Taking DSGE models to the policy environment

Pedro Alvarez-Lois[a], Richard Harrison[a*], Laura Piscitelli[a], Alasdair Scott[a†]

[a] Bank of England, Threadneedle St., London EC2 R 8AH, United Kingdom.

First version: 11 November 2005 This version: 18 August 2006

Abstract

Considerable effort has been directed towards developing dynamic macroeconomic theory and to taking DSGE models to the data. Comparatively little appears to have been given to taking such models to the policy environment, especially to how projections from such models will be communicated to policy-makers. We argue that this issue deserves more attention.

To put our claim in context, we describe the 'core/non-core' approach used in the Bank of England's Quarterly Model, and compare it with two others: the 'measurement error' approach of Ireland and the 'shocks-in-parameters' approach utilised by Smets and Wouters and others. Using a mock forecast scenario for illustration, we argue that each of these approaches would present model users with difficulties in communicating with policy-makers. We conjecture that it is this problem, not theory or fitting the data, that currently hinders more widespread influence of DSGE-type models on policy-making.

Keywords: DSGE models

JEL classification: C51, C53

1. Introduction

The aim of this essay is to discuss some issues that arise in taking dynamic general equilibrium models to economic policy environments. For a model to be useful – in particular, for preparing economic projections that will guide policy decisions – it needs to do more than offer elegant theory and pass goodness of fit tests: attention should be paid to how judgments will be factored in and how its outputs can be communicated to policy-makers. Our case is that there is still a lot of progress yet to be made in this direction.

This argument might be obvious to some but surprising to others. To give the discussion a context, we introduce the 'core/non-core' (CNC) concept used in the Bank of England's Quarterly Model (BEQM)[1] and compare it to two other approaches that have recently been advocated for taking DSGE models to data: the 'measurement error' approach advocated by Peter Ireland and the 'shocks in parameters' approach, prominent examples of which are found in the work of Frank Smets and Raf Wouters. We then illustrate some issues that could arise when these approaches are applied to macroeconomic projections. We finish with our assessment of how well macroeconomic model designs are suited for 'real world' policy environments, and some conjectures on how to move forward.

*Corresponding author: richard.harrison@bankofengland.co.uk

†Our thanks to James Bell, Alex Brazier, Spencer Dale, Iain de Weymarn, Marco del Negro, Philip Evans, Jeff Fuhrer, Michael Grady, Michel Juillard, Frank Schorfheide for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect views held by the Bank of England.

[1] A full write-up can be found in Harrison et al. (2005).

2. Models and the policy environment

It seems uncontroversial to ask that policy models provide a solid, internally-consistent and transparent theoretical framework for analysis of macroeconomic shocks and their implications for policy.[2] Commensurately, the model needs to match the data, so that the Committee can have some confidence in the multipliers that are affecting the projections.

But the modeller-forecaster faces many other challenges: official data do not necessarily match model concepts very well; identities do not always hold in practice; data are noisy, and subject to frequent and possibly large revisions; the model will be misspecified, while policymakers might have different views on how the economy works; and the model needs to be tractable and reliable when producing projections in real time. This suggests that, even if the model performed well by accepted measures of 'fit', the model needs to be tractable, especially with regards to the application of judgements, and reliable, as many simulations will need to be performed in short period of time.

2.1. The core/non-core approach

To illustrate one response to these requirements, we describe the approach taken in BEQM. This was constructed in two distinct parts: a theoretical 'core' model, and 'non-core' equations that include additional variables and dynamics not modelled formally in the core. These two parts form the full model that is used for producing projections and facilitating the direct application of judgement.

Formally, the core model is a tightly-specified DSGE-type model. Using Y^* to denote endogenous variables in the core model, and ω to denote the forcing variables (shocks), we have

$$f\left(Y_{t-1}^{*}, Y_{t}^{*}, Y_{t+1}^{*}, \{\omega_{t+i}\}_{i=0}^{\infty}\right) = 0, \tag{1}$$

so that the solution for any given period is consistent with predetermined values of Y^*, expectations of Y^*, and the whole forecast of forcing variables, ω.

We then have to relate Y^* to actual data, denoted by Y. There will be some misspecification, reflected in an error term, ϵ_t:

$$Y_t = Y_t^* + \epsilon_t. \tag{2}$$

or $\Delta Y_t = \Delta Y_t^* + \Delta\epsilon_t$. If this error were simply a VAR(1), $\epsilon_t = \Phi\epsilon_{t-1} + \iota_t$, then we would have $\Delta Y_t = \Delta Y_t^* + (\Phi - I)\,\epsilon_{t-1} + \iota_t$,which, knowing (2), implies

$$\Delta Y_t = \Delta Y_t^* + (I - \Phi)\left(Y_{t-1}^* - Y_{t-1}\right) + \iota_t, \tag{3}$$

where Φ represents the VAR parameters.

Equation (3) is a VECM linking Y_t and Y_t^*. Once a history of shocks, ω, has been provided for (1), the parameters in Φ can be estimated, so that (3) provides

[2] This often implies that policy models need to be quite broad: in an inflation-targeting central bank, for example, a model dedicated solely to inflation forecasts is insufficient, as policymakers need to understand why inflation is expected to rise or fall.

a way of measuring how well Y_t^* tracks Y_t. This implies that the cointegrating vectors implicit from the theory in (1) should be a good match to those in the data, Y, so that ϵ is a stationary series.

One might have some ideas about what is missing in (1), and so one could add other variables to (3):

$$\Delta Y_t = \Delta Y_t^* + (1 - \Phi)\left(Y_{t-1}^* - Y_{t-1}\right) + \Psi Z_t + \iota_t. \tag{4}$$

Indeed, one can imagine testing whether variables in Z are significant, which might help to see in what ways the core model (1) could be developed to "reduce" the size of the error ι. The only restrictions are that variables in Z have no effect on the steady state, so that the projected path for a variable in Y will converge to the path generated by the core model in the long run, and that stock-flow, accounting and other consistency conditions are enforced.

2.2. An illustration with a small model

Consider Hansen's (1985) real business cycle model with indivisible labour. The representative agent maximises utility defined over consumption and hours worked:

$$\max \sum_{i=0}^{\infty} \beta^i \left(\ln C_{t+i}^* - \gamma H_{t+i}^*\right), \tag{5}$$

such that

$$K_{t+i+1}^* = (1-\delta) K_{t+i}^* + A_{t+i}^* \left(K_{t+i}^*\right)^{\alpha} \left(\mu_{t+i} H_{t+i}^*\right)^{1-\alpha} - C_{t+i}^*, \tag{6}$$

where C^* denotes consumption, H^* hours worked, K^* the capital stock (measured at the beginning of the discrete period) and A^* productivity, all in levels. Other characters include β, the household discount factor; γ, the leisure preference weight; δ, depreciation; α, the exponent on capital in production; and μ, the trend growth rate of labour productivity.

As is well known, the solution to (5) subject to (6) yields a intertemporal condition,

$$C_t^* = C_{t+1}^* \left[\beta \left(\alpha A_{t+1}^* \left(K_{t+1}^*\right)^{\alpha-1} \left(\mu_{t+1} H_{t+1}^*\right)^{1-\alpha} + 1 - \delta\right)\right]^{-1}, \tag{7}$$

and an intratemporal condition

$$C_t^* = \frac{1-\alpha}{\gamma} A_t^* \mu_t \left(\frac{K_t^*}{\mu_t H_t^*}\right)^{\alpha}. \tag{8}$$

Equations (6), (7) and (8) form a simultaneous block, which we could solve in nonlinear levels directly (assuming perfect foresight, as here), and from which we can recursively calculate a number of variables of interest, such as investment,

$$I_t^* = K_{t+1}^* - (1-\delta) K_t^*; \tag{9}$$

output,

$$Y_t^* = A_t^* (K_t^*)^{\alpha} (\mu_t H_t^*)^{1-\alpha}; \tag{10}$$

and a competitive market real interest rate,

$$R_t^* = 1 + \alpha A_t^* \left(\frac{K_t^*}{\mu_t H_t^*}\right)^{\alpha-1} - \delta. \tag{11}$$

Equations (6) to (11) would be the equivalent of the structural core model (1), yielding paths for K^*, C^*, H^*, I^*, Y^*, and R^*.

Assume now that we want to make *ad hoc* adjustments to the consumption and investment paths. Following the transformations in section 2.1., we could specify error correcting equations following (4) for actual consumption and investment:

$$\Delta C_t = \Delta C_t^* + \left(1 - \Phi^C\right)\left[C_{t-1}^* - C_{t-1}\right] + \Psi^C Z_t + \iota_t^C \tag{12}$$

$$\Delta I_t = \Delta I_t^* + \left(1 - \Phi^I\right)\left[I_{t-1}^* - I_{t-1}\right] + \Psi^I Z_t + \iota_t^I. \tag{13}$$

For example, we might include some measure of 'confidence' effects, say from a survey, in Z in order to improve the dynamic fit of the model in consumption and investment spaces. These would have to be normalised to zero or specified as changes, so as to ensure that C converged to C^* and I converged to I^* Similarly, we could specify an adjustment for hours worked:

$$\Delta H_t = \Delta H_t^* + \left(1 - \Phi^H\right)\left[H_{t-1}^* - H_{t-1}\right] + \Psi^H Z_t + \iota_t^H, \tag{14}$$

and under the same assumptions H would return to H^*.

At this stage, the non-core part of the system is incomplete. First, we have to predict Z to form projections for C, I, and H. Second, we need to pay attention to accounting consistency. For stock-flow consistency, we would want to duplicate the perpetual inventory consistency condition, (9):

$$K_{t+1} = (1 - \delta) K_t + I_t, \tag{15}$$

which will be sufficient to ensure that K eventually converged to K^*.[3] This value would also allow us consistently to define corresponding values for a production measure of output, Y^P, given the series for (exogenous) productivity A^* and *non-core* capital K and *non-core* hours H:

$$Y_t^P = A_t^* (K_t)^{\alpha} (\mu_t H_t)^{1-\alpha}. \tag{16}$$

We could also apply the same process to define the non-core real interest rate:

$$R_t = 1 + \alpha A_t^* \left(\frac{K_t}{\mu_t H_t}\right)^{\alpha-1} - \delta. \tag{17}$$

These also have long-run limits $\lim_{t\to\infty} R_t = R_t^*$ and $\lim_{t\to\infty} Y_t^P = Y_t^*$. Total expenditure is defined as $Y_t = C_t + I_t$,which also ensures that $\lim_{t\to\infty} Y_t = Y_t^P = Y_t^*$.[4]

[3] It is sufficient so long as depreciation, δ, is greater than zero and less than one.

[4] However, $Y_t \neq Y_t^P$ in the short run, which is a consequence of the *ad hoc* additions that override the saddlepath dynamics of the structural core model. The hybrid approach of Ireland (2004) that we consider below also has the same implication.

3. Some other methods of taking theory to data

We consider the generalised 'measurement error' (GME) approach advocated by Peter Ireland and the 'shocks in parameters' (SIP) approach associated with the work of Frank Smets and Raf Wouters and others. Our intention here is not to find out which approach is 'best' (such as by fit). Instead, we want to find out how these methods would be different in an operational sense. In particular, we want to know: what would each method imply for the way that forecast issues are handled, and would the same economic 'stories' come through to policy-makers?

As is well known, a typical DSGE model has a linear rational expectations solution that can be written, using lower case to denote log deviations from the steady state, as

$$f_t = B_{fs} s_t \tag{18}$$

$$s_t = B_{ss} s_{t-1} + B_{s\varepsilon} \varepsilon_t \tag{19}$$

where f denotes forward-looking endogenous (or 'jump') variables; s denotes the state variables, which includes both endogenous state variables and exogenous (forcing) variables; ε is the vector of shocks associated with the forcing variables; and B is a matrix of coefficients, such that B_{xy} denotes the elasticity of x with respect to y. If we stack (18) and (19), we have

$$\begin{bmatrix} f_t \\ s_t \end{bmatrix} = \begin{bmatrix} 0 & B_{fs} B_{ss} \\ 0 & B_{ss} \end{bmatrix} \begin{bmatrix} f_{t-1} \\ s_{t-1} \end{bmatrix} + \begin{bmatrix} B_{fs} B_{s\varepsilon} \\ B_{s\varepsilon} \end{bmatrix} \varepsilon_t. \tag{20}$$

If $y_t = \begin{bmatrix} f'_t, & s'_t \end{bmatrix}'$, we have the VAR representation of the DSGE model:

$$y_t = B_{yy} y_{t-1} + B_{y\varepsilon} \varepsilon_t. \tag{21}$$

The number of variables in y – the variables the modeller is trying to match – is typically larger than the number of stochastic processes in ε. In this case, the the model would be stochastically singular if the researcher attempted to take it directly to the data y using a maximum likelihood estimator, because a subset of variables f is deterministically related to the state variables s.

3.1. 'Generalised Measurement Error'

Ireland (2004) uses additive innovations in the observation equation (18) to reconcile the predictions of a DSGE model with the data. These innovations are given a VAR structure and are therefore silent about what drives the discrepancy between theory and data. Because the errors are not restricted to be orthogonal – structure is deliberately introduced to allow for cross- and auto-correlation – the term 'measurement error' is, strictly speaking, a misnomer. We offer the label Generalised Measurement Error (GME) instead.

To represent this approach, we start with a state equation like (19) and augment (18) with a sufficient number of 'measurement errors' to get around the stochastic singularity problem:

$$f_t = B_{fs} s_t + B_{fu} u_t. \tag{22}$$

The process u is assumed to be a VAR(1):

$$u_t = B_{uu}u_{t-1} + B_{u\zeta}\zeta_t. \tag{23}$$

This allows for cross correlation between the elements of u; however, the measurement errors ζ and the shocks to the forcing processes ε are assumed to be orthogonal.

Stacking these equations with the state equation (19), with the assumption that $B_{fu} = I$ and $B_{u\zeta} = I$, we have

$$\begin{bmatrix} f_t \\ s_t \\ u_t \end{bmatrix} = \begin{bmatrix} 0 & B_{fs}B_{ss} & B_{uu} \\ 0 & B_{ss} & 0 \\ 0 & 0 & B_{uu} \end{bmatrix} \begin{bmatrix} f_{t-1} \\ s_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} B_{fs}B_{s\varepsilon} & I \\ B_{s\varepsilon} & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \zeta_t \end{bmatrix}. \tag{24}$$

With the same definition of y_t as before, the VAR representation of this model is:

$$\begin{bmatrix} y_t \\ u_t \end{bmatrix} = \begin{bmatrix} B_{yy} & B_{yu} \\ 0 & B_{uu} \end{bmatrix} \begin{bmatrix} y_{t-1} \\ u_{t-1} \end{bmatrix} + \begin{bmatrix} B_{y\varepsilon} & B_{y\zeta} \\ 0 & I \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \zeta_t \end{bmatrix}. \tag{25}$$

From this it can be seen that the marginal difference compared to (21) is the addition of the measurement errors, u.

3.2. 'Shocks in parameters'

We take the work of Smets and Wouters (2003) as representative of this approach, in which a subset of parameters is treated as potentially time-varying. This approach is motivated, in part, by the singularity problem. (21) is augmented by making some parameters subject to shocks. In practice, this means adding new forcing processes x to the model so that some parameters become state variables. We assume that these new forcing processes do not depend on the existing state variables in (19), so we can write

$$x_t = B_{xx}x_{t-1} + B_{x\eta}\eta_t. \tag{26}$$

The observation equation is now

$$f_t = B_{fs}s_t + B_{fx}x_t, \tag{27}$$

reflecting the impact of the extra forcing processes, and the state equation is correspondingly expanded to

$$s_t = B_{ss}s_{t-1} + B_{sx}x_{t-1} + B_{s\varepsilon}\varepsilon_t + B_{s\eta}\eta_t. \tag{28}$$

Stacking these elements, we have

$$\begin{aligned} \begin{bmatrix} f_t \\ s_t \\ x_t \end{bmatrix} &= \begin{bmatrix} 0 & B_{fs}B_{ss} & B_{fs}B_{sx} + B_{fx}B_{xx} \\ 0 & B_{ss} & B_{sx} \\ 0 & 0 & B_{xx} \end{bmatrix} \begin{bmatrix} f_{t-1} \\ s_{t-1} \\ x_{t-1} \end{bmatrix} \\ &\quad + \begin{bmatrix} B_{fs}B_{s\varepsilon} & B_{fs}B_{s\eta} + B_{fx}B_{x\eta} \\ B_{s\varepsilon} & B_{s\eta} \\ 0 & B_{x\eta} \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix}, \end{aligned} \tag{29}$$

so that if we define $y_t = \begin{bmatrix} f_t', & s_t' \end{bmatrix}'$, as before, we have the VAR representation

$$\begin{bmatrix} y_t \\ x_t \end{bmatrix} = \begin{bmatrix} B_{yy} & B_{yx} \\ 0 & B_{xx} \end{bmatrix} \begin{bmatrix} y_{t-1} \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} B_{y\varepsilon} & B_{y\eta} \\ 0 & B_{x\eta} \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix}. \tag{30}$$

This shows that the marginal difference compared to (21) is the addition of new state variables and the increase in the number of shocks. But the coefficient matrices B_{yy} and $B_{y\varepsilon}$ are the same as those in (21).

3.3. Core/non-core revisited

Equations (18) and (19) constitute the 'core' model, so we define $y_t^* = \begin{bmatrix} f_t', & s_t' \end{bmatrix}'$. This implies that we can write the core model as

$$y_t^* = B_{y^*y^*} y_{t-1}^* + B_{y^*\varepsilon}\varepsilon_t. \tag{31}$$

We also have a set of auxiliary variables z that we first saw in (4). These are assumed to follow a VAR on past values of auxiliary variables z and past values of observed variables y:

$$z_t = B_{zz} z_{t-1} + B_{zy} y_{t-1} + B_{z\upsilon}\upsilon_t. \tag{32}$$

These variables are combined in the non-core side of the model, which we can represent as[5]

$$y_t = B_{yy} y_{t-1} + B_{yy_0^*} \cdot y_t^* + B_{yy_1^*} \cdot y_{t-1}^* + B_{yz} z_t + B_{y\xi}\xi_t. \tag{33}$$

Given the processes in (31), (32) and (33), the autoregressive representation of the core/non-core approach can then be written as follows:

$$\begin{bmatrix} y_t^* \\ y_t \\ z_t \end{bmatrix} = \begin{bmatrix} B_{y^*y^*} & 0 & 0 \\ \left(B_{yy_0^*} B_{y^*y^*} + B_{yy_1^*}\right) & \left(B_{yy} + B_{yz} B_{zy}\right) & B_{yz} B_{zz} \\ 0 & B_{zy} & B_{zz} \end{bmatrix} \begin{bmatrix} y_{t-1}^* \\ y_{t-1} \\ z_{t-1} \end{bmatrix} + \begin{bmatrix} B_{y^*\varepsilon} & 0 & 0 \\ B_{yy_0^*} B_{y^*\varepsilon} & B_{y\xi} & B_{yz} B_{z\upsilon} \\ 0 & 0 & B_{z\upsilon} \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \xi_t \\ \upsilon_t \end{bmatrix}. \tag{34}$$

Remembering that y^* represents the same DSGE model as in (18) and (19), we see that the marginal difference compared to (21) is twofold: the auxiliary variables z designed to 'soak up' the gap between y and y^*, and their combination in the equation for y. One interpretation of the core/non-core approach is therefore that the forecast of each variable using (33) is a weighted combination of the predictions from a theoretically-tight model and the atheoretic elements of extra persistence and auxiliary variables. While the representation (34) looks complicated, in practice it is likely to be quite sparse and restricted.

These approaches all share a basic DSGE model as a 'core'. In each case, something is added to allow the model to match the data, y, better. But the

[5] Here $B_{yy_0^*}$ denotes the elasticity of y with respect to contemporaneous y^*, whereas $B_{yy_1^*}$ denotes the elasticity of y with respect to the first lag of y^*.

nature of what is added differs. In the GME approach, it is the measurement error, u. In SIP, it is the extra state variables, x. CNC can have some of the flavour of all of these methods, but differs from the others in the role of auxiliary variables, z.[6]

4. Comparing the approaches using mock forecasts

We conduct mock forecasts on UK data. The object of the exercise is not to explore goodness of fit, nor to show that one method produces more accurate forecasts than another. Instead, we want to see whether the model's structure – that is, the choice of GME, SIP or CNC approaches – would make a difference to the way judgements are handled and economic stories are communicated to policymakers. For this exercise, the basic theoretical structure we use in this section is the Hansen model from (2.2.).

4.1. Model-consistent data

We use a sample of UK data running from 1978:1 to 2005:2, yielding 110 observations. Consumption and investment data are real, quarterly, and seasonally-adjusted, from the national accounts. To be consistent with the theory, we define an artificial output series as $Y = C + I$, where C is consumption data and I is investment data. The hours data, H, are from the Office of National Statistics (ONS) whole economy hours worked series. Source codes are listed in Appendix A.

We convert data into per capita terms by dividing by a measure of working age population from the ONS, which we denote N. We then take natural logarithms, defining $\tilde{y} = \log(Y/N)$, $\tilde{c} = \log(C/N)$, $\tilde{\imath} = \log(I/N)$, and $\tilde{h} = \log(H/N)$. (These data are plotted in levels in Appendix A.) Expenditure data are then detrended using a constant growth rate (μ, estimated as the trend growth in $\tilde{y}$ by OLS). Deviations from the steady-state level are then calculated. This defines the data series $y_t = \tilde{y}_t - t\log\mu - \bar{y}$, $c_t = \tilde{c}_t - t\log\mu - \bar{c}$, $i_t = \tilde{\imath}_t - t\log\mu - \bar{\imath}$, and $h_t = \tilde{h}_t - \bar{h}$, where we denote the (deterministic) steady-state value of a variable x by $\bar{x}$.

4.2. Parameterising the models

4.2.1. The generalised measurement error version

In the GME approach, we define the observation vector

$$f_t = \begin{bmatrix} y_t & c_t & h_t \end{bmatrix}', \tag{35}$$

[6] There is also the approach of del Negro and Schorfheide (eg, 2004), which begins with a SIP or GME version of the theoretical model. (In practice, there is room for some overlap: a modeller could use both measurement equation errors and shocks in parameters in order to fit the data.) Their approach can be thought of as a way of finding the optimal weights by which to combine the predictions of a DSGE model and an atheoretic model, using Bayesian shrinkage.

where y is output, c is consumption, and h is hours. The state vector is

$$s_t = \begin{bmatrix} k_t & a_t \end{bmatrix}', \tag{36}$$

where k is (unobserved) capital, and a is productivity. The matrices B_{fs}, B_{ss} and $B_{s\varepsilon}$ in (24) will depend on β, γ, α, and δ.

We add three measurement errors to the observation equation[7], so that the error structure $u_t = B_{uu}u_{t-1} + \zeta_t$ is a three-by-three VAR(1)

$$\begin{aligned} \begin{bmatrix} u_t^i \\ u_t^c \\ u_t^h \end{bmatrix} &= \begin{bmatrix} \rho_{ii} & \rho_{ic} & \rho_{ih} \\ \rho_{ci} & \rho_{cc} & \rho_{ch} \\ \rho_{hi} & \rho_{hc} & \rho_{hh} \end{bmatrix} \begin{bmatrix} u_{t-1}^i \\ u_{t-1}^c \\ u_{t-1}^h \end{bmatrix} + \begin{bmatrix} \zeta_t^i \\ \zeta_t^c \\ \zeta_t^h \end{bmatrix}, \\ \operatorname{cov}(\zeta_t) &= \begin{bmatrix} \sigma_i^2 & \lambda_{ci}\sigma_i\sigma_i & \lambda_{hi}\sigma_h\sigma_i \\ \lambda_{ci}\sigma_i\sigma_i & \sigma_c^2 & \lambda_{hc}\sigma_h\sigma_c \\ \lambda_{hi}\sigma_h\sigma_i & \lambda_{hi}\sigma_h\sigma_i & \sigma_h^2 \end{bmatrix}. \end{aligned} \tag{37}$$

We also have one structural shock, to productivity. This structure implies that we need values for 23 parameters: β, γ, α, δ and $\bar{a}$; the persistence of the productivity shock, ρ; the variance of the productivity shock, σ, and ρ_{ii}, ρ_{ic}, ρ_{ih}, ρ_{ci}, ρ_{cc}, ρ_{ch}, ρ_{hi}, ρ_{hc}, ρ_{hh}, σ_i^2, σ_c^2, σ_h^2, λ_{ic}, λ_{ih}, and λ_{ch}.

4.2.2. The shocks-in-parameters version

To implement the SIP approach, the maximisation problem described in section (2.2.) is modified to include a consumption preference shock and a labour supply shock,

$$\max \sum_{i=0}^{\infty} \beta^i \left(\theta_{t+i} \ln C_{t+i} - \gamma_{t+i} H_{t+i}\right), \tag{38}$$

where we have introduced a term θ_t and γ_t is now allowed to vary over time.[8] As with the technology shock, these follow AR(1) processes:

$$\ln \theta_t = \rho_\theta \ln \theta_{t-1} + \varepsilon_t^\theta \tag{39}$$

$$\ln \gamma_t = \left(1 - \rho_\gamma\right) \bar{\gamma} + \rho_\gamma \ln \gamma_{t-1} + \varepsilon_t^\gamma. \tag{40}$$

The budget constraint (6) is unchanged.

[7] In this and Ireland's (2004) implementation, there are more measurement errors than necessary, so the model is over-identified.

[8] Of course, we could use any structural shocks, so long as they are not perfectly colinear. These shocks are only illustrative.

We use the same measurement vector as in the GME version, (35), but the state vector is larger than (36), so that the state equation follows the process

$$\begin{bmatrix} k_t \\ a_t \\ \theta_t \\ \gamma_t \end{bmatrix} = \begin{bmatrix} b_{kk} & b_{ka} & b_{k\theta} & b_{k\gamma} \\ 0 & \rho & 0 & 0 \\ 0 & 0 & \rho_\theta & 0 \\ 0 & 0 & 0 & \rho_\gamma \end{bmatrix} \begin{bmatrix} k_{t-1} \\ a_{t-1} \\ \theta_{t-1} \\ \gamma_{t-1} \end{bmatrix} + \begin{bmatrix} b_{k\varepsilon} & b_{k\varepsilon^\theta} & b_{k\varepsilon^\gamma} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \varepsilon_t^\theta \\ \varepsilon_t^\gamma \end{bmatrix}, \qquad \mathrm{cov}\,(\varepsilon) = \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \sigma_\theta^2 & 0 \\ 0 & 0 & \sigma_\gamma^2 \end{bmatrix}. \tag{41}$$

The SIP structure means we need values for 11 parameters: β, γ, α, δ and $\bar{a}$; and ρ, ρ_θ, ρ_γ, σ, σ_θ, and σ_γ in (41).

4.2.3. The core/non-core version

To implement the CNC version, we need values for the core model's parameters, β, γ, α, δ and $\bar{a}$, and then parameters for the non-core equations. We have used a more general specification than is shown in section 2.1., specifically:

$$c_t = \phi^c c_t^* + (1 - \phi^c)\, c_{t-1} + \psi^c z_t + c_t^{res} \tag{42}$$

$$i_t = \phi^i i_t^* + \left(1 - \phi^i\right) i_{t-1} + i_t^{res}, \tag{43}$$

$$y_t = c_t + i_t \tag{44}$$

and

$$h_t = \phi^h h_t^* + \left(1 - \phi^h\right) h_{t-1} + h_t^{res}, \tag{45}$$

where res denotes a residual for the given non-core equation; in each case, ϕ and ψ are scalars. For simplicity, we only include one example of a z variable: real house prices in the consumption equation (to capture "credit" effects), which necessitates an equation to forecast house prices (and expands the observation vector to include house price data). Together, this structure implies that we need to estimate 13 parameters.

4.2.4. Estimating the parameters

The different designs can be regarded as different ways by which to adapt the theory to the data, by allowing for more richness in the model's dynamics. To focus on these differences, we impose the same values for the parameters for β, δ, γ, α and $\bar{a}$. We take a "calibrationist" approach, by relating steady-state expressions for endogenous variables to mean values of the data.

The remaining parameters were estimated using Bayesian maximum likelihood via the Kalman filter and the state space representations discussed in sections (3.1. to 3.3.). Prior values for the persistence, ρ, and the standard deviation, σ, of the technology shock come from OLS estimates of a constructed Solow residual (see

Appendix B). As we wanted to be agnostic about the marginal effects of the extra dynamic structure, our priors were set such that the modal parameters implied that the extra dynamics in each approach had zero persistence and very small variances. (Appendix B shows the results of estimating the models.)

4.3. A comparison of forecasts

4.3.1. Variance decompositions

A first step to understanding the effects of these different approaches is to examine variance decompositions. For example, the technology shock, ε, that is common to all three approaches explains the data to differing degrees. It accounts for nearly 100 per cent of the variation of consumption in the SIP version, but only about 80 per cent in the CNC case and closer to 10 per cent in the GME version (Table 1).

Table 1: variance contributions of technology shock (per cent)

	GME	SIP	CNC
Output	28.65	97.53	82.12
Consumption	10.40	98.63	83.40
Investment	72.15	76.73	44.63
Hours	7.75	31.67	89.10

Clearly, these approaches matter for identification of shocks. It would be tempting, for example, to attribute business cycles primarily to technology shocks, based on the results of the SIP model; there is little role for the preference and labour supply shocks in expenditures, and even hours are substantially accounted for by movements in TFP. But in the GME version, the consumption measurement error dominates output, consumption and hours. There is a significant role for the investment residual in the CNC version. (Variance decompositions for each version are summarised in Appendix B3.)

4.3.2. Contributions to historical fit

The next question we ask is how the different versions explain history. This is an important part of communicating with policymakers, as they will often want to examine what can be said about past episodes, with an eye to how policy should react to similar episodes in the future. We show the contributions to fitted values from the shocks driving the each version of the model – that is, we examine the MA representation of the fitted values. In addition to the TFP shock and the extra components from the different approaches, there is also a contribution from initial conditions.[9] Some examples are seen in Figure 1, which plots contributions to output according to the three different approaches.

[9] To produce the decomposition, we use the VAR representation

$$Y_t = B_{YY} Y_{t-1} + B_{Y\varepsilon} \varepsilon_t$$

The contributions are often offsetting. For example, output in recent years is accounted for by positive productivity shocks and negative consumption measurement errors in the GME case; by positive labour supply and productivity shocks and negative preference shocks in the SIP case,[10] and negative productivity shocks but positive investment residuals in the CNC case.

In all three cases, TFP shocks are procyclical. But the scale of the contribution differs quite markedly. In the CNC and SIP versions, for example, around three quarters of the -0.12 fall in detrended output from its peak in 1989Q2 to its trough in 1994Q2 is explained by productivity shocks. But in the GME version, the proportion is closer to one half (Table 2).

Table 2: net contributions to fall in detrended output from 1989Q2 to 1994Q2

GME	y_0	ε	u^i	u^c	u^h
	0.0029	-0.0578	-0.0305	-0.0468	0.011
SIP	y_0	ε	ε^γ	ε^θ	
	0.0007	-0.0870	0.0125	-0.0473	
CNC	y_0	ε	i^{res}	c^{res}	h^{res}
	-0.0007	-0.1084	0.0003	0.0013	-0.0135

Clearly, the different approaches are imposing different identification structures on the data, with large differences in implicit stories. This aspect of identification therefore appears to be just as important as conventional issues concerning inference.[11]

4.3.3. Contributions to the projections

Figure 1 also shows projections from the three versions of the Hansen model. All the projections show offsetting effects, such as the gradual decline of the effects of TFP and consumption measurement errors in the GME case. The projections are qualitatively different: the effects of TFP shocks are negative in SIP and CNC but positive in GME, for example.

In levels, the projections appear smooth and persistent. But in growth space, all versions show strong trend reversion, with little or no extrapolation of recent below trend growth (Figure 2). Interestingly, house price effects do seem to be important in shaping the profile for output growth in the CNC projection.

to calculate

$$Y_t = \sum_{i=0}^{t-1} (B_{YY})^i B_{Y\varepsilon}\varepsilon_{t-i} + (B_{YY})^t Y_0.$$

But the true VAR representation of the model is infinite; because we have only a finite sample, initial conditions matter. These can take some time to decay away – if an element of Y_0 is large and/or the loading in $(B_{YY})^t$ is close to one, then the effects of initial conditions will persist.

[10] The extremely persistent effects of initial conditions are because the implied capital stock is estimated to be some distance from its steady state level.

[11] For the latter, see Canova and Sala (2005).

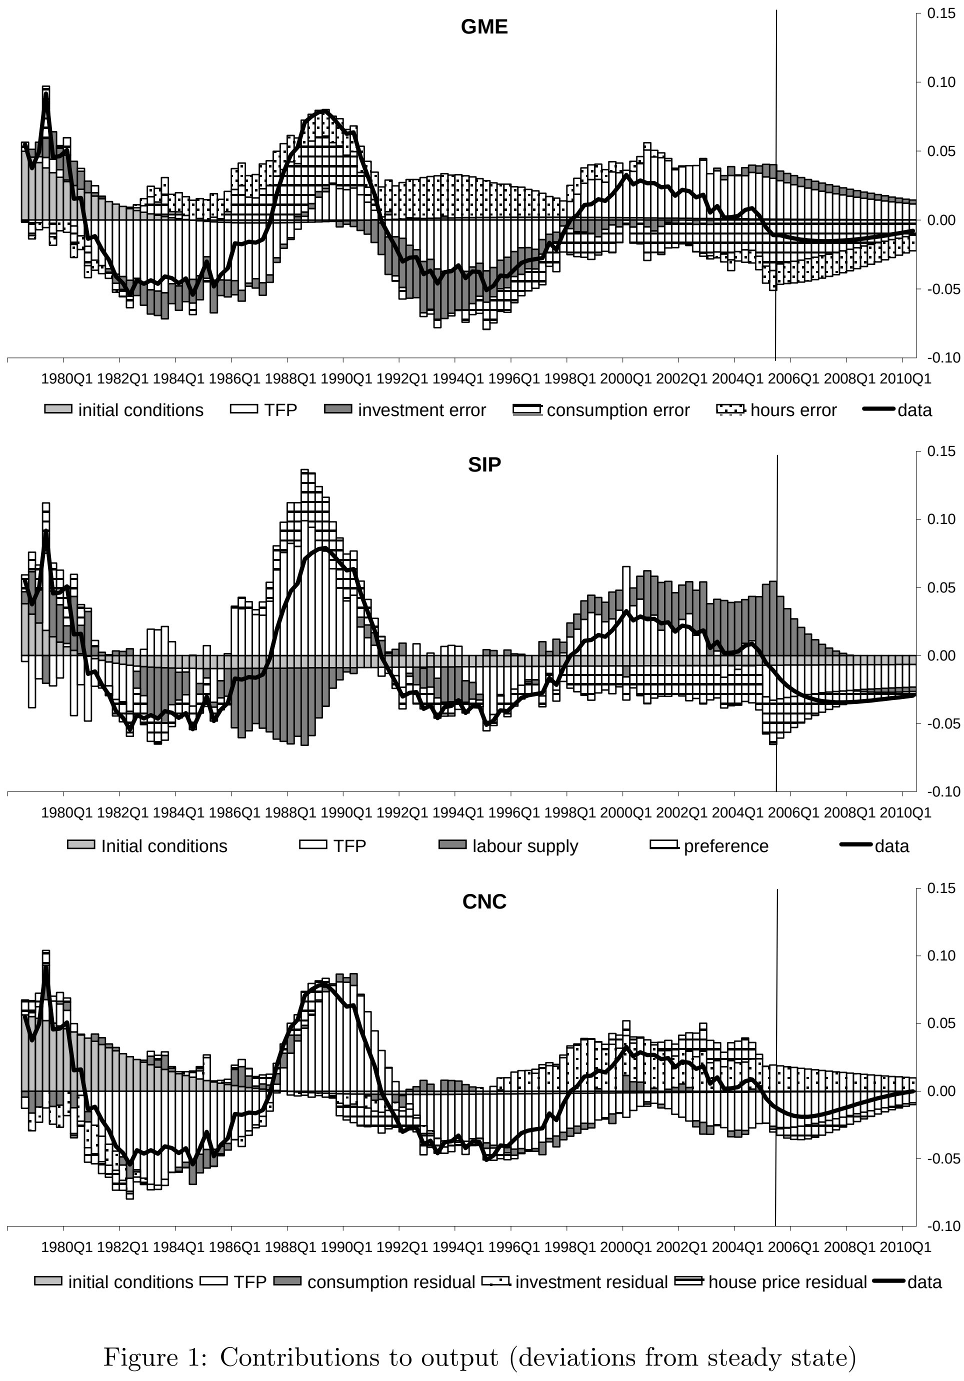

Figure 1: Contributions to output (deviations from steady state)

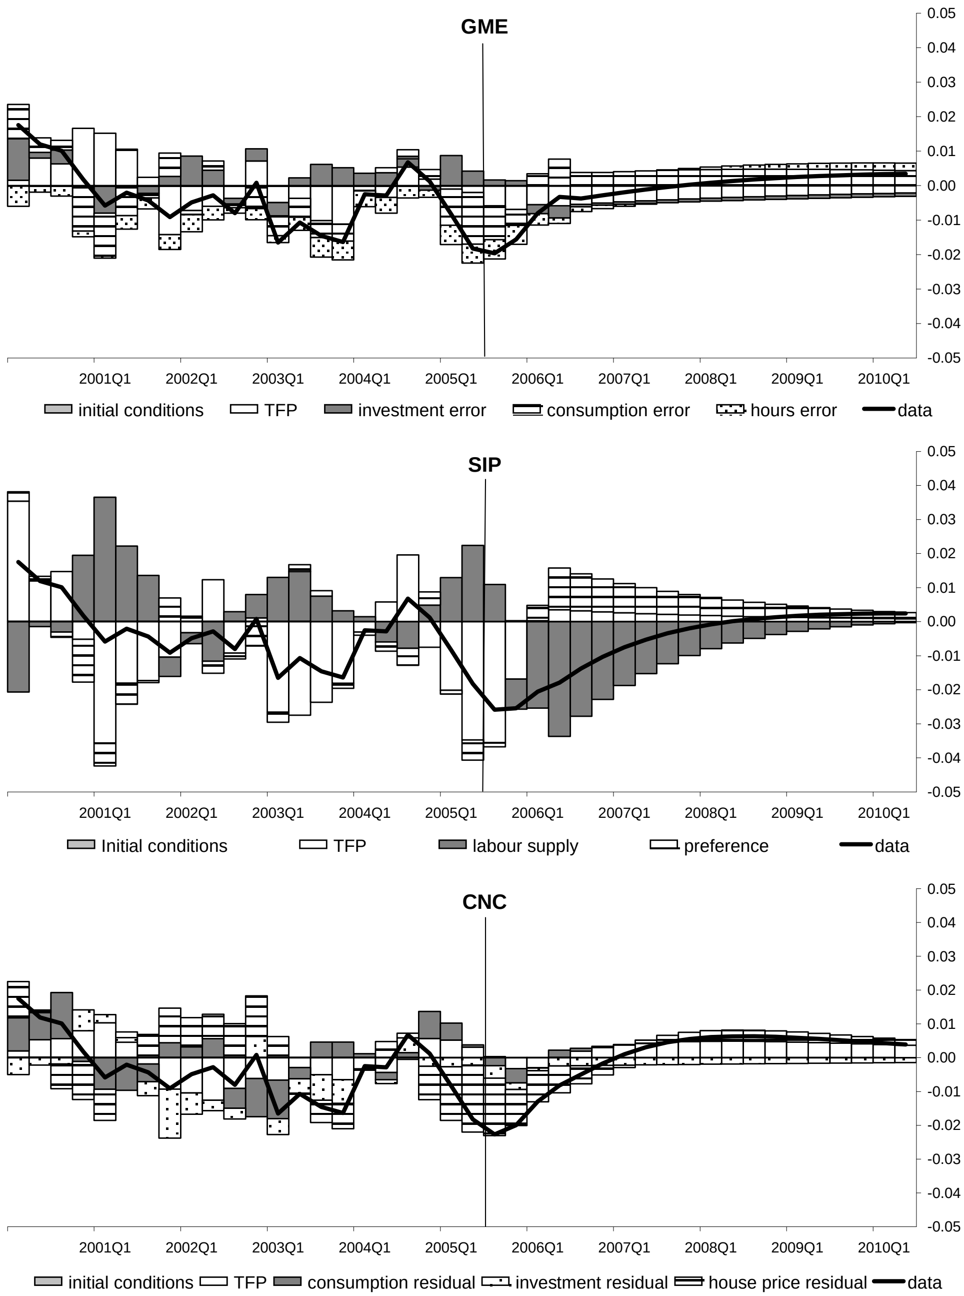

Figure 2: Contributions to output growth projections (deviations from steady state)

5. Model design and economic policy: where are we?

We are now in the position to ask the most important questions of this essay. First, are the differences between the approaches merely mechanical, or would they make a meaningful difference in the way economic stories are communicated to policymakers? Second, if no single method is unambiguously superior than any other, how should we be looking to move forward?

5.1. Do these approaches affect communication with policymakers?

Policymakers want clear and coherent economic stories about the condition of the economy, what pressures have driven it there, how these pressures could play out, and what such scenarios imply for their policy targets. In this context, a story means a mapping from what they can observe to what they care about as policymakers. For example, a policymaker is typically not interested just in an observation that margins have risen, but in what *caused* the rise. We see three aspects to this: how these models would explain recent outturns in the data, how they would incorporate judgements, and how they would tell stories about agents' expectations.

5.1.1. Explaining the recent data

We see a clear difference between SIP and GME approaches on the one hand and the CNC approach on the other. The first two would account for recent movements in the data by the addition of flexible but unobserved time series processes. Our experience is that policymakers are uncomfortable with reliance on unobserved factors, if only because they have to explain their decisions to policy watchers in terms of developments that all people can observe.

This influenced the CNC approach and the use of the z variables. In that case, the staff member would be able to tell a story about how much weight to put on a purist, textbook explanation, and how much to put on short-run factors that, while ad hoc, have exhibited plausible correlations.[12] On the other hand, we can easily run into problems if we are tempted to put structural stories on the auxiliary variables, or expect to find cross-correlations that might not be there.

5.1.2. Incorporating judgement

In practice, it is unlikely that staff or policymakers would be happy with judgement-free forecasts. Our case is that none of the approaches would handle judgements in a completely satisfactory way.

The simple reason is that different judgements are motivated by different reasons. In some cases, something is known to be missing from the underlying economic theory contained in the model, and so judgements need to be applied in a coherent and systematic way so as to compensate for this. For example, there might be useful

[12] We note that the DSGE-VAR approach has some of the same flavour. See, for example, the exercise in del Negro and Schorfheide (2003, pp44-47).

information from other, more specialised models; or the policymakers might simply have different views on how the economy works. On the other hand, there are also occasions when judgements are required to deal with idiosyncratic problems, such as data discrepancies, or the forecaster might simply want to "tune" the projection to incorporate the advice from sectoral experts or achieve a consensus projection.

For example, in a strict SIP application, the modeller could change the demand for consumption goods using a preference shock. The shock would then affect all other variables that depended on the marginal utility of consumption. This clearly has the advantage of consistency. But it might prove difficult to find a unique set of parameter shocks that could deal with, say, an idiosyncratic data issue, without affecting the implicit story for other variables.

In the GME case, the opposite applies: idiosyncratic fixes are easy to apply. Similarly, the non-core equations in the CNC case can simply be inverted to produce any path. But sometimes a more structural interpretation might be desired, and the modeller might be tempted to put more structure on the non-core equations than is justified.

5.1.3. Expectations

A crucial aspect of telling stories involves talking about agents' expectations. In this respect, there are clear differences: only the SIP approach automatically tells consistent stories about expectations. By contrast, agents in the GME and CNC approaches don't "see" the effects of the ad hoc elements and factor those in their decisions. This is clearest in the representation of the GME approach: there the adjustment is completely orthogonal and separable from the DSGE "core". The same applies to the CNC approach; in effect, we are saying that agents follow text-book theory and react to what *should* happen, rather than what actually eventuates each period.[13]

5.2. Towards a hybrid of hybrids

Our experiences with BEQM show us that the model users naturally migrate to a de facto hybrid. They use measurement error, in the sense of shocks to accounting identities in the non-core equations, to deal with data inconsistencies and revisions. They use the shocks to "behavioural" non-core equations when they simply want to "fine tune" the projection at a late stage. They might use shocks to auxiliary variables to try to tell stories about marginal effects that the core model does not capture. Finally, they treat some structural parameters in the core model as potentially time-varying. This appears to be a combination of SIP, GME and CNC approaches. In this section, we attempt to formalise this.

The key part of our proposal is simply to include z variables in the stochastic processes for parameters. For example, the process for the laboour supply

[13] This inconsistency also creates problems when using the projections model for policy experiments, such as optimal control exercises. Hence, there is still a barrier between small, highly stylised models that can be used for policy experiments and richer, larger models used to inform projections, which was surely a motivation to explore the DSGE paradigm in the first place.

parameter γ could be changed to

$$\ln \gamma_t = (1-\rho^\gamma)\ln \bar{\gamma} + \rho^\gamma \ln \gamma_{t-1} + \psi z_{t-1} + \varepsilon_t^\gamma,$$

where, as usual, the z variables should be stationary. In effect, one would try to see whether the unobserved "shock" could be correlated to some other observed factors.[14] In this example, we could include measures of wage bargaining power or other indicators of labour market pressure to provide a better account for cyclical variations in wages and employment. We feel this would provide a much better basis for communication with policymakers than simply accounting for the variation by an AR(1) process.[15]

The state equation would now include direct effects from the auxiliary variables:

$$\begin{bmatrix} s_t \\ x_t \end{bmatrix} = \begin{bmatrix} B_{ss} & B_{sx} \\ 0 & B_{xx} \end{bmatrix} \begin{bmatrix} s_{t-1} \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} B_{sz} \\ B_{xz} \end{bmatrix} z_t + \begin{bmatrix} B_{s\varepsilon} & B_{s\eta} \\ 0 & B_{x\eta} \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix},$$

and the modeller would also provide some process for z in order to project it forward.[16] In addition, we could use generalised measurement errors to cope with idiosyncratic issues, such as with data discrepancies and predictable revisions, implying the measurement equation

$$f_t = \begin{bmatrix} B_{fs} & B_{fx} \end{bmatrix} \begin{bmatrix} s_t \\ x_t \end{bmatrix} + B_{fu} u_t,$$

where u follows some process, such as the VAR(1) in the GME exposition in section 3.1..

Note that we only allow z variables to affect states, because we want expectations to be formed consistently with respect to the predictable effects of these ad hoc variables. They could also be allowed to have effects on the non-predetermined variables directly, but then they would not be expectationally consistent (in the same way that GME and CNC approaches are not).

We conjecture that this approach would offer flexibility, using measurement errors to cope with idiosyncratic issues; consistent treatment, in the sense that other, structural shocks that would feed everywhere theory suggested they should; coherence, in the sense that, because there is one model, expectations are consistent; and the ability to communicate better stories by relating shocks to what we can easily observe.

---

[14] There are links here to Holland and Scott (1998), who construct measures of shocks and see whether they are correlated with macro variables. (They find some correlation.) It is similar in spirit to the informal analysis in Smets and Wouters (2004), who examine the implied shock series from their SIP model and attempt to relate them to the discussions of the conjuncture contained in ECB monthly reports.

[15] We note that Boivin and Giannoni (2005) have independently made a similar point recently, for example: "measures of labour productivity, oil prices, or commodity prices may all be noisy indictors of productivity containing independent information that could be exploited... this could be the case for all exogenous shocks in the model."

[16] We need to provide such a process so that it is correctly factored into agents' decisions when they form their rational expectations.

6. Summary

This essay began by noting that, while much attention has been given to taking DSGE models to the data, little seems to have been directed towards taking such models to the policy environment. Perhaps it has been assumed that the former is sufficient for the latter. We have tried to show why we think this is not the case. We have stressed the importance of communicating clear stories to policymakers, who will often have to justify their decisions by telling stories themselves. In particular, we think that the current reliance of DSGE models on unobserved but highly persistent extrinsic dynamics creates a credibility gap in the eyes of policymakers.

We investigated three approaches that offer a structure that could be used to take DSGE models to the policy environment. We saw that the different methods imposed different structures on the data, so that conclusions about what shocks accounted for history and what would drive projections could be very different. This aspect of identification has received little attention. Moreover, each approach would have different implications for projections, especially as concern the imposition of judgement and the treatment of expectations. In particular, we have argued that the structures would have a meaningful impact on the way projections are communicated. We have argued that there is no clear 'winner': no single method, taken literally and used exclusively, is problem-free.

We proposed a new structure that attempts to maintain expectational consistency, with the flexibility to deal with idiosyncratic, 'non-economic' issues. In particular, we feel that mapping shocks onto observables would greatly aid communication. These are merely conjectures at this stage, but we hope that our essay will prompt discussion, debate, and experimentation.

References

Boivin, Jean, and Marc Giannoni (2005), "DSGE models in a data-rich environment." manuscript.

Canova, Fabio and Luca Sala (2005), "Back to square one: identification issues in DSGE models" manuscript.

del Negro, Marco and Frank Schorfheide (2003), "Take your model bowling: forecasting with general equilibrium models." *Federal Reserve Bank of Atlanta Economic Review*, 35-50.

del Negro, Marco and Frank Schorfheide (2004), "Priors from general equilibrium models for VARs" *International Economic Review* 45(2), 643-73.

Harrison, Richard, Kalin Nikolov, Meghan Quinn, Gareth Ramsay, Alasdair Scott and Ryland Thomas (2005), *The Bank of England Quarterly Model* London, Bank of England.

Holland, Alison and Andrew Scott (1998), "The determinants of UK business cycles" *The Economic Journal* 108, 1067-1092.

Ireland, Peter N. (2004), "A method for taking models to the data" *Journal of Economic Dynamics and Control* 28(6), 1205-26.

Smets, Frank R. and Raf Wouters (2003), "An estimated dynamic stochastic general equilibrium model of the Euro area" *Journal of the European Economic Association* 1(5), 1123-75.

A Model-consistent data

Series	ONS code	Descriptor
Population, N	MGSL.Q	LFS: Population aged 16+ UK, thousands (NSA)
Total hours, H	YBUS.Q	LFS: Total actual weekly hours worked UK, millions (SA)
Consumption, C	ABJR.Q	Household final consumption expenditure CVM (SA) £m
Investment, I	NPEL.Q	Gross Fixed Capital Formation, Business Investment CVM (SA) £m

For the CNC approach we also need house price data, which were constructed as follows: we create a nominal house price index as the average of the Nationwide and Halifax house price indices from 1983Q1. Before 1983Q1 we use data from the Nationwide index only and splice this to the average series. The nominal house price data are then divided by the consumption deflator (ABJQ.Q/ABJR.Q) to obtain a real house price measure. This was then further detrended by trend growth μ.

B Estimation

B1. DSGE model parameters and priors

The parameters $\alpha, \beta, \delta, \gamma, \mu, a$ determine the steady state of the model. We calibrated these parameters to match the sample averages of the data using the steady-state versions of the model equations, as derived in Ireland (2004). We began by setting $\beta = 0.99$, which is a standard assumption. We then set $\delta = 0.0115$, which is the average of the implied depreciation rate using annual ONS data for the nominal (whole economy) capital stock.[17] We then estimate a log-linear time trend for per capita output, which delivers a point estimate of $\mu = 1.0069$. These assumptions allow us to set the remaining parameters using the steady-state

[17] Specifically, we compute the average of NQAE.A/(1000*CIXM.A) for 1978–2004 and convert to a quarterly rate.

relationships in the model. Specifically, we set

$$\alpha = \frac{\mu/\beta - 1 + \delta}{\mu - 1 + \delta}\overline{I/Y},$$

where $\overline{I/Y}$ is the sample average of the investment:output ratio. We set

$$\gamma = \frac{1-\alpha}{\overline{H/N}}\left[1 - \frac{\alpha\left(\mu - 1 + \delta\right)}{\mu/\beta - 1 + \delta}\right]^{-1},$$

where $\overline{H/N}$ is the sample average of per capita hours. Finally, we set the value of $\bar{a}$ to match the sample average of detrended per capita output. Following the notation of Section 4, output data are $\exp\left(\tilde{y}_t - t\log\mu\right)$, which, denoting the sample average as $\overline{\exp\left(y\right)}$, allows us to find $\bar{a}$ by solving

$$\overline{\exp\left(y\right)} = \bar{a}^{1/(1-\alpha)}\left(\frac{\alpha}{\mu/\beta - 1 + \delta}\right)^{\frac{\alpha}{1-\alpha}}\frac{1-\alpha}{\gamma}\left[1 - \frac{\alpha\left(\mu - 1 + \delta\right)}{\mu/\beta - 1 + \delta}\right]^{-1}.$$

We also need to set priors for the parameters determining the dynamics of the DSGE model: ρ and the properties of the productivity shock ε. We do so by constructing a Solow residual as follows. In logs, the production function is

$$y_t = a_t + \left(1 - \alpha\right)h_t + \alpha k_{t-1},$$

and the capital accumulation identity means that

$$\begin{aligned} k_{t-1} &= \left(1-\delta\right)\frac{k_{t-2}}{\mu} + \frac{i_{t-1}}{\mu} \\ &= \mu^{-1}\sum_{j=0}^{t-1}\left(\frac{1-\delta}{\mu}\right)^{j} i_{t-j} + \left(\frac{1-\delta}{\mu}\right)^{t} k_0, \end{aligned}$$

where k_0, the initial capital stock is unknown. We construct a data series s_t

$$s_t = y_t - \left(1-\alpha\right)h_t - \frac{\alpha}{\mu}\sum_{j=0}^{t-1}\left(\frac{1-\delta}{\mu}\right)^{j} i_{t-j}$$

and run the following regression:

$$s_t = \kappa_1 + \kappa_0\left(\frac{1-\delta}{\mu}\right)^{t-1} + u_t.$$

We interpret the residuals u_t as the deviations of productivity from steady state. An AR(1) regression of the residuals delivers estimates for the parameter ρ, its standard error and also an estimate of the standard error of the productivity shock. These estimates are used to set the priors for these parameters.

The parameter values implied by this approach are summarised in the table below:

Parameter	Value
β	0.99
δ	0.0115
μ	1.0069
α	0.1919
γ	47.818
a	30.469
ρ	0.8476
s.e.(ρ)	0.0513
σ	0.0098

B2. Estimation results

For each approach, we used Bayesian maximum likelihood to estimate the parameters. The mode of the posterior distribution was computed numerically before conducting Markov Chain Monte Carlo methods (using the Metropolis Hastings algorithm) to estimate the posterior distribution. In each case, we performed a total of 50,000 replications in four chains. The first half of each chain was burned. The average acceptance rates for the Metropolis Hastings algorithm for GME, SIP and CNC were 0.44, 0.280 and 0.36 respectively.

In setting our priors, we assumed that the additional shocks that were added to the RBC model were serially uncorrelated with small variance.[18] This corresponds to the prior that the RBC model is a good representation of the data so that additional dynamics are of little importance, and is chosen to provide some symmetry across the prior assumptions used for each approach.

For each approach, we estimate the parameters governing the dynamics of the producitivity shock: ρ and the standard deviation of ε (σ). Each approach adds additional dynamic parameters to be estimated. For the GME approach, we need to estimate the parameters of the 'measurement error' process described in equation (37). For SIP, we need to estimate the parameters governing the shocks to preferences (ρ_θ, σ_θ) and labour supply (ρ_γ, σ_γ) shown in equations (39), (40) and (41). For the CNC approach, we need to estimate the parameters of the non-core equations, which are specified as follows:

$$\begin{aligned}
c_t &= \phi^c c_t^* + (1-\phi^c)\, c_{t-1} + \psi^c\,(lph_t - lph_{t-1}) + c_t^{res} \\
h_t &= \phi^h h_t^* + \left(1-\phi^h\right) h_{t-1} + h_t^{res} \\
i_t &= \phi^i i_t^* + \left(1-\phi^i\right) i_{t-1} + i_t^{res} \\
y_t &= c_t + i_t \\
\Delta lph_t &= \psi^{lph}\Delta lph_{t-1} - \phi^{lph}\left(lph_{t-1} - \overline{lph}\right) + lph_t^{res}
\end{aligned}$$

where we use Δ to denote the first difference operator. Solutions from the core model are denoted with '*' and non core residuals by the *res* superscript.

The estimation results from each approach are presented in Tables B1–B3 below.

[18] We set the variances of the shocks to 1% of the sample variances of the data series they are

Table B1: Parameter estimates for GME approach

Parameter		Prior				Posterior		
	Type*	Mean	Std dev	Mode	St err	5%	Mean	95%
Persistence of technology shock, ρ	B(0,1)	0.8476	0.0513	0.9645	0.0094	0.9529	0.9665	0.9863
VAR term in measurement error, ρ_{cc}	B(-1,1)	0	0.25	0.6366	0.1252	0.5558	0.7075	0.8931
VAR term in measurement error, ρ_{ch}	B(-1,1)	0	0.25	-0.3067	0.1484	-0.4119	-0.1333	0.2207
VAR term in measurement error, ρ_{ci}	B(-1,1)	0	0.25	-0.4373	0.1358	-0.561	-0.2984	-0.0035
VAR term in measurement error, ρ_{hc}	B(-1,1)	0	0.25	-0.0732	0.0531	-0.1202	0.0582	0.3321
VAR term in measurement error, ρ_{hh}	B(-1,1)	0	0.25	0.7943	0.0669	0.6831	0.7667	0.8798
VAR term in measurement error, ρ_{hi}	B(-1,1)	0	0.25	-0.1893	0.0621	-0.2573	-0.0623	0.1787
VAR term in measurement error, ρ_{ic}	B(-1,1)	0	0.25	-0.4377	0.1316	-0.5597	-0.3246	0.0132
VAR term in measurement error, ρ_{ih}	B(-1,1)	0	0.25	-0.1110	0.1368	-0.3016	-0.0521	0.1685
VAR term in measurement error, ρ_{ii}	B(-1,1)	0	0.25	0.5655	0.1471	0.3368	0.6377	0.9433
Std dev technology shock, σ	IG	0.0098	10	0.0026	0.0003	0.0022	0.0026	0.003
Std dev consumption measurement error, σ_c	IG	0.003	10	0.0097	0.0008	0.009	0.0103	0.0118
Std dev hours measurement error, σ_h	IG	0.003	10	0.0040	0.0004	0.0012	0.0032	0.0046
Std dev investment measurement error, σ_i	IG	0.013	10	0.0103	0.0024	0.0067	0.0143	0.0235
Measurement error correlation, λ_{ch}	B(-1,1)	0	0.25	-0.0682	0.1150	-0.1765	0.0886	0.5579
Measurement error correlation, λ_{ci}	B(-1,1)	0	0.25	-0.3116	0.2004	-0.5292	-0.2556	0.0728
Measurement error correlation, λ_{hi}	B(-1,1)	0	0.25	-0.5253	0.1609	-0.7117	-0.2955	0.0886

* B(x,y)=beta distribution defined on support (x,y); N=Normal; IG=Inverted Gamma (standard deviation reports degrees of freedom)

Table B2: Parameter estimates for SIP approach

Parameter		Prior				Posterior		
	Type*	Mean	Std dev	Mode	St err	5%	Mean	95%
Persistence of technology shock, ρ	B(0,1)	0.8476	0.0513	0.9964	0.0000	0.9964	0.9964	0.9964
Persistence of preference shock, ρ_θ	B(-1,1)	0	0.25	0.8919	0.0162	0.8548	0.8836	0.9139
Persistence of labour supply shock, ρ_θ	B(-1,1)	0	0.25	0.8676	0.0210	0.8201	0.8584	0.8985
Std dev technology shock, σ	IG	0.0098	10	0.0098	0.0006	0.009	0.0101	0.0112
Std dev preference shock, σ_θ	IG	0.003	10	0.0080	0.0006	0.0071	0.008	0.0089
Std dev labour supply shock, σ_h	IG	0.003	10	0.0045	0.0003	0.004	0.0046	0.0051

* B(x,y)=beta distribution defined on support (x,y); N=Normal; IG=Inverted Gamma (standard deviation reports degrees of freedom)

Table B3: Parameter estimates for CNC approach

Parameter		Prior				Posterior		
	Type*	Mean	Std dev	Mode	St err	5%	Mean	95%
Persistence of technology shock, ρ	B	0.8476	0.0513	0.8898	0.0248	0.8686	0.9037	0.9435
Non-core consumption attractor, ϕ^c	B(0,2)	1	0.25	0.5639	0.1140	0.3734	0.5734	0.7528
Non-core effect of house prices on consumption, ψ^c	B(0,2)	1	0.25	0.2677	0.0435	0.2005	0.2732	0.3431
Non-core hours attractor, ϕ^h	B(0,2)	1	0.25	0.0891	0.0105	0.0744	0.0938	0.1124
Non-core investment attractor, ϕ^i	B(0,2)	1	0.25	0.0332	0.0055	0.0319	0.0372	0.0435
AR coefficient in house price equation, ψ^{lph}	B(0,2)	1	0.25	0.6869	0.0652	0.6043	0.7076	0.8078
ECM coefficient in house price equation, ϕ^{lph}	B(0,2)	1	0.25	0.0371	0.0089	0.0319	0.0418	0.0510
Long-run level of log real house prices, $\overline{lph}$	N	-0.1964	0.0175	-0.1976	0.0232	-0.2328	-0.1949	-0.1595
Std dev technology shock, σ	IG	0.0098	10	0.0117	0.0015	0.0093	0.0115	0.0139
Std dev consumption residual, σ_c	IG	0.003	10	0.0078	0.0007	0.0067	0.0079	0.0092
Std dev hours residual, σ_h	IG	0.003	10	0.0032	0.0004	0.0027	0.0033	0.0039
Std dev investment residual, σ_i	IG	0.013	10	0.0293	0.0020	0.0267	0.0302	0.0338
Std dev house price residual, σ_{lph}	IG	0.175	10	0.0206	0.0000	0.0206	0.0206	0.0206

* B(x,y)=beta distribution defined on support (x,y); N=Normal; IG=Inverted Gamma (standard deviation reports degrees of freedom)

B3. Variance decompositions

Variance decompositions using the modal parameter estimates from each approach are presented in Tables B4–B6.

Table B4: Variance decomposition for GME approach

Variable	ε	ζ^c	ζ^h	ζ^i
Output	28.65	60.53	2.42	8.39
Consumption	10.40	75.48	2.99	11.12
Investment	72.15	21.10	1.73	5.02
Hours	7.75	62.07	14.97	15.22
Consumption measurement error	0.00	84.24	3.34	12.42
Hours measurement error	0.00	67.28	16.22	16.50
Investment measurement error	0.00	75.74	6.22	18.04

Table B5: Variance decomposition for SIP approach

Variable	ε	ε^γ	ε^θ
Output	97.53	1.75	0.72
Consumption	98.63	0.34	1.03
Investment	76.73	23.27	0.00
Hours	31.67	46.53	21.8

Table B6: Variance decomposition for CNC approach

Variable	ε	c^{res}	i^{res}	h^{res}	lph^{res}
Output	82.12	2.88	10.05	0.00	4.95
Consumption	83.40	6.10	0.00	0.00	10.50
Investment	44.63	0.00	55.37	0.00	0.00
Hours	89.10	0.00	0.00	10.9	0.00
Real house prices	0.00	0.00	0.00	0.00	100.00

intended to match.