$	$\beta_{\gamma y}$	β_{α_0}	β_{y^*}	β_α	$\beta_{z\pi}$	$\beta_{z\omega}$	Prob.
1	1	0	1	1	1	1	1	0.169
1	1	0	1	1	1	0	1	0.132
1	1	0	1	0	1	1	1	0.116
1	1	0	1	0	1	0	1	0.088
1	0	0	1	1	1	1	1	0.045
1	0	1	1	1	1	1	1	0.038
1	0	1	1	1	1	0	1	0.034
1	0	0	1	0	1	1	1	0.029
1	0	0	1	1	1	0	1	0.025
1	0	1	1	0	1	1	1	0.025

Source: Author’s analysis.

As can be seen from the table, the ten most likely models account for the majority (67.5 per cent) of the posterior model probabilities, but no one model

dominates the system. This shows both the considerable degree of uncertainty with respect to individual model selection, and the benefit that accrues from averaging across the set of models. This result supports our modelling approach, and remains consistent with the suite-of-models approach that emphasises the reduction in estimate uncertainty from using multiple techniques to decompose macroeconomic variables into trends and cycles.

The most likely model supported by the system is the one in which: the unemployment gap is present in both the output and wage inflation equations; the output gap is absent from the wage inflation equation; output has a time-varying growth rate and a permanent level shock component; and innovations to both the wage inflation and output equations follow a Student's t -distribution. This model is closely followed by a variant in which shocks to inflation follow a Gaussian distribution, rather than a Student's t -distribution. Together, these two models account for less than 30 per cent of the posterior probability, suggesting that model uncertainty is not insubstantial. Hence, the results from our estimates of interest in the remainder of the paper are obtained using a weighted average across all model specifications, rather than only taking estimates from the most likely model.

6.3 Estimation of Model Parameters

In this sub-section, we discuss the parameter estimates from the stochastic model specification search exercise, calculated by averaging across all models in our model space. We focus on relationships among the wage inflation gap, the unemployment gap and the output gap, how persistent they are, and how often large shocks affect the wage inflation and output equations.

Table 4 reports the posterior means and 2.5 and 97.5 percentiles for the main parameters of interest from the model. The posterior mean of γ^u is estimated to be -0.20 , consistent with a downward-sloping wage Phillips curve. This suggests that a 1 percentage point unemployment gap is associated with an wage inflation rate that is 0.2 percentage points below trend. Consistent with the posterior estimate of β_{γ^y} , the coefficient on the output gap in the wage inflation equation is not estimated to be significant: a 1 percentage point output gap is associated with a rate of wage inflation that is only 0.02 percentage points above trend. However, the estimate of θ^u is both large and significant: a 1 percentage point increase in the unemployment gap leads to a decrease in the output gap of 2.08 percentage points (equivalent to an Okun coefficient of 0.48). This result is in line with the model of Ball et al. (2013), which estimate an Okun coefficient of 0.46 for Ireland.

Turning to the auto-regressive coefficients in the model, the AR(1) coefficient for the wage inflation gap (τ^π) is positive, with the estimate of 0.4 suggesting that the wage inflation gap is weakly persistent. Similarly, the unemployment gap is found to be strongly persistent; both AR coefficients are significant, with the sum of the coefficients estimated to be 0.98. Neither of the AR coefficients on the cyclical output growth term is found to be significant.

Finally, the degrees of freedom parameters associated with the Student's t -distributions of the innovations to the wage inflation and output equations are estimated to be 25.7 and 8.2, respectively. This suggests that rare shocks to wage inflation and output are larger than those allowed under more restrictive Gaussian assumptions. Relatively, the distribution of wage inflation shocks more closely approximates a Gaussian distribution than the distribution of output shocks (as supported by the posterior probability estimates of $\beta_{z\pi}$ and $\beta_{z\omega}$ in Table 2), but a Student's t -distribution with $df = 26$ still possesses important modal and leptokurtic differences from a Gaussian distribution.

Table 4: Posterior Parameter Means and Percentiles, Model Average Estimates

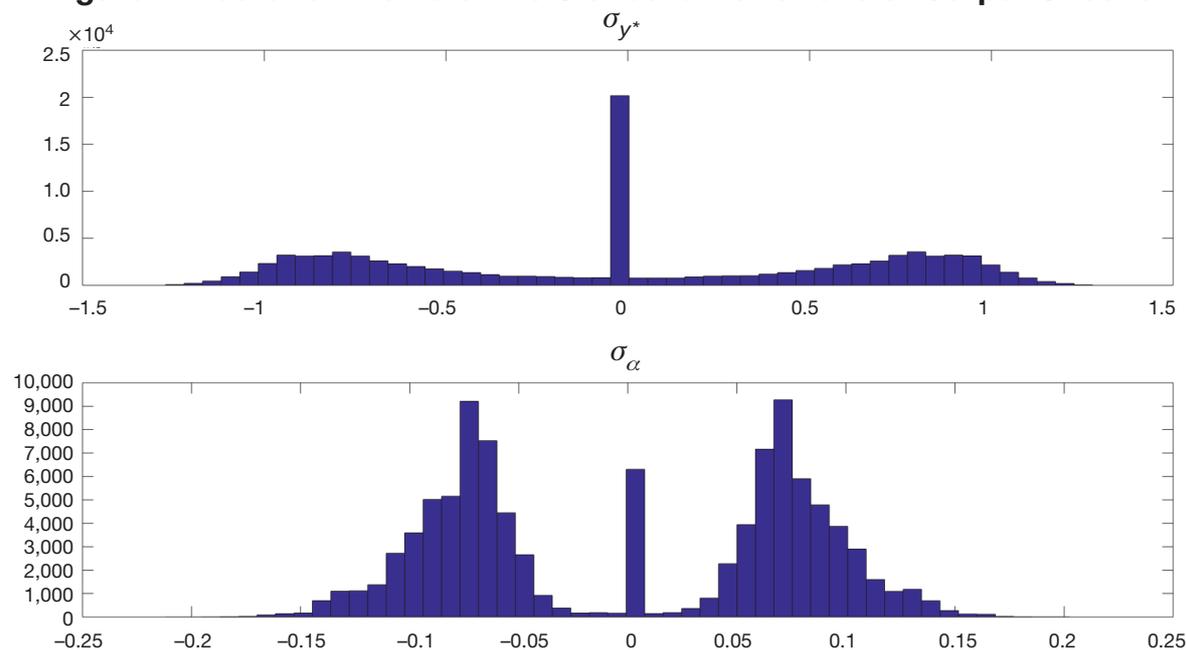
<i>Parameter</i>	<i>Posterior Mean</i>	<i>2.5 Percentile</i>	<i>97.5 Percentile</i>
γ^u	-0.20	-0.55	0.00
θ^u	-2.08	-2.73	-1.41
γ^y	0.02	-0.02	0.15
τ^π	0.40	0.12	0.67
τ_1^ω	0.04	-0.43	0.47
τ_2^ω	0.18	-0.14	0.47
τ_1^u	1.88	1.74	1.96
τ_2^u	-0.89	-0.98	-0.76
ν_π	25.74	4.06	48.68
ν_ω	8.16	2.17	40.65

Source: Author's analysis.

Reinforcing the importance of our distributional assumptions, Figure 1 presents histograms of the posterior draws of the standard deviations of the shocks to potential output σ_{y^*} and potential output growth σ_α . As can be seen from the first chart, the distribution of σ_{y^*} is highly leptokurtic, displaying a point mass at 0, with evidence of two additional modes in either tail of the distribution. The distribution of σ_α is trimodal, with mass in both the left and right tail, as well as a point mass around 0. These results provide further evidence to support the inclusion of both permanent level shocks to, and a time-varying growth rate of, potential output, supporting the posterior model probabilities attached to β_{y^*} and β_α .

6.4 Evaluation of the Output Gap and Trend Output Growth

In this sub-section, we report estimates of the output gap and potential output from our model-averaged specification. We also report estimates from two other sources: the IMF World Economic Outlook database, and the DG-ECFIN AMECO database. To assess the benefits of our Bayesian model averaging (BMA) approach, we perform a historical analysis of the trend and cycle decomposition of output from

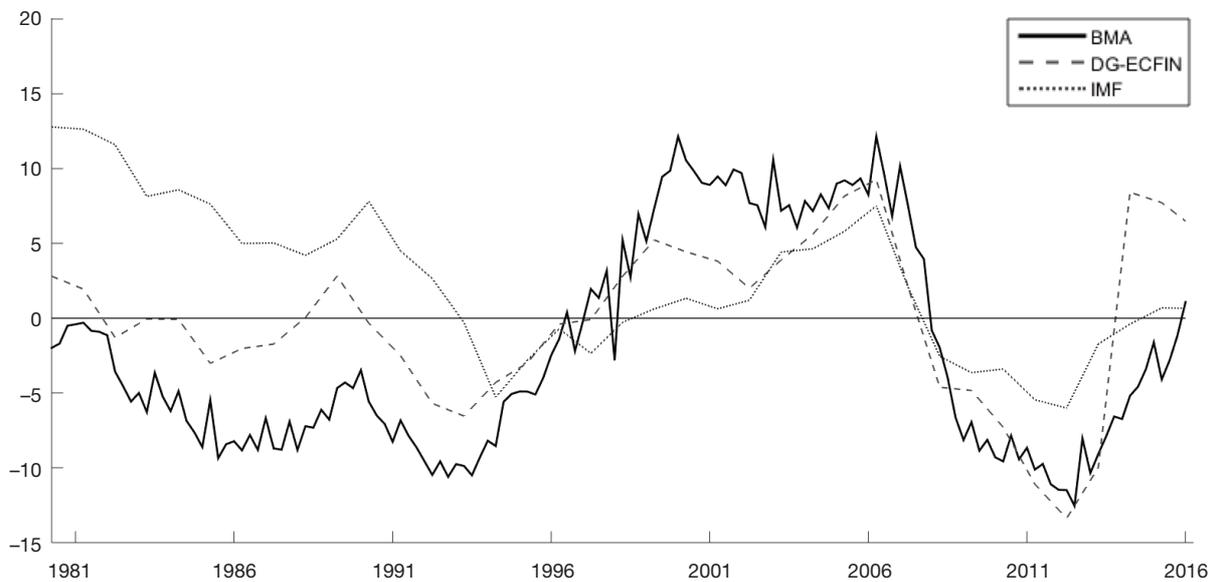
Figure 1: Posterior Draws of the Standard Deviations of Output Shocks

Source: Author's analysis.

our BMA model versus the IMF and DG-ECFIN models. While an in-sample forecasting analysis would be the preferred way to compare the models, we do not have the necessary real-time estimates of output (either MDD or GDP), access to the IMF and DG-ECFIN models, or a relevant series exogenous to all three models to perform an alternative variables forecast.

Figure 2 presents all three estimates of the output gap series. While the three output gaps generally follow a similar time path, with local maxima around 1990, 2000 and 2007, and local minima around 1994 and 2012, the magnitudes of the gaps differ markedly across the three estimates. From the beginning of the sample, there appears to be considerable disparity between the IMF model, and the DG-ECFIN and BMA models. Economic underperformance characterised most of the 1980s and early 1990s in Ireland, with high unemployment, low levels of saving and investment, and currency over-valuation hindering the export sector. However, the IMF model suggests that the Irish output gap was positive for this entire period, with a negative output gap only emerging in 1994, at the beginning of the Celtic Tiger period. In contrast, both the DG-ECFIN and our BMA estimates show more persistent negative output gaps across the early part of the sample, with large negative gaps estimated under both models before the expansionary Celtic Tiger phase.

For the first half of the Celtic Tiger period, each model suggests that there was an upward trend in the output gap, with all models showing positive gaps by 2000. However, the DG-ECFIN model peaks before the others, despite the fact that wage growth was still rising and unemployment was falling at its fastest rate over the

Figure 2: Estimates of the Output Gap for Ireland

Source: Author's analysis.

sample period. Similarly, while both our model and the DG-ECFIN model estimates peak in 2001, the magnitudes of the output gap estimates differ substantially. While our model suggests that the effects of the Celtic Tiger, the strongest period of rapid economic expansion in the history of the Irish State, pushed the output gap to previously unobserved levels, the IMF model estimates the output gap to be 1.3 per cent, smaller than any value observed during the period of economic stagnation in the 1980s.

Following these peak points, the time path of output gap estimates is broadly similar among all three models. Each shows a decline in the output gap of 1-2 years, before following a strong upward path, culminating in a local maximum in 2007. From 2007 to 2009, all models show a sharp, precipitous decline in the output gap, with estimates turning negative within this period. Again, both the DG-ECFIN model and our BMA model show comparable time-paths and gap estimates, in contrast to the IMF model which relatively underestimates the severity of the output gap decline. This decline is further exacerbated by the Sovereign Debt Crisis, with all models estimating a local minimum in 2012. Post-crisis, all models show an upward movement in the estimated series, with positive output gaps prevailing at the end of the sample period. However, as the DG-ECFIN model uses GDP as its output series, the distortionary effects of MNE activity are observed in the output gap estimates from 2014, with the model attributing part of the large positive spike in output to the cyclical component of output. This effectively limits the usefulness of the DG-ECFIN model in determining the appropriate stance of macroeconomic policy. Both the IMF and BMA models suggest a relatively less rapid adjustment in the output gap.

Overall, our model attributes a greater share of variance in the data to output gap movements, rather than potential output changes. Consequently, we do not suffer from the conjectural secular stagnation issue raised by Rogoff (2016). For example, the peak of the economic expansion in 2007 is estimated to be 11.8 per cent under our trivariate model average; in contrast, the comparable DG-ECFIN estimate is 9.25 per cent, while the IMF estimate is smaller again at 7.5 per cent. Similarly, peak recession effects in response to the GFC are estimated at -11.9 per cent in our model, versus -13.3 per cent and -6 per cent in the EC-DGFIN and IMF series. These results may reflect lower amplitude and smaller periodicity in our estimation of trend output growth, likely due to our BMA model incorporating stochastic volatility and Student's t -distributional assumptions. If our model attributes a lower share of negative output growth to changes in permanent factors during recessions, a greater share will be attributed to the transitory component of output growth, leading to a larger output gap (in absolute terms) during these periods. As a result, our BMA model shows more plausible GFC and SDC declines, and post-crisis improvements, in the output gap, with estimates rising from strongly negative positions in 2013 to a closed (marginally positive) gap at the end of the sample. This narrative is consistent with the the post-crisis data on core inflation in Ireland.

Figure 3 presents the estimates of trend output growth from 1980 to 2016. As is immediately apparent from the graph, the DG-ECFIN model imposes a high degree of periodicity on its potential output series. Aside from the low likelihood of potential output following such a smooth curvature over time, the imposition of this structure also causes a number of turning point issues. All three models identify similar levels of potential output at the beginning of the sample, with values rising over the 1980s to provide identical estimates by 1990. However, while the DG-ECFIN model follows its smooth growth path, both the IMF and BMA models show lower potential output growth rates for the early part of the 1990s than the latter period; a decomposition that seems more consistent with the effects of the Celtic Tiger. Similarly, while the DG-ECFIN curve suggests that potential output growth peaked in 1997 (before the full effects of the Celtic Tiger were observed), the BMA estimates suggest the growth rate did not peak until 2000.

Furthermore, the DG-CFIN model suggests a smooth, consistent decline in the growth rate of potential output between 1997 and 2009, at odds with both the large-scale decline in unemployment in the late 1990s and the immediacy of the crisis effects observed in the manufacturing, construction and financial sectors in 2008. Our BMA estimates show a much slower rate of decline in potential output over the 2000-2007 period when compared to the 2007-2012 period, consistent with prior assumptions. And while both the IMF and BMA models suggest that potential output growth reached its lowest point during the sovereign debt crisis in 2012, the DG-ECFIN estimate suggests that this occurred in 2009, and that potential output growth increased substantially during the SDC. Finally, the BMA and IMF model

Figure 3: Estimates of Potential Output for Ireland

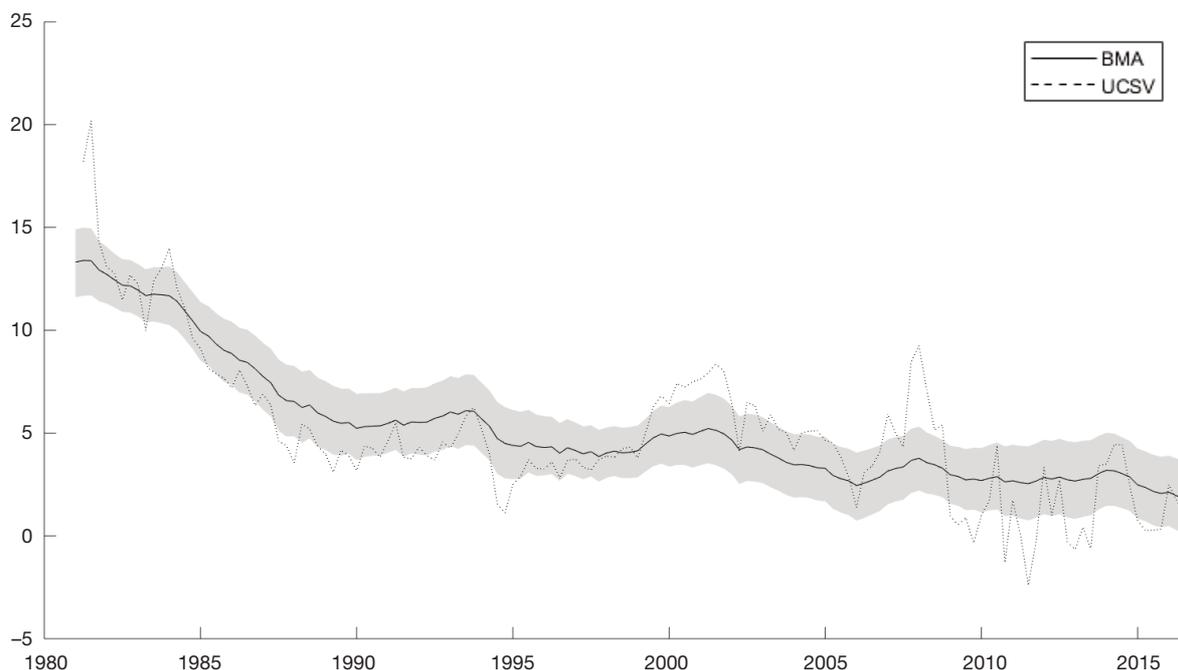
Source: Author's analysis.

estimates are broadly similar at the end of the sample, while the DG-ECFIN model estimate is considerably larger, most likely due to the distortionary MNE activity present in the GDP series.

Overall, both the IMF and DG-ECFIN estimates of potential output show a stronger degree of periodicity than our BMA model, with greater peaks and troughs over the sample period. However, as has been suggested by Jarociński and Lenza (2018), the imposition of these high-amplitude, strongly periodic structures in standard output gap models for the euro area Member States estimates a post-crisis, secular stagnation-like decline in trend growth and a small output gap, which is inconsistent with developments in consumer price inflation and other macroeconomic variables.

6.5 Estimation of Trend Inflation, NAWRU and Stochastic Volatility

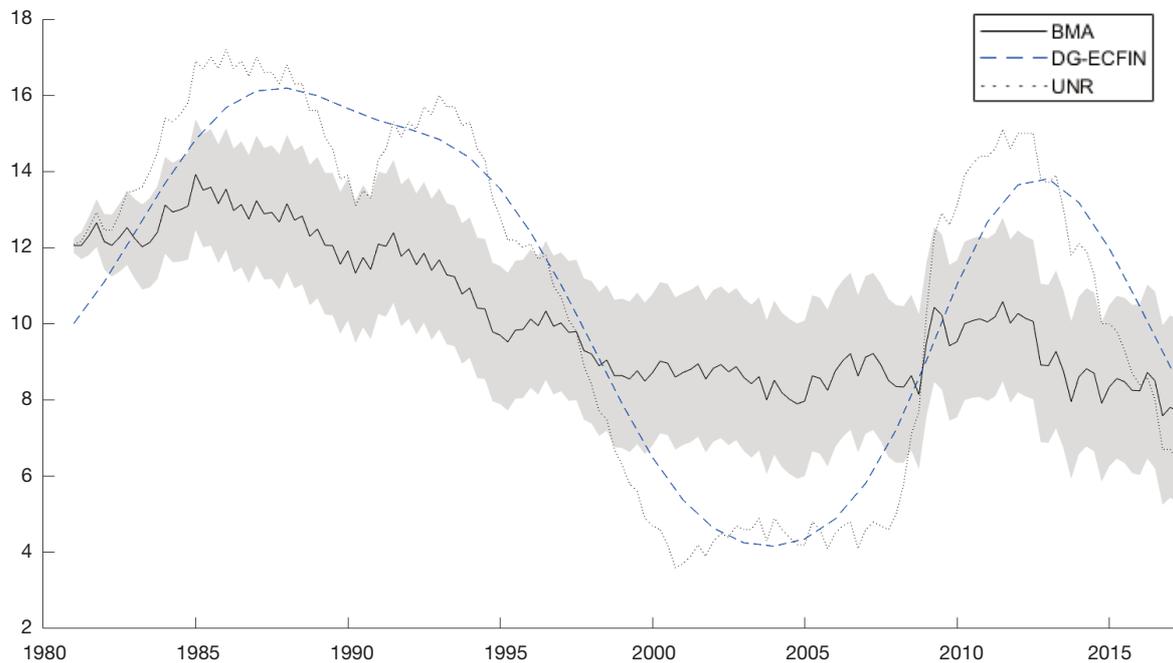
Turning to the wage inflation and unemployment series, Figure 4 presents the trend wage inflation estimates and sample error bands from our averaged model. We calculate the error bands as the 10 per cent and 90 per cent quantiles of the posterior draws of π^* . As there are no publicly available comparators of trend wage inflation for the Irish economy, we evaluate our BMA estimates against the unobserved components with stochastic volatility (UCSV) model of Stock and Watson (2007). The UCSV model is a univariate model that assumes the wage inflation gap exhibits no persistence and embeds the principle that the gap is itself governed by a separate process for stochastic volatility. Innovations to trend wage inflation and the wage inflation gap, plus the shocks to stochastic volatility, are assumed to be serially and mutually uncorrelated.

Figure 4: Estimates of Trend Wage Inflation for Ireland

Source: Author's analysis.

As can be seen from Figure 4, both series broadly follow a similar time path, although there are some characteristic differences. Unsurprisingly, the UCSV model results in more extreme values for the estimates of trend wage inflation across the sample period. This can, in part, be attributed to the univariate nature of the UCSV model, versus our multivariate model, where both output and unemployment can potentially explain some of the movements in the wage inflation series. For example, there is a lower trend level of wage inflation in the UCSV series post-2009, than the trend estimated by our BMA model. Thus, our BMA model suggests that, in order to explain both the high unemployment rate and a wage inflation rate that is higher than predicted by the wage Phillips curve relationship, trend wage inflation must be higher.

Figure 5 presents the estimates of the NAWRU from our process, plotted against the DG-ECFIN estimate of the NAWRU and the raw unemployment rate series. As can be seen from the graph, the NAWRU estimated by the BMA model has declined steadily since 1985, from a peak of 13.9 per cent in Q1 1985 to 7.7 per cent in Q4 2016. Given the proximity of the output gap to zero at the end of our sample (Figure 2), it is unsurprising that our estimate of the NAWRU intersects with the prevailing unemployment rate in 2016, indicating that the excess unemployment that had developed in Ireland since Q1 2009, as a result of the GFC and SDC, had effectively dissipated by 2016.

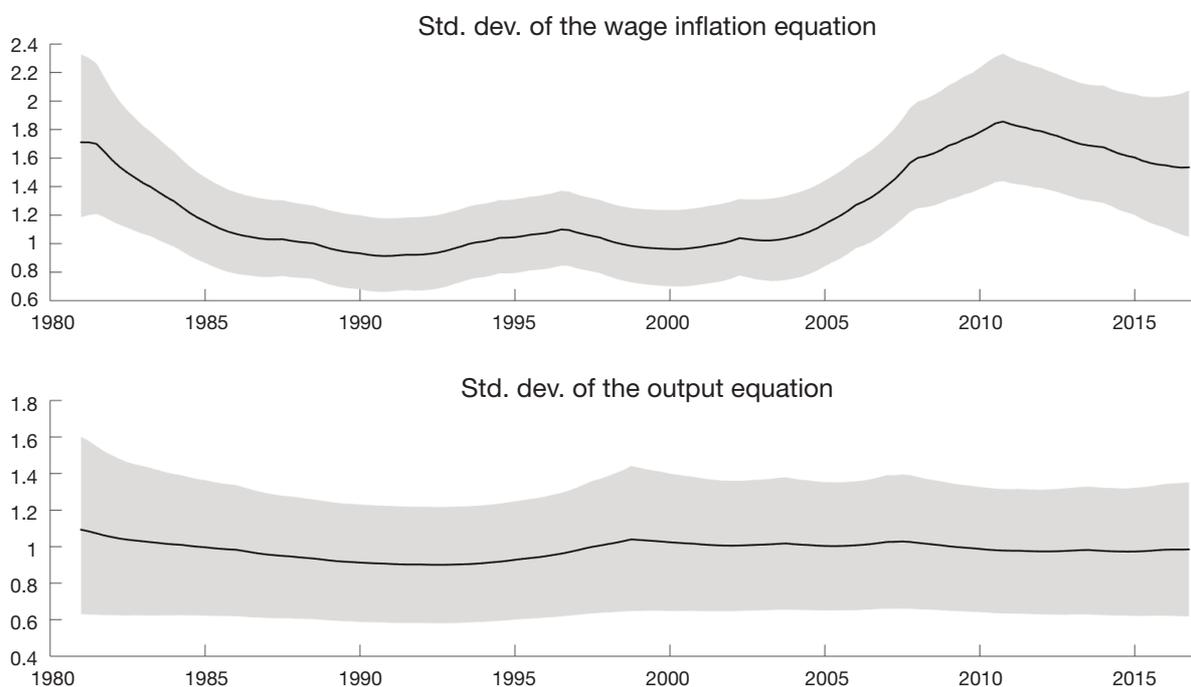
Figure 5: Estimates of the Trend Unemployment Rate for Ireland

Source: Author's analysis.

In contrast, the DG-ECFIN estimate of the NAWRU shows a number of issues related to the use of the HP-filter in identifying latent variables. From 2006, the DG-ECFIN estimate of the NAIRU starts to rise, before crisis effects are observed in the data. This problem results from using a statistical filter that conditions on future observations of the data to obtain a latent trend value. Similarly, while the DG-ECFIN NAWRU declines between 1995 and 2004, the rate of decline starts to slow in 1999, despite the continued decline in unemployment until 2001. Again, this effect is an artefact of using a filter that conditions on future values of the series to conduct a trend/cycle decomposition.

Finally, for this section, we plot the estimates of stochastic volatility from the wage inflation and output equations of our model in Figure 6. As can be seen from the upper panel, with respect to their volatility, wage inflation innovations appear to be strongly episodic. From a high starting point in 1980, the volatility of wage inflation shocks declines during the 1980s and early 1990s, before rising in the late 1990s until the advent of the SDC in 2011, and declining from this point to the end of sample in 2016. In contrast, the lower panel of Figure 6 suggests that stochastic volatility is not as strong a feature in output shocks, declining from 1980 to the mid-1990s, rising moderately in the mid-to-late 1990s, before remaining at a relatively constant level over the remainder of the sample period. However, given that the search exercise identified output shocks to follow Student's t -distributions, the argument that rare, large shocks drive Irish business cycles remains supported.

Figure 6: Estimates of Stochastic Volatility in Wage Inflation and Output Shocks



Source: Author's analysis.

VII CONCLUSION

This paper attempts to overcome the inherent uncertainty in estimating the output gap via semi-structural equations and univariate filtering, by employing a Bayesian model averaging approach. Using a multivariate model with a trivariate set of equations, our approach uses a weighted average of estimates from a broad set of potential specifications of output, wage inflation and unemployment gaps. Choosing across the set of variables included in the system of equations, the specifications for modelling our potential output variable, and the distribution of innovations to our wage inflation and output equations, we identify 256 potential variants of our system and construct a suitable model space in which our specification search can be conducted.

Our model averaging suggests that the unemployment gap is a strong predictor of both the output gap and the wage inflation gap; but conditional on the unemployment gap, the output gap does not provide any information on the wage inflation gap. Our results indicate that potential output growth declined between 2001 and 2012, from a peak of 5.9 per cent to a low of 2 per cent. However, potential output has risen since this point, with stronger growth recorded at the end of our sample period in 2016. Similarly, the output gap declined significantly between 2008 and 2012, but has strengthened to the point where the gap was above the zero line at the end of 2016.

Furthermore, we find evidence in the data to suggest that innovations to wage inflation and output are better characterised by Student's t -distributions, rather than the standard assumption of Gaussian errors. This indicates that tail events occur more frequently than suggested by the Gaussian distribution, which places relatively more weight on median outcomes. This Gaussian characterisation may not be reflective of the true state of the world, particularly given Ireland's exposure to large macroeconomic shocks as a small open economy. Further research on incorporating Student's t -distributed errors (and stochastic volatility) into forecasting models of the Irish economy may be warranted.

Comparing our results to output gap and potential output estimates from DG-ECFIN and the IMF, our BMA model produces estimates of potential output that are less periodic, and lower in amplitude. As a result, we do not observe conjectural secular stagnation issues from the global financial crisis (IMF model), allowing for output gap estimates that are consistent with the argument that weak growth post-crisis reflects cyclical, although persistent, sources of fluctuations. This narrative is consistent with the data on core inflation in Ireland. As we do not use filters that condition on future values of a series, our potential output series does not suffer from turning point issues (DG-ECFIN model), where recessionary or expansionary effects alter the path of the series before they occur. Our output gap estimates show the cyclical component of output to be stronger than those observed in the IMF and DG-ECFIN models, with larger gaps (in absolute terms) identified in recessionary and expansionary periods. This is likely due to the less-periodic estimates of potential output, and the incorporation of Student's t -distributed error terms into our model, highlighting the importance of rare shocks in driving the Irish business cycle.

REFERENCES

- Alberola, Enrique, Ángel Estrada and Daniel Santabárbara, 2013. "Growth Beyond Imbalances. Sustainable Growth Rates and Output Gap Reassessment", Working Papers 1313, Banco de España, October 2013.
- Álvarez, Luis J. and Ana Gómez-Loscos, 2018. "A Menu on Output Gap Estimation Methods", *Journal of Policy Modeling*, Vol. 40(4):827-850.
- Ball, Laurence M., Daniel Leigh and Prakash Loungani, 2013. "Okun's Law: Fit at Fifty?". NBER Working Papers 18668, National Bureau of Economic Research, Inc, January 2013.
- Basistha, Arabinda and Charles R. Nelson, 2007. "New Measures of the Output Gap based on the Forward-Looking New Keynesian Phillips Curve", *Journal of Monetary Economics*, Vol. 54, No. 2, pp. 498-511, March 2007.
- Basistha, Arabinda and Richard Startz, 2008. "Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach", *The Review of Economics and Statistics*, Vol. 90, No. 4, pp. 805-811, November 2008.
- Berger, Tino and Julia Richter, 2017. "Estimating Finance-Neutral Output Gaps". Working Paper, Gottingen University Department of Economics.

- Berger, Tino, Gerdie Everaert and Hauke Vierke, 2016. "Testing for Time Variation in an Unobserved Components Model for the U.S. Economy", *Journal of Economic Dynamics and Control*, Vol. 69(C), pp. 179-208.
- Bergin, Adele and John FitzGerald, 2014. "The Structural Balance for Ireland", *Quarterly Economic Commentary: Special Articles*.
- Beveridge, Stephen and Charles R. Nelson, 1981. "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle", *Journal of Monetary Economics*, Vol. 7, No. 2, pp. 151-174.
- Borio, Claudio, Piti Disyatat and Mikael Juselius, 2017. "Rethinking Potential Output: Embedding Information about the Financial Cycle", *Oxford Economic Papers*, Vol. 69, No. 3, pp. 655-677.
- Byrne, Stephen and Martin D. O'Brien, 2017. "Understanding Irish Labour Force Participation", *The Economic and Social Review*, Vol. 48, No. 1, pp. 27-60.
- Chan, Joshua C. C., Gary Koop and Simon M. Potter, 2016. "A Bounded Model of Time Variation in Trend Inflation, NAIRU and the Phillips Curve", *Journal of Applied Econometrics*, Vol. 31, No. 3, pp. 551-565.
- Chan, Joshua C. C. and Angelia L. Grant, 2017. "Measuring the Output Gap using Stochastic Model Specification Search". CAMA Working Papers 2017-02, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- Chan, Joshua C.C. and Ivan Jeliazkov, 2009. "Efficient Simulation and Integrated Likelihood Estimation in State Space Models", *International Journal of Mathematical Modelling and Numerical Optimisation*, Vol. 1, Nos. 1-2, pp. 101-120.
- Chiu, Ching-Wai (Jeremy), Haroon Mumtaz and Gábor Pintér, 2016a. "Bayesian Vector Autoregressions with Non-Gaussian Shocks", CREMFi Discussion Papers 5, CREMFi, School of Economics and Finance, QMUL.
- Chiu, Ching-Wai (Jeremy), Haroon Mumtaz and Gábor Pintér, 2016b. "VAR Models with Non-Gaussian Shocks", Discussion Papers 1609, Centre for Macroeconomics (CFM).
- Chiu, Ching-Wai (Jeremy), Haroon Mumtaz and Gábor Pintér, 2017. "Forecasting with VAR Models: Fat Tails and Stochastic Volatility", *International Journal of Forecasting*, Vol. 33, No. 4, pp. 1124-1143.
- Clark, Peter K., 1987. "The Cyclical Component of U.S. Economic Activity", *The Quarterly Journal of Economics*, Vol. 102, No. 4, pp. 797-814.
- Coibion, Olivier and Yuriy Gorodnichenko, 2015. "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation", *American Economic Journal: Macroeconomics*, Vol. 7, No. 1, pp. 197-232.
- Coibion, Olivier, Yuriy Gorodnichenko and Mauricio Ulate, 2017. "The Cyclical Sensitivity in Estimates of Potential Output", NBER Working Papers 23580, National Bureau of Economic Research, Inc.
- Conefrey, Thomas, Martina Lawless and Suzanne Linehan, 2014. "Developments in the Irish Labour Market during the Crisis: What Lessons for Policy?", *Journal of the Statistical and Social Inquiry Society of Ireland*, Vol. XLIV, pp. 18-39.
- Connor, Gregory, Thomas Flavin and Brian O'Kelly, 2012. "The US and Irish Credit Crises: Their Distinctive Differences and Common Features", *Journal of International Money and Finance*, Vol. 31, No. 1, pp. 60-79.
- Cross, Jamie and Aubrey Poon, 2016. "Forecasting Structural Change and Fat-Tailed Events in Australian Macroeconomic Variables", *Economic Modelling*, Vol. 58(C), pp. 34-51.
- Cúrdia, Vasco, Marco Del Negro and Daniel L. Greenwald, 2014. "Rare Shocks, Great Recessions", *Journal of Applied Econometrics*, Vol. 29, No. 7, pp. 1031-1052.
- Darvas, Zsolt and András Simon, 2015. "Filling the Gap: Open Economy Considerations for More Reliable Potential Output Estimates". Working Papers 10401, Bruegel, October 2015.

- de la Fuente, Ángel, 2009. "A Mixed Splicing Procedure for Economic Time Series". Working Papers 415, Barcelona Graduate School of Economics, December 2009.
- Eisenstat, Eric, Joshua C.C. Chan and Rodney W. Strachan, 2016. "Stochastic Model Specification Search for Time-Varying Parameter VARs", *Econometric Reviews*, Vol. 35, Nos. 8-10, pp. 1638-1665.
- FitzGerald, John, 2018. "National Accounts for a Global Economy: the Case of Ireland". Trinity Economics Papers TEP0418, Trinity College Dublin, Department of Economics.
- Frühwirth-Schnatter, Sylvia and Helga Wagner, 2010. "Stochastic Model Specification Search for Gaussian and Partial non-Gaussian State Space Models", *Journal of Econometrics*, Vol. 154, No. 1, pp. 85-100.
- Fueki, Takuji, Ichiro Fukunaga, Hibiki Ichiue and Toyochiro Shirota, 2016. "Measuring Potential Growth with an Estimated DSGE Model of Japan's Economy", *International Journal of Central Banking*, Vol. 12, No. 1, pp. 1-32.
- Galí, Jordi, 2011. "The Return Of The Wage Phillips Curve", *Journal of the European Economic Association*, Vol. 9, No. 3, pp. 436-461.
- Garnier, Christine, Elmar Mertens and Edward Nelson, 2015. "Trend Inflation in Advanced Economies", *International Journal of Central Banking*, Vol. 11, No. 4, pp. 65-136.
- Gerlach, Stefan, Reamonn Lydon and Rebecca Stuart, 2016. "Unemployment and Inflation in Ireland: 1926-2012", *Cliometrica*, Vol. 10, No. 3, pp. 345-364.
- Geweke, J. and Amisano, G., 2011. "Hierarchical Markov Normal Mixture Models with Applications to Financial Asset Returns", *Journal of Applied Econometrics*, Vol. 26, No. 1, pp. 1-29.
- Gordon, Robert J., 2013. "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery", NBER Working Papers 19390, National Bureau of Economic Research, Inc, August 2013.
- Grassi, Stefano and Tommaso Proietti, 2014. "Characterising Economic Trends by Bayesian Stochastic Model Specification Search", *Computational Statistics & Data Analysis*, Vol. 71(C), pp. 359-374.
- Hardouvelis, Gikas A., 1988. "Inflationary Bias and Openness". Research Paper 8807, Federal Reserve Bank of New York.
- Hardouvelis, Gikas A., 1992. "Monetary Policy Games, Inflationary Bias, and Openness", *Journal of Economic Dynamics and Control*, Vol. 16, No. 1, pp. 147-164.
- Harvey, Andrew C. and Albert Jaeger, 1993. "Detrending, Stylized Facts and the Business Cycle", *Journal of Applied Econometrics*, Vol. 8, No. 3, pp. 231-247, July-Sept 1993.
- Harvey, Andrew C., Thomas M. Trimbur and Herman K. Van Dijk, 2007. "Trends and Cycles in Economic Time Series: A Bayesian Approach", *Journal of Econometrics*, Vol. 140, No. 2, pp. 618-649.
- Honohan, Patrick, 1992. "The Link between Irish and UK Unemployment", *Quarterly Economic Commentary: Special Articles*, 1992(1-Spring).
- Jarociński, Marek and Michele Lenza, 2018. "An Inflation-Predicting Measure of the Output Gap in the Euro Area", *Journal of Money, Credit and Banking*, Vol. 50, No. 6, pp. 1189-1224.
- Justiniano, Alejandro and Giorgio E. Primiceri, 2008. "The Time-Varying Volatility of Macroeconomic Fluctuations", *The American Economic Review*, Vol. 98, No. 3, pp. 604-641.
- Kamber, Günes, James Morley and Benjamin Wong, 2016. "Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter". BIS Working Papers 584, Bank for International Settlements, September 2016.
- Kiley, Michael T., 2013. "Output Gaps", *Journal of Macroeconomics*, Vol. 37, pp. 1-18.
- Kim, Chang-Jin, Pym Manopimoke and Charles R. Nelson, 2014. "Trend Inflation and the Nature of Structural Breaks in the New Keynesian Phillips Curve", *Journal of Money, Credit and Banking*, Vol. 46, Nos. 2-3, pp. 253-266.

- Kim, Sangjoon, Neil Shephard and Siddhartha Chib, 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models", *Review of Economic Studies*, Vol. 65, No. 3, pp. 361-393.
- Konuki, Tetsuya, 2010. "Estimating Potential Output and the Output Gap in Slovakia", *Eastern European Economics*, Vol. 48, No. 2, pp. 39-55, March 2010.
- Lane, Philip R., 1997. "Inflation in Open Economies", *Journal of International Economics*, Vol. 42, Nos. 3-4, pp. 327-347.
- Lane, Philip R., 2014. "International Financial Flows and the Irish Crisis". *CESifo Forum*, Vol. 15, No. 2, pp. 14-19, April 2014.
- Lane, Philip R. and Francis Ruane, 2006. "Globalisation and the Irish Economy", The Institute for International Integration Studies Discussion Paper Series, IIIS, March 2006.
- Lane, Philip R. and Gian Maria Milesi-Ferretti, 2012. "External Adjustment and the Global Crisis", *Journal of International Economics*, Vol. 88, No. 2, pp. 252-265.
- Luo, Sui and Richard Startz, 2014. "Is it One Break or Ongoing Permanent Shocks that Explains US Real GDP?", *Journal of Monetary Economics*, Vol. 66(C), pp. 155-163.
- Meyler, Aidan, 1999. "The Non-Accelerating Inflation Rate of Unemployment (NAIRU) in a Small Open Economy: The Irish Context". Research Technical Papers 5/RT/99, Central Bank of Ireland, December 1999.
- Morley, James and Jeremy Piger, 2012. "The Asymmetric Business Cycle", *The Review of Economics and Statistics*, Vol. 94, No. 1, pp. 208-221.
- Morley, James C., Charles R. Nelson and Eric Zivot, 2003. "Why are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP so Different?", *The Review of Economics and Statistics*, Vol. 85, No. 2, pp. 235-243.
- Orlandi, Fabrice, Werner Roeger and Anna Thum-Thysen, 2018. "The Return of the European Wage Phillips Curve". European Economy - Discussion Papers 2015 – 085, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission, September 2018.
- Planas, Christophe, Alessandro Rossi and Gabriele Fiorentini, 2008. "Bayesian Analysis of the Output Gap", *Journal of Business & Economic Statistics*, Vol. 26, pp. 18-32.
- Rogoff, Kenneth, 2016. *Debt Supercycle, Not Secular Stagnation*,. Cambridge: MIT Press, pp. 19-28, .
- Romer, David, 1993. "Openness and Inflation: Theory and Evidence", *The Quarterly Journal of Economics*, Vol. 108, No. 4, pp. 869-903.
- Slevin, Geraldine, 2001. "Potential Output and the Output Gap in Ireland", Research Technical Papers 5/RT/01, Central Bank of Ireland, September 2001.
- Stella, Andrea and James H. Stock, 2012. *A State-Dependent Model for Inflation Forecasting*. International Finance Discussion Papers 1062, Board of Governors of the Federal Reserve System (US).
- Stock, James H. and Mark W. Watson, 2007. "Why has US Inflation Become Harder to Forecast?", *Journal of Money, Credit and Banking*, Vol. 39(s1), pp. 3-33.
- Walsh, Brendan, 2000. "Cyclical and Structural Influences on Irish Unemployment", *Oxford Economic Papers*, Vol. 52, No. 1, pp. 119-145.
- Walsh, Brendan, 1999. "The Persistence of High Unemployment in a Small Open Labour Market: The Irish Case", in Barry F. (ed.), *Understanding Ireland's Economic Growth*, pp. 193-226. London: Palgrave Macmillan.
- Watson, Mark W., 1986. "Univariate Detrending Methods with Stochastic Trends", *Journal of Monetary Economics*, Vol. 18, No. 1, pp. 49-75.

APPENDIX

The process we use to obtain the estimates in this paper is a mixed Markov Chain Monte Carlo algorithm, taking elements from the samplers proposed by Frühwirth-Schnatter and Wagner (2010), Chan and Jeliaskov (2009) and Kim *et al.* (1998). Posterior draws are obtained through the following multi-step procedure:

1. **Sample $\boldsymbol{\pi}^*$ from its full conditional distribution.** Since our wage inflation and trend wage inflation equations define a linear Gaussian state-space model, we can use standard algorithms to sample $\boldsymbol{\pi}^*$. To derive the conditional density, we can express the wage inflation equation in (1) as

$$\mathbf{W}_{\tau^\pi} \boldsymbol{\pi} = \mathbf{W}_{\tau^\pi} \boldsymbol{\pi}^* + \beta_{\gamma^u} \gamma^u (\mathbf{u} - \mathbf{u}^*) + \beta_{\gamma^y} \gamma^y (\mathbf{y} - \mathbf{y}^*) + \boldsymbol{\varepsilon}^\pi$$

where

$$\mathbf{W}_{\tau^\pi} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -\tau^\pi & 1 & 0 & \cdots & 0 \\ 0 & -\tau^\pi & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -\tau^\pi & 1 \end{pmatrix}$$

and $\boldsymbol{\varepsilon}^\pi \sim N(\mathbf{0}, \mathbf{S}_\pi)$. If $\beta_{z^\pi} = 1$, then $\mathbf{S}_\pi = \text{diag}(z_1^\pi e^{h_1^\pi}, z_2^\pi e^{h_2^\pi}, \dots, z_T^\pi e^{h_T^\pi})$. Otherwise, $\mathbf{S}_\pi = \text{diag}(e^{h_1^\pi}, e^{h_2^\pi}, \dots, e^{h_T^\pi})$. Since $|\mathbf{W}_{\tau^\pi}| = 1 \forall \tau^\pi$, the conditional log-likelihood for $\boldsymbol{\pi}$ is given by

$$-\frac{1}{2} (\boldsymbol{\pi} - \boldsymbol{\pi}^* - \boldsymbol{\Delta}_{\boldsymbol{\pi}^*})' \mathbf{W}_{\tau^\pi}' \mathbf{S}_\pi^{-1} \mathbf{W}_{\tau^\pi} (\boldsymbol{\pi} - \boldsymbol{\pi}^* - \boldsymbol{\Delta}_{\boldsymbol{\pi}^*})$$

where $\boldsymbol{\Delta}_{\boldsymbol{\pi}^*} = \mathbf{W}_{\tau^\pi}^{-1} (\beta_{\theta^u} \theta^u (\mathbf{u} - \mathbf{u}^*) + \beta_{\theta^y} \theta^y (\mathbf{y} - \mathbf{y}^*))$. We can rewrite the trend wage inflation rate from (2) as

$$\mathbf{W} \boldsymbol{\pi}^* = \tilde{\boldsymbol{\eta}}_{\boldsymbol{\pi}^*} + \boldsymbol{\varepsilon}^{\boldsymbol{\pi}^*}$$

where $\boldsymbol{\varepsilon}^{\boldsymbol{\pi}^*} \sim N(\mathbf{0}, \mathbf{S}_{\boldsymbol{\pi}^*})$, $\mathbf{S}_{\boldsymbol{\pi}^*} = \text{diag}(V_{\boldsymbol{\pi}^*}, \sigma_{\boldsymbol{\pi}^*}^2, \dots, \sigma_{\boldsymbol{\pi}^*}^2)$, $\tilde{\boldsymbol{\eta}}_{\boldsymbol{\pi}^*} = (\boldsymbol{\pi}_0^*, 0, \dots, 0)'$, and

$$\mathbf{W} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -1 & 1 \end{pmatrix}$$

As $|\mathbf{W}| = 1$ the log prior density of $\boldsymbol{\pi}^*$ can be represented as

$$-\frac{1}{2}(\boldsymbol{\pi}^* - \boldsymbol{\eta}_{\pi^*})' \boldsymbol{W}' \boldsymbol{S}_{\pi^*}^{-1} \boldsymbol{W} (\boldsymbol{\pi}^* - \boldsymbol{\eta}_{\pi^*})$$

where $\boldsymbol{\eta}_{\pi^*} = \boldsymbol{W}^{-1} \tilde{\boldsymbol{\eta}}_{\pi^*}$. Through regression analysis, it can be shown that $\boldsymbol{\pi}^*$ follows the conditional distribution $\boldsymbol{\pi}^* \sim N(\hat{\boldsymbol{\pi}}^*, \boldsymbol{P}_{\pi^*}^{-1})$, where

$$\begin{aligned} \boldsymbol{P}_{\pi^*} &= \boldsymbol{W}' \boldsymbol{S}_{\pi^*}^{-1} \boldsymbol{W} + \boldsymbol{W}'_{\tau\pi} \boldsymbol{S}_{\pi}^{-1} \boldsymbol{W}_{\tau\pi} \\ \hat{\boldsymbol{\pi}}^* &= \boldsymbol{P}_{\pi^*}^{-1} (\boldsymbol{W}' \boldsymbol{S}_{\pi^*}^{-1} \boldsymbol{W} \boldsymbol{\eta}_{\pi^*} + \boldsymbol{W}'_{\tau\pi} \boldsymbol{S}_{\pi}^{-1} \boldsymbol{W}_{\tau\pi} (\boldsymbol{\pi} - \boldsymbol{\Delta}_{\pi^*})) \end{aligned}$$

As \boldsymbol{W} and $\boldsymbol{W}_{\tau\pi}$ are both banded matrices, the precision matrix \boldsymbol{P}_{π^*} is also a banded matrix, and $\boldsymbol{\pi}^*$ can be sampled efficiently using a standard precision sampler.

2. **Sample \boldsymbol{u}^* from its full conditional distribution.** Information on \boldsymbol{u}^* is contained in three of the four structural equations of our model (1), plus the state equation of unemployment (3). As in Step 1, if $\beta_{z\omega} = 1$, define $\boldsymbol{S}_{\omega} = \text{diag}(z_1^{\omega} e^{h_1^{\omega}}, z_2^{\omega} e^{h_2^{\omega}}, \dots, z_T^{\omega} e^{h_T^{\omega}})$. Otherwise, $\boldsymbol{S}_{\omega} = \text{diag}(e^{h_1^{\omega}}, e^{h_2^{\omega}}, \dots, e^{h_T^{\omega}})$. We can rewrite the four equations as

$$\begin{aligned} \boldsymbol{W}_{\tau\pi} (\boldsymbol{\pi} - \boldsymbol{\pi}^*) &= \beta_{\gamma^u} \boldsymbol{\gamma}^u (\boldsymbol{u} - \boldsymbol{u}^*) + \beta_{\gamma^y} \boldsymbol{\gamma}^y (\boldsymbol{y} - \boldsymbol{y}^*) + \boldsymbol{\varepsilon}^{\pi} \\ \boldsymbol{y} - \boldsymbol{y}^* &= \beta_{\theta^u} \boldsymbol{\theta}^u (\boldsymbol{u} - \boldsymbol{u}^*) + \boldsymbol{\omega} \\ \boldsymbol{W}_{\tau u} (\boldsymbol{u} - \boldsymbol{u}^*) &= \boldsymbol{\varepsilon}^u \\ \boldsymbol{W}_{\tau\omega} \boldsymbol{\omega} &= \boldsymbol{\varepsilon}^{\omega} \\ \boldsymbol{W} \boldsymbol{u}^* &= \boldsymbol{W} \boldsymbol{\eta}_{u^*} + \boldsymbol{\varepsilon}^{u^*} \end{aligned}$$

where $\boldsymbol{\varepsilon}^{\omega} \sim N(\mathbf{0}, \boldsymbol{S}_{\omega})$, $\boldsymbol{\varepsilon}^u \sim N(\mathbf{0}, \boldsymbol{S}_u)$, $\boldsymbol{S}_u = \text{diag}(\sigma_u^2, \sigma_u^2, \dots, \sigma_u^2)$, $\boldsymbol{\varepsilon}^{u^*} \sim N(\mathbf{0}, \boldsymbol{S}_{u^*})$, $\boldsymbol{S}_{u^*} = \text{diag}(V_{u^*}, \sigma_{u^*}^2, \dots, \sigma_{u^*}^2)$, $\boldsymbol{\eta}_{u^*} = \boldsymbol{W}^{-1}(u_0^*, 0, \dots, 0)$ and

$$\boldsymbol{W}_{\tau u} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -\tau_1^u & 1 & 0 & 0 & \cdots & 0 \\ -\tau_2^u & -\tau_1^u & 1 & 0 & \cdots & 0 \\ 0 & -\tau_2^u & -\tau_1^u & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -\tau_2^u & -\tau_1^u & 1 \end{pmatrix}$$

$$\boldsymbol{W}_{\tau\omega} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -\tau_1^{\omega} & 1 & 0 & 0 & \cdots & 0 \\ -\tau_2^{\omega} & -\tau_1^{\omega} & 1 & 0 & \cdots & 0 \\ 0 & -\tau_2^{\omega} & -\tau_1^{\omega} & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -\tau_2^{\omega} & -\tau_1^{\omega} & 1 \end{pmatrix}$$

As with $\boldsymbol{\pi}^*$, it can be shown that the conditional distribution of the NAWRU can be represented as $\boldsymbol{u}^* \sim N(\hat{\boldsymbol{u}}, \boldsymbol{P}_{\boldsymbol{u}^*}^{-1})$, where

$$\begin{aligned} \boldsymbol{P}_{\boldsymbol{u}^*} &= \boldsymbol{W}' \boldsymbol{S}_{\boldsymbol{u}^*}^{-1} \boldsymbol{W} + \boldsymbol{W}'_{\tau u} \boldsymbol{S}_u^{-1} \boldsymbol{W}_{\tau u} + (\beta_{\theta u} \theta^u)^2 \boldsymbol{W}'_{\tau \omega} \boldsymbol{S}_\omega^{-1} \boldsymbol{W}_{\tau \omega} + (\beta_{\gamma^u} \gamma^u)^2 \boldsymbol{S}_\pi^{-1} \\ \hat{\boldsymbol{u}}^* &= \boldsymbol{P}_{\boldsymbol{u}^*}^{-1} (\boldsymbol{W}' \boldsymbol{S}_{\boldsymbol{u}^*}^{-1} \boldsymbol{W} \boldsymbol{\eta}_{\boldsymbol{u}^*} + \boldsymbol{W}'_{\tau \pi} \boldsymbol{S}_u^{-1} \boldsymbol{W}_{\tau u} \boldsymbol{u} - \beta_{\theta u} \theta^u \boldsymbol{W}'_{\tau \omega} \boldsymbol{S}_\omega^{-1} \boldsymbol{W}_{\tau \omega} (\boldsymbol{y} - \boldsymbol{y}^* - \beta_{\theta u} \theta^u \boldsymbol{u}) \\ &\quad - \beta_{\gamma^u} \gamma^u \boldsymbol{S}_\pi^{-1} (\boldsymbol{W}_{\tau \pi} (\boldsymbol{\pi} - \boldsymbol{\pi}^*) - \beta_{\gamma^y} \gamma^y (\boldsymbol{y} - \boldsymbol{y}^*) - \beta_{\theta u} \theta^u \boldsymbol{u})) \end{aligned}$$

As with $\boldsymbol{P}_{\boldsymbol{\pi}^*}$ the precision matrix $\boldsymbol{P}_{\boldsymbol{u}^*}$ is also a banded matrix, so \boldsymbol{u}^* can be sampled efficiently using a standard precision sampler.

3. **Sample $\tilde{\boldsymbol{y}}^*$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distributions and randomly permute the signs of $\tilde{\boldsymbol{y}}^*$ and σ_{y^*} .** To derive the conditional distribution of $\tilde{\boldsymbol{y}}^*$, we can rewrite the non-centered parameterisation of the trend output equation as

$$\boldsymbol{y}^* = y_0^* \boldsymbol{I}_1 + \beta_{\alpha_0} \alpha_0 \boldsymbol{I}_T + \beta_{y^*} \sigma_{y^*} \tilde{\boldsymbol{y}} + \beta_\alpha \sigma_\alpha \tilde{\boldsymbol{A}}$$

where $\boldsymbol{I}_1 = (1, 1, \dots, 1)'$, $\boldsymbol{I}_T = (1, 2, \dots, T)'$, and $\tilde{\boldsymbol{A}} = (\tilde{\alpha}_1, \tilde{\alpha}_1 + \tilde{\alpha}_2, \dots, \sum_{t=1}^T \tilde{\alpha}_t)'$. Re-writing the structural equations for the wage inflation and output gaps in terms of $\tilde{\boldsymbol{y}}$,

$$\begin{aligned} \boldsymbol{W}_{\tau \pi} \boldsymbol{\pi} &= \boldsymbol{W}_{\tau \pi} \boldsymbol{\pi}^* + \beta_{\gamma^u} \gamma^u (\boldsymbol{u} - \boldsymbol{u}^*) + \beta_{\gamma^y} \gamma^y (\boldsymbol{y} - \boldsymbol{\xi}_{\tilde{\boldsymbol{y}}^*} - \beta_{y^*} \sigma_{y^*} \tilde{\boldsymbol{y}}^*) + \boldsymbol{\varepsilon}^\pi \\ \boldsymbol{y} &= \boldsymbol{\xi}_{\tilde{\boldsymbol{y}}^*} + \beta_{y^*} \sigma_{y^*} \tilde{\boldsymbol{y}} + \beta_{\theta u} \theta^u (\boldsymbol{u} - \boldsymbol{u}^*) + \boldsymbol{\omega} \end{aligned}$$

where $\boldsymbol{\xi}_{\tilde{\boldsymbol{y}}^*} = y_0^* \boldsymbol{I}_1 + \beta_{\alpha_0} \alpha_0 \boldsymbol{I}_T + \beta_\alpha \sigma_\alpha \tilde{\boldsymbol{A}}$, $\boldsymbol{\varepsilon}^\pi \sim N(\mathbf{0}, \boldsymbol{S}_\pi)$ and

$\boldsymbol{\omega} \sim N(\mathbf{0}, (\boldsymbol{W}'_{\tau \omega} \boldsymbol{S}_\omega^{-1} \boldsymbol{W}_{\tau \omega})^{-1})$. The state equation controlling trend output growth can be restated as

$$\boldsymbol{W} \tilde{\boldsymbol{y}} = \boldsymbol{\varepsilon}^{\tilde{\boldsymbol{y}}^*}$$

where $\boldsymbol{\varepsilon}^{\tilde{\boldsymbol{y}}^*} \sim N(\mathbf{0}, \boldsymbol{I}_1)$. Using the same derivation as for the NAWRU, it can be shown that $\tilde{\boldsymbol{y}}^* \sim N(\hat{\boldsymbol{y}}^*, \boldsymbol{P}_{\tilde{\boldsymbol{y}}^*}^{-1})$, where

$$\begin{aligned} \boldsymbol{P}_{\tilde{\boldsymbol{y}}^*} &= \boldsymbol{W}' \boldsymbol{W} + (\beta_{\gamma^y} \gamma^y \beta_{y^*} \sigma_{y^*})^2 \boldsymbol{S}_\pi^{-1} + (\beta_{y^*} \sigma_{y^*})^2 \boldsymbol{W}'_{\tau \omega} \boldsymbol{S}_\omega^{-1} \boldsymbol{W}_{\tau \omega} \\ \hat{\boldsymbol{y}}^* &= \boldsymbol{P}_{\tilde{\boldsymbol{y}}^*}^{-1} (-\beta_{\gamma^y} \gamma^y \beta_{y^*} \sigma_{y^*} \boldsymbol{S}_\pi^{-1} (\boldsymbol{W}_{\tau \pi} (\boldsymbol{\pi} - \boldsymbol{\pi}^*) - \beta_{\gamma^y} \gamma^y (\boldsymbol{y} - \boldsymbol{\xi}_{\tilde{\boldsymbol{y}}^*}) - \beta_{\theta u} \theta^u (\boldsymbol{u} - \boldsymbol{u}^*)) \\ &\quad + \beta_{y^*} \sigma_{y^*} \boldsymbol{W}'_{\tau \omega} \boldsymbol{S}_\omega^{-1} \boldsymbol{W}_{\tau \omega} (\boldsymbol{y} - \boldsymbol{\xi}_{\tilde{\boldsymbol{y}}^*} - \beta_{\theta u} \theta^u (\boldsymbol{u} - \boldsymbol{u}^*))) \end{aligned}$$

As the $\boldsymbol{P}_{\tilde{\boldsymbol{y}}^*}$ precision matrix is banded, we can sample $\tilde{\boldsymbol{y}}^*$ efficiently using standard precision sampler. To derive the conditional distribution of $\Phi_{y^*} = (y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ the prior on Φ_{y^*} is given by $\Phi_{y^*} \sim N(\Phi_{y^*,0}, \boldsymbol{V}_{\Phi_{y^*}})$, where

$\Phi_{y^*,0} = (y_{0,0}^*, \alpha_{0,0}, 0, 0)$ and $V_{\Phi_{y^*}} = \text{diag}(V_{y^*}, V_{\alpha_0}, V_{\sigma_{y^*}}, V_{\sigma_\alpha})$, Letting $X_{y^*} = (I_1, \beta_{\alpha_0} I_T, \beta_{y^*} \tilde{y}, \beta_\alpha \tilde{A})$, our expression for trend output can be re-expressed as $y^* = X_{y^*} \Phi_{y^*}$. Our structural equations for the output and wage inflation gaps then become

$$\begin{aligned} W_{\tau\pi}(\pi - \pi^*) &= \beta_{\gamma^y} \gamma^y (y^* - X_{y^*} \Phi_{y^*}) + \beta_{\gamma^u} \gamma^u (u - u^*) + \varepsilon^\pi \\ y &= X_{y^*} \Phi_{y^*} + \beta_{\theta^u} \theta^u (u - u^*) + \omega \end{aligned}$$

Again, through regression analysis, it can be shown that Φ_{y^*} follows the conditional distribution $\Phi_{y^*} \sim N(\hat{\Phi}_{y^*}, P_{\Phi_{y^*}}^{-1})$, where

$$\begin{aligned} P_{\Phi_{y^*}} &= V_{\Phi_{y^*}}^{-1} + (\beta_{\gamma^y} \gamma^y)^2 X_{y^*}' S_\pi^{-1} X_{y^*} + X_{y^*}' W_{\tau\omega}' S_\omega^{-1} W_{\tau\omega} X_{y^*} \\ \hat{\Phi}_{y^*} &= P_{\Phi_{y^*}}^{-1} (V_{\Phi_{y^*}}^{-1} \Phi_{y^*,0} - \beta_{\gamma^y} \gamma^y X_{y^*}' S_\pi^{-1} (W_{\tau\pi}(\pi - \pi^*) - \beta_{\gamma^y} \gamma^y y - \beta_{\theta^u} \theta^u (u - u^*)) \\ &\quad + X_{y^*}' W_{\tau\omega}' S_\omega^{-1} W_{\tau\omega} (y - \beta_{\theta^u} \theta^u (u - u^*))) \end{aligned}$$

As outlined in Frühwirth-Schnatter and Wagner (2010), the signs of \tilde{y}^* and σ_{y^*} are not identifiable. Hence, to improve the efficiency of the sampler, the signs of both are randomly permuted. Letting U be a random variable that takes values in $\{-1, 1\}$ with equal probabilities, we take the current draws of \tilde{y}^* and σ_{y^*} and return $U\tilde{y}^*$ and $U\sigma_{y^*}$.

4. **Sample $\tilde{\alpha}$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distributions and randomly permute the signs of $\tilde{\alpha}$ and σ_α .** Given that $W\tilde{A} = \tilde{\alpha}$, we can sample \tilde{A} and transform the draw to get $\tilde{\alpha}$. We can rewrite our definition of trend output from Step 3 as

$$y^* = \xi_{\tilde{A}} + \beta_\alpha \sigma_\alpha \tilde{A}$$

where $\xi_{\tilde{A}} = y_0^* I_1 + \beta_{\alpha_0} \alpha_0 I_T + \beta_{y^*} \sigma_{y^*} \tilde{y}$. From this equation, we can re-write our structural wage inflation and output gap equations in terms of \tilde{A} , such that

$$\begin{aligned} W_{\tau\pi}(\pi - \pi^*) &= \beta_{\gamma^u} \gamma^u (u - u^*) + \beta_{\gamma^y} \gamma^y (y - \xi_{\tilde{A}} - \beta_\alpha \sigma_\alpha \tilde{A}) + \varepsilon^\pi \\ y &= \xi_{\tilde{A}} + \beta_\alpha \sigma_\alpha \tilde{A} + \beta_{\theta^u} \theta^u (u - u^*) + \omega \end{aligned}$$

where $\varepsilon^\pi \sim N(\mathbf{0}, S_\pi)$ and $\omega \sim N(\mathbf{0}, (W_{\tau\omega}' S_\omega^{-1} W_{\tau\omega})^{-1})$. The state equation for $\tilde{\alpha}$ is given by

$$W\tilde{\alpha} = \varepsilon^{\tilde{\alpha}} \quad \varepsilon^{\tilde{\alpha}} \sim N(\mathbf{0}, I_1)$$

which implies that $\tilde{A} = W^{-1}\tilde{\alpha} \sim N(\mathbf{0}, ((W^2)'W^2)^{-1})$. It can be shown that $\tilde{A} \sim N(\hat{\tilde{A}}, P_{\tilde{A}}^{-1})$, with

$$\begin{aligned} \mathbf{P}_{\tilde{\mathbf{A}}} &= (\mathbf{W}^2)' \mathbf{W}^2 + (\beta_{\gamma^y} \gamma^y \beta_{\alpha} \sigma_{\alpha})^2 \mathbf{S}_{\pi}^{-1} + (\beta_{\alpha} \sigma_{\alpha})^2 \mathbf{W}'_{\tau\omega} \mathbf{S}_{\omega}^{-1} \mathbf{W}_{\tau\omega} \\ \hat{\mathbf{A}} &= \mathbf{P}_{\tilde{\mathbf{A}}}^{-1} (-\beta_{\gamma^y} \gamma^y \beta_{\alpha} \sigma_{\alpha} \mathbf{S}_{\pi}^{-1} (\mathbf{W}_{\tau\pi} (\boldsymbol{\pi} - \boldsymbol{\pi}^*) - \beta_{\gamma^y} \gamma^y (\mathbf{y} - \boldsymbol{\xi}_{\tilde{\mathbf{A}}}) - \beta_{\gamma^u} \gamma^u (\mathbf{u} - \mathbf{u}^*)) \\ &\quad + \beta_{\alpha} \sigma_{\alpha} \mathbf{W}'_{\tau\omega} \mathbf{S}_{\omega}^{-1} \mathbf{W}_{\tau\omega} (\mathbf{y} - \boldsymbol{\xi}_{\tilde{\mathbf{A}}} - \beta_{\theta^u} \theta^u (\mathbf{u} - \mathbf{u}^*))) \end{aligned}$$

As $\mathbf{P}_{\tilde{\mathbf{A}}}$ is a band matrix, it can also be sampled efficiently using a precision sampling algorithm. Then, $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_{\alpha})$ can be sampled as before. Finally, we randomly permute the signs of $\tilde{\boldsymbol{\alpha}}$ and σ_{α} as per the same reasons given in Step 3.

5. **Sample $(\beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u})$ and $(\tau^{\pi}, \gamma^u, \gamma^y, \theta^u)$ jointly, by first sampling $(\beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u})$ marginally of $(\tau^{\pi}, \gamma^u, \gamma^y, \theta^u)$ then sampling $(\tau^{\pi}, \gamma^u, \gamma^y, \theta^u)$ from their full conditional distribution.**

Firstly, we must derive the full conditional distribution of $(\tau^{\pi}, \gamma^u, \gamma^y, \theta^u)$. Defining $\Phi_{\pi} = (\tau^{\pi}, \gamma^u, \gamma^y)'$, it is obvious that Φ_{π} and θ^u are conditionally independent given other parameters and states, and we can therefore sample them in turn. Rewriting the structural wage inflation gap equation as

$$\boldsymbol{\pi} - \boldsymbol{\pi}^* = \mathbf{X}_{\pi} \Phi_{\pi} + \boldsymbol{\varepsilon}^{\pi}$$

where $\boldsymbol{\varepsilon}^{\pi} \sim N(\mathbf{0}, \mathbf{S}_{\pi})$ and

$$\mathbf{X}_{\pi} = \begin{pmatrix} 0 & \beta_{\gamma^u}(u_1 - u_1^*) & \beta_{\gamma^y}(y_1 - y_1^*) \\ \pi_1 - \pi_1^* & \beta_{\gamma^u}(u_2 - u_2^*) & \beta_{\gamma^y}(y_2 - y_2^*) \\ \vdots & \vdots & \vdots \\ \pi_{T-1} - \pi_{T-1}^* & \beta_{\gamma^u}(u_T - u_T^*) & \beta_{\gamma^y}(y_T - y_T^*) \end{pmatrix}$$

The full conditional distribution of Φ_{π} can be expressed as $\Phi_{\pi} \sim N(\hat{\Phi}_{\pi}, \mathbf{P}_{\Phi_{\pi}}^{-1})$, where

$$\mathbf{P}_{\Phi_{\pi}} = V_{\Phi_{\pi}}^{-1} + \mathbf{X}'_{\pi} \mathbf{S}_{\pi}^{-1} \mathbf{X}_{\pi} \quad \hat{\Phi}_{\pi} = \mathbf{P}_{\Phi_{\pi}}^{-1} (V_{\Phi_{\pi}} \Phi_{\pi,0} + \mathbf{X}'_{\pi} \mathbf{S}_{\pi}^{-1} (\boldsymbol{\pi} - \boldsymbol{\pi}^*))$$

Similarly, the conditional distribution of θ^u can be expressed as $\theta^u \sim N(\hat{\theta}^u, P_{\theta^u}^{-1})$, with

$$\begin{aligned} P_{\theta^u} &= V_{\theta^u}^{-1} + \beta_{\theta^u}^2 (\mathbf{u} - \mathbf{u}^*)' \mathbf{W}'_{\tau\omega} \mathbf{S}_{\omega}^{-1} \mathbf{W}_{\tau\omega} (\mathbf{u} - \mathbf{u}^*) \\ \hat{\theta}^u &= P_{\theta^u}^{-1} (V_{\theta^u}^{-1} \theta_0^u + \beta_{\theta^u} (\mathbf{u} - \mathbf{u}^*)' \mathbf{W}'_{\tau\omega} \mathbf{S}_{\omega}^{-1} \mathbf{W}_{\tau\omega} (\mathbf{y} - \mathbf{y}^*)) \end{aligned}$$

To derive the conditional distribution of $(\beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u})'$ marginally of the coefficients Φ_{π} and θ^u , our prior is the product of the individual indicators Bernoulli distributions

$$p(\beta_{\gamma^u}, \beta_{\gamma^y}, \beta_{\theta^u}) = \prod_{i=\gamma^u, \gamma^y, \theta^u} p_i^{\beta_i} (1-p_i)^{1-\beta_i}$$

From here, it can be shown that the joint density function of $\boldsymbol{\pi}$ and \mathbf{y} marginally of Φ_π and θ^u is proportional to

$$p(\boldsymbol{\pi}, \mathbf{y} | \beta_{\gamma^y}, \beta_{\gamma^u}, \beta_{\theta^u}, \cdot) \propto |\mathbf{P}_{\Phi_\pi}|^{-\frac{1}{2}} e^{-\frac{1}{2} \hat{\Phi}'_{\pi} P_{\Phi_\pi} \hat{\Phi}_\pi} |\mathbf{P}_{\theta^u}|^{-\frac{1}{2}} e^{-\frac{1}{2} (\hat{\gamma}^y)^2 P_{\theta^u}}$$

The conditional distribution of the indicators can now be evaluated marginally of Φ_π and θ^u . Finally, $(\beta_{\gamma^u}, \beta_{\gamma^y}, \beta_{\theta^u})$ can be sampled via the inverse-transform method.

6. **Sample $(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha)$ and $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ jointly, by sampling $(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha)$ marginally of $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ then sampling $(y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)$ from their full conditional distribution.** The full conditional distribution of $\Phi_{y^*} = (y_0^*, \alpha_0, \sigma_{y^*}, \sigma_\alpha)'$ is $\Phi_{y^*} \sim N(\hat{\Phi}_{y^*}, \mathbf{P}_{\Phi_{y^*}}^{-1})$ which was derived in Step 4 of the process. Using a similar derivation to that used in Step 5, it can be shown that the distribution of $(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha)$ marginally of Φ_{y^*} is given by

$$p(\beta_{\alpha_0}, \beta_{y^*}, \beta_\alpha | \boldsymbol{\pi}, \mathbf{y}, \cdot) \propto \prod_{i=\alpha_0, y^*, \alpha} p_i^{\beta_i} (1-p_i)^{1-\beta_i} |\mathbf{P}_{\Phi_{y^*}}|^{-\frac{1}{2}} e^{-\frac{1}{2} \hat{\Phi}'_{y^*} P_{\Phi_{y^*}} \hat{\Phi}_{y^*}}$$

7. **Sample $(\beta_{z^\omega}, \beta_{z^\pi}), \mathbf{z}^\omega$ and \mathbf{z}^π jointly, by first sampling $(\beta_{z^\omega}, \beta_{z^\pi})$ marginally of \mathbf{z}^ω and \mathbf{z}^π , then drawing \mathbf{z}^ω and \mathbf{z}^π from their full conditional distributions.** From our modelling assumptions, the joint conditional density of $\boldsymbol{\varepsilon}^\pi = (\varepsilon_1^\pi, \varepsilon_2^\pi, \dots, \varepsilon_T^\pi)'$ is given by

$$p(\boldsymbol{\varepsilon}^\pi | \mathbf{h}^\pi, \beta_{z^\pi} = 0) = \prod_{t=1}^T (2\pi e^{h_t^\pi})^{-\frac{1}{2}} e^{-\frac{1}{2} e^{-h_t^\pi} (\varepsilon_t^\pi)^2}$$

$$p(\boldsymbol{\varepsilon}^\pi | \mathbf{h}^\pi, \beta_{z^\pi} = 1) = \left(\frac{\Gamma\left(\frac{1+v_\pi}{2}\right)}{\sqrt{v_\pi \pi} \Gamma\left(\frac{v_\pi}{2}\right)} \right)^T \prod_{t=1}^T e^{-\frac{1}{2} h_t^\pi} \left(1 + \frac{1}{v_\pi} e^{-h_t^\pi} (\varepsilon_t^\pi)^2 \right)^{-\frac{1+v_\pi}{2}}$$

The conditional density of β_{z^π} marginally of \mathbf{z}^π can be represented as

$$p(\beta_{z^\pi} | \boldsymbol{\varepsilon}^\pi, \mathbf{h}^\pi) \propto p_{z^\pi}^{\beta_{z^\pi}} (1-p_{z^\pi})^{1-\beta_{z^\pi}} p(\boldsymbol{\varepsilon}^\pi | \mathbf{h}^\pi, \beta_{z^\pi})$$

and a similar expression can be derived for β_{z^ω} . Given, $\beta_{z^\pi}, z_1^\pi, z_2^\pi, \dots, z_T^\pi$ are independent Inverse-Gamma random variables,

$$(z_t^\pi | v_\pi, \beta_{z^\pi} = 0) \sim IG\left(\frac{v_\pi}{2}, \frac{v_\pi}{2}\right)$$

$$(z_t^\pi | \nu_\pi, \beta_{z^\pi} = 1) \sim IG\left(\frac{1 + \nu_\pi}{2}, \frac{e^{-h_t^\pi} (\varepsilon_t^\pi)^2 + \nu_\pi}{2}\right)$$

The same sampling procedure and distributional form can then be applied to $z_1^\omega, z_2^\omega, \dots, z_T^\omega$.

8. **Sample τ^u and τ^ω from their full conditional distributions.** As τ^u and τ^ω are conditionally independent (given other parameters), they can be sampled jointly. Define X_{τ^u} and X_{τ^ω} such that

$$X_{\tau^u} = \begin{pmatrix} u_0 - u_0^* & u_{-1} - u_{-1}^* \\ u_1 - u_1^* & u_0 - u_0^* \\ \vdots & \vdots \\ u_{T-1} - u_{T-1}^* & u_{T-2} - u_{T-2}^* \end{pmatrix} \quad X_{\tau^\omega} = \begin{pmatrix} \omega_0 & \omega_{-1} \\ \omega_1 & \omega_0 \\ \vdots & \vdots \\ \omega_{T-1} & \omega_{T-1} \end{pmatrix}$$

Our structural equations for the unemployment gap and the cyclical component of output growth can then be represented as

$$\begin{aligned} \mathbf{u} - \mathbf{u}^* &= X_{\tau^u} \boldsymbol{\tau}^u + \boldsymbol{\varepsilon}^u \\ \boldsymbol{\omega} &= X_{\tau^\omega} \boldsymbol{\tau}^\omega + \boldsymbol{\varepsilon}^\omega \end{aligned}$$

where $\boldsymbol{\varepsilon}^u \sim N(\mathbf{0}, \sigma_u^2 \mathbf{I}_1)$ and $\boldsymbol{\varepsilon}^\omega \sim N(\mathbf{0}, \mathbf{S}_\omega)$. Thus, the conditional distributions of $\boldsymbol{\tau}^u$ and $\boldsymbol{\tau}^\omega$ can be represented as

$$\begin{aligned} \boldsymbol{\tau}^u &\sim N(\hat{\boldsymbol{\tau}}^u, \mathbf{P}_{\tau^u}^{-1}) | (\boldsymbol{\tau}^u \in \mathbf{R}) \\ \boldsymbol{\tau}^\omega &\sim N(\hat{\boldsymbol{\tau}}^\omega, \mathbf{P}_{\tau^\omega}^{-1}) | (\boldsymbol{\tau}^\omega \in \mathbf{R}) \end{aligned}$$

where \mathbf{R} is the stationary region of the parameter-space, and

$$\mathbf{P}_{\tau^u} = V_{\tau^u}^{-1} + \frac{1}{\sigma_u^2} X_{\tau^u}' X_{\tau^u} \quad \hat{\boldsymbol{\tau}}^u = \mathbf{P}_{\tau^u}^{-1} \left(V_{\tau^u}^{-1} \boldsymbol{\tau}_0^u + \frac{1}{\sigma_u^2} X_{\tau^u}' (\mathbf{u} - \mathbf{u}^*) \right)$$

$$\mathbf{P}_{\tau^\omega} = V_{\tau^\omega}^{-1} + X_{\tau^\omega}' \mathbf{S}_\omega^{-1} X_{\tau^\omega} \quad \hat{\boldsymbol{\tau}}^\omega = \mathbf{P}_{\tau^\omega}^{-1} \left(V_{\tau^\omega}^{-1} \boldsymbol{\tau}_0^\omega + X_{\tau^\omega}' (\mathbf{S}_\omega^{-1} \boldsymbol{\omega}) \right)$$

9. **Sample h^π and h^ω from their full conditional distributions.** Using the auxiliary mixture sampler of Kim *et al.* (1998), we can sample the log-volatilities h^π and h^ω . Under this sampler, if $\beta_{z^\omega} = 1$, define $v_t^\omega = \log\left(\frac{(\varepsilon_t^\omega)^2}{z_t^\omega}\right)$; otherwise, let $v_t^\omega = \log((\varepsilon_t^\omega)^2)$. We can then generate the series $v_1^\omega, v_2^\omega, \dots, v_T^\omega$ and pass it through the auxiliary mixture sampler as data.

10. **Sample the variances $\sigma_{u^*}^2, \sigma_{\pi^*}^2, \sigma_u^2, \sigma_{h\pi}^2$ and $\sigma_{h\omega}^2$ using an independence-chain Metropolis-Hastings step.** As each of the variances are conditionally independent, given the other parameters of the model, they can be sampled independently. As σ_u^2 possesses a standard inverse gamma distribution, it can be easily sampled as

$$(\sigma_u^2 | \mathbf{u}, \mathbf{u}^*, \tau^u) \sim IG\left(v_{u,0} + \frac{T}{2}, S_{u,0} + \frac{1}{2} \sum_{t=1}^T (\varepsilon_t^u)^2\right)$$

As the other variances are non-standard (due to their Gamma priors), they must be sampled via a Metropolis-Hastings step with an Inverse-Gamma proposal density. To sample $\sigma_{u^*}^2$, we obtain a candidate draw, s^2 , from

$$IG\left(\frac{T}{2} - 1, \sum_{t=2}^T \frac{(u_{t^*} - u_{t-1}^*)^2}{2}\right)$$

Given the current draw $\sigma_{u^*}^2$, we accept the candidate draw s^2 with probability

$$\min\left\{1, \exp\left(-\frac{1}{2V_{\sigma_{u^*}^2}}(s^2 - \sigma_{u^*}^2)\right)\right\}$$

Otherwise, we keep $\sigma_{u^*}^2$. The same process is then used to sample the other variances.

11. **Sample ν_ω and ν_π using an independence-chain Metropolis-Hastings step.** As the conditional distributions of ν_ω and ν_π are non-standard, they must also be sampled using the Metropolis-Hastings step. Both parameters can be sampled using the same approach. To sample ν_ω , we maximize the value $\log(p(\nu_\omega | z^\omega))$ using the Newton-Raphson technique to obtain the mode ($\hat{\nu}_\omega$) and the negative Hessian evaluated at the mode (\hat{P}_{ν_ω}). Then, we can implement an independence-chain Metropolis-Hastings step with proposal density $N(\hat{\nu}_\omega, P_{\nu_\omega}^{-1})$.