



BANK OF ENGLAND

Working Paper No. 471

The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models

Stephen Burgess, Emilio Fernandez-Corugedo,
Charlotta Groth, Richard Harrison, Francesca Monti,
Konstantinos Theodoridis and Matt Waldron

May 2013



BANK OF ENGLAND

Working Paper No. 471

The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models

Stephen Burgess,⁽¹⁾ Emilio Fernandez-Corugedo,⁽²⁾ Charlotta Groth,⁽³⁾
Richard Harrison,⁽⁴⁾ Francesca Monti,⁽⁵⁾ Konstantinos Theodoridis⁽⁶⁾
and Matt Waldron⁽⁷⁾

Abstract

This paper introduces the Bank of England's new forecasting platform and provides examples of how it can be applied to practical forecasting problems. The platform consists of four components: COMPASS, a structural central organising model; a suite of models, used to fill in the gaps in the economics of COMPASS and provide cross-checks on the forecast; MAPS, a macroeconomic modelling and projection toolkit; and EASE, a user interface. The platform has been in use since the end of 2011 in support of production of the projections produced for the Monetary Policy Committee's quarterly *Inflation Reports*. In this paper we provide a full description of COMPASS, including discussion of its estimation and its properties. We also illustrate how the suite of models can be used to mitigate some of the trade-offs inherent in building a projection with a central organising model such as COMPASS, and discuss the role of the suite in addressing problems of model misspecification.

Key words: Forecasting, macro-modelling, misspecification.

JEL classification: E17, E20, E30, E40, E50.

(1) Bank of England. Email: stephen.burgess@bankofengland.co.uk

(2) IMF. Email: EFernandez-Corugedo@imf.org

(3) Zurich Insurance Group. Email: charlotta.groth@zurich.com

(4) Bank of England. Email: richard.harrison@bankofengland.co.uk

(5) Bank of England. Email: francesca.monti@bankofengland.co.uk

(6) Bank of England. Email: konstantinos.theodoridis@bankofengland.co.uk

(7) Bank of England. Email: matthew.waldron@bankofengland.co.uk (corresponding author)

The bulk of the forecasting platform described in this paper was developed during a project that took place between the autumn of 2009 and the end of 2011. The authors would like to acknowledge the contributions of David Bradnum, Stephen Elliott, Sujeevan Kanageswaran, Kate Monaghan and Anish Patel in developing the IT user interface, Yaser Al-Saffar and Hitendra Shah in providing the business analysis, Angela Middlemiss for managing the project and Martin Andreasen, Christoph Görtz, Alex Haberis and David Reifschneider for their help and advice at various stages of the project. Prior to the start of that project, the Bank conducted an internal review of its previous macroeconomic forecasting platform, BEQM. The authors would like to thank Dario Caldara, Anna Lipińska, Tim Taylor and Gregory Thwaites for their work on that review. Finally, the authors would also like to thank Charlie Bean, Spencer Dale and Robert Woods for helpful comments on a draft of this paper. This paper describes research in progress at the Bank of England and has been published to elicit comments and to further debate. It was finalised on 16 May 2013.

The Bank of England's working paper series is externally refereed.

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

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

Summary

Since autumn 2011 the Monetary Policy Committee (MPC) has used a new forecasting platform to help put together its quarterly economic forecasts. The MPC's judgement is paramount when agreeing their forecasts, but the process also relies on a range of economic models. The new forecast platform includes a central organising model (called COMPASS¹), an enhanced suite of forecasting models, and new IT tools to assist the forecast process. This paper provides detailed documentation of each of these components of the platform and has been published to elicit comments and further debate.

COMPASS is a “New Keynesian” general equilibrium model and shares many features with similar models in use at other central banks and policy institutions. Prices and wages are assumed to be sticky, so monetary policy affects output and employment in the short to medium term. Expectations of future events, including the actions of monetary policy makers, can also affect current output and inflation. COMPASS provides the basic set of relationships that articulate core macroeconomic mechanisms and provides a disciplining framework by ensuring that forecasts are internally consistent. COMPASS itself only provides forecasts for fifteen variables: “key” macroeconomic series such as GDP, inflation, interest rates, trade, wages and consumption.

COMPASS is smaller and simpler than previous central models used at the Bank of England. This makes it easier to estimate and to use, enabling Bank staff to produce timely updates to the MPC's forecast in the weeks ahead of an *Inflation Report*. But it also implies some sacrifice of detailed economic structure. To compensate for that, the suite of models is very much an equal partner in the new forecasting platform. The suite contains over 50 separate models, covering a huge range of different frameworks and ways of thinking about the economy. Different models can be selected from the suite, depending on what insight is required. The suite provides the means to cross-check the projections in COMPASS, expand the forecast to cover more variables, and challenge the key judgements in the forecast.

This paper offers various illustrations of how the suite of models can be used to inform the forecast. Although COMPASS does not include an explicit role for a banking sector, there are several models in the suite that can be used to consider the impact of credit on the economy, and so explore the effects of an impaired banking sector. The forecast platform can be used to estimate the underlying shocks driving the economy and that can be a useful framework to interpret recent events. It is also possible to use the platform to explore the impact of different paths for monetary policy on the economy.

The forecasting platform is likely to evolve over time. The parameter values in COMPASS will be re-estimated on a regular basis, and the structure of the model may be modified as Bank staff learn more about its performance. The Bank's vision for the suite of models is also a dynamic one: models should be added or removed as economic modelling progresses and also as the questions facing policymakers change.

¹ The Central Organising Model for Projection Analysis and Scenario Simulation.

Contents

1	Introduction	1
2	Motivation and design	3
2.1	Forecasting at the Bank of England	3
2.2	The role of the forecasting platform	4
2.3	Design principles	5
3	An overview of the forecasting platform	9
3.1	COMPASS	9
3.2	The suite of models	10
3.3	MAPS	10
3.4	EASE	11
4	COMPASS	12
4.1	The general modelling approach	12
4.2	The model	13
4.2.1	Supply	14
4.2.2	Price and wage setting	16
4.2.3	Private domestic demand	18
4.2.4	Interactions with the rest of the world	19
4.2.5	Fiscal and monetary policy	20
4.2.6	Forcing processes and shocks	22
4.3	Estimation	23
4.3.1	Data and measurement equations	24
4.3.2	Priors	27
4.3.3	Posterior parameter estimates	33
4.4	Model properties	34
4.4.1	A monetary policy shock	34
4.4.2	A labour augmenting productivity shock	36
4.4.3	A domestic risk premium shock	37
5	The suite of models	39
5.1	Reasons to employ a suite of models	39
5.2	Models which articulate economic shocks and channels missing from COM- PASS	40
5.2.1	Models with energy	40
5.2.2	Models and tools for understanding the impact of the financial sector	41
5.3	Models which expand the scope of the forecast	42
5.3.1	The Post-Transformation Model (PTM)	43
5.3.2	The Balance Sheet Model (BSM)	46
5.4	Models which generate alternative forecasts	47
6	The IT infrastructure	51
6.1	EASE	51
6.2	MAPS	52
6.2.1	Modelling framework	52
6.2.2	Estimation	54

6.2.3	Model Analysis	55
6.2.4	Projection and simulation	57
6.2.5	Decompositions	59
6.2.6	Expectations & policy analysis	61
6.2.7	Non-Linear Backward-Looking models in MAPS	61
7	Forecasting with misspecified models and the role of judgement	63
7.1	The ‘misspecification algorithm’	63
7.1.1	Understanding the economics of the misspecification	64
7.1.2	Quantifying the effects of misspecification	65
7.1.3	Incorporating the quantitative effects of misspecification	65
7.2	Alternative approaches	67
8	The forecasting platform in action	69
8.1	Introducing data news	70
8.1.1	Does the identification of shocks look sensible given the data?	73
8.1.2	How should the forecast be changed in light of the data news?	75
8.2	Incorporating effects of VAT changes	78
8.3	Incorporating effects of financial frictions	84
8.3.1	Financial frictions and credit spreads	85
8.3.2	The suite models	87
8.3.3	Quantifying the effects of financial shocks	88
8.3.4	Mimicking the effects of financial shocks using COMPASS	92
8.4	Incorporating policy changes	95
9	Conclusions	99
	Bibliography	101

1 Introduction

Towards the end of 2011, staff at the Bank of England adopted a new central organising model, COMPASS, to assist with the production of forecasts presented by the Monetary Policy Committee (MPC) in their quarterly *Inflation Reports*. This replaced the previous central organising model, BEQM, which had been in use since 2003. An enhanced and updated suite of models was introduced alongside COMPASS, and the models were all supported by new IT infrastructure. The purpose of this paper is to document the new models and IT tools, and to demonstrate how they are used in practice to inform the judgemental forecasts made by the MPC.

The new forecasting platform recognises more explicitly the importance of the suite of models and the costs of operating large, intractable models. Relative to previous central organising models at the Bank, COMPASS is both smaller and simpler, with the aim of making it more straightforward for Bank staff to use the model to aid the MPC's discussions and to articulate the narrative of the MPC's forecast. Of course, all economic models are misspecified, and COMPASS is no exception, but the decision to use a smaller central organising model places a greater onus on the suite of models in being able to address known misspecifications, and in providing cross-checks on the forecast. One important aim of this paper is to explain our approach to dealing with misspecification of the central model and to illustrate how the suite of models can be used to try to mitigate it. The new IT infrastructure is particularly important in that regard because it provides much more comprehensive support for multiple models.

The paper is structured as follows. Section 2 explains the motivation for creating a new forecasting platform in more detail, and places it in the context of the wider forecast process at the Bank of England. Sections 3 to 6 document the individual components of the new forecasting platform: COMPASS; the suite of models; MAPS and EASE. The remainder of the paper (Sections 7 and 8) explains how the models and tools can be used to support judgemental forecasting. In particular, we document our approach to dealing with misspecification, and provide concrete examples of how the models and tools can be used in practice.

COMPASS, the new central organising model, is described in Section 4. COMPASS serves three main purposes: to be the main organising framework for the construction of the forecast; to analyse and explain the forecast; and to assess the sensitivity of the forecast to alternative assumptions. COMPASS is an open economy, New Keynesian DSGE model, estimated on UK data using Bayesian methods. It shares many features with similar models at other central banks. Wages and prices are assumed to be sticky, and so monetary policy can influence real variables such as output and employment over short to medium horizons, but not in the long run. And expectations of monetary policy actions are an important determinant of current output and inflation. A full derivation of COMPASS and a complete set of impulse responses are provided in accompanying appendices.

The suite of models is documented in Section 5. Because the suite is diverse and contains a large number of models, many of which are documented in past Bank publications, we do not seek to describe every model. Rather, we provide examples of models within three broad categories: models which articulate economic shocks and channels which are omitted from COMPASS; models which expand the scope of the forecast, by producing forecasts for variables not in COMPASS itself; and models which offer cross-checks by generating alternative forecasts for variables which are in COMPASS.

Section 6 documents the new IT infrastructure, MAPS and EASE. MAPS is a modelling toolkit which supports all of the models described in the paper. It offers two broad classes of functionality: model analysis, to estimate and interrogate the properties of models; and projection, which allows the construction of forecasts using those models, including the imposition of judgement. Given the importance of judgement in the forecast, a detailed description of the toolkit for imposing judgement is provided in an accompanying appendix. EASE is a user interface which provides access to all of the models and tools. It supports the staff's workflow in updating the projections and producing analysis as inputs to key MPC meetings.

Section 7 explains our approach to dealing with problems of model misspecification. In general, there are three steps involved: first, to understand the economics of the misspecification; second, to quantify it; and third, to find a suitable method to incorporate a quantitative correction into the judgemental forecast organised using COMPASS.

Section 8 provides concrete examples of how COMPASS, the suite of models and the associated IT tools can be used together to address a selection of problems commonly encountered in forecasting. We focus on the following: the updating of an MPC forecast for new and revised data; the use of suite models to correct for known misspecifications in COMPASS; and the application of conditioning paths to the forecast. One of the known misspecifications we consider is the absence of financial frictions in COMPASS. We demonstrate how the suite of models can be used to quantify the impact of financial sector shocks in a variety of ways.

As this paper makes clear, macroeconomic models play a crucial supporting role in the MPC's forecast process. They provide a framework for organising the forecast, and important insights which can be fed into discussions of the forecast with the MPC. However, the production of a forecast is not an exercise in feeding data into a model, or even a set of models. MPC members and Bank staff are acutely aware of the strengths and limitations of macroeconomic models, and the judgement of policymakers remains paramount when setting monetary policy and agreeing forecasts for the quarterly *Inflation Report*. The MPC's projections are ultimately made by the MPC, not by economic models.

2 Motivation and design

This section explains the motivation behind the creation of the new forecasting platform, with reference to the process that it supports. The design of the new platform flows from a desire to deliver a forecasting architecture that best supports that process.

2.1 Forecasting at the Bank of England

Each quarter, in accordance with section 18 of the Bank of England Act 1998, Bank of England staff produce an *Inflation Report* on behalf of the Monetary Policy Committee (MPC). Among the key charts in each report are the ‘fan charts’, which represent “the MPCs best collective judgement about the most likely paths for inflation and output, and the uncertainties surrounding those central projections.”¹

Forecasting at the Bank of England therefore has two key characteristics: first, the forecasts are ‘owned’ by the MPC; second, they are expressed as probability distributions, since a full assessment of the outlook has to capture risks and uncertainties. In constructing their forecast, the MPC has the following objectives in mind:

- To discuss the economic outlook and come to a view on the balance of risks to economic activity and inflation.
- To come to a view on the appropriate response of monetary policy in light of the discussion of the economics of the forecast and the uncertainty around it.
- To communicate the outlook to the public in a manner that promotes transparency and accountability.

It is discussion of the economics of the forecast, including the balance of risks, that underpins those objectives, not a desire to maximise the accuracy of their point forecasts per se. The internal process through which the staff provide inputs to the MPC’s forecast discussions is tailored to those objectives. Bean and Jenkinson (2001) describe the internal process that supports the production of the forecast. An important feature is a high level of engagement from the MPC, taking place through a sequence of meetings in the weeks leading up to the production of the *Inflation Report*. At each stage of the process, MPC judgements are discussed and incorporated into the forecasts.

While the structure of the forecast process described in Bean and Jenkinson (2001) remains broadly unchanged, the tools used by the staff to implement that process have evolved over time.² Bean and Jenkinson (2001) note that “A central tool in the production of these forecasts is a relatively standard macroeconomic model (MM)”.³ The MM was replaced in 2003 by the Bank of England Quarterly Model (BEQM) (see Harrison et al. (2005)). And from the November 2011 *Inflation Report*, the forecast process has been supported by the forecasting platform described in this paper.

¹This text appears in the foreword of each *Inflation Report*.

²There may be some changes to the forecast process as the Bank implements some of the recommendations in the Stockton Review – see the discussion towards the end of Section 2.3.

³The MM is described in Bank of England (1999) and Bank of England (2000).

2.2 The role of the forecasting platform

The ‘platform’ used for the production of the quarterly forecasts consists of a set of tools used by the staff to support the MPC’s discussions. Economic models form an important part of that toolkit.

The Bank’s long-standing approach to forecasting has consistently recognised the strengths and weaknesses of macroeconomic models. In the foreword to the 1999 volume *‘Economic models at the Bank of England’*, Governor Eddie George wrote:⁴

The Bank’s use of economic models is pragmatic and pluralist. In an ever-changing economy, no single model can possibly assimilate in a comprehensible way all the factors that matter for policy. Forming judgements about those factors, and their implications for policy, is the job of the Committee, not something that can be abdicated to models or even to modellers. But economic models are indispensable tools in that process.

This view is reiterated in Bean and Jenkinson (2001, p438):⁵

All economic models are highly imperfect reflections of the complex reality that is the UK economy and at best they represent an aid to thinking about the forces affecting economic activity and inflation. The MPC is acutely aware of these limitations.

So the economic models used in the forecast process play a supporting role, rather than a starring one. The forecasting platform used by the staff provides a way to organise the contributions from a range of economic models. The types of contributions that different models can provide are wide-ranging and include:

- Elucidating the economic mechanisms that might be determining the behaviour of particular macroeconomic variables.
- Assessing the quantitative effects of particular shocks or events.
- Identifying which types of economic shocks best explain the current state of the economy.
- Quantifying the sensitivity of any of the answers above to different assumptions about the underlying structure of the economy.
- Exploring the policy implications of particular shocks or events.

Casual inspection of the list above reveals that it would be extremely difficult for a single economic model to deliver every item, consistent with the Bank’s long-standing use of a ‘suite’ of economic models.⁶

The decision to build a new forecasting platform was motivated in part by rapid advances in the tools available to estimate and analyse the outputs of models, enabled by

⁴See Bank of England (1999).

⁵It is also evident in the documentation of BEQM: “The new macroeconomic model [BEQM] is by no means the only input into the forecasting and policy processes.” (Harrison et al., 2005, p151).

⁶While the considerations listed here are clearly of importance to the MPC, there are additional requirements to ensure that the platform is practically useable. For example, the staff require that judgement can be applied to the model efficiently in order to be able to construct the forecast within the required timetable.

advances in computing power.⁷ This progress went hand in hand with development and implementation of new forecasting models in other central banks and policy institutions, reflecting concerted efforts in many central banks to use structural economic models at the heart of their policy and forecast processes.⁸

2.3 Design principles

The philosophy behind the new forecasting platform is summarized succinctly by George Box (Box and Draper, 1987, p424): “Essentially, all models are wrong, but some are useful”. As highlighted in the preceding section, any models that support the forecast process will be misspecified. Nevertheless, models can provide useful insights into the discussions that help the MPC to produce each forecast. The key challenge, therefore, is to ensure that the forecasting platform helps the Bank’s staff to extract the most useful insights from the wide range of models at its disposal.

From the perspective of producing the best statistical forecasts, a popular approach is to combine the insights from many models by taking a weighted average of their forecasts.⁹ Indeed, the staff produce forecasts from a set of econometric models optimized for forecast performance.¹⁰ These forecasts are used as cross-checks on the MPC’s forecast.

However, as noted in Section 2.1, a primary purpose of the *Inflation Report* is for the MPC to present a narrative describing its best collective judgement of the forces influencing the current state of the economy and the alternative paths it might take over the future. Models optimised for statistical forecasting performance rarely provide a clear story of *why* they produce the forecasts they do.

Partly for this reason, the new forecasting platform consists of a “central” forecasting model, surrounded by a suite of other models and tools. The purpose of the central model is to provide an organising framework to help frame the discussions of the key forces shaping the current state of the economy and how they might affect the forecast. The surrounding suite of models and tools provide ways to cross-check, interrogate and adjust the forecast, particularly in the areas in which the central model is more likely to be deficient.

As already noted, the process of producing the *Inflation Report* forecasts using a “central organising model”, surrounded by other models and tools is very much a continuation of the approach taken at the Bank of England for many years.¹¹ However, the forecasting platform described in this paper more explicitly recognises the role that the suite of models has to play. In particular, the IT infrastructure (described in Section 6) that staff at the Bank use to produce and analyse forecasts has been designed with a

⁷It should be noted that the decision to build the new forecasting platform predates the financial crisis. The process of building the new platform involved significant investment in developing new tools (including IT systems), which necessarily took time to undertake.

⁸Examples include the ‘g3’ model introduced by the Czech National Bank (Andrle et al., 2009); the RBNZ’s KITT model (Beneš et al., 2009); the NEMO model developed at Norges Bank (Brubakk et al., 2006); the Riksbank’s RAMSES model (Adolfson et al., 2007); the ECB’s NAWM (Christoffel et al., 2008); the EDO model of the Federal Reserve Board of Governors (Edge et al., 2007; Chung et al., 2010); the Bank of Canada’s ToTEM (Murchison and Rennison, 2006).

⁹There are of course, different ways to weight the forecasts. See, for example, Kapetanios et al. (2006) and Kapetanios et al. (2007).

¹⁰See Kapetanios et al. (2008).

¹¹In the late 1990s and early 2000s, the central organising model was the Medium-Term Macroeconomic Model (MTMM, see Bank of England, 1999, 2000)). From 2003, it was the Bank of England Quarterly Model (BEQM, see Harrison et al., 2005).

suite of models in mind, so this aspect of the design of the forecasting platform marks a material improvement over previous ones used in the Bank.

In retaining the long-standing generic design of the forecasting platform, the most recent improvements represent an evolution rather than a revolution. However, the new platform has been designed to incorporate more explicitly the fact that the central model is misspecified and consequently to improve the way that additional insights can be applied to it.

To deliver those benefits, the design of the central organising model must be chosen carefully. This involves assessing the trade-offs between the model's strengths and weaknesses in a number of dimensions. As noted above, policymakers would ideally like to use models for a wide range of purposes. In discussing the design of BEQM, Harrison et al. (2005, p 12) note a trade-off between the 'empirical coherence' and 'theoretical coherence' of different model types.¹² Though different types of model may indeed exhibit different degrees of empirical and theoretical coherence, our view is that the trade-offs faced when designing a central organising model in fact involve many more dimensions than this. Moreover, these trade-offs are unlikely to be as continuous and smooth as the simple trade-off characterised by Harrison et al. (2005) and Pagan (2003).

Our view is that there are typically trade-offs between the performance of a potential central organising model in the following dimensions:

- **Theoretical foundations.** The behaviour of the central organising model should be consistent with the theory underpinning policymakers' views of the monetary transmission mechanism.
- **Empirical fit.** The central organising model should be able to explain the macroeconomic data well.
- **Tractability.** The central organising model should be easy to use, easy to understand, reliable and robust.
- **Flexibility.** It should be possible to examine easily the implications of alternative economic assumptions (eg the implications of different parameter values) on the behaviour of the central organising model.
- **Comprehensiveness.** The central organising model should provide adequate coverage of the key economic mechanisms and variables required to support policymakers' discussions.

Fully exploring the precise nature of these trade-offs is beyond the scope of this paper. However, it may be useful to highlight some obvious examples. A large model is likely to contain a comprehensive coverage of the range of macroeconomic variables relevant to policy and forecast discussions. However, it is also likely that a large model will be less tractable (as defined above) than a smaller model. The tradeoff between 'empirical coherence' and 'theoretical coherence' stressed by Harrison et al. (2005) and Pagan (2003) clearly appears here too.

The design of the new forecasting platform reflects our assessment of the nature of the trade-offs listed above. Our assessment led us to build the platform around a central

¹²For example, vector autoregression (VAR) models provide a good fit to the data ('empirical coherence') and dynamic stochastic general equilibrium (DSGE) models dictate that the model's predictions accord with the underlying theoretical assumptions ('theoretical coherence').

organising model of a particular type: a New Keynesian Dynamic Stochastic General Equilibrium (DSGE) model, similar to those implemented over recent years in other central banks. This judgement was based on a number of considerations:

1. New Keynesian DSGE models incorporate a well-understood baseline description of some key elements of the monetary transmission mechanism that policymakers agree are important. For example, monetary policy can affect activity and inflation in the short to medium term because of rigidities in setting nominal prices and wages, as discussed in Monetary Policy Committee (1999). New Keynesian DSGE models also incorporate the notion that expectations and the stabilising role of monetary policy are important for understanding the economy, two important elements of the “consensus view” of monetary policy discussed by Bean (2007).
2. Tools for using DSGE models and analysing their outputs are now well established.¹³ Including these tools within the forecasting platform helps to make the operation of the central DSGE model tractable and its outputs more quickly interpreted.
3. It is now possible to estimate the parameters of relatively large DSGE models. The estimation process provides important information about the model’s parameter values and the extent to which the model fits the data. Recently developed DSGE models have been shown to fit the data no worse than model types that are primarily designed for empirical coherence, like VARs.¹⁴
4. Many other central banks and international institutions use New Keynesian DSGE models to support their forecast and policy processes. This makes it easier to share results with colleagues at other central banks.

Given recent arguments that the current generation of New Keynesian DSGE models are ill-suited to analysing the causes and consequences of financial crises, using a model of this type as the central organising model may seem surprising, particularly given that the model does not include a financial sector.¹⁵ While there is a debate over the validity of these critiques,¹⁶ it is undoubtedly true that the events surrounding the financial crisis have posed significant challenges to all existing modelling approaches, not just DSGE.

If the forecasting platform relied on a single model, this would be an acute concern. After all, the only thing we know with certainty about *any* model is that it is wrong. However, the forecasting platform is designed to help staff to address the fact that the

¹³See, for example Del Negro and Schorfheide (forthcoming), Gómez et al. (2009) and Beneš et al. (2009).

¹⁴See Smets and Wouters (2007). Such results do not imply that the forecasts of a DSGE model cannot be improved if cross-equation restrictions are relaxed (see Del Negro and Schorfheide (2004)). Regardless of the relative forecast performance of these models, Edge et al. (2010) note that the *absolute* forecast performance of DSGE models and their competitors is poor. In terms of their ability to forecast individual variables, like GDP and inflation, these models typically fail to beat simple univariate statistical models (see Kapetanios et al. (2008)).

¹⁵Our judgement is that the benefits of adding a financial sector to COMPASS would be outweighed by the costs of the added complexity. It is possible that we will come to a different view in the future as this rapidly developing area of the academic literature advances. Some alternative approaches to capturing financial frictions in the forecast are discussed in Section 8.3.

¹⁶Our sense is that there is still some difference between academic presentation of DSGE models and their practical use in central banks. Staff and policymakers are acutely aware of the problems associated with model misspecification. For a review of the use of DSGE models at central banks see Tovar (2008).

central organising model is misspecified. In particular, the choice of a relatively simple and well-understood central organising model makes it easier to identify where and how the model is misspecified. The first step in that process is to understand how the model works. In subsequent steps, one needs to consider how to adjust the outputs of the central organising model to address the misspecification. As we will see in Section 7, many of the tools that have been developed to analyse model misspecification are designed for models with a DSGE (or similar) structure. It is therefore easier to apply these tools to the central organising model if it has a DSGE structure. A crucial element of the design of the forecasting platform is the suite of models and the infrastructure that supports their use. The models in the suite are essential in helping the staff to adjust the outputs of the central organising model in order to help to account for misspecification. Section 5 discusses the role and use of the suite of models in more detail and Section 8.3 uses the suite to demonstrate how financial shocks can be incorporated into the forecast.

It is important to note that the design of the forecasting platform needs to take the institutional framework into account.¹⁷ As noted in Section 2.1, key elements of the forecasting process at the Bank of England include the ownership of the final forecast by the MPC, their intensive involvement in the production of the forecasts and the fact that the published forecasts in the *Inflation Report* are judgemental.

We believe that the new forecasting platform will support the internal processes well, though it is possible that other designs may actually support it at least as well. In the process of developing the new platform, we considered an alternative design in which the central organising model would be an even simpler, empirically-based model, supported by an estimated DSGE model as a crucial element of a supporting suite of models.

Of course, in practice, it was only possible to implement and operate a forecasting platform of a single design (it was infeasible to build and compare the performance of two alternative platforms in ‘live’ use). Reflecting that, a continual review process is in place to assess the performance of the new forecasting platform. Importantly, the infrastructure has been designed to be sufficiently flexible to allow such changes to the nature of the central organising model if required. Moreover, such a change would not alter the general approach to macroeconomic forecasting at the Bank of England described in this paper.

The same is true of the recommendations in the Stockton Review of the MPC’s forecasting capability, which was published in October 2012 and to which the Bank published a response in May 2013.¹⁸ The review concluded that the forecast process and the forecasting tools are fundamentally sound, but that there was scope for improvement in increasing the transparency of external communications and enhancing the role for analysis of monetary policy strategy in the forecast process. In its response, the Bank’s Executive stated that it intends to implement many of the recommendations in the review. The forecasting platform described in this paper supports the sort of changes the Bank intends to make. In particular, COMPASS and the IT infrastructure provide a platform for analysis of monetary policy strategy that would not have been possible in BEQM.¹⁹

¹⁷See, for example, Laxton et al. (2009).

¹⁸The Stockton review was conducted by David Stockton, a former Director of Research and Statistics at the Board of Governors of the Federal Reserve System. It formed part of three reviews commissioned by the Court of the Bank of England, covering the Bank’s performance during the financial crisis. See <http://www.bankofengland.co.uk/about/Pages/courtreviews/default.aspx> for the reviews in full and see <http://www.bankofengland.co.uk/publications/Pages/news/2013/051.aspx> for the Bank’s response.

¹⁹The ‘core/non-core’ structure of BEQM made analysis of monetary policy very difficult because it decoupled agents’ expectations of policy, which were articulated in the ‘core’, from the resulting outcomes, which were the product of the ‘core’ and the ‘non-core’.

As discussed in Section 6.2.6, providing support for analysis of monetary policy strategy will be a key part of future development for the tools that support the forecast process.

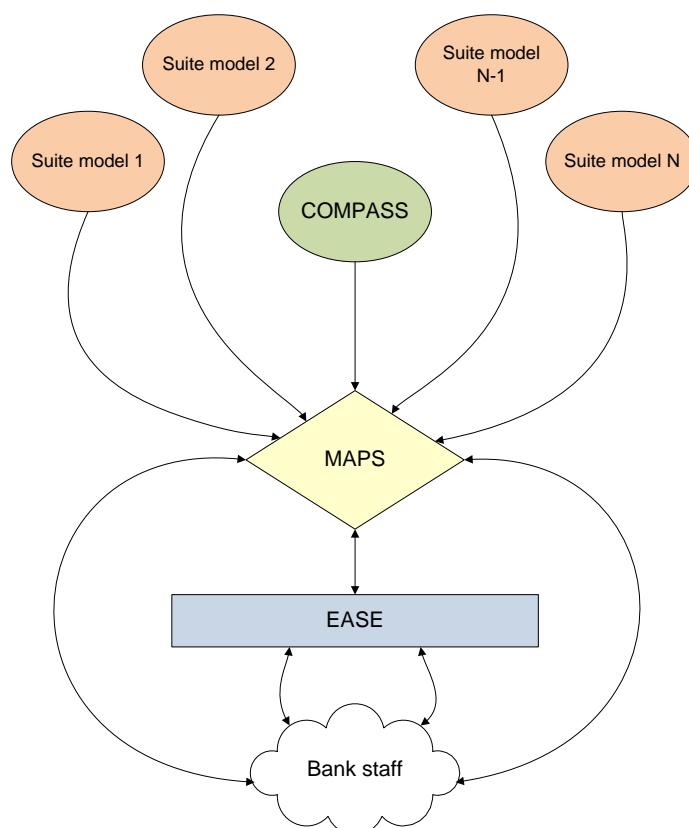
3 An overview of the forecasting platform

This section outlines the components of the forecasting platform. The aim is to explain each component at a high level, by way of introduction to Sections 4, 5 and 6.

Figure 1 provides a schematic representation of the forecasting platform. The platform has four main components, which we discuss in turn:

1. The central organising model, **COMPASS**
2. The suite of models
3. A modelling toolkit, **MAPS**
4. A user interface, **EASE**

Figure 1: Overview of the forecasting platform



3.1 COMPASS

COMPASS is the *Central Organising Model for Projection Analysis & Scenario Simulation*. As its name suggests, COMPASS is intended to serve three key purposes: to be the main organising framework for the construction of the forecast; to analyse and explain

the forecast (projection analysis); and to construct experiments to assess the sensitivity of the forecast to alternative assumptions (scenario simulation).

COMPASS is an open economy New Keynesian DSGE model, sharing many features with antecedents at other central banks.²⁰ Because prices and wages are assumed to be sticky, monetary policy can influence demand and hence output and employment in the short and medium term. In the long term, however, output is determined by technology and the supply of factors of production. One implication of these properties is that there is no long-run trade-off between inflation and output (or growth). Another implication is that expectations of future monetary policy actions can have important effects on current output and inflation.

3.2 The suite of models

COMPASS is a relatively small and simple central model, and this makes it relatively easy to use and understand. But, like all models, it is misspecified. Suite models can be used to help to overcome the main misspecifications and to inform staff and MPC judgements made to the forecast produced by COMPASS.

The models in the suite can be divided into three broad classes, according to their main purpose:

- (a) Models which articulate economic shocks and channels which are omitted from COMPASS (see Section 5.2);
- (b) Models which expand the scope of the forecast, by producing forecasts for additional variables not included in COMPASS (see Section 5.3);
- (c) Models which generate alternative forecasts for variables which are in COMPASS. These models play an important role in validating and adjusting the output from the central organising model (see Section 5.4).

3.3 MAPS

MAPS is the *Model Analysis & Projection System*. MAPS has been designed as a relatively general modelling language. A critical feature of MAPS is that it works with many models, rather than just the central organising model.²¹ The ability to use the same tools with a broad range of models greatly reduces the costs of actively using a suite of models in the production of forecast analysis.

MAPS supports two broad classes of functionality. ‘Model analysis’ comprises functions to estimate and interrogate the properties of compatible models. The ‘projection system’ provides tools that underpin the ability to construct forecasts, impose judgement and analyse the properties of forecasts produced by compatible models, including COMPASS.

This design enables the staff to construct new functions and capabilities more quickly and, as shown in Figure 1, independently of the user interface, EASE. This provides the flexibility to conduct new experiments to support forecast analysis and to extend the toolkit in the future.

²⁰COMPASS is perhaps closest in structure to RAMSES model developed at the Riksbank (Adolfson et al., 2007) and the ECB’s New Area Wide Model (Christoffel et al., 2008).

²¹By contrast, the infrastructure supporting the MTMM and BEQM worked solely with those models.

3.4 EASE

EASE is the *Economic Analysis & Simulation Environment*, a new IT user interface. EASE has been designed to provide very straightforward access to the new models and tools at our disposal. Importantly, EASE allows users to apply the same tools to COMPASS and suite models, via a common user interface and the results from these models can be easily charted and compared. EASE supports the staff's workflow in the production of inputs to key MPC meetings, increasing the efficiency of standard forecast operations. This frees up staff time to analyse the key economic questions posed by the forecast.

4 COMPASS

COMPASS is the *Central Organising Model for Projection Analysis & Scenario Simulation*. As its name suggests it plays three key roles. First, as described in Section 2.3, COMPASS sits at the centre of the forecasting platform and is used as the main organising framework for constructing the MPC's projections. Second, COMPASS can be used as a tool for analysing those projections in terms of the most likely pattern of shocks that supports them. Third, COMPASS can be used to examine the effects of alternative sets of assumptions (or 'scenarios') on those projections, including alternative assumptions about the response of monetary policy.

In order to fulfil these roles efficiently, COMPASS is designed to be relatively small and simple. As emphasised in Section 5, COMPASS therefore abstracts from a wide range of important economic mechanisms, which we incorporate using insights from the suite of models. Most macroeconomic models in use at central banks evolve over time and COMPASS will be no exception. There are two reasons for that. First, the model parameters will be re-estimated in the light of new and revised data series on a regular basis (most likely once per year). Second, the structure of the model itself is likely to evolve as we learn more about its performance.²²

The rest of this section is organised as follows. Section 4.1 discusses the high level features of the modelling approach, describing the generic features of DSGE models and how these models can be related to macroeconomic data. Section 4.2 provides a high-level description of COMPASS, focusing on the key parts of the model in turn. Of course, because COMPASS is a general equilibrium model, its overall behaviour is a result of the interaction of all components.

A more detailed derivation, with a comprehensive list of all model equations, is provided in Appendix A.²³ In Section 4.3 we describe how the parameters of the model are estimated using Bayesian methods. We present a summary of some of the model's properties in Section 4.4. A more comprehensive discussion of model properties can be found in Appendix B.²⁴

4.1 The general modelling approach

In this section we explain the general features of our modelling approach. Our approach is relatively standard in the context of recently developed models within the same class of dynamic stochastic general equilibrium models. The model is a system of behavioural equations that derive from decisions made by optimising economic agents (for example, households and firms). Agents' optimisation problems are typically dynamic (for example, a household chooses consumption to satisfy a lifetime or intertemporal budget constraint), which means that expectations of future outcomes have important implications for agents' decisions. More specifically, the model consists of a set of first order conditions from agents' optimisation problems together with the budget constraints and market clearing conditions that define an equilibrium of the model as a whole. The behaviour of the model is determined by parameters describing preferences, technologies

²²For example, the Model Development Team in the Monetary Analysis area of the Bank (responsible for maintenance and development of the forecasting platform) intends to investigate alternative specifications for the world block of COMPASS as part of the planned autumn 2013 re-estimation.

²³Available at: <http://www.bankofengland.co.uk/publications/Pages/workingpapers/2013/wp471.aspx>

²⁴Available at: <http://www.bankofengland.co.uk/publications/Pages/workingpapers/2013/wp471.aspx>

and constraints. Frictions such as sticky prices are introduced through the appropriate specification of objective functions and constraints. Because the model is derived from explicit optimisation problems, the number of parameters is relatively small compared with the number of variables.

The model is stochastic in the sense that exogenous random shocks to preferences, technologies and constraints will affect agents' decisions. In the absence of shocks, the model settles on a balanced growth path where all variables are growing at constant, but possibly different, rates, reflecting exogenous population and technology trends.²⁵ Shocks push the variables in the model away from the balanced growth path temporarily, with the speed at which they return to the balanced growth path governed by the persistence of the shocks and the strength of the model's propagation mechanisms (which in turn depends on the specific frictions in the model).

As noted, the intertemporal nature of the optimisation problems means that expectations of future events can have important effects on current decisions. The default treatment of expectations in COMPASS follows the conventional approach of assuming that expectations are 'model consistent' (often also referred to as 'rational expectations'). This means that agents' expectations of the future paths of all variables coincide with the future paths of those variables produced by COMPASS in the absence of future unanticipated shocks. While this standard assumption represents a convenient benchmark, it has some very strong implications (for example, agents' forecast errors are uncorrelated with actual out-turns). The MAPS toolkit described in Section 6.2.6 includes tools for analysing versions of COMPASS in which expectations are formed by alternative means.

We follow a conventional solution approach and solve a log-linear approximation to the model equations. To do this, we first de-trend the variables in the model by scaling them relative to the exogenous processes that generate growth in population and technology. This delivers a set of model equations in terms of stationary (de-trended) variables. In the absence of shocks, these stationary variables will return to a steady state characterised by a constant value for those variables. We approximate the model equations by taking a log-linear approximation around this steady state. A set of 'measurement equations' are used to relate the stationary model variables to observable data. We discuss these equations in more detail in Section 4.3.1.

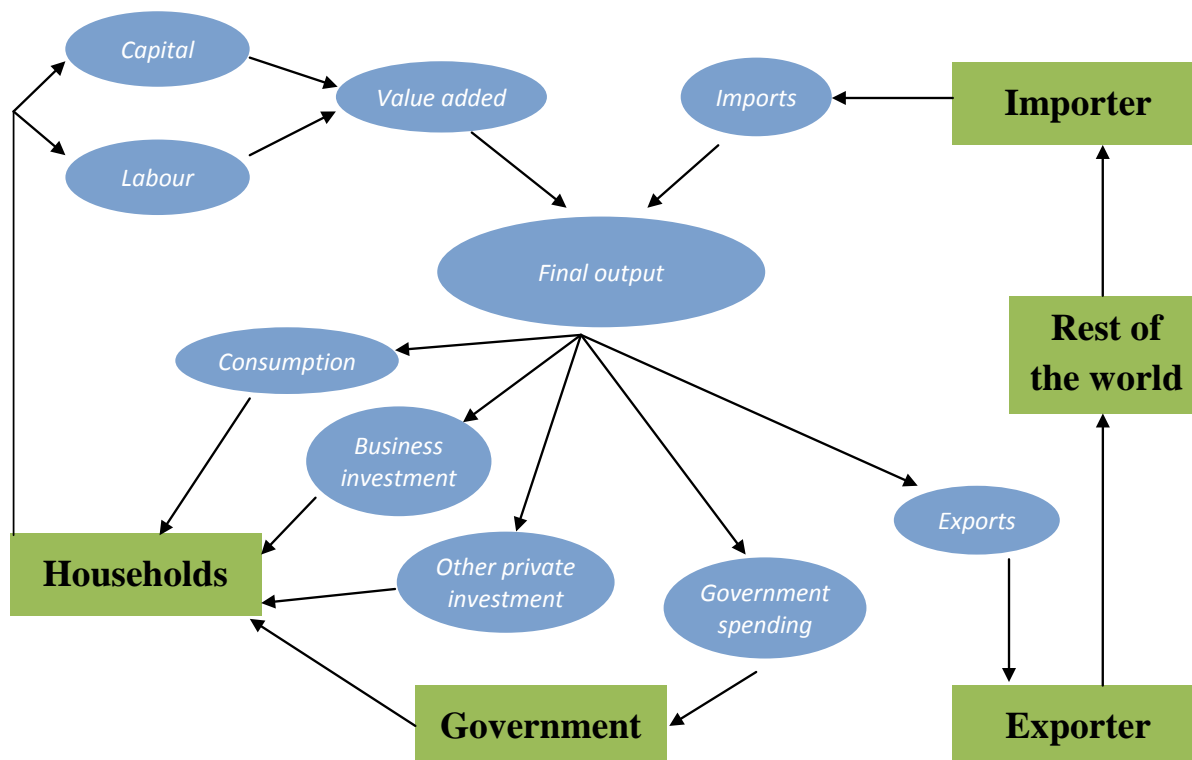
4.2 The model

COMPASS comprises five types of economic agents: households, firms, the government, the rest of the world and the monetary policy maker. A very stylised representation of the interactions between the different sectors is shown in Figure 2, where boxes represent sectors, and the arrows represent flows of goods and services around the economy.

In the remainder of this section, we discuss the key equations that describe these interactions. As explained above, COMPASS is log-linearised around stationary variables which are measured relative to the balanced growth path. The equations in this section are therefore written in terms of log-deviations of the relevant variables from their steady state values. For example c_t denotes the log-deviation of consumption from steady state at date t . This is defined as $c_t \equiv \log C_t - \log C$ where C_t is the level of de-trended consumption and C is the steady state value of de-trended consumption. The detrended

²⁵Some variables may have a growth rate of zero (so that they are stationary) or negative (for example, the relative price of a particular expenditure component may fall over time reflecting higher productivity growth in that sector).

Figure 2: Overview of COMPASS



value of consumption is defined as $C_t \equiv \frac{\tilde{C}_t}{\tilde{\chi}_t^Z}$, where \tilde{C}_t is the level of consumption at date t and $\tilde{\chi}_t^Z$ is the stochastic trend of consumption.²⁶

4.2.1 Supply

The production function for domestic value added (GDP) is a Cobb-Douglas production function:

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

where v is value added, l is labour input (total hours worked), k is the stock of capital and $\hat{\varepsilon}^{TFP}$ denotes exogenous movements in total factor productivity. The assumption that output is produced using a Cobb-Douglas production function implies that it is optimal for firms to keep the expenditure shares of labour and capital in value added (or GDP) at a constant level in the long run. The parameter $0 < \alpha_L < 1$ is the share of labour in the production function.²⁷

The efficiency of value-added firms can fluctuate due to *temporary* changes in total factor productivity, $\hat{\varepsilon}_t^{TFP}$. The efficiency of labour in production of value-added is increasing over time due to exogenous technological progress. This gives rise to a

²⁶See Section A.2 & Section A.4 in Appendix A for full details of the de-trending and log-linearisation process.

²⁷Previous (unpublished) versions of COMPASS also contained variable capital utilisation (with associated adjustment costs). We experimented with two forms: one in which higher capital utilisation is costly because of higher depreciation of physical capital (or ‘wear and tear’ – see Harrison et al. (2005) for an example of this approach); and another in which capital utilisation decisions impose a direct cost on the resources available for consumption and other uses (see, for example, Christiano et al. (2005) and Smets and Wouters (2007) for examples of this approach). Under both specifications, the model’s estimated dynamics were similar with and without variable capital utilisation so, consistent with our general design principle of tractability, we dropped it from the model.

common stochastic trend $\tilde{\chi}_t^Z$, the logarithm of which follows a unit root with drift, $\log \tilde{\chi}_t^Z = \log \Gamma^Z + \log \tilde{\chi}_{t-1}^Z + \gamma_t^Z$, where γ_t^Z is stochastic (see Section 4.2.6 for a definition). Note that $\tilde{\chi}_t^Z$ does not appear in the log-linearised equation because the model variables are expressed in detrended percentage (or log) deviations from the model's balanced growth path: they have been de-trended by the technology term.

Firms hire labour and capital to produce value added output. Cost minimisation gives rise to a demand for capital relative to labour input as a function of the relative factor prices:

$$k_{t-1} - \gamma_t^Z - l_t = w_t - r_t^K \quad (2)$$

where w is the real wage and r^K is the price that firms pay to rent one unit of capital services.²⁸ A rise in the rental price of capital relative to the wage reduces the demand for capital services. Cost minimisation also implies that the real marginal cost of producing a unit of value added is a weighted average of the factor prices:²⁹

$$mc_t^V = (1 - \alpha_L) r_t^K + \alpha_L w_t - \hat{\varepsilon}_t^{TFP} \quad (3)$$

Domestic value added is combined with imports by final output producers to produce final output. The production function for final output producers is also Cobb-Douglas:

$$z_t = \alpha_V v_t + (1 - \alpha_V) m_t \quad (4)$$

where z is final output and m are imported intermediates. Cost minimisation by final output producers gives rise to an import demand function that depends negatively on the price of imports relative to domestic value added prices:

$$m_t - v_t = p_t^V - p_t^M - \hat{\varepsilon}_t^M \quad (5)$$

where the demand for imports includes a stochastic disturbance, $\hat{\varepsilon}_t^M$, which shifts the relative demand for imports for reasons unrelated to the price of imports relative to value added.

Final output can be used for consumption (by households and government), investment or exports. This means that:

$$\tilde{Z}_t = \tilde{Z}_t^C + \tilde{Z}_t^G + \tilde{Z}_t^I + \tilde{Z}_t^{I^O} + \tilde{Z}_t^X \quad (6)$$

where \tilde{Z} is the total production of final output and \tilde{Z}^C , \tilde{Z}^G , \tilde{Z}^I , \tilde{Z}^{I^O} , and \tilde{Z}^X represent the quantities of final output allocated towards household consumption, government consumption, capital investment, 'other' investment and exports, respectively. The use of the '~' character above each variable indicates that the variable has not been detrended.

Perfectly competitive retailers buy the final-output good from final-output producers, and transform it into consumption, investment, government spending or export goods using a simple linear technology. For example, for consumption goods:

$$\tilde{C}_t = \tilde{\chi}_t^C \tilde{Z}_t^C \quad (7)$$

²⁸All prices in the stationary version of the model are *relative* prices (and are measured relative to the price of final output). The real wage is also detrended by the level of labour augmenting productivity (to ensure that it is stationary).

²⁹Given the Cobb-Douglas nature of the production function, the marginal cost of value-added can also be expressed as the labour share in value-added, $mc_t^V = w_t + l_t - v_t - p_t^V$.

where \tilde{C} denotes (non-detrended) consumption and $\tilde{\chi}_t^C$ captures productivity in the consumption retail sector. There are analogous equations for investment, government and export goods which permit different productivities in the different sectors. The $\tilde{\chi}_t^C$, $\tilde{\chi}_t^G$, $\tilde{\chi}_t^I$, $\tilde{\chi}_t^{IO}$ and $\tilde{\chi}_t^X$ variables all follow deterministic trends.³⁰

This approach permits different trend growth rates across sectors without requiring an explicit treatment of a multi-sectoral supply side. When equation (6) is detrended and log-linearised the resulting equation is:

$$z_t = \omega_{CZ}c_t + \omega_{GZ}g_t + \omega_{IZ}i_t + \omega_{IO}i_t^O + \omega_{XZ}x_t \quad (8)$$

where ω_{CZ} denotes the steady-state share of consumption expenditure in final output, and so on.

4.2.2 Price and wage setting

Firms in the value added and final output sectors operate under monopolistic competition. They set their prices as a mark-up over the marginal cost of production. It is costly to change prices, and firms take these adjustment costs into account when setting prices. This means that the mark-up depends on the current and expected future changes in the inflation rate.

These assumptions give rise to pricing equations of the form:

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

and:

$$\pi_t^Z = \hat{\mu}_t^Z + \frac{1}{\phi_Z(1 + \beta\Gamma^H\xi_Z)}mc_t^Z + \frac{\xi_Z}{1 + \beta\Gamma^H\xi_Z}\pi_{t-1}^Z + \frac{\beta\Gamma^H}{1 + \beta\Gamma^H\xi_Z}\mathbb{E}_t\pi_{t+1}^Z \quad (10)$$

where π_t^V denotes inflation in value-added prices and π_t^Z is inflation in final output prices. Expected inflation one-period ahead is denoted by $\mathbb{E}_t\pi_{t+1}^V$ and $\mathbb{E}_t\pi_{t+1}^Z$. Exogenous fluctuations in producers' desired margins are captured by the shocks $\hat{\mu}_t^V$ and $\hat{\mu}_t^Z$. The shocks to desired margins are assumed to be entirely transitory (they have no persistence at all) and are intended to proxy for effects such as temporary changes in market conditions or taxes applied to spending or production.³¹

The generic form of the two pricing equations is the same, reflecting the fact that the same type of nominal rigidities affect both types of producer. The pricing equations show that inflation tends to rise when real marginal costs rise (as firms pass on higher costs in the form of price increases) or when expected inflation rises (as firms raise their price today anticipating future price rises), or when past inflation rises (since price adjustment costs are assumed to depend on past inflation). The sensitivity of current inflation to these variables will depend on the parameters that govern the nominal rigidities. For instance, if the parameter that governs the degree of indexation to past value-added inflation is set to zero ($\xi_V = 0$), the equation reduces to a purely forward-looking equation for value added pricing, where past inflation does not influence current price-setting. And value-added firms will tend to increase prices by less in the face of higher costs if the parameter governing the extent of nominal rigidities, ϕ_V , is larger.

³⁰Note that the productivity trend in the consumption retail sector is normalised so that $\tilde{\chi}_t^C \equiv \tilde{\chi}_t^Z$ and that imports share the same trend as exports. See A.2 in Appendix A for discussion.

³¹The terms β and Γ^H represent households' discount factor and population growth respectively, where $\beta\Gamma^H \simeq 1$.

Final output inflation and CPI inflation are equivalent, because consumption taxes such as VAT are not included in the model.³² Since marginal costs in the final output sector are a weighted average of imported and value added (domestic) prices, CPI inflation is a function of both domestic and imported prices.

Importers purchase output produced overseas and sell it to domestic producers. Import prices are set in domestic currency as a mark-up over the cost that importers pay for the world export good on world markets. Again, we assume that prices are costly to adjust so that the mark-up of import prices over costs depends on changes in the rate of import price inflation:

$$\pi_t^M = \hat{\mu}_t^M + \frac{p_t^{XF} - q_t - p_t^M}{\phi_M (1 + \beta\Gamma^H \xi_M)} + \frac{\xi_M}{1 + \beta\Gamma^H \xi_M} \pi_{t-1}^M + \frac{\beta\Gamma^H}{1 + \beta\Gamma^H \xi_M} \mathbb{E}_t \pi_{t+1}^M \quad (11)$$

where π^M is import price inflation, p^{XF} is the foreign currency price of world exports and q is the real exchange rate, so that $p^{XF} - q - p^M$ represents the domestic currency price of world exports relative to imports.³³ An exchange rate appreciation – which corresponds to a *rise* in q – tends to lead to a fall in the domestic-currency price of world export goods, which in turn tends to reduce import price inflation. Price adjustment costs means that, in the short run, the pass-through from exchange rate movements to the price of imports is incomplete, but as prices adjust, pass-through is complete in the long run.

Exporters are monopolistically competitive and have some price-setting power on world markets. Export prices are set in foreign currency as a mark-up over exporters' marginal costs, given by the price of domestic final output expressed in foreign currency, subject to costs of adjusting their prices. The price adjustment costs imply that the mark-up is a function of changes in export price inflation:

$$\pi_t^{EXP} = \hat{\mu}_t^X + \frac{q_t - p_t^{EXP}}{\phi_X (1 + \beta\Gamma^H \xi_X)} + \frac{\xi_X}{1 + \beta\Gamma^H \xi_X} \pi_{t-1}^{EXP} + \frac{\beta\Gamma^H}{1 + \beta\Gamma^H \xi_X} \mathbb{E}_t \pi_{t+1}^{EXP} \quad (12)$$

where π_t^{EXP} is export price inflation expressed in foreign currency, q is the real exchange rate, and p_t^{EXP} is the foreign currency relative export price. An appreciation of the domestic exchange rate leads to a rise in marginal cost measured in foreign currency. Analogous to the treatment of import prices, the pass-through from exchange rate movements to export prices expressed in foreign currency is incomplete in the short run, due to the price adjustment costs, but is complete in the long run.

Similar costs are assumed to apply to adjusting nominal wages, which gives rise to a labour supply relationship in which wages are set as a mark-up over the expected cost of supplying labour. The mark-up depends on the costs of adjusting wage inflation:

$$\pi_t^W = \hat{\mu}_t^W + \frac{\hat{\epsilon}_t^L + \epsilon_L l_t + \frac{\epsilon_C (c_t^\circ - \psi_C c_{t-1}^\circ)}{1 - \psi_C} - w_t}{\phi_W (1 + \beta\Gamma^H \xi_W)} + \frac{\xi^W}{1 + \beta\Gamma^H \xi^W} \pi_{t-1}^W + \frac{\beta\Gamma^H}{1 + \beta\Gamma^H \xi^W} \mathbb{E}_t \pi_{t+1}^W \quad (13)$$

where π_t^W denotes wage inflation and w is the real wage measured in terms of final output (or, equivalently, consumption). The cost of supplying labour is increasing in the amount

³²Section 8.2 uses an extended version of COMPASS that includes a consumption tax rate to show how the effects of changes in VAT can be incorporated into a forecast built using COMPASS.

³³It is assumed that p^{XF} measures the world price of exported goods relative to the world price of final output goods.

of labour supplied, since households value leisure.³⁴ As the cost of supplying labour increases relative to the real wage, a higher wage is demanded, which tends to increase wage inflation.

4.2.3 Private domestic demand

The model includes two types of households. ‘Constrained’ or ‘rule-of-thumb’ households have no access to financial markets and cannot save: they spend their current labour income on consumption. The remaining ‘unconstrained’ or ‘optimising’ households can accumulate assets, which allows them to smooth their consumption over time. Unconstrained households also dislike changing the level of consumption abruptly over time (so-called ‘habit formation’). This means that they smooth the rate of change of consumption as well as the level, giving rise to more inertia in consumption. These unconstrained households save in two types of assets: physical capital, rented out to firms, and deposits, which are invested in domestic bond and foreign bonds by a ‘portfolio packager’ on households’ behalf. The portfolio packagers pay households a stochastic return, funded by nominal returns yielded on domestic and foreign bonds.³⁵

Unconstrained households choose a consumption path to maximise lifetime utility. The consumption path is optimal in the sense that the marginal utility lost by giving up a unit of consumption in the current quarter is equal to the expected marginal utility gained by consuming the proceeds of the additional saving in the following quarter. When combined with the consumption of constrained households, this optimality condition can be represented in terms of an ‘Euler equation’ for aggregate consumption:³⁶

$$\begin{aligned}
c_t = & \frac{1}{1 + \psi_C + \epsilon_\beta(1 - \psi_C)\epsilon_C^{-1}} [\mathbb{E}_t c_{t+1} + \psi_C c_{t-1}] \\
& - \frac{\omega_o(1 - \psi_C)}{(1 + \psi_C)\epsilon_C + \epsilon_\beta(1 - \psi_C)} [r_t - \mathbb{E}_t \pi_{t+1}^Z + \hat{\epsilon}_t^B - \mathbb{E}_t \gamma_{t+1}^Z] \\
& + (1 - \omega_o) \frac{wL}{C} \left[w_t + l_t - \frac{\mathbb{E}_t w_{t+1} + \mathbb{E}_t l_{t+1} + \psi_C (w_{t-1} + l_{t-1})}{1 + \psi_C + \epsilon_\beta(1 - \psi_C)\epsilon_C^{-1}} \right] \quad (14)
\end{aligned}$$

The first line of equation (14) shows that consumption is a function of expected future consumption and lagged consumption. Expected future consumption matters because unconstrained households are forward looking. Lagged consumption matters when the parameter governing habit formation $0 \leq \psi_C < 1$ is strictly positive.³⁷ The second line shows that consumption is sensitive to the real interest rate, $r_t - \mathbb{E}_t \pi_{t+1}^Z$, which represents the (opportunity) cost of current consumption relative to one-period ahead consumption (i.e. the cost of borrowing or the return to saving) adjusted for the effects

³⁴The return to supplying labour depends negatively on the level of consumption, since leisure and consumption are substitutes and the marginal utility of consumption decreases as consumption rises. The relevant measure of consumption is that of ‘optimising’ households, c^o , discussed in more detail in Section 4.2.3.

³⁵We assume that households hold the capital stock directly rather than claims to the profits of firms that accumulate capital (ie equity). In the absence of frictions that cause the firm and household to value capital differently, the behaviour of the model is identical under both assumptions. Assuming that households hold the capital stock directly is the simpler approach.

³⁶This equation is derived by combining equations (A.276), (A.287) and (A.288) from Appendix A.

³⁷If $\psi_C = 0$, lagged consumption does not have a direct effect on current consumption.

of a risk premium shock, $\hat{\varepsilon}_t^B$.^{38,39} The effect of changes in the adjusted real interest rate on consumption is larger if the share of unconstrained households in the economy, $0 < \omega_o \leq 1$, is higher and if the elasticity of intertemporal substitution, $\epsilon_C^{-1} > 0$, is larger.⁴⁰ The third line of equation (14) shows that consumption is influenced by contemporaneous labour income, with the importance of these effects increasing with the share of constrained households $(1 - \omega_o)$.

As described above, capital is an input to value-added production. The value of the capital stock is determined by the discounted value of the stream of returns that one unit of capital is expected to generate:

$$tq_t = \frac{1 - \delta_K}{r^K + (1 - \delta_K)} \mathbb{E}_t tq_{t+1} - (r_t - \mathbb{E}_t \pi_{t+1}^Z + \hat{\varepsilon}_t^B) + \frac{r^K}{r^K + (1 - \delta_K)} \mathbb{E}_t r_{t+1}^K \quad (15)$$

where tq_t is ‘‘Tobin’s Q’’, the value of one unit of capital, $r_t - \mathbb{E}_t \pi_{t+1}^Z + \hat{\varepsilon}_t^B$ is the real interest rate adjusted for the risk premium shock, $\mathbb{E}_t r_{t+1}^K$ is the expected rental rate of capital and δ_K is the rate at which capital depreciates.

Investment in physical capital is subject to adjustment costs. So, while additional investment is induced by increases in the value of capital, tq , this effect is tempered by the desire to smooth the change in investment:

$$i_t = \frac{\beta \Gamma^H}{1 + \beta \Gamma^H} (i_{t+1} + \gamma_{t+1}^Z) + \frac{1}{1 + \beta \Gamma^H} (i_{t-1} - \gamma_t^Z) + \frac{1}{(1 + \beta \Gamma^H) (\Gamma^H \Gamma^Z \Gamma^I)^2} \left(\frac{tq_t}{\psi_I} + \hat{\varepsilon}_t^I \right) \quad (16)$$

where i is investment, $\hat{\varepsilon}_t^I$ is a shock to the investment adjustment cost function and $\psi_I \geq 0$ is the parameter determining the size of investment adjustment costs. As $\psi_I \rightarrow 0$, the investment rate becomes closer to being costlessly adjustable. The term $(\Gamma^H \Gamma^Z \Gamma^I)^2$ appears because $\Gamma^H \Gamma^Z \Gamma^I$ is the steady state growth rate of investment and the adjustment cost is specified so that the cost of adjusting investment is quadratic in the deviation of investment from its trend growth rate.⁴¹

4.2.4 Interactions with the rest of the world

Exporters supply whatever quantity of exports is demanded by the rest of the world. Export demand is a simple function of world output and the price of exports relative to the world export price:

$$x_t = z_t^F + \hat{\varepsilon}_t^{\kappa^F} - \epsilon_F \left(p_t^{EXP} - p_t^{XF} \right) \quad (17)$$

³⁸An adjustment is also made for the expected one-period ahead growth of the trend in final output (or consumption), $\mathbb{E}_t \gamma_{t+1}^Z$. Given the estimation described in Section 4.3, this term is equal to zero.

³⁹The risk-free real interest rate and the rate that households actually face. However, it is purely exogenous and does not arise endogenously as the consequence of a financial friction – see Section 8.3 for a description of a suite model that does contain financial frictions.

⁴⁰The equation also contains a parameter ϵ_β , which determines the degree of ‘over-discounting’ in the model. This mechanism is included to ensure that the net foreign asset position in the model is stationary.

⁴¹And the presence of β and Γ^H reflects the same discounting as in the Phillips curves described above – see footnote 31 on page 16.

where x is exports, z^F is world output, p^{EXP} is domestic export prices expressed in foreign currency and p^{XF} is world export prices. The exogenous process $\hat{\varepsilon}_t^{KF}$ represents shocks to foreign preferences for domestically-produced exports, so that export demand may fluctuate for reasons that are not related to the overall level of world demand or the relative price of domestic exports.

Optimal portfolio allocation between foreign and domestic bonds specifies that the returns on domestic and foreign bonds are equalised when measured in a common currency. So the real exchange rate satisfies an uncovered interest parity (UIP) condition:

$$q_t = \mathbb{E}_t q_{t+1} + (r_t - \mathbb{E}_t \pi_{t+1}^Z) - \hat{\varepsilon}_t^{BF} \quad (18)$$

where q_t is the real exchange rate (measured as the price of domestic consumption relative to that of foreign consumption, so that an increase represents a domestic *appreciation*), $r_t - \mathbb{E}_t \pi_{t+1}^Z$ is the *ex ante* real interest rate⁴² and $\hat{\varepsilon}_t^{BF}$ is an exogenous shock to the UIP condition, capturing, for example, movements in the real interest rate in the rest of the world (which is assumed constant in the model), or in the risk premium.⁴³

The domestic economy is assumed to be small with respect to the rest of the world, in the sense that world output and world prices are not affected by developments in the domestic economy. They are modelled as simple exogenous AR(1) processes:⁴⁴

$$z_t^F = \omega_t^F + \rho_{ZF} z_{t-1}^F + \hat{\varepsilon}_t^{ZF} \quad (19)$$

$$p_t^{XF} = \rho_{PXF} p_{t-1}^{XF} + \hat{\varepsilon}_t^{PXF} \quad (20)$$

where ω^F is a term that controls the speed at which the rest of the world (and hence demand for domestic exports) inherits permanent labour augmenting productivity (LAP) shocks realised in the domestic economy. Cointegration of world output with domestic output is a condition that ensures balanced growth in the model. The ω^F term is included as a device to decouple the speed with which world output responds to LAP shocks from the persistence of world output shocks.⁴⁵

4.2.5 Fiscal and monetary policy

Real government expenditure is assumed to follow a simple autoregressive process:

$$g_t - g_{t-1} + \gamma_t^Z = (\rho_G - 1) g_{t-1} + \hat{\varepsilon}_t^G \quad (21)$$

⁴²As detailed in the section beginning on page A7 of Appendix A the domestic risk premium shock arises in the returns paid to households by a ‘portfolio packager’. The returns received by that portfolio packager on their investments in domestic and foreign bonds are not subject to this shock and so it does not feature in the UIP condition.

⁴³Earlier (unpublished) versions of COMPASS contained a net foreign asset term in the UIP condition. This assumption is commonly used to ensure that the model returns to a unique steady-state net foreign asset position following temporary shocks. As mentioned in footnote 40 on page 19, the version of COMPASS in this paper instead uses over-discounting of consumption (see Schmitt-Grohe and Uribe (2003) for a comparison of this type of approach with several alternatives).

⁴⁴Many small open economy DSGE models model foreign variables as vector autoregressions (VARs) (see, for example, Adolfson et al., 2007; Christoffel et al., 2008; Harrison and Oomen, 2010). We choose a simpler approach, noting that we are likely to explore alternative specifications in future.

⁴⁵It is defined as: $\omega_t^F = -\gamma_t^Z + (1 - \zeta_{\omega_F}) \omega_{t-1}^F$, where $0 < \zeta_{\omega_F} \leq 1$ is the parameter that controls the speed with which world demand ‘inherits’ permanent domestic labour productivity shocks. An alternative (and arguably more natural) approach would be for the domestic economy to inherit permanent productivity shocks from the rest of the world. That would require a much richer structure for the world block of the model in order to be able to identify those shocks.

where g is real government expenditure and ϵ^G captures the effects of exogenous shocks to government spending. The equation is written in an ‘error correction’ format. It states that the growth of government spending (measured in ‘actual’ – i.e. non-detrended – units) responds negatively to the deviation of *de-trended* government spending from steady state. A positive shock to labour augmenting productivity, γ_t^Z , pushes de-trended government spending down in the short run (as government spending does not expand one-for-one with the improvement in productivity). Over the longer term, de-trended government spending rises until it has adjusted to the new, higher level of productivity. The parameter $0 < \rho_G < 1$ controls the speed at which government spending responds to changes in labour augmenting productivity and the persistence of government spending shocks.

Government spending is financed by lump-sum taxes levied on unconstrained households, though in principle the government may run budget deficits or surpluses by increasing or reducing the stock of debt that it issues to private agents. The lump-sum (non-distortionary) tax levied on households is used as the fiscal instrument to ensure that government debt is stabilised. The use of a non-distortionary tax gives rise to a ‘Ricardian equivalence’ result that (for a given path of government spending) the path of government debt (and hence the fiscal rule for the lump-sum tax used to control it) has no effect on the decisions of private agents. This result allows us to simplify the model by assuming that the lump-sum tax is used to balance the government’s budget each period. The simplified specification of fiscal policy in COMPASS means that it is ill-suited for analysis of changes in fiscal policy. Instead, the effects of changes in fiscal policy are quantified and analysed using the suite.⁴⁶

The domestic short-term nominal interest rate is set by the central bank. A simple reaction function – a so-called Taylor rule – is used to specify how the nominal interest rate responds to key macroeconomic variables:

$$r_t = \theta_R r_{t-1} + (1 - \theta_R) \left[\theta_\Pi \left(\frac{1}{4} \sum_{j=0}^3 \pi_{t-j}^Z \right) + \theta_Y \hat{y}_t \right] + \hat{\epsilon}_t^R \quad (22)$$

where r is the quarterly policy rate, π_t^Z is the deviation of quarterly consumer price inflation from target and \hat{y}_t is a measure of the output gap. The reaction function allows for interest rate smoothing so that the current interest rate depends on the rate set in the previous quarter. The measure of the output gap used is the difference between value-added, v_t , and the level of value-added that would prevail if all prices (including wages) were perfectly flexible – a so-called ‘flexible-price’ output gap.⁴⁷ A common alternative approach is to define the output gap as the difference between output and the level of output implied by the value-added production function with inputs measured at ‘trend’ levels. In COMPASS, the two approaches yield very similar results indeed (both in terms of the estimation and the resulting model properties), reflecting that the two measures of the output gap are quantitatively very similar over the estimation sample.⁴⁸

⁴⁶At the time of writing, changes in fiscal policy are incorporated into the MPC’s forecast using a range of empirical estimates or ‘multipliers’, which vary according to the particular fiscal instrument in use. Bank staff are also engaged in building a set of fiscal suite models, which will allow a more structural analysis of fiscal issues in future.

⁴⁷Specifically, the measure of potential output in this case is computed using a version of COMPASS in which all prices and wages are flexible and temporary fluctuations to desired mark-ups are ignored.

⁴⁸The definition of the ‘trend’ level of output used in the production function approach is $(1 - \alpha_L) k_{t-1} + \hat{\epsilon}_t^{TFP}$.

Note that in practice the *Inflation Report* forecasts have been published under the assumption that Bank Rate and government spending follow a measure of market expectations of Bank Rate and the government's published spending plans respectively. These 'conditioning' assumptions and other issues related to changes in monetary policy are discussed in Section 8.4.

4.2.6 Forcing processes and shocks

All movements in the endogenous variables away from the steady state are ultimately caused by exogenous, random shocks. In COMPASS we assume that each shock is drawn from a Gaussian (or 'normal') distribution with a mean of zero and variance equal to 1.⁴⁹ The shocks in the model are uncorrelated with each other and the realisations of a particular shock are uncorrelated over time. Shocks in COMPASS are labelled using the symbol η .

Shocks affect the model via so-called forcing processes, which typically impart a degree of persistence to the effect of the shock. Forcing processes are typically modelled as AR(1) processes, with the symbol ρ used to denote the AR coefficient. We use parameters denoted with the symbol σ to denote the standard deviations of the *forcing processes*. In cases where the forcing processes have persistence, the coefficient on the shock η in each forcing process equation is defined as a function of both a persistence parameter, ρ , and a standard deviation scaling parameter, σ . This ensures that the standard deviation of each forcing process is a function of the relevant σ only.⁵⁰

A complete list of the forcing processes is given below:

$$\begin{aligned}
\hat{\varepsilon}_t^{ZF} &= (1 - \rho_{ZF}^2)^{1/2} \sigma_{ZF} \eta_t^{ZF} \\
\hat{\varepsilon}_t^{PXF} &= (1 - \rho_{PXF}^2)^{1/2} \sigma_{PXF} \eta_t^{PXF} \\
\hat{\varepsilon}_t^B &= \rho_B \hat{\varepsilon}_{t-1}^B + (1 - \rho_B^2)^{1/2} \sigma_B \eta_t^B \\
\hat{\varepsilon}_t^I &= \rho_I \hat{\varepsilon}_{t-1}^I + (1 - \rho_I^2)^{1/2} \sigma_I \eta_t^I \\
\hat{\varepsilon}_t^G &= (1 - \rho_G^2)^{1/2} \sigma_G \eta_t^G \\
\hat{\varepsilon}_t^{\kappa^F} &= \rho_{\kappa^F} \hat{\varepsilon}_{t-1}^{\kappa^F} + (1 - \rho_{\kappa^F}^2)^{1/2} \sigma_{\kappa^F} \eta_t^{\kappa^F} \\
\hat{\varepsilon}_t^M &= \rho_M \hat{\varepsilon}_{t-1}^M + (1 - \rho_M^2)^{1/2} \sigma_M \eta_t^M \\
\hat{\varepsilon}_t^{IO} &= (1 - \rho_{IO}^2)^{1/2} \sigma_{IO} \eta_t^{IO} \\
\hat{\varepsilon}_t^R &= \sigma_R \eta_t^R \\
\hat{\varepsilon}_t^{BF} &= \rho_{BF} \hat{\varepsilon}_{t-1}^{BF} + (1 - \rho_{BF}^2)^{1/2} \sigma_{BF} \eta_t^{BF} \\
\hat{\varepsilon}_t^L &= \rho_L \hat{\varepsilon}_{t-1}^L + (1 - \rho_L^2)^{1/2} \sigma_L \eta_t^L \\
\hat{\varepsilon}_t^{LAP} &= \rho_{LAP} \hat{\varepsilon}_{t-1}^{LAP} + (1 - \rho_{LAP}^2)^{1/2} \sigma_{LAP} \eta_t^{LAP} \\
\hat{\varepsilon}_t^{TFP} &= \rho_{TFP} \hat{\varepsilon}_{t-1}^{TFP} + (1 - \rho_{TFP}^2)^{1/2} \sigma_{TFP} \eta_t^{TFP}
\end{aligned}$$

⁴⁹See Section 6.2.1 for a discussion of the general modelling framework.

⁵⁰This is useful for separate identification of the standard deviations and persistence parameters in the estimation, as well as in setting the priors for the estimation.

$$\begin{aligned}
\hat{\mu}_t^X &= \sigma_{\mu^X} \eta_t^{\mu^X} \\
\hat{\mu}_t^M &= \sigma_{\mu^M} \eta_t^{\mu^M} \\
\hat{\mu}_t^W &= \sigma_{\mu^W} \eta_t^{\mu^W} \\
\hat{\mu}_t^Z &= \sigma_{\mu^Z} \eta_t^{\mu^Z} \\
\hat{\mu}_t^V &= \sigma_{\mu^V} \eta_t^{\mu^V}
\end{aligned}$$

Note that forcing processes are denoted $\hat{\varepsilon}$ with the exception of forcing processes for firms' desired price mark-ups, which are denoted $\hat{\mu}$ and which have no persistence reflecting a strong prior that desired mark-ups should not fluctuate persistently away from their long-run levels. Note also that the value-added production function in equation (1) implies that growth in the unit root process for trend technological progress can be defined as: $\gamma_t^Z = \alpha_L \hat{\varepsilon}_t^{LAP} + (1 - \alpha_L) \gamma_{t-1}^Z$.

4.3 Estimation

We use Bayesian maximum likelihood techniques to estimate the parameters of the model, using the MAPS toolkit described in Section 6.2.2.⁵¹ As this approach has become commonplace in the estimation of large-scale DSGE models, including those recently developed at central banks, our discussion of the methodology is brief.⁵²

We denote the vector of parameters to be estimated as θ^{est} . Estimation of these parameters by Bayesian maximum likelihood proceeds in two steps. First, we specify prior distributions for the parameters. This prior information can be combined with the value of the likelihood function implied by the the model and the available data to form the posterior. We use the notation Y_T to denote the data for the set of macroeconomic time series used to estimate the model (described in Section 4.3.1). The likelihood function can be evaluated using the Kalman filter and the (linear) state space representation of the rational expectations solution of the DSGE model.

From Bayes' rule, the posterior distribution satisfies:

$$p(\theta^{est} | Y_T) \propto p(\theta^{est}) p(Y_T | \theta^{est}) \quad (23)$$

where $p(\theta^{est})$ represents the prior distribution for the estimation parameters and $p(Y_T | \theta^{est})$ is the value of the likelihood function conditional on a particular value for the parameter vector. The modal parameter vector can be found by numerically maximising the right-hand side of (23). To approximate the posterior distribution, we use Markov Chain Monte Carlo methods (specifically, the random walk Metropolis-Hastings algorithm) to generate simulated draws from the posterior density of θ^{est} . Specifically, we simulate 12 separate chains of 1.25 million draws, from which we discard the first 250,000 draws and then 'thin' the remainder to leave 100,000 draws in each of the 12 chains. We then combine those chains to deliver a simulated posterior distribution of 1.2 million draws. In order to obtain point estimates of the parameters for the central organising model used to produce the forecast, we compute the mean values of the parameters across that combined set of draws.

⁵¹More specifically, the model is comprised of the equations outlined in Section A.4 of Appendix A and the data mapping described below.

⁵²An and Schorfheide (2007) provide an excellent review of these methods applied to DSGE models.

4.3.1 Data and measurement equations

To estimate the model, we need to specify how the variables in the model are related to UK data. The fact that COMPASS equations are written in terms of de-trended variables means that we need to account for this de-trending process. This sub-section describes the data we use and explains how those data are related to COMPASS variables.

We use data for fifteen macroeconomic variables (described below).⁵³ The sample period is 1993Q1–2007Q4.⁵⁴ This is a relatively short sample, but is motivated by a desire to ensure that the estimation period excludes large changes in monetary policy regime. An alternative approach would be to use a longer sample period and account for the changes in monetary regime, though this creates other issues.⁵⁵ We exclude the recent financial crisis from our sample period to avoid this episode having a disproportionate effect on the properties of the data and to avoid using the most recent data that is subject to larger revisions.

The data series that feed into the model are called ‘raw observables’, because they represent the raw data prior to any transformation or de-trending.⁵⁶ These raw observables are transformed into appropriate units using a set of ‘data transformation equations’, which de-trend the raw observables in the appropriate way to convert them into units that correspond to the units in which COMPASS variables are measured (typically 100 times a logarithmic deviation from steady state). The raw observables, their data sources and data transformation equations are shown in Table 1.

The information in Table 1 shows how raw data is mapped to model observable units. As an example consider the data transformation equation for the import price deflator:

$$\text{dlnpmdef}_t = 100\Delta \ln \text{pmdef}_t - \Pi_t^{*,\text{tt}} - \Pi_t^{m,\text{tt}} \quad (24)$$

This equation defines a variable dlnpmdef_t that can be mapped into COMPASS variables measured as (100 times) the logarithm of the stationary variable. To arrive at dlnpmdef_t , we take the first difference of the logarithm of the raw data for the import deflator and multiply it by 100.⁵⁷ Then we subtract two ‘time-varying trends’, which are introduced to correct for deterministic movements in the trend of the import deflator over our sample that are not captured by the trends within the model itself.

In this example, the time-varying trends correct for the low-frequency decline in CPI inflation associated with the transition to inflation targeting over the early (and in the training) part of our sample, $\Pi_t^{*,\text{tt}}$, and the fact that import prices have fallen less rapidly

⁵³As is common in models like COMPASS, the number of data series used (treated as ‘observable’) is much smaller than the total number of endogenous variables in the model. In part, this reflects the fact that some of the equations in DSGE models are identities (for example, a market clearing condition specifying that total expenditure equals total output). But it also reflects that many of the endogenous variables do not have a measured counterpart in the data (eg TFP).

⁵⁴Data from 1987Q3–1992Q4 is used as a ‘training sample’ to initialise the Kalman filter.

⁵⁵Harrison and Oomen (2010) use UK data from the 1960s and account for regime shifts by using deterministic breaks in the trends for nominal variables. That exercise demonstrates the difficulties involved in producing a set of consistently measured UK macroeconomic data, even for a small number of series. Recent work by Cúrdia and Finocchiaro (2013) proposes a method for incorporating regime shifts explicitly into the estimation of structural models.

⁵⁶The term ‘observables’ comes from the state space modelling literature. We use state space methods to estimate the model parameters using the Kalman filter. For more details, see Section 6.2.2.

⁵⁷Scaling by 100 is simply a normalisation. It means that the units of variables in COMPASS correspond (approximately) to percentage deviations from steady state, which is useful as a starting point in analysis of the model’s outputs when preparing the forecast.

than predicted by the supply side structure of the model, $\Pi_t^{m,tt}$. The CPI inflation target trend is calculated relative to the current inflation target of 2% per annum for CPI inflation, so that $\Pi_t^{*,tt} = 0$ over most of our sample.⁵⁸ The trend $\Pi_t^{m,tt}$ is computed as the average difference between the observed growth rate of import prices in the data and the rate predicted by the supply side structure of the model adjusted for the CPI inflation trend.⁵⁹ The same approach is used for the other time-varying trends in Table 1. That is, they are used to ensure that the mean of the transformed observables over the estimation sample is consistent with the assumptions embodied in the model. These trends are: $\Pi_t^{x,tt}$, dlnxkp_t^{tt} , dlnmkp_t^{tt} , $\Pi_t^{x,tt}$ & dlnyf_t^{tt} . The underlying economic rationale for these time-varying trends is that it is difficult to match the secular increase in global trade as a share of activity using standard supply side assumptions. This is a feature of the data in many countries (a similar de-trending approach is taken by Adolfson et al. (2007) using euro area data).

Having defined observable data in appropriate units, a set of ‘measurement equations’ is used to relate the transformed observables to COMPASS variables. These are also shown in Table 1. For example, the measurement equation associated with import price inflation, dlnpmdef_t is:

$$\text{dlnpmdef}_t = \pi_t^M + 100 \ln \frac{\Pi^*}{\Gamma^X} + \sigma_{me,pm} me_t^{pm} \quad (25)$$

which tells us that the detrended (log difference of the) import deflator is mapped into the model variable π_t^M (the deviation of quarterly import price inflation from steady state) plus the steady-state trend growth in import prices implied by the model, plus an *iid* measurement error. The steady-state growth of import prices in the model is determined by the inflation target, Π^* , and the relative deterministic growth rate of technological progress in the import (and export) sector, Γ^X (see the discussion in Section A.2 of Appendix A for more details). Measurement error, me_t^{pm} , is included to cope with the fact that import prices may be particularly prone to measurement issues.⁶⁰ We also include measurement error for investment growth, wage growth, hours worked, export prices and import and export volumes.

⁵⁸Over the earlier part of our sample, $\Pi_t^{*,tt}$ takes positive values, based on the implicit inflation target implied by official policy announcements (following the approach used in Batini and Nelson (2001)).

⁵⁹The supply-side structure of the model implies that relative import prices should decline at the rate of relative productivity in the export retail sector, Γ^X . Since it is not possible to calibrate simultaneously the trend growth rates of GDP, final output and all the expenditure components, the approach taken is to calibrate the relative trend growth of exports and imports, Γ^X , to deliver target trend growth rates for final output, Γ^Z , and GDP, Γ^V . That means that $\ln \Gamma^X = \frac{\alpha_V}{1-\alpha_V} (\ln \Gamma^Z - \ln \Gamma^V)$.

⁶⁰As is the case for the structural shocks, measurement errors are treated as *iid* with standard normal distributions. The parameter $\sigma_{me,pm}$ is used to scale the measurement error and determines its standard deviation.



Table 1: Observables, data transformation and measurement equations

Variable	Description	ONS code	Data transformation equation	Measurement equation
gdppk	Real GDP	ABMI	$d\ln gdppk_t \equiv 100\Delta \ln gdppk_t$	$d\ln gdppk_t = \Delta v_t + \gamma_t^Z + 100 \ln \left(\Gamma^Z \Gamma^H (\Gamma^X)^{-\frac{1-\alpha_V}{\alpha_V}} \right)$
ckp	Real consumption	ABJR+HAYO	$d\ln ckp_t \equiv 100\Delta \ln ckp_t$	$d\ln ckp_t = \Delta c_t + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H)$
ikkp	Real business investment	NPEL ^a	$d\ln ikkp_t \equiv 100\Delta \ln ikkp_t$	$d\ln ikkp_t = \Delta i_t + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H \Gamma^I) + \sigma_I^{me} me_t^I$
gonskp	Real government spending	NMRY+DLWF ^b	$d\ln gonskp_t \equiv 100\Delta \ln gonskp_t$	$d\ln gonskp_t = \Delta g_t + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H \Gamma^G)$
xkp	Real exports	IKBK ^c	$d\ln xkp_t \equiv 100\Delta \ln xkp_t - d\ln xkp_t^{tt}$	$d\ln xkp_t = \Delta x_t + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H \Gamma^X) + \sigma_X^{me} me_t^X$
mkp	Real imports	IKBL ^c	$d\ln mkp_t \equiv 100\Delta \ln mkp_t - d\ln mkp_t^{tt}$	$d\ln mkp_t = \Delta m_t + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H \Gamma^X) + \sigma_M^{me} me_t^M$
pxdef	Export deflator	IKBH/IKBL ^c	$d\ln pxdef_t \equiv 100\Delta \ln pxdef_t - \Pi_t^{*,tt} - \Pi_t^{x,tt}$	$d\ln pxdef_t = \Delta p_t^{EX} - \Delta q_t + \pi_t^Z + 100 \ln \frac{\Pi^*}{\Gamma^X} + \sigma_{PX}^{me} me_t^{PX}$
pmdef	Import deflator	IKBI/IKBL ^c	$d\ln pmdef_t \equiv 100\Delta \ln pmdef_t - \Pi_t^{*,tt} - \Pi_t^{m,tt}$	$d\ln pmdef_t = \pi_t^M + 100 \ln \frac{\Pi^*}{\Gamma^X} + \sigma_{PM}^{me} me_t^{PM}$
awe	Nominal wage per capita	KAB9	$d\ln awe_t \equiv \Delta \ln awe_t - \Pi_t^{*,tt}$	$d\ln awe_t = \Delta w_t + \gamma_t^Z + \pi_t^Z + 100 \ln (\Gamma^Z \Pi^*) + \sigma_W^{me} me_t^W$
cpisa	Seasonally adjusted CPI	D7BT ^d	$d\ln cpisa_t \equiv 100\Delta \ln cpisa_t - \Pi_t^{*,tt}$	$d\ln cpisa_t = \pi_t^C + 100 \ln \Pi^*$
rga	Bank Rate	In-house ^e	$robs_t \equiv 100 \ln \left(1 + \frac{rga_t}{100} \right)^{\frac{1}{4}} - \Pi_t^{*,tt}$	$robs_t = r_t + 100 \ln R$
eer	Sterling ERI	In-house ^f	$d\ln eer_t \equiv 100\Delta \ln eer_t$	$d\ln eer_t = \Delta q_t - \pi_t^Z$
hrs	Total hours worked	YBUS	$d\ln hrs_t \equiv 100\Delta \ln hrs_t$	$d\ln hrs_t = \Delta l_t + 100 \ln \Gamma^H + \sigma_L^{me} me_t^L$
yf	World output	In-house ^g	$d\ln yf_t \equiv 100\Delta \ln yf_t - d\ln yf_t^{tt}$	$d\ln yf_t = \Delta z_t^F + \gamma_t^Z + 100 \ln (\Gamma^Z \Gamma^H)$
pxfdef	World export deflator	In-house ^h	$d\ln pxfdef_t \equiv 100\Delta \ln pxfdef_t - \Pi_t^{x,tt}$	$d\ln pxfdef_t = \Delta p_t^{XF} + 100 \ln \frac{\Pi^*}{\Gamma^X}$

^aFollowing the 2011 ONS Blue Book publication, this series was unavailable from 1997 onwards. Bank staff constructed the series back to 1987 by projecting the pre-2011 Blue Book growth rates of business investment from 1987-1996 backwards from the published level of business investment in 1997Q1.

^bData for government investment (DLWF) prior to 1997 was constructed by Bank staff in the same way as for business investment – see footnote *a*.

^cIn-house adjustments are made to take account of the effects of MTIC fraud. See the box on pp 22-23 of the August 2006 Inflation Report for more details.

^dThis CPI series is seasonally adjusted by Bank staff using the X-12 method after accounting for the effect of VAT changes. The resulting series is converted to a quarterly frequency by averaging the index numbers in each month of each quarter.

^eAvailable from the Bank's Statistical Interactive Database with the code IUQABEDR.

^fComputed as a weighted average of individual bilateral exchange rates where the weights are determined by shares in UK trade. See, for example, Lynch and Whitaker (2004). Available from the Bank's Statistical Interactive Database with the code XUQBK67.

^gConstructed as a measure of world trade, taking the average of imports across countries weighted by those countries respective shares in UK trade.

^hConstructed as the average of export prices across countries weighted by those countries respective shares in UK trade.

4.3.2 Priors

In principle, the vector θ^{est} could include all of the parameters in COMPASS. However, we split the model parameters into two broad groups. The first group is mainly comprised of parameters that are important in determining the steady state of the model, but which have little or no influence over its dynamic properties. The second group contains parameters that predominantly influence the dynamic behaviour of the model with little or no effect on its steady state. We calibrate the first set of parameters and so hold them fixed when estimating the dynamic parameters, θ^{est} , using Bayesian methods. This is common practice when estimating models with a relatively large number of parameters.⁶¹

Calibration of ‘steady state’ and other parameter values

As noted above, we calibrate a subset of the parameters in the model, including a set of parameters that primarily influence the steady state. Table 2 shows the values of those calibrated parameters.⁶² Their calibration reflects the following rationale:

- Parameters that govern the steady state shares of the expenditure components in final output (eg the share of consumption in final output ω_{CZ}) are calibrated to match the average ratio of each expenditure component to total final expenditure in the estimation sample (1993Q1–2007Q4).⁶³
- Parameters that govern the steady state growth rates of hours worked, the expenditure components and final output are calibrated in the following way: the steady state growth rate of hours worked, Γ^H , is calibrated to match the average growth of total hours worked in the estimation sample; given that, the steady state growth rate of final output per capita (or in this case per hour), Γ^Z , and the steady state growth rate of investment per capita relative to final output, Γ^I , are calibrated so that steady state consumption and business investment growth match the averages over the sample; given those, the steady state growth rate of imports and exports per capita relative to final output, Γ^X , is calibrated so that steady state growth of value-added matches the average growth rate of GDP in the estimation sample; the steady state growth rate of government spending relative to final output, Γ^G , is backed out by residual to ensure that the weighted sum of the growth rate of the expenditure components in the final output accounting identity delivers the calibrated value for steady state final output growth.⁶⁴

⁶¹Note that Canova and Sala (2009) show that this approach can generate potential identification problems.

⁶²Note also that we make two inconsequential steady state normalisations. We normalise steady state hours worked to 1, backing out a value for the parameter ν_L to deliver that normalisation, and we normalise the steady state real exchange rate to be consistent with steady state mark-ups in the import and export sectors greater than 1, backing out a value for the parameter κ_F to deliver that. See Section A.3 of Appendix A for a derivation of the steady state.

⁶³Although Table 2 includes all of these share parameter values, the share of consumption in final output is in fact backed out by residual (and so it technically not treated as a parameter) reflecting that the shares of the expenditure components in final output must sum to 1.

⁶⁴This turns out to be very similar to the average of the sample. That is not surprising because differences can only arise due to differences in the growth rates of components of final expenditure that are, in effect, treated as a residual in COMPASS like dwellings investment and stockbuilding, which are captured by I^O and which are assumed to grow at the same rate as final output. Since these have a

- The inflation target, Π^* is calibrated to an annual rate of 2% in line with the MPC's inflation target. Given that and given the calibration of Γ^Z , the household discount factor, β , (which determines their rate of time preference) is calibrated so that the steady state nominal interest rate, R , matches the average over the sample of just under 5.5%.⁶⁵
- The labour share (in terms of value added), $\omega_{LV} = \frac{\alpha_L}{\mu^V}$, is calibrated to match the average share of wages and salaries (adjusted for employer contributions) in nominal GDP measured at basic prices in the estimation sample. And the value-added share (in terms of final output), $\omega_{VZ} = \frac{\alpha_V}{\mu^Z}$, is calibrated so that the share of imports in final output (equal $1 - \omega_{VZ}$) matches the sample average of the share of imports in total final expenditure at basic prices.⁶⁶
- The steady state final output price mark-up, μ^Z , is calibrated to a low value of 1.005 to ensure that the national accounting structure in COMPASS comes close to matching the ONS national accounting structure.⁶⁷
- The capital depreciation rate, δ_K , is calibrated to match the value from the Bank of England's previous quarterly forecast model, BEQM.
- The parameter that governs the rate at which the rest of the world inherits domestic LAP shocks, ζ_{ω_F} (see Section 4.2.4 for details), is calibrated to 0.9. The parameter that determines the over-discounting of consumption to ensure a well-defined net foreign asset position in steady state, ϵ_β , is calibrated so that it has very little effect on model properties (since it is a device to implement the technical requirement that the model returns to steady state following temporary shocks) with $\beta_{factor} = \frac{\epsilon_\beta(1-\psi_C)}{\epsilon_C} = 0.01$ (see A5 of Appendix A for a brief discussion of the over-discounting approach).
- We calibrate the persistence in the LAP forcing process, ρ_{LAP} , to zero reflecting a strong prior that permanent shocks to productivity should not be persistent. The parameters of the TFP and value-added mark-up forcing processes are also calibrated, because COMPASS is 'over-identified', meaning it contains more shocks than observable variables, as discussed below.

Over-identification is useful in the context of a misspecified model because it affords more flexibility around the economic rationale for the imposition of forecast judgements and is a reflection of our general approach to dealing with misspecification of COMPASS (see Section 7). However, separately identifying all of the forcing process parameters could be problematic given that some of them have fairly similar economic effects, so

small weight in final output, their growth rates in the data would have to be drastically different to that of final output to drive a significant wedge between the growth rate of government spending in the data and that backed out by residual in the model.

⁶⁵Note that this is adjusted for the inflation trend term discussed above, $\Pi^{*,tt}$.

⁶⁶We proxy import volumes at basic prices with data for import volumes at market prices.

⁶⁷As noted by Basu and Fernald (2002), the presence of imperfect competition when intermediate inputs are used in production drives a wedge between measured GDP and true value added because the marginal product of intermediates does not equal their price. Since COMPASS includes intermediate imports in final output production and imperfect competition in the market for final output, that effect is present. However, by calibrating the mark-up in the final output sector to be very small, the effects of imperfect competition on national income accounting measures become negligible.

we choose to calibrate a subset of them.⁶⁸ The calibration ensures that these shocks do not have much impact on the properties of the model, but that they are still effective instruments for the imposition of forecast judgements.⁶⁹

Table 2: Calibrated parameters

Parameter	Description	Value
ω_{CZ}	Steady state share of consumption in final output	0.5031
ω_{IZ}	Steady state share of business investment in final output	0.0845
ω_{I^oZ}	Steady state share of ‘other’ investment in final output	0.0370
ω_{GZ}	Steady state share of government spending in final output	0.1662
ω_{XZ}	Steady state share of exports in final output	0.2092
ω_{MZ}	Steady state share of imports in final output	0.2197
Γ^H	Trend population growth	1.002
Γ^Z	Trend productivity growth	1.007
Γ^I	Trend investment growth relative to final output growth	1.0036
Γ^X	Trend export growth relative to final output growth	1.0025
Γ^G	Trend government spending growth relative to final output	0.995
Π^*	Inflation target	1.005
β	Household discount factor	0.9986
ω_{LV}	Steady state labour share	0.6774
ω_{VZ}	Steady state value added share	0.7599
μ^Z	Steady state final output price mark-up	1.005
δ_K	Capital depreciation rate	0.0077
ζ_{ω_F}	Speed at which rest of the world inherits LAP shocks	0.9
β_{factor}	‘Over-discounting’ factor	0.01
σ_{μ^V}	Standard deviation of value added price mark-up shock	0.05
σ_{TFP}	Standard deviation of TFP shock	0.05
ρ_{TFP}	Persistence of TFP forcing process	0.9
ρ_{LAP}	Persistence of LAP forcing process	0

Priors for ‘dynamic’ parameters

The parameters that mainly affect the dynamic properties of the model are further subdivided into two groups. The first set of parameters are more directly related to the frictions that affect agents’ decisions within the model (for example, the degree of price stickiness and the size of investment adjustment costs). The second set of parameters are those that control the variances of the measurement errors and shocks and the persistence of the forcing processes. Tables 3 & 4 detail the distribution used for each of the priors. The rest of this sub-section discusses the rationale for those priors.

The overall strategy for setting the priors for the structural parameters is to set the means of the prior distributions using a mixture of estimates from empirical studies

⁶⁸For example, the value-added mark-up shock has a similar effect to the final output mark-up shock. See Appendix B.

⁶⁹This is achieved by setting their standard deviations to be relatively small, but still large enough that there are no numerical issues in applying forecast judgements through inversion – see Appendix C for a description.

(applied to the UK economy where possible) and, where those are not readily available, using ‘standard’ priors from other studies (again, applied to the UK where possible).

- We set the priors for the parameters of the Taylor rule in a standard way. The means of the priors for θ_{Π} , θ_Y & θ_R , are set to 1.5, 0.125 & 0.8.
- We set the means of the priors for the domestic price adjustment costs, which govern the stickiness of domestic prices in responses to changes in the relevant measures of domestic marginal costs, to be roughly in line with evidence from Greenslade and Parker (2012), who report that firms change prices at an average frequency of around three quarters. We then set the mean of the prior for wage adjustment costs to imply a slightly longer average duration of around 4-5 quarters based on a prior that rigidities in nominal wage setting are greater than those in price setting (consistent with the findings of the papers surveyed by Harrison and Oomen (2010)). Evidence for the stickiness of UK import and export prices is much harder to come by. We choose to set the prior means of the export and import price adjustment between those of the wage and domestic price adjustment costs.
- We set the priors for each of the price indexation parameters in an identical way, allowing for 90% of the probability mass to lie between 0.14 and 0.38 in all cases.
- The mean for the prior on habit formation is set to 0.7, roughly in the middle of a range of estimates for the United Kingdom (Carroll et al., 2011). The standard deviation of the prior is set so that 90% of the probability mass roughly covers the range of estimates from the same study and also encompasses another UK study (Banerjee and Batini, 2003). The mean for the prior on investment adjustment costs is set to 2. This is somewhat lower than the estimates (for models estimated using US and Euro area data) surveyed by Harrison and Oomen (2010), which would imply a value closer to 5. However, most of the papers considered by Harrison and Oomen (2010) used models in which costs of adjusting capital utilisation were relatively low. Since COMPASS does not include variable capital utilisation (see footnote 27), we reduce our prior for the costs of adjusting physical capital accordingly.
- We follow Harrison and Oomen (2010) and set the mean of the inverse of the intertemporal elasticity of substitution, ϵ_C , to 1.5, with a standard deviation that allows for some mass towards a log specification for utility (a commonly-used specification) and for some mass towards an elasticity of substitution of one half ($\epsilon_C = 2$).
- We set the mean of the share of optimising households, ω_O , to 0.7 with a variance implying that roughly 90% of the mass of the distribution lies between 0.6 and 0.8 in line with the range reported in an empirical study for the share of UK households displaying excess sensitivity to their current incomes (Benito and Mumtaz, 2006).



Table 3: Priors and posteriors for estimated parameters

Parameter	Description	Prior			Posterior			
		Distribution	Mean	Std	Mean	Std	5%	95%
θ_R	Monetary policy rule interest rate smoothing	Beta	0.8	0.05	0.8336	0.02423	0.7913	0.8704
θ_Π	Monetary policy rule inflation response	Normal	1.5	0.125	1.497	0.1154	1.309	1.689
θ_Y	Monetary policy rule output gap response	Beta	0.125	0.075	0.1512	0.06654	0.05497	0.272
ϕ_Z	Final output price adjustment cost	Gamma	7	1.5	8.702	1.815	5.971	11.92
ϕ_V	Value added price adjustment cost	Gamma	7	1.5	6.672	1.478	4.461	9.257
ϕ_M	Import price adjustment cost	Gamma	10	1.5	9.467	1.395	7.31	11.9
ϕ_X	Export price adjustment cost	Gamma	10	1.5	10.54	1.442	8.306	13.02
ϕ_W	Nominal wage adjustment cost	Gamma	14	1.5	15.02	1.524	12.61	17.61
ξ_Z	Indexation of final output prices	Beta	0.25	0.075	0.2079	0.06249	0.1136	0.3182
ξ_V	Indexation of value added prices	Beta	0.25	0.075	0.2363	0.07256	0.127	0.3642
ξ_M	Indexation of import prices	Beta	0.25	0.075	0.1745	0.0543	0.09364	0.2711
ξ_X	Indexation of export prices	Beta	0.25	0.075	0.2441	0.0615	0.1489	0.3503
ξ_W	Indexation of nominal wages	Beta	0.25	0.075	0.2446	0.07177	0.1355	0.3706
ψ_C	Habit formation parameter	Beta	0.7	0.15	0.2188	0.06825	0.1134	0.3375
ψ_I	Investment adjustment cost	Gamma	2	0.4	2.998	0.4258	2.347	3.741
ϵ_C	Coefficient of relative risk aversion	Gamma	1.5	0.2	1.218	0.14	0.9984	1.459
ϵ_L	Labour supply elasticity	Gamma	2	0.3	1.921	0.2783	1.488	2.404
ϵ_F	Price elasticity of world demand for UK exports	Gamma	0.75	0.1	0.5198	0.06152	0.4223	0.6248
ω_o	Share of optimising households	Beta	0.7	0.05	0.7831	0.03849	0.7166	0.843
ρ_B	Persistence of risk premium forcing process	Beta	0.75	0.1	0.7058	0.04334	0.6318	0.7744
ρ_I	Persistence of investment adjustment forcing process	Beta	0.75	0.1	0.6977	0.04873	0.6125	0.7727
ρ_G	Persistence of government spending forcing process	Beta	0.9	0.05	0.91	0.01466	0.8839	0.9317
ρ_{IO}	Persistence of other investment forcing process	Beta	0.75	0.1	0.5036	0.07382	0.376	0.6187
ρ_{κ^F}	Persistence of export preference forcing process	Beta	0.75	0.1	0.7152	0.04771	0.6309	0.7862
ρ_M	Persistence of import preference forcing process	Beta	0.75	0.1	0.7082	0.04396	0.6323	0.776
ρ_L	Persistence of labour supply forcing process	Beta	0.75	0.1	0.7799	0.1012	0.5921	0.9187



Table 4: Priors and posteriors for estimated parameters (continued)

Parameter	Description	Prior			Posterior			
		Distribution	Mean	Std	Mean	Std	5%	95%
ρ_{BF}	Persistence of UIP shock forcing process	Beta	0.75	0.1	0.8723	0.02064	0.8374	0.9053
ρ_{PX^F}	Persistence of world export price forcing process	Beta	0.9	0.05	0.911	0.01523	0.8839	0.9337
ρ_{ZF}	Persistence of world output forcing process	Beta	0.9	0.05	0.9016	0.01731	0.8709	0.927
σ_B	St dev of risk premium shock	Gamma	0.5	0.1	0.5376	0.07988	0.4178	0.6783
σ_I	St dev of investment adjustment shock	Gamma	1.9	0.1	2.23	0.09481	2.075	2.387
σ_G	St dev of government spending shock	Gamma	3	0.1	3.061	0.09816	2.901	3.225
σ_{IO}	St dev of other investment shock	Gamma	14	0.5	13.98	0.4892	13.19	14.8
σ_{κ^F}	St dev of export preference shock	Gamma	2.2	0.1	2.283	0.09822	2.123	2.447
σ_M	St dev of import preference shock	Gamma	2.2	0.1	2.266	0.09488	2.112	2.424
σ_{LAP}	St dev of LAP growth shock	Gamma	0.35	0.05	0.4686	0.03604	0.4115	0.53
σ_L	St dev of labour supply shock	Gamma	0.75	0.1	0.7658	0.1033	0.6039	0.9438
σ_R	St dev of monetary policy shock	Gamma	0.1	0.05	0.08806	0.008425	0.07538	0.1029
σ_{BF}	St dev of UIP shock	Gamma	0.65	0.1	0.5797	0.06477	0.4798	0.692
σ_{μ^Z}	St dev of final output markup shock	Gamma	0.1	0.05	0.2237	0.02242	0.1896	0.2632
σ_{μ^W}	St dev of wage markup shock	Gamma	0.3	0.05	0.3472	0.05061	0.2651	0.4315
σ_{μ^M}	St dev of import markup shock	Gamma	1.3	0.1	1.174	0.08113	1.044	1.311
σ_{μ^X}	St dev of export markup shock	Gamma	1.3	0.1	1.421	0.08418	1.286	1.563
σ_{PX^F}	St dev of world export price shock	Gamma	1.6	0.1	1.701	0.08828	1.56	1.85
σ_{ZF}	St dev of world output shock	Gamma	2.5	0.1	2.525	0.09749	2.368	2.687
σ_i^{me}	St dev of investment measurement error	Gamma	0.35	0.1	0.5042	0.1724	0.266	0.8228
σ_X^{me}	St dev of export measurement error	Gamma	0.18	0.055	0.4142	0.06212	0.3191	0.5235
σ_M^{me}	St dev of import measurement error	Gamma	0.18	0.055	0.1848	0.05885	0.1008	0.2921
σ_L^{me}	St dev of hours measurement error	Gamma	0.045	0.013	0.1859	0.02339	0.1456	0.2226
σ_W^{me}	St dev of wage measurement error	Gamma	0.125	0.0275	0.3152	0.03644	0.2558	0.3755
σ_{PM}^{me}	St dev of import price measurement error	Gamma	0.34	0.075	0.4321	0.07488	0.3143	0.5605
σ_{PX}^{me}	St dev of export price measurement error	Gamma	0.34	0.075	0.348	0.08526	0.2235	0.5007

The parameters associated with the forcing processes can be split into two: those that govern their persistence and those that govern their standard deviations. For the parameters governing the persistence of the forcing processes, we try to select priors that are consistent with the notion that most of the persistence in the model variables is generated by internal propagation mechanisms rather than the persistence of the forcing processes. In practice, this means that most of our priors for the autoregressive parameters are set to values that imply a short half-life for the forcing process: the means of our priors imply a half-life of around 2.5 quarters.⁷⁰ We set a higher prior mean for the parameters governing the persistence of the world output, world export price and government spending forcing processes, ρ_{PX^F} , ρ_{ZF} & ρ_G because the corresponding endogenous variables are modelled as exogenous and so their persistence is driven entirely by those parameters.

For the parameters governing the variance of the shocks, we use a minimum distance estimation to set the means of the prior distributions to match the variances of the observables in the model with their counterparts in the data as closely as possible, conditional on the values of the calibrated parameters and the means of all the other priors.⁷¹

Finally, the priors for the standard deviations of the measurement errors are set conservatively to ensure that not ‘too much’ of the variance of the observables concerned is explained by measurement error. Specifically, the priors for the measurement errors on imports, exports, hours worked and investment are set so that there is an upper bound that no more than 10% of the variance of those variables is driven by measurement error, while the priors for the measurement errors on wage growth and the trade deflators are set so that no more than 25% of their respective variances is driven by measurement error.

4.3.3 Posterior parameter estimates

Tables 3 & 4 report the results of the Bayesian estimation of COMPASS using the priors detailed above in the form of moments and percentiles from the posterior distributions of the estimated parameters. The following results are particularly noteworthy:

- The posterior means of the parameters governing indexation of inflation to lagged inflation (the ξ family) for final output, value-added, imports, exports and wages are less than 0.25 and so lower than the prior means.
- Both the posterior mean for habit formation, ψ_C , and the coefficient of relative risk aversion, ϵ_C , are materially lower than the prior means, implying a surprisingly low degree of habit formation (given the empirical evidence discussed above) and a point specification of the utility function that is close to a log function (i.e.

⁷⁰The half-life of a process is defined as the number of periods it takes for the process to fall to half its value as measured in the period of the shock. For an AR(1) process $y_t = \rho y_{t-1} + \eta_t$, the half-life is equal to $\frac{\ln 0.5}{\ln \rho}$.

⁷¹It was not possible to match all of the variances in the data. Reflecting that, we took a judgemental approach to identification and held the standard deviations of the LAP, labour supply, risk premium and export price mark-up shocks fixed in the minimum distance estimation and chose not to try to match GDP, hours and export price growth. Note that all standard deviations in the estimated model (other than those of the TFP and value-added mark-up shocks that we calibrated) are informed by a combination of the prior and the likelihood in the full Bayesian estimation. Note also that the minimum distance estimation was conducted conditional on the means of all the other priors, so it is possible that there are parameter combinations that would improve the fit of the *prior* variances in the model to the data.

a relatively high intertemporal elasticity of substitution). One interesting result arising from informal analysis of the sensitivity of the estimation results (to different priors) was that there seems to be a trade-off between the share of rule-of-thumb households, ω_o , and the degree of habit formation. The economic intuition for that is that the consumption of rule-of-thumb households inherits the stickiness in wages, so a higher share of rule-of-thumb households can impart sluggishness to consumption in much the same way as higher habit formation.

- The investment adjustment cost parameter's posterior mean is materially higher than the prior mean. In addition, the standard deviation of the investment adjustment cost posterior distribution is wider than the prior distribution (which is the case for very few of the parameters). This parameter does not seem to be particularly well identified, consistent with its posterior distribution being quite sensitive to the specification of the priors.
- The posterior mean for the price elasticity of demand for UK exports is materially lower than the prior mean, implying a relatively low degree of price elasticity.
- On the whole, the means of the posteriors for the persistence parameters in the forcing processes for the shocks (the ρ family of parameters) are lower than the prior means (consistent with not 'too much' of the persistence in the endogenous variables being driven by those parameters).
- The posterior means of the standard deviations of the shocks and measurement errors are, in general, slightly larger than the prior means.

4.4 Model properties

This section illustrates some of the properties of the estimated model by showing the impulse responses to a selection of shocks. For a complete set of impulse responses, see Appendix B. When presenting the impulse responses, we take advantage of the fact that the Bayesian estimation results described above provide us with an estimate for the (joint) posterior distribution of the model's parameters. We use this information to compute the posterior distributions of the impulse responses, which we characterise in the figures below as a 'swathe' of impulse responses.⁷² We consider in turn a shock to the monetary policy reaction shock function, a shock to labour augmenting productivity and a domestic risk premium shock.

4.4.1 A monetary policy shock


Figure 3 summarises the effects of a one standard deviation monetary policy shock.⁷³ Given the parameter estimates, a one standard deviation shock raises the policy rate by around 30 basis points on impact. The real interest rate rises in response because of the presence of nominal rigidities (sticky prices and wages). The increase in the real interest rate has two important effects. First, it acts to depress private domestic demand as the

⁷²Constructing these charts is straightforward. We take 10,000 draws from the posterior distribution of the model parameters. For each draw we compute the impulse response function. We collect together the sample of impulse responses thus generated and compute the range lying between the percentiles of interest.

⁷³The shock is a surprise increase in η^R by 1.

increased return to saving induces households to postpone consumption and investment in physical capital. The second effect of a persistent increase in the real interest rate is that the real exchange rate appreciates (via the UIP condition), reducing export demand slightly.⁷⁴ Taken together, these effects reduce GDP and final output.

Figure 3: Impulse responses to a monetary policy shock



The charts show responses of COMPASS variables to a one standard deviation shock. Responses are measured in percentage changes from steady state or percentage points (pp) where specifically indicated. The x-axis shows the number of quarters following the shock. The shaded area shows the range between the 5-th percentile and the 95-th percentile of the distribution of responses associated with the posterior parameter uncertainty. Dashed lines show the responses when parameter values are set to the mean of the posterior distribution.


The fall in output reduces the demand for labour and capital and so the relative prices of these factors of production fall. This acts to lower the marginal cost of producing value added. Taken together with a reduction in import price inflation (generated by the exchange rate appreciation), these effects reduce the cost of producing final output and hence CPI inflation. Despite the presence of interest rate smoothing in the monetary policy reaction function, the policy rate returns to steady state relatively quickly, reflecting an endogenous response to weaker inflation and a negative output gap.

⁷⁴The fall in exports is relatively small because export prices are sticky in foreign currency, so exchange rate movements take some time to be reflected in the foreign currency price of exports.

4.4.2 A labour augmenting productivity shock

Figure 4 plots the results of a temporary one standard deviation shock to the *growth rate* of labour augmenting productivity (LAP).⁷⁵ This shock temporarily raises the growth rate of the economy's supply capacity. Although the effect on growth is temporary, this shock has a permanent effect on the levels of variables along the balanced growth path. This is because the shock permanently increases the *level* of output that can be produced using any combination of factor inputs.

Figure 4: Impulse responses to a labour augmenting productivity shock



The charts show responses of COMPASS variables to a one standard deviation shock. Responses are measured in percentage changes from steady state or percentage points (pp) where specifically indicated. The x-axis shows the number of quarters following the shock. The shaded area shows the range between the 5-th percentile and the 95-th percentile of the distribution of responses associated with the posterior parameter uncertainty. Dashed lines show the responses when parameter values are set to the mean of the posterior distribution.

Value added prices (measured relative to final output prices) fall, reflecting increased productivity and falling production costs. The reduction in costs leads to a small fall in CPI inflation. The fall in costs is partially offset by an initial rise in relative import prices associated with a depreciation in the real exchange rate, resulting from the monetary policy maker reducing the nominal interest rate in response to weaker inflation. Total

⁷⁵The shock is a surprise increase in η^{LAP} by 1.

hours fall initially because sticky prices and wages prevent demand from increasing as much as supply capacity. This means that it is possible to meet the increased demand without using additional factors of production.⁷⁶

The permanent nature of this shock means that the dynamics of adjustment are relatively prolonged, as the stocks of physical capital and financial assets take some time to accumulate to their new, permanently higher, levels. Until this adjustment is complete, expenditure flows remain away from steady state. The response of the exchange rate to these stock-flow adjustments is relatively small because we assume that the stochastic trend in the domestic economy is co-integrated with the stochastic trend in the rest of the world. This means that foreign demand for exports gradually increases in response to the increase in domestic productivity and so delivering the increase in exports to the higher balanced growth path does not require a large exchange rate depreciation.

4.4.3 A domestic risk premium shock

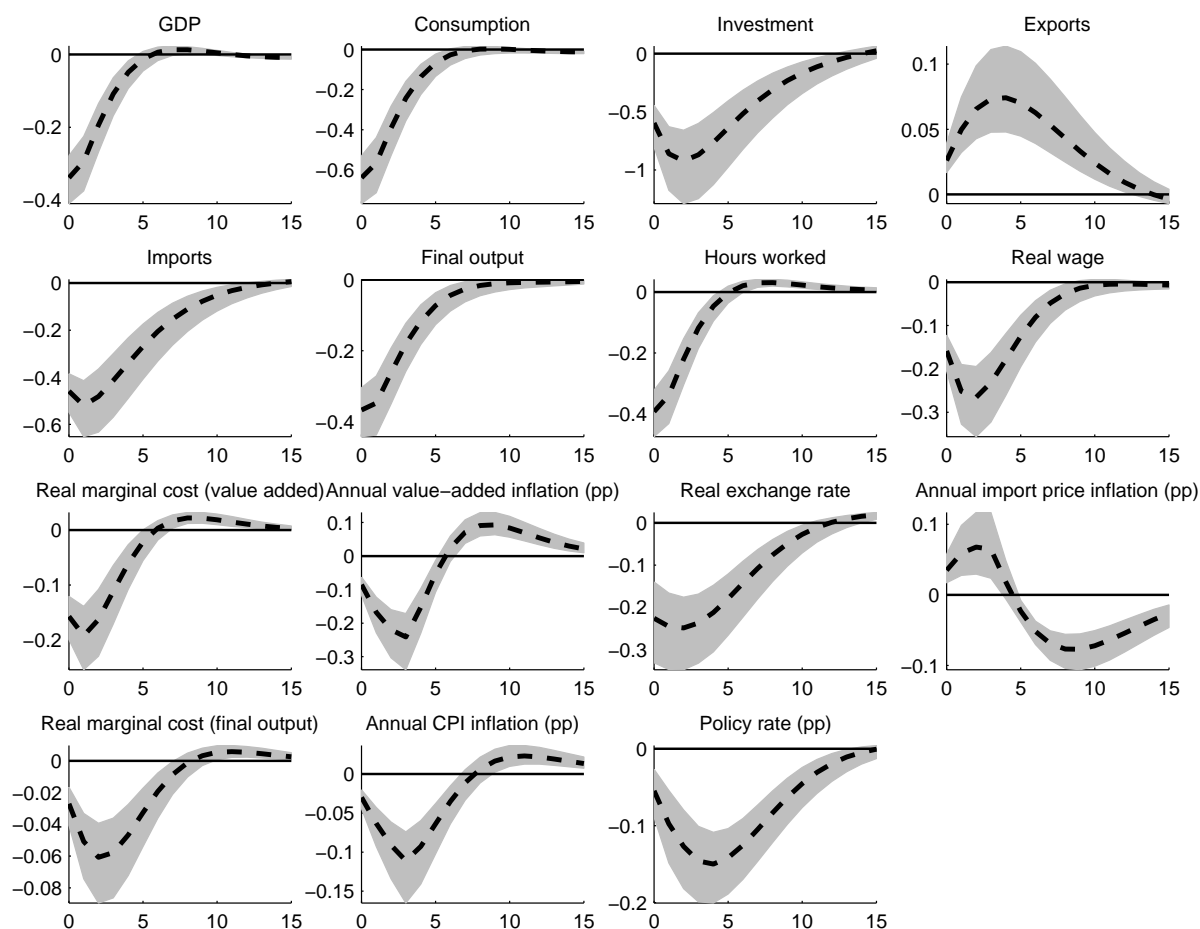
Figure 5 summarises the key results of a one standard deviation shock to the domestic risk premium.⁷⁷ A positive domestic risk premium shock temporarily increases the effective cost of borrowing for a given real interest rate. This induces households to postpone consumption and investment spending. The short-run reduction in private domestic demand reduces final output and GDP.

The fall in production lowers firms' demand for factor inputs, including imports. This reduces the prices of domestic factors of production and the marginal cost of producing value added. As a result, value added inflation declines, leading to a fall in CPI inflation. CPI inflation falls despite a rise in import price inflation brought about by a depreciation of the exchange rate. The depreciation reflects a reduction in interest rates as the monetary policy reaction function responds to weaker inflation and activity. The depreciation is sufficient to generate a small increase in exports, which partially offsets the reduction in private domestic demand brought about by the shock.

⁷⁶If prices and wages were flexible, factor inputs would increase because aggregate demand would expand more than proportionally to the increase in productive capacity. This is because the persistence of the shock gives rise to positive wealth effects in this instance.

⁷⁷The shock is a surprise increase in η^B by 1.

Figure 5: Impulse responses to a domestic risk premium shock



The charts show responses of COMPASS variables to a one standard deviation shock. Responses are measured in percentage changes from steady state or percentage points (pp) where specifically indicated. The x-axis shows the number of quarters following the shock. The shaded area shows the range between the 5-th percentile and the 95-th percentile of the distribution of responses associated with the posterior parameter uncertainty. Dashed lines show the responses when parameter values are set to the mean of the posterior distribution.

5 The suite of models

This section explains the motivation for using a suite of models alongside the central organising model, and describes the content of the suite in more detail. The suite of models provides Bank of England staff with the means to cross-check, interrogate and adjust the forecast. Section 8 demonstrates, using some detailed case studies, how the suite can be used alongside COMPASS to support the production of the MPC’s forecast.

5.1 Reasons to employ a suite of models

The suite can be used to manage some of the trade-offs inherent in using any central organising model. COMPASS is a relatively small and simple central model, and this makes it easier to use and understand. But, like any model, it is misspecified. Suite models can be selected to help to overcome the main misspecifications, once those have been identified. They can therefore be used to adjust a forecast, by imposing appropriate judgements back on to the central organising model.

The Bank of England has always used a plurality of models when producing its forecasts (see Whitley (1997) and Section 1.1 of Bank of England (2000)). This can reduce the danger of relying too heavily on any one model or paradigm and adds richness and breadth to the MPC’s discussions of the forecast.

Using a suite of models alongside a smaller central model also provides flexibility. Models can be withdrawn from the suite if they are not needed, or introduced quickly if a new question needs to be addressed. A crucial feature of the suite is that it is dynamic, evolving over time as economic modelling techniques advance and the economic environment itself changes.

Some models in the suite will always form a key part of the MPC’s forecast process, while others may be used on a more occasional basis. Recognising this, the suite of models distinguishes an “inner” suite and a “wider” suite. The models in the inner suite typically have more resources devoted to their maintenance.

In addition, as discussed in Section 4 on page 12, COMPASS is likely to evolve over time as the staff learn more about its strengths and weaknesses. As part of that process, suite model features may be added to COMPASS in the future. Indeed, one way of assessing the suitability of a proposed COMPASS extension is to build it into the suite in the first instance and then add it to COMPASS later on if the benefits (in terms of the economics it provides) are judged to outweigh the costs (in terms of operating and understanding a larger central model).

Models in the suite can be divided into three broad classes, according to their main purpose:

- (a) Models which articulate economic shocks and channels which are omitted from COMPASS (Section 5.2).
- (b) Models which expand the scope of the forecast, by producing forecasts for additional variables not included in COMPASS (Section 5.3).
- (c) Models which generate alternative forecasts for variables which are in COMPASS. These models play an important role in cross-checking the output from the central organising model (Section 5.4).

The remainder of Section 5 describes these three classes of model in more detail.

5.2 Models which articulate economic shocks and channels missing from COMPASS

As noted in Section 2.3, the one thing we surely know about any model is that it is misspecified, and COMPASS is no exception. In particular, the decision to use a small, tractable organising model at the centre of the forecast platform means that many economic channels are actually excluded by design. A range of ancillary models and tools is therefore needed to articulate those “omitted” channels, and to understand how the MPC’s forecast might need to be adjusted to take them into account. The first class of suite models does this.

For example, all of the following economic channels are missing from COMPASS:

- The model does not account explicitly for energy as an input to production or consumption. But the impact of changes in energy prices on marginal cost and inflation is likely to be substantially different from the impact of changes in the prices of other goods and services.
- It contains no (explicitly modelled) financial frictions, and no role for a financial sector intermediating funds between different sectors of the economy.
- Fiscal policy is only modelled in a very simple way. Government spending is assumed not to add to household utility, distortionary taxes play little role, and households behave in a way which is Ricardian (see Section 4.2.5).
- There is only a single, short-term, interest rate. This means that COMPASS is not well suited to analysing the effects of unconventional policies such as asset purchases (Benford et al. (2009)).
- There is no mechanism via which firms can vary the intensity of their labour input, so firms do not have the option of ‘hoarding’ (under-utilising) labour (see equation (1)). This could be material because under-utilisation of labour is one possible explanation for the puzzling resilience of UK employment relative to output over the past five years (see Faccini and Hackworth (2010) and Hughes and Saleheen (2012)).
- The model assumes that the final output price of firms is the same as that paid by consumers. Indirect taxes, such as VAT and duties, which may drive a wedge between those prices, are ignored.

But in each of the above cases, although the channels are not included in COMPASS, they can nonetheless be incorporated in the MPC’s forecast by using suitable suite models and supporting tools. The rest of this section explores two examples in more detail: a model of energy; and the treatment of frictions arising from the financial sector. We also return to these issues (notably financial frictions, asset purchases and VAT) in Section 8.

5.2.1 Models with energy

For reasons of tractability, energy is not modelled separately in COMPASS.⁷⁸ But the UK economy has been subject to a number of shocks in recent years from changes in

⁷⁸That was also the case in BEQM, which was much larger.

energy prices, and they have had important consequences for inflation and monetary policy. There are several models available in the suite which incorporate energy channels within a general equilibrium framework. Those can be used to calibrate the effects of a movement in energy prices on key forecast variables. Those responses can then be mapped back on to the central organising model using appropriate shocks.

Harrison et al. (2011) show how a workhorse DSGE model can be expanded to include energy channels. Theirs is a small open economy model with real and nominal rigidities such as investment adjustment costs and sticky prices and wages. Petrol and utilities are modelled separately from non-energy goods and services, and enter the first order conditions for households and firms. Oil and gas prices are set in the global market and are exogenous to the UK economy, though the UK is assumed to have endowments of oil and gas and can trade freely in these, as well as non-energy goods and services. The model can be used to study policy responses, possible real wage resistance, and the effect of both temporary and permanent shocks. Millard (2011) takes a similar model and estimates its parameters on UK data using Bayesian techniques similar to those described in Section 6.2.2. He demonstrates how model-based decompositions can be used to trace the impact of energy price shocks through the supply chain.

In a similar spirit, the suite also includes a version of COMPASS which has been extended to include energy price channels. The assumptions used in this model are much simpler than in Harrison et al. (2011) and Millard (2011). However, in keeping with the Bank's vision for how the suite of models should operate, Bank staff may develop more sophisticated versions as the suite evolves over time. The current version of "COMPASS with Energy" inherits its balanced growth path from COMPASS, and the additional parameters are calibrated. Final output consists of non-energy goods and energy, so energy enters the optimisation problems of both firms and households. The production of final output is assumed to be Leontief in energy and non-energy inputs, so firms are unable to substitute labour and capital for energy or vice versa; the same approach is used in the household utility function. Although a simplification, the Leontief approach is probably a reasonable approximation to reality, at least in the short run. Both the UK's endowment of energy and the world relative price of energy are assumed to be exogenous. Shocks to energy prices affect firms' marginal costs and the price of final output but, as is the case more generally, the wider macroeconomic effects depend ultimately on the stance of monetary policy. The monetary policymaker can choose either to accommodate or to offset the inflationary impact of energy price shocks, depending on the circumstances. Section 8.2 demonstrates how Compass with Energy can be used to account for the contribution of changes in energy prices over the recent past to inflation, conditional on a particular assumption about the response of monetary policy.

5.2.2 Models and tools for understanding the impact of the financial sector

At the time of writing there is no canonical model in the academic literature which articulates all of the effects of the financial sector on the wider economy.⁷⁹ Bank staff have therefore used a pragmatic approach to capturing the effects of the financial channels which are missing from COMPASS. A range of suite models and ancillary tools are used to calibrate these effects and to apply judgement to the MPC's forecast. But this remains an active area of academic research and an important priority for future suite development.

⁷⁹Of course, this is an active research area, with a number of recent contributions: see for example Gertler and Kiyotaki (2010) and Gertler and Karadi (2011).

The suite includes various extensions of COMPASS to incorporate financial sector channels. One of these, which we describe here by way of illustration, introduces credit spreads into COMPASS. These drive a wedge between the official policy rate and the effective marginal interest rates faced by households and firms. The household rate enters the consumption equation (14) in both the sticky and flexible price models (in the same way as the risk premium shock), so a rise in credit spreads has a similar effect to a negative demand shock. On the production side, a working capital channel is included: firms have to borrow to pay for their labour and capital in advance of sales, so a higher credit spread increases their marginal cost (see, for example, Fernandez-Corugedo et al. (2011)). This means that a shock which increases spreads faced by firms leads to higher inflation and a fall in output. The model also allows for a monetary policy response to credit spread shocks.

The Bank’s forecast process regularly includes judgements applied via a “credit spread adjustment”. This is a summary statistic which measures by how much the effective marginal interest rates faced by household and firms differ from the official policy rate.⁸⁰ A wide range of information is fed into producing the adjustment, including data on a portfolio of lending and deposit rates, data from household surveys, and judgements about the outlook for conditions in the banking sector. By considering the implications of a change in spreads on key model variables, it is possible to map these changes back into the central organising model and thereby incorporate judgements about the outlook for financial conditions into the MPC’s forecasts.

A range of other models can be used to cross-check these judgements. The suite also includes a structural VAR (described in Barnett and Thomas (forthcoming)) which traces through the implications of shocks to credit supply and credit demand for key macro-economic variables. Different types of credit shock are identified using sign restrictions: a negative credit supply shock is assumed to lead to a rise in credit spreads as well as a fall in lending quantities, whereas a negative demand shock leads to a fall in both spreads and quantities. Barnett and Thomas find that negative shocks to credit supply are expected to reduce the level of GDP, but also to push up on inflation, though the effect of the latter is not statistically significant. They also find that Bank Rate has, over the past, tended to offset movements in credit spreads. Since the impact of an adverse credit supply shock on inflation is small, such shocks could also be associated with a reduction in potential supply. An example of the analysis of credit spreads using suite models (including the Barnett and Thomas model) and the application of their effects to COMPASS is provided in Section 8.3.

5.3 Models which expand the scope of the forecast

A second consequence of using a smaller central model is that a number of important variables are excluded. For example, COMPASS includes one measure of labour input (total hours worked), but does not produce forecasts of the employment or unemployment rate. The second class of suite model seeks to fill in these gaps and thereby expand the scope of the forecast. While these extra variables are not used as direct inputs to the COMPASS model equations, they serve as diagnostics on the output of COMPASS, and can be used to motivate judgements.

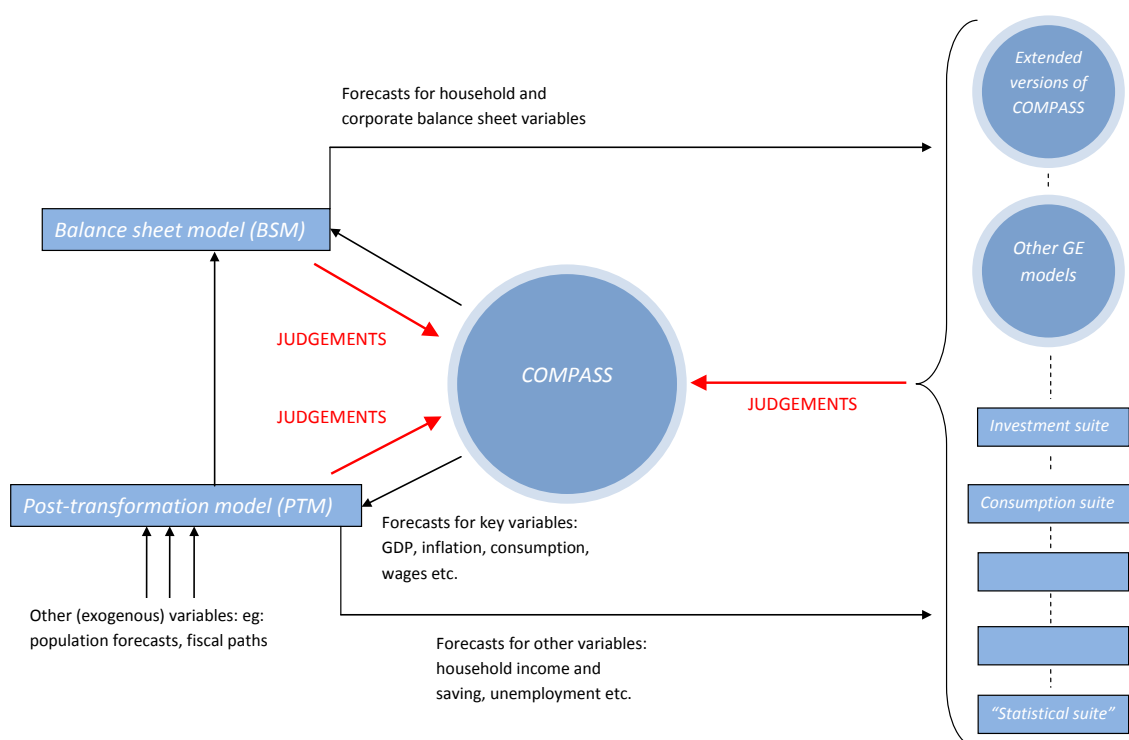
The current suite has two models in this class. The first is the “Post-Transformation Model” (PTM) which is run alongside COMPASS throughout the Bank’s forecast process.

⁸⁰This is described in more detail in Section 8.3.

The outputs from this model form a core of important variables which are regularly used in discussions with the MPC. The second is the “Balance Sheet Model” (BSM), which extends the forecast to encompass an even wider range of variables, with a special focus on money, lending, and corporate and household balance sheets.

It is of course imperative that these extra forecasts are consistent with those being generated directly from COMPASS. To ensure this, forecasts from COMPASS are fed into the suite models as inputs, along with other exogenous variables where necessary. For example, if the COMPASS forecast for hours worked is updated, labour market variables in the PTM will be updated automatically. Figure 6 illustrates how forecasts from COMPASS are fed into these two suite models. It is also possible to use forecasts from the PTM and BSM to run other suite models, such as the investment and consumption models described in Section 5.4.

Figure 6: Stylised diagram of interaction between COMPASS and models in the suite



5.3.1 The Post-Transformation Model (PTM)

The PTM is a backward-looking, recursive model which serves as an add-on to COMPASS. Unlike COMPASS, it is nonlinear, which facilitates the specification of equations in terms of actual published data.⁸¹ The current version contains 150 variables, which fit loosely into four categories:

- (a) Labour market: The PTM uses the COMPASS forecast for total hours, along with other exogenous variables, to produce forecasts for average hours, employment in heads, the participation rate, the unemployment rate and other labour market series;

⁸¹For example, the saving rate is a nonlinear function of consumption and income. Most, though not all, of the data come from ONS sources.

- (b) GDP components: COMPASS does not produce forecasts for every expenditure component of GDP. The PTM produces disaggregated forecasts for components which are not separately identified in COMPASS, chiefly stockbuilding and housing investment.
- (c) Household income: COMPASS can be used to project the income that households might receive from wages and salaries. But households also receive income from transfers, dividends, interest and rent, as well as paying taxes. These are difficult to model in a tractable DSGE setting, but a full assessment of household consumption, saving and income needs to take into account these additional variables and associated economic channels.
- (d) Fiscal conditioning paths: The PTM also produces projections for nominal government spending (with separate identification of consumption, investment, paybill and net transfers). The PTM cannot easily be used to study the impact of fiscal policy on agents' behaviour. However, forming a view about the flow of funds between sectors is an important part of cross-checking the forecast. The PTM produces forecasts for net lending by each sector (households, companies, government and rest of the world).

Example 1: Forecasting unemployment

The PTM uses the hours worked forecast from COMPASS as the primary input to generate a projection for the unemployment rate. The following set of equations and identities, which is fairly typical of the rest of the PTM, show how this is carried out in the current version of the model:

$$AVHRS_t = AVHRS_{t-1} \left\{ 1 + \frac{1}{100} \begin{bmatrix} -0.095 + 0.49 \left(100 \times \frac{HRS_t}{HRS_{t-1}} - 1 \right) \\ -0.24 \left(100 \times \frac{AVHRS_{t-1}}{AVHRS_{t-2}} - 1 \right) \\ -0.46 (AVHRS_{t-1} - AVHRS_{t-1}^{trend}) \\ + AVHRS_t^{res} \end{bmatrix} \right\} \quad (26)$$

$$EAGGHDS_t = \frac{HRS_t}{AVHRS_t} \quad (27)$$

$$ER_t = \frac{EAGGHDS_t}{NHDS_t} \quad (28)$$

$$LP_t = LP_{t-1} + (LP_t^{trend} - LP_{t-1}^{trend}) - 0.31(LP_{t-1} - LP_{t-1}^{trend}) + 0.15(ER_t - ER_t^{trend}) + LP_t^{res} \quad (29)$$

$$UNEMP_t = 1 - \frac{ER_t}{LP_t} \quad (30)$$

where HRS is total hours worked, $AVHRS$ is average hours worked, $EAGGHDS$ is employment in heads, ER is the employment rate (as a share of the labour force), $NHDS$ is the size of the labour force, LP is the participation rate, $UNEMP$ is unemployment as a share of the labour force and $AVHRS^{res}$ and LP^{res} are residuals which map the respective variables to the data and through which judgement can be applied. The variables LP^{trend} , and $AVHRS^{trend}$ are trend paths based on structural factors such as demographic change, and estimates of them are produced using other models in the suite.

The employment trend ER^{trend} in this specification of the equation is based on a measure generated using a Hodrick-Prescott filter.

Equations (26) and (29) are essentially error correction models, allowing for slow adjustment of average hours and participation towards their long-run paths. However, the equations also allow for variations over the cycle: when overall hours worked are high, average hours and participation would also be expected to pick up, all else equal. Figure 7 shows the fit of Equation (26) over the past.

Figure 7: Fit of the average hours equation (26) in the PTM



Notes: The chart shows the data for whole economy average hours per worker per week and the fitted values from the post-transformation model equation for average hours discussed in the text.

Example 2: ‘Residual’ GDP components

The PTM includes a simple model which forecasts stockbuilding⁸², $SDELSKP$, similar to that in Bank of England (2000):

$$\begin{aligned}
 SDELSKP_t &= SKP_t - SKP_{t-1} \\
 s_t - s_{t-1} &= -0.00052 + 0.18(y_t - y_{t-1}) + 0.25(y_{t-1} - y_{t-2}) \\
 &\quad + 0.28(y_{t-2} - y_{t-3}) - 0.19(s_{t-1} - y_{t-1} - sgdphp_{t-1}) + s_t^{res} \quad (31)
 \end{aligned}$$

where $s_t = \log(SK P_t)$ gives the level of inventories, $y_t = \log(GDP K P_t)$ and $sgdphp_t$ is a trend for the ratio of inventories to GDP, estimated using a Hodrick-Prescott Filter

⁸²Excluding the alignment adjustment.

(see Elder and Tsoukalas (2006)). Stockbuilding is assumed to be driven by short-term movements in GDP, with the level of stocks gravitating towards a certain share of GDP in the medium to long term.

Housing investment, $IHKP$, is projected using a range of separate suite models. Since, in the ONS national accounts,

$$\begin{aligned} GDPK_t = & CKP_t + IKKP_t + GONSKP_t + XKP_t - MKP_t \\ & + SDELSKP_t + IHKP_t + RES_t \end{aligned} \quad (32)$$

where RES_t just picks up components such as the alignment adjustment and statistical discrepancy⁸³, and since the other components are modelled in COMPASS, those two forecasts are sufficient to complete the GDP expenditure breakdown.

There is no guarantee that the suite-implied forecasts of $SDELSKP$ and $IHKP$ will necessarily be consistent with the implied residual from the GDP expenditure breakdown in COMPASS, because the two accounting structures are different.⁸⁴ This is an example of how suite models can be used as a cross-check on forecasts produced using COMPASS. Staff can apply judgements to the COMPASS forecasts for GDP or its components if Equation (32) in the PTM indicates that an inconsistency may be present.

5.3.2 The Balance Sheet Model (BSM)

The BSM (see Benito et al. (2001)) can be thought of as an extension of the PTM, producing forecasts for even more variables, all consistent with the MPC judgements embodied in the central model. Like the PTM, the BSM is backward-looking, recursive and nonlinear, and adds around another 140 variables to the scope of the MPC's forecast. Although these variables are all inter-related, they can be classified into five broad groups:

- (a) Effective interest rates faced by households and firms, using the yield curve and staff analysis on credit spreads;
- (b) Components of household and corporate borrowing, based on the macroeconomic outlook encapsulated in the MPC's forecast;
- (c) Aggregate and sectoral money balances;
- (d) Metrics of balance sheet health for each sector, such as net wealth, capital gearing, income gearing and the debt-to-income ratio;
- (e) Other important variables from the financial and income accounts, such as dividend, tax and interest flows, accumulation of liquid assets and disposable income for each sector.

An example of a typical BSM forecasting equation is given below. This shows how forecasts for real consumption, prices and unemployment, taken from COMPASS and

⁸³These are typically small. The term RES_t also includes residuals generated by the chain-linking process. As data become older, this residual can become very large. However, in the base year and thereafter (which is most relevant for forecasting), the contribution from such effects will be zero.

⁸⁴See footnote 67 on page 28 for a brief discussion.

the PTM, can be used to produce a forecast for household unsecured credit. The current specification is:

$$\begin{aligned} \log(CREDH_t) = & \log(CREDH_{t-1}) + \log(CCP_t) - \log(CCP_{t-1}) \\ & - 3.36(UNEMP_t - UNEMP_{t-1}) + CREDH_t^{res} \end{aligned} \quad (33)$$

where $CREDH$ is the break-adjusted stock of unsecured lending to households, $CCP \equiv PCDEF \times CKP$ is nominal consumption, $UNEMP$ is the unemployment rate (see Section 5.3.1) and $CREDH^{res}$ is a residual which can be used to impose judgement. A more complete equation listing can be found in Benito et al. (2001).

The BSM provides an important cross-check on the MPC's forecast. The profiles from the BSM are used regularly as an input to MPC discussions, and judgements about financial conditions and the nominal environment are often fed back into the projections in COMPASS.

5.4 Models which generate alternative forecasts

Even in those cases where a variable is modelled in COMPASS, and the model articulates some economic channels through which it can respond to shocks, the model will still be misspecified in some way. As a result, it would be unwise to rely on a single model. The final class of suite models comprises those which might shed light on such misspecifications, by offering an alternative view. Of course, the suite models themselves will be misspecified, so Bank staff pay careful attention to the assumptions underlying each model, and the economic channels which are (and, as importantly, are not) articulated in them.

One set of models within this class are statistical forecasting models, such as those described in Kapetanios et al. (2008). They derive a wide range of models for GDP and inflation, from simple benchmark models such as univariate time series equations through to Bayesian VARs and large factor models. There are currently around 15 models in the "statistical suite"; they are normally used to produce judgement-free forecasts which act as a cross-check on the MPC's projections.

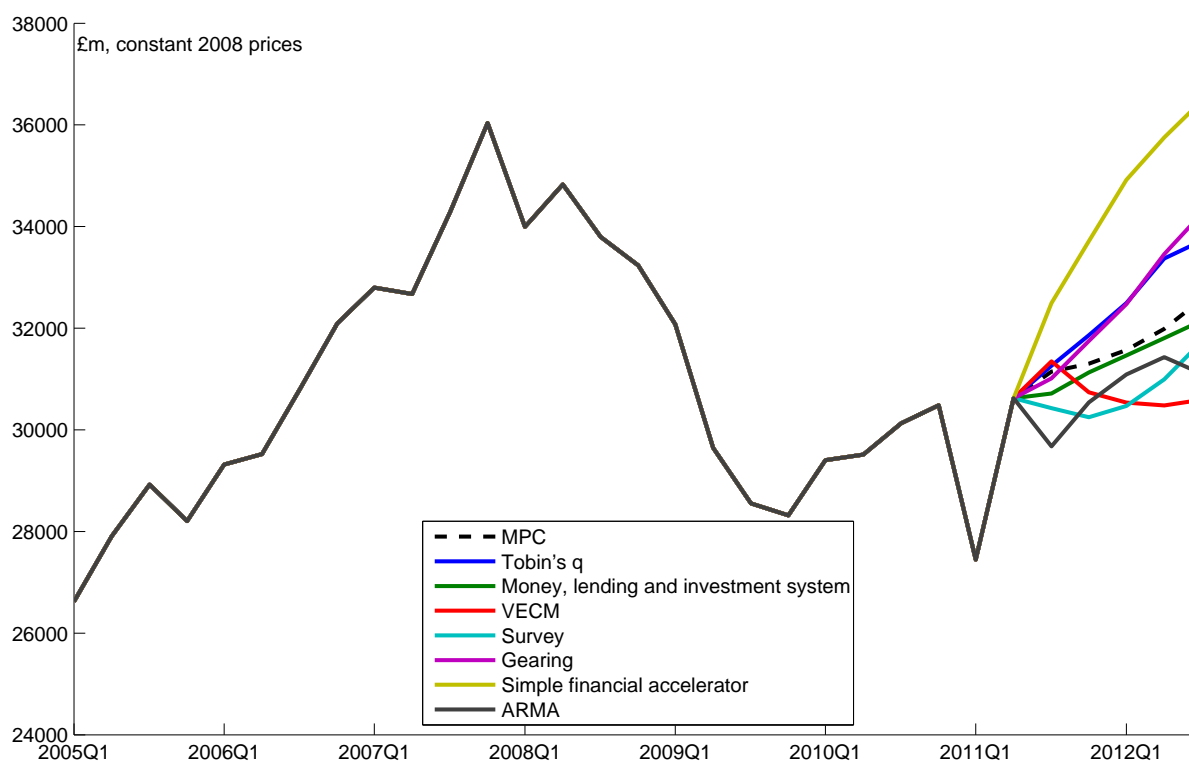
The list below describes some of the other models in this class. The majority use simple econometric relationships, such as error correction models (ECMs) (see Davidson et al. (1978)), to produce alternative forecasts of some of the key variables in COMPASS. These more traditional models have undoubted strengths: they are usually simple to understand, can quickly identify potential inconsistencies in COMPASS-based forecasts, and in many cases have an established track record in the Bank's forecast process. However, they also have limitations when compared with more structural models. They are not designed to produce joint forecast densities for the complete set of COMPASS observable variables, which makes direct comparison problematic. Moreover, in some cases, they can only produce conditional forecasts, taking some variables from COMPASS and other suite models as inputs. As a result, their forecasts may not be fully independent of all the judgements captured in the central organising model.

- Consumption models: The suite contains several 'Keynesian' consumption functions, which model household spending as a function of current labour income. These can be augmented with other factors, such as financial wealth, housing wealth, unemployment and interest rates. Because the treatment of the relationship between wealth and consumption in COMPASS is a simple one, the use of suite

models enables a richer treatment of these interactions. There is also a model which expresses the household saving rate as a function of credit conditions, the unemployment rate and the ratio of household wealth to income, inspired by a model in Carroll et al. (2012) estimated on US data.

- Suite of investment models: Although business investment is included within the current central organising model, modelling investment in a DSGE framework is challenging, and this is one dimension in which COMPASS could be badly specified. The suite of investment models provides an important cross-check, expanding the set of explanatory variables, and using a variety of functional forms. There are seven models currently in the investment suite, which are described below. Figure 8 shows a comparison of the forecasts from these seven models with the MPC's modal forecast for business investment consistent with the November 2011 *Inflation Report*.
 1. ARMA model: A simple baseline model, expressing business investment as a function of lag dynamics (see Box et al. (1970));
 2. Simple financial accelerator model: An ECM which assumes that in the long run, the level of investment depends on the level of GDP, the capital stock and the cost of capital, but that in the short run, financial channels such as firms' cash flow, interest payments and net financial assets play an important role.
 3. Gearing model: An ECM, which assumes that in the long run, investment is determined by GDP, the cost of capital and the "gearing disequilibrium": the extent to which firms' debt levels are away from a "target" level determined by tax incentives and the risks of distress (Bunn and Young (2004));
 4. Money, lending and investment system: A three-equation VECM which jointly models business investment, non-financial companies' money holdings and M4 lending to non-financial companies. A range of other explanatory variables are included, such as spare capacity and firms' retained earnings. See Brigden and Mizen (1999);
 5. Tobin's Q model: A model for the ratio of investment to the capital stock, which in the long run depends on a proxy for Tobin's Q, the value of the firm (see Kapetanios et al. (2006));
 6. Survey model: This uses the investment intentions balances in the BCC Quarterly Economic Survey to project investment in the year ahead;
 7. VECM: A four-equation system embodying two assumed long-run relationships. One relates investment to the size of the capital stock; the other is based on a profit-maximising condition and links the capital-output ratio to the real cost of capital. See Ellis and Price (2004).
- Simple wage equations: The suite includes a number of models for studying the labour market, many grounded in the tradition of Layard et al. (1991). Some of these are single equation models which describe how nominal wages vary with productivity, slack in the economy and inflation expectations; others are systems of equations which articulate more complex interactions between wage determination and price-setting.

Figure 8: Forecasts from the individual investment suite models in November 2011



- Alternative models of inflation: COMPASS is primarily designed for understanding the factors determining inflation over the medium term. Over shorter horizons, other models may have a comparative advantage. The suite includes a detailed supply chain model, which describes how changes in commodity and other input prices are passed through to final consumer prices, and other models which provide a more detailed account of firms' cost structures and pricing decisions. A wide range of judgement-based and statistical tools are also available for forecasting inflation over the first year of the forecast.⁸⁵
- Suite of trade models: Stand-alone models exist to forecast export and import volumes, and prices. For example, export volumes can be modelled as a function of world trade and relative export prices (that is, UK export prices relative to world export prices expressed in sterling terms). Import volumes can be modelled as a function of UK total final expenditure (suitably weighted for import intensity) and relative import prices.⁸⁶ Although inspired by the same theory as that embodied in COMPASS, the estimation of these equations can relax some of the restrictions applied in COMPASS to derive different, and possibly richer, dynamics.⁸⁷

⁸⁵For example, the first two quarters of the MPC's inflation forecast are heavily guided by the Staff's "Short Term Inflation Forecast" (STIF). This is a bottom-up forecast of inflation which projects components of the CPI basket based on a variety of inputs, such as intelligence from the Bank's Agents, commodity prices, specific information about known forthcoming price changes, and simple statistical models.

⁸⁶See, for example, equations (6.2.18) and (6.2.19) in Bank of England (2000).

⁸⁷For example, exports in COMPASS have a unit world demand income elasticity – see equation (A.96) – and imports are part of a Cobb-Douglas production function for final output and so have a unit price elasticity – see equation (5).

In situations where the suite contains several alternative forecasts for a single variable, such as inflation, one option available to Bank staff is to use weighting techniques to combine the individual forecasts. This is often done using forecasts from the statistical suite, but is less common for other combinations of models. A more typical approach is to produce forecasts using several models for a given variable and to understand the economics of why they differ, so that the MPC can use that information to decide whether or not to make further adjustments to the forecast produced using COMPASS.

6 The IT infrastructure

The IT infrastructure was created jointly with the rest of the forecast platform and so was designed to help meet the overall vision described in Section 2. In particular, the design paid close attention to the following considerations. First, the IT infrastructure was designed to support a range of different models. This reduces the costs of using the suite of models and, without it, the suite of models approach would not be viable. Second, it was designed to support the forecast process efficiently with the aim of maximising the amount of time available to Bank staff for analysis of the forecast as inputs to the MPC's forecast discussions.

The infrastructure is comprised of two components: a user interface called Economic Analysis and Simulation Environment (EASE) and a modelling toolbox called Model Analysis and Projection System (MAPS). The rest of this section describes EASE and MAPS in more detail, starting with EASE.

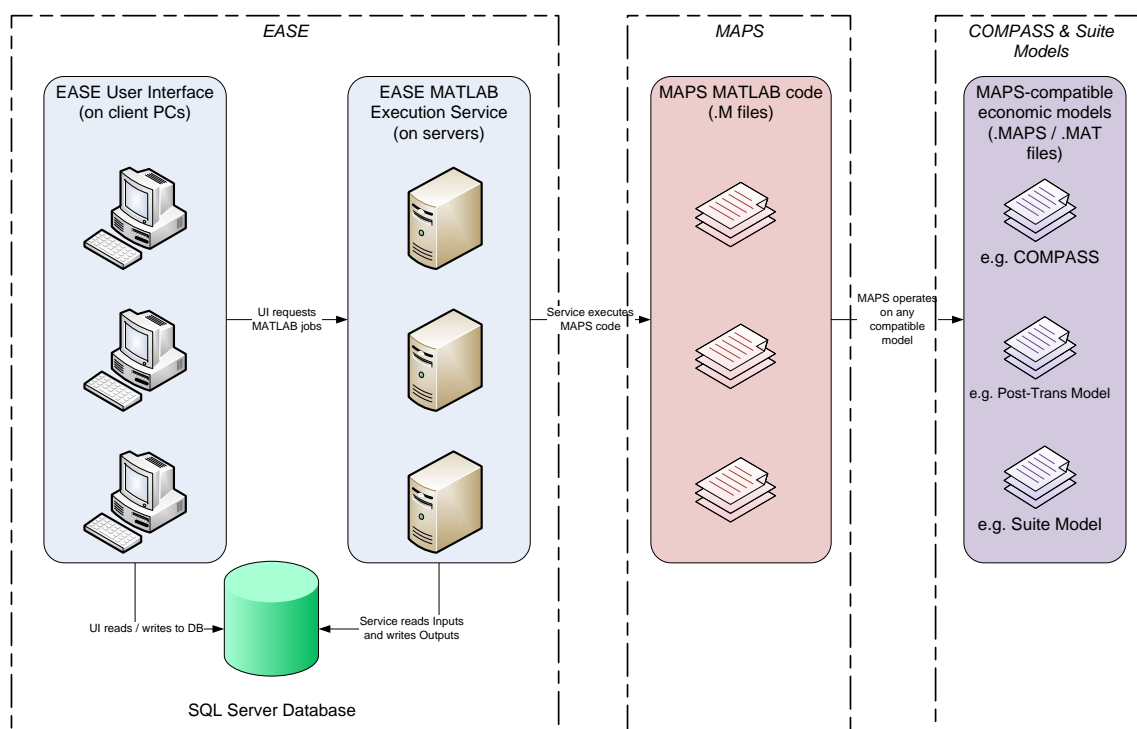
6.1 EASE

Figure 9 depicts the EASE architecture at a high level. The figure can be read from left to right (and back again). EASE is a desktop application that Bank staff use to build and analyse the forecast (see Section 8 for a description of the forecast process). It is a rich Windows Forum application, built in C# on the Microsoft.NET platform, and provides an easy-to-use interface for the underlying forecasting software, MAPS.⁸⁸ Instructions are passed from the EASE client to an execution service on a server, which manages the acquisition of the necessary data from a SQL Server database and executes the relevant function in MAPS, connecting to Matlab via COM automation. The results are stored in the same database, and are then accessible through the EASE client for visualisation and interrogation. The range of visualisation and decomposition tools is a key part of the EASE design, making use of various MAPS tools (which are described below) and a range of time series manipulation functions that form part of the Bank's in-house Time Series Viewer software. Finally, as emphasised above, the design of MAPS and EASE is such that identical operations can be carried out on multiple models, provided that those models meet the criteria for inclusion in the system, which is explained in Section 6.2.1 below.

The key guiding principle for the design of EASE was that it should support the forecast process efficiently. In particular, it provides a straightforward work-flow for the processes of combining a model with data to produce a projection, adjusting that projection to incorporate judgement (see Section 6.2.4) and visualising the result in a number of different ways. It also supports the efficient use of a wide range of suite models by providing a single store for all suite models used in the forecast process and by making it easy to transfer data from one model to another, in order to support applications of the type described in Section 8.2 and Section 8.3. The rest of this section describes the modelling framework and functionality in MAPS that underpins EASE.

⁸⁸Unlike MAPS and other parts of the forecasting platform, EASE was built and is maintained by IT staff at the Bank rather than economists.

Figure 9: Overview of EASE architecture



6.2 MAPS

MAPS is a MATLAB toolkit built and maintained by economists at the Bank. As described above, MAPS contains the underlying functionality necessary for the construction of a forecast using EASE. MAPS is also designed to be operable independently of EASE, thereby supporting additional parts of the forecast process not covered by EASE like estimation of COMPASS or suite models. MAPS is similar to other MATLAB-based modelling toolboxes like DYNARE (Adjemian et al., 2011), IRIS (see Beneš et al. (2009)) and YADA (see Christoffel et al. (2008)), but is specifically designed to support the forecast process at the Bank with the flexibility to be extended and adapted by Bank staff in the future. Indeed, since the project to build EASE was completed, MAPS has been extended in a number of ways. The rest of this sub-section describes the main modelling framework embodied in MAPS and the functionality provided for models that fit into that framework.

6.2.1 Modelling framework

MAPS is predominantly designed for models in the Linear State Space (LSS) class (see, for example, Hamilton (1986)). More precisely, MAPS provides a broad range of functionality (as described below) for any model that can be written as follows, where x is a vector of stationary, mean-zero model variables (which could be partitioned into ‘predetermined’ and ‘non-predetermined’ variables), z is a vector of disturbances or shocks, Y is a vector of model observables and ω is a vector of measurement errors. Both the disturbances, z and measurement errors, ω , are assumed to be independent of each other and across time

(iid) with standard normal distributions.⁸⁹

$$x_t = Bx_{t-1} + \Phi z_t \quad (34)$$

$$Y_t = D + Gx_t + V\omega_t \quad (35)$$

The vector of model variables, x , includes variables that are unobservable (i.e. that do not load into the vector of observables, Y). For example, COMPASS contains over eighty variables, but only fifteen of those load into the observables. Given data for the observables, Y , and the state space model described by equations (34) and (35), these unobservable variables can be uncovered using the Kalman filter and smoother.⁹⁰ Although the above modelling framework appears restricted by the lag order of the state transition equation, systems with higher order lags can be nested in this framework by adding an appropriate set of lag identities into the state transition equations.

The LSS class is a fairly broad class of linear model that encompasses several popular model types. One of those types is linearised DSGE models, like COMPASS, and like several of the suite models described in Section 5. Such models have the following structural equations in MAPS.⁹¹

$$H_B x_{t-1} + H_C x_t + H_F E_t x_{t+1} = \Psi z_t \quad (36)$$

As long as a standard set of conditions is met (see, for example, Blanchard and Kahn (1980)), this set of second-order difference equations can be solved numerically under the assumption of rational expectations to yield the state transition equation in equation (34) and combined with a set of measurement equations as in (35) to form a linear state space model. There are several widely available numerical procedures for solving rational expectations models, including Blanchard and Kahn (1980), Klein (2000) and Sims (2002). Following many other central banks, MAPS employs the AiM procedure developed at the Federal Reserve Board (Anderson and Moore, 1985), which has proved to be reliable and efficient.

Other models that fall into the LSS class and that are supported by MAPS include VARs, structural VARs (like the one used in Section 8.3), Bayesian VARs, and DSGE-VARs. Crucially, once in state space form, much of the same functionality described below can be used identically regardless of the underlying type of model being used.

The MAPS LSS modelling framework also allows for an optional set of data transformation equations which facilitate a flexible detrending of data into stationary, model observable units. For example, one of these data transformation equations in COMPASS transforms GDP, which is (chain-volume) measured in £bns, into quarterly GDP growth (see Table 1 on page 26 for the complete set of data transformation equations in COMPASS).

These data transformation equations also allow for the optional inclusion of a set of deterministic time-varying trends, which is useful when one or more data series is not compatible with assumptions made about trends within the state space model itself. For example, as discussed in Section 4.3, an inflation trend is employed in COMPASS to account for the fact that the mean level of inflation was higher prior to the advent of

⁸⁹This standardisation is inconsequential. Non-unitary standard deviations for the shocks and measurement errors can be entered into the loadings in the matrices Φ and V .

⁹⁰See, for example, Hamilton (1994) for discussion and derivation of the Kalman filter.

⁹¹Consistent with the discussion above, this setup encompasses models with higher order lags and leads because those models can be written compactly with appropriate lag and lead identities.

inflation targeting.⁹² Importantly, these data transformation equations can be inverted, so MAPS (and, by extension, EASE) produces forecasts for the observables in their untransformed space, as well as in model space. This allows staff at the Bank to apply any desired transformation to the underlying forecast data using a broad set of pre-existing time series functions (see Section 6.1 for brief discussion). This approach underpins a range of tools in MAPS and EASE designed to facilitate efficient analysis of forecast outputs.

6.2.2 Estimation

As discussed in Section 2, an important advance in applied structural modelling over the past few years is the ability to estimate relatively large models like COMPASS as systems. The linear state space model outlined in equations (34) and (35) yields a likelihood function that relates the underlying structural parameters of the model, in a vector θ ,⁹³ to a sample of model observable data, $p(Y_T|\theta)$. In principle, the likelihood function, which can be evaluated using the Kalman filter, could be used to estimate the parameters of an LSS model – so-called maximum likelihood. But estimation of medium to large-scale structural models like COMPASS and some of the models that form part of the suite is fraught with identification problems.⁹⁴ Reflecting that, Bayesian estimation has become popular both at central banks and in academia as a practical means for overcoming these problems.

Bayesian estimation is based around Bayes’s rule outlined in equation (23) on page 23. The essence of Bayesian estimation is the combination of ‘prior’ information in the form of prior beliefs about the distributions of the parameters being estimated with the likelihood function, which updates those beliefs given the data. Given the widespread use of Bayesian methods in macroeconomics and for the sake of brevity, discussion of Bayesian methodology is not included here. For a formal discussion of Bayesian methods see, for example, Geweke (2005). And for a discussion of applied Bayesian estimation in the context of DSGE models see, for example, An and Schorfheide (2007).

MAPS includes a fairly standard toolkit for Bayesian estimation of linear state space models. These tools support the estimation of COMPASS, as outlined in Section 4.3, and other linear state space models that form part of the suite discussed in Section 5. The fundamental goal of Bayesian estimation is to characterise the posterior distribution. As a precursor to that, it is often useful to first estimate the parameter vector that maximises the posterior density (the mode). MAPS includes functionality for posterior optimisation, supported by a range of alternative optimisation routines, including all those in MATLAB’s optimisation and global optimisation toolboxes, Chris Sims ‘csminwel’ routine (first used in Leeper and Sims (1994)) and a version of a routine called Covariance Matrix Adaptation Evolutionary Strategy (see Hansen and Kern (2004)).⁹⁵ The toolkit also includes functionality for simulation of the posterior distribution using

⁹²In principle, these data transformation equations could also be used to implement popular detrending methods like Hodrick-Prescott filtering.

⁹³That is, those that form part of the original, structural form as in equation (36) in the context of a forward-looking model. For example, one of the structural parameters of COMPASS is the intertemporal elasticity of substitution – see equation (14) on page 18.

⁹⁴See, for example, Iskrev (2010).

⁹⁵We have found the CMA-ES routine to be quite robust and reliable for estimation of large-scale models. See Andreasen (2010) for a comparison of the CMA-ES routine with a version of Simulated Annealing in the maximum likelihood estimation of a DSGE model.

Markov Chain Monte Carlo (MCMC) simulation and, more specifically, the random walk Metropolis-Hastings algorithm (see Schorfheide (2000)). The typical strategy employed for estimation in MAPS is to use posterior optimisation to find an initial condition for a posterior simulation of multiple chains (distributed across a server cluster) containing multiple draws.⁹⁶ This toolkit and approach was used to produce the estimation of COMPASS as described in Section 4.3.

Finally, the Bayesian estimation toolkit also includes a set of simulation diagnostic functions (chain convergence, auto-correlations and Gelman-Rubin statistics), an auto-generated report function which produces a PDF report containing the results and diagnostics of the estimation, and functions that extend the model analysis tools to include posterior parameter uncertainty such as the function used to produce impulse responses that take into account posterior uncertainty around the parameter estimates for COMPASS in Section 4.4.

As well as Bayesian estimation, MAPS also contains a toolkit for Minimum Distance Estimation of a linear state space models parameters to come closest to achieving a (reference) set of model properties.⁹⁷ This functionality was used in the COMPASS estimation described in Section 4.3 to set the prior means for the shock standard deviations. It could also be useful in (quickly) estimating suite models designed for a particular purpose (such as the quantification of a particular shock), where full-information system estimation may not be necessary.

6.2.3 Model Analysis

A pre-requisite for understanding forecasts produced using any model is an understanding of the model itself. MAPS provides a range of standard tools for doing that. In particular, MAPS provides functionality for impulse responses, ‘fix’ responses, model observable autocovariances and forecast error variance decompositions.⁹⁸

Impulse responses are a standard tool to aid understanding of the economic mechanisms embodied in a structural model. For example, Section 4.4 shows the responses of a range of endogenous variables in COMPASS to three different unanticipated shocks. Impulse responses are a good starting point for understanding a model because they isolate its response to single realisations of the shocks, which are the exogenous drivers of the model. But, as discussed in Section 6.2.4 below, a fundamental part of forecasting is to apply judgement by ‘fixing’ the path of a particular endogenous variable. Fix response functionality in MAPS helps to bridge the gap between impulse responses and forecasting by providing a controlled environment in which ‘fixes’ can be applied to a model and the consequences assessed. For example, impulse response functionality could be used to assess the response of activity and inflation to a one-off monetary policy shock, while fix response functionality could be used to assess the response of activity and inflation to a sequence of monetary policy shocks which raise the level of interest rates by 1pp for 12 quarters.

The linear state space model described by equations (34) and (35) implies a set of covariances among the endogenous variables. In particular, if the model is covariance stationary (i.e. if all eigenvalues of B lie inside the unit circle) then the unconditional

⁹⁶A number of these draws are usually burned and the resulting chains are then thinned.

⁹⁷This could be used to implement GMM estimation. See Christiano et al. (2005) for example.

⁹⁸The decompositions described in Section 6.2.5 can also be useful as part of the suite of model analysis tools.

covariance of the model variables can be defined as:

$$P = BPB' + \Phi\Phi' \quad (37)$$

from this, it is straightforward to define the covariances of the model observables, Y as:

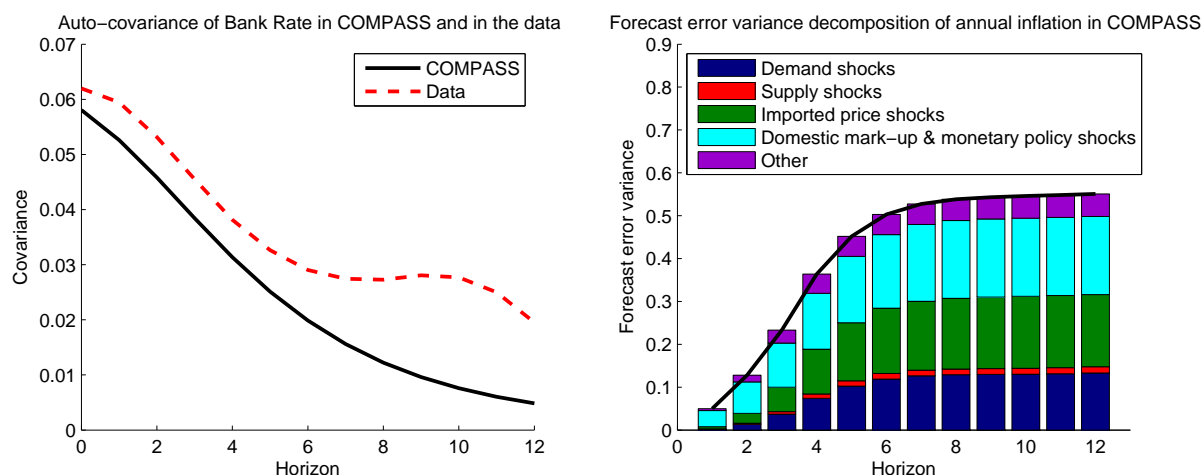
$$E[Y_t Y_t'] = GPG' + VV' \quad (38)$$

or any auto-covariance as (where $h \geq 1$):⁹⁹

$$E[Y_t Y_{t-h}'] = GB^h PG' \quad (39)$$

These model observable autocovariances can be readily compared to equivalent measures from a sample of observable data such as that used to estimate an LSS model (as described in Section 6.2.2) to provide a comparison of how well the model matches various moments in the data. As an example, the left-panel of Figure 10 shows a comparison of the auto-covariance of the interest rate in COMPASS with that in the data over the estimation sample (1993Q1-2007Q4).

Figure 10: Examples of autocovariance & forecast error variance decomposition model analysis outputs using COMPASS



Notes: The left panel shows auto-covariances of Bank Rate at various horizons from the estimated COMPASS model described in Section 4.3 compared to that in the data over the estimation sample (1993Q1-2007Q4). The right panel shows a forecast error variance decomposition of annual inflation from the estimated COMPASS model. The contributions of the shocks have been grouped together in the same way as in Figure 20 on page 82.

Forecast error variance decompositions provide a useful decomposition of the variances of a model's forecast errors at any desired horizon into contributions from each of the shocks of the model. They can be derived from the definition of a forecast error at any horizon, h :

$$x_{t+h} - E_t x_{t+h} = \sum_{s=1}^h B^{h-s} \Phi z_{t+s} \quad (40)$$

and it follows that the variance of the forecast error is defined as:

$$E[(x_{t+h} - E_t x_{t+h})(x_{t+h} - E_t x_{t+h})'] = \sum_{s=1}^h B^{h-s} \Phi \Phi' (B^{h-s})' \quad (41)$$

⁹⁹Noting that the measurement errors are by assumption uncorrelated across time.

which relates forecast errors to the structural disturbances. A decomposition of the forecast error variance can be computed using the above definition by adding the relevant columns of the matrix Φ one-by-one. As an example, the right-hand panel of Figure 10 shows a forecast error variance decomposition of inflation from COMPASS.

These tools are fairly common to many of the modelling toolboxes referenced in the introduction to this section. What is perhaps less common is the range of tools in MAPS for efficient visualisation of a model's properties. These tools help to support one of the key elements of the suite approach for the forecasting platform, which is to 'use models efficiently'. They include a range of visualisation functions, including a Graphical User Interface for visualisation of a model's impulse responses, comparison of responses across multiple models and exploration of how responses change when a model's parameters change. They also include a set of functions which auto-generate PDF reports containing user-defined sets of the model analysis outputs described above (in the form of charts and tables).

6.2.4 Projection and simulation

The starting point for projection is the Kalman filter and smoother, which uncover estimates for the model variables, x_T , given data for the observables, Y_T , and the state space model described by equations (34) and (35). With those estimates in hand, it is then straightforward to use equation (34) to project them forward to produce a plain-vanilla LSS model projection for any horizon, h , as follows:¹⁰⁰

$$x_{T+h} = B^h x_T \quad (42)$$

As discussed at length in Section 8, judgement is a key part of the forecast process, so a plain-vanilla projection is merely a starting point for further analysis. Sections 7 and 8 discuss the rationale for applying judgement to a forecast but, from a mechanical perspective and therefore the functionality in MAPS, the precise reason for a judgement is irrelevant. Fundamentally, all judgements are implemented in MAPS by assigning values to the structural disturbances of the model over the forecast horizon. In addition to the standard set of unanticipated disturbances that form part of state space models as in equation (34), MAPS also allows for forecast judgements (in forward-looking models) to be implemented using anticipated disturbances. This provides greater flexibility over the economic assumptions underpinning particular judgements and also facilitates analysis of expectations and policy as described in Section 6.2.6. Indeed, the default assumption for judgements applied to COMPASS in constructing the MPC's forecast is that they are treated as anticipated.¹⁰¹

¹⁰⁰And, given this projection, it is also straightforward to recover a consistent projection for the observables using equation (35) (and the fact that measurement errors are zero in a forecast).

¹⁰¹Note that it is now necessary to assign a horizon up to which anticipated shocks affect the forecast (since, in principle, anticipated shocks into the infinite future matter). As Section 8 suggests, that horizon, H , is usually set to twelve quarters to construct the MPC's forecast. Note also that equation (43) implies that the standard deviations of the anticipated disturbances are identical to the unanticipated disturbances. This is an assumption borne of convenience and, in principle, it would be possible to associate the anticipated shocks with a different Φ matrix. In practice, for estimated models like COMPASS it would be hard to do so because the estimated version of the model is based on equations (34) and (35) in which all disturbances are unanticipated. Identifying the anticipated shocks in addition to those would require the model to be augmented with data representing agents' expectations in each time period, posing problems for the measurement of those expectations and for the practical task of estimation given the resulting size of the model. See Del Negro and Schorfheide (forthcoming), for example.

This means that a projection in MAPS is made up of the following, where a distinction is drawn between unanticipated, u , and anticipated disturbances, a :

$$x_{T+h} = B^h x_T + \sum_{s=1}^h B^{h-s} \Phi u_{T+s} + \sum_{s=1}^{h-1} B^{h-s} \sum_{j=s+1}^H F^{j-s} \Phi a_{T+j|T+s} + \sum_{s=h+1}^H F^{s-h} \Phi a_{T+s|T+h} \quad (43)$$

where:¹⁰²

$$F = -(H_C + H_F B)^{-1} H_F \quad (44)$$

In words, an h -step ahead judgemental projection is the sum of the initial conditions projected forward h periods, the projected impact of unanticipated shocks realised over the first h periods, the projected impact of anticipated shocks realised over the first $h - 1$ periods, and the impact of future anticipated shocks. The same equation can be expressed recursively as follows:

$$x_{T+h} = Bx_{T+h-1} + \Phi u_{T+h} + \sum_{s=h+1}^H F^{s-h} \Phi a_{T+s|T+h} \quad (45)$$

As discussed above, forecast judgements are embodied in values for the structural disturbances, a and u . It follows that the most obvious way of applying judgement to a forecast is by changing the values for those disturbances. But, as is made clear in Section 8, judgements are typically applied to endogenous variables directly, rather than indirectly via the shocks. For example, it would be more natural to policy makers to apply judgement to GDP growth rather than to a productivity shock, not least because the units of GDP growth are more familiar. Reflecting that, MAPS contains a powerful and flexible toolkit for the imposition of judgement. In particular, the judgement toolkit allows for judgements to be implemented to the shocks or to forecast paths for endogenous variables, often known as ‘fixes’ or ‘conditioning’. If judgement is applied via conditioning of endogenous variables it is necessary to ‘invert’ those conditioning paths to find the values for a (chosen) sub-set of the disturbances that support them. The MAPS inversion technology is flexible and allows for judgements to be implemented using a mix of anticipated and unanticipated disturbances, identifying them either exactly, or using one of two identification schemes in cases where the number of shocks (instruments) used is greater than the number of conditioning paths (targets), or minimising the sum of squared deviations from the targets in cases where the number of conditioning paths is greater than the number of shock instruments.¹⁰³ Appendix C outlines the MAPS inversion algorithm in detail.¹⁰⁴

¹⁰² H_C and H_F are defined in equation (36).

¹⁰³The MAPS toolkit is built around what is known as hard conditioning, whereby the conditioning information is imposed with certainty. An extension to the toolkit would be to allow for a degree of uncertainty around the conditioning information, known as soft conditioning and as discussed in Maih (2010).

¹⁰⁴As is made clear in the appendix, MAPS makes two simplifying assumptions relative to those implied by equation (43). First, for convenience MAPS allows for anticipated shocks to occur contemporaneously. These shocks are identical to unanticipated shocks and so, in effect, the vector of unanticipated disturbances that underpin the judgemental projection is partitioned into two. Second, the identification (which is based on a single set of conditioning paths) relies on there being a single set of anticipated shocks over the forecast such that the same set of shocks is anticipated in any arbitrary period of the forecast, j , in all periods up to and including that $T + 1 \dots j$.

In addition to those tools, which allow the forecaster to impose judgements on point forecasts, MAPS also includes prototype tools for the imposition of judgement on entire forecast densities like those published as fan charts in the *Inflation Report* (see Section 2.1 for context).

Finally, MAPS and EASE include a ‘simulation mode’ in which all of the tools described above can be used to explore scenarios. The advantage of simulation mode is that it does not require data, allowing forecast scenarios to be explored more efficiently (noting that, given the inherent linearity of the modelling framework described above, it is valid to add marginal simulations as marginal judgement to a projection).¹⁰⁵ For example, both the VAT scenarios in Section 8.2 and credit scenarios in Section 8.3 were computed in ‘simulation mode’.

6.2.5 Decompositions

As briefly discussed in Section 6.1, a key objective of the IT infrastructure is to make it as easy as possible for forecasters to understand the models being used and the results of using them. Reflecting that, MAPS contains a wide-ranging and flexible decompositions toolkit for shock-based decompositions, hybrid decompositions, equation-based decompositions and ‘flexible’ decompositions.¹⁰⁶

Shock-based decompositions provide a decomposition of an endogenous variable into the contributions of current and past realisations of the exogenous disturbances. The basis for shock-based decompositions can be derived directly from backward induction of the state transition equation (34):

$$x_t = B^t x_0 + \sum_{s=1}^t B^{t-s} \Phi z_s \quad (46)$$

and a similar expression can be derived for forecasts of the endogenous variables taking into account anticipated shocks used to implement forecast judgements (and where the unanticipated shocks over the forecast have been merged with those estimated over the past):

$$x_{T+h} = B^{T+h} x_0 + \sum_{s=1}^{T+h} B^{T+h-s} \Phi u_s + \sum_{s=1}^{h-1} B^{h-s} \sum_{j=s+1}^H F^{j-s} \Phi a_{T+j|T+s} + \sum_{s=h+1}^H F^{s-h} \Phi a_{T+s|T+h} \quad (47)$$

In either case, it is then straightforward to produce a decomposition of the endogenous variables into initial conditions and shocks by adding each shock into the relevant equation one at a time and recording the effect. For example, Figure 20 shows a decomposition of annual inflation into the shocks filtered using COMPASS. And, as discussed in Section 8.1, a shock-based account of the changes in the data forms a part of the forecast process.¹⁰⁷

¹⁰⁵In simulation mode, the vector of initial conditions for the projection x_T are set to zero, which is their steady state. Simulations explore deviations from that steady state base.

¹⁰⁶All of these are accessible via EASE.

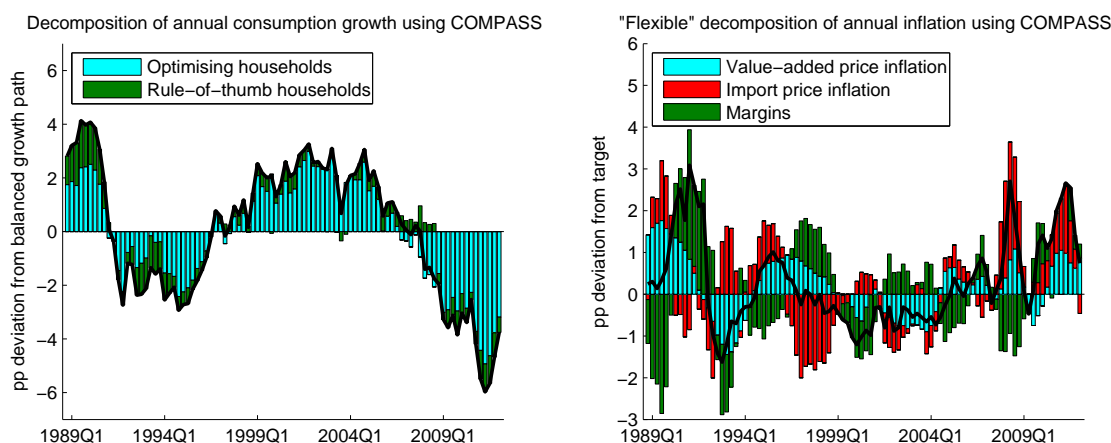
¹⁰⁷A key tool in this part of the process is a PDF document produced by MAPS, which reports decompositions of the changes in the forecasts for the endogenous variables into the shocks and then decompositions of the changes in the shocks into the news in the data. This helps to build an understanding of what has driven news in the projection, which acts as a useful starting point for further analysis.

While shock-based decompositions provide a useful starting point for thinking about what has been driving data, they can only be interpreted as structural if the model is the true data generating process. As discussed in Section 7, all models are misspecified, so the shocks cannot be interpreted as truly structural and are in fact likely to be correlated. In particular, shocks that are not captured by a model are likely to show up in particular correlations among the shocks that are included in a model. In cases where those missing shocks can be identified, it is possible to try to correct for the misspecification. For example, Section 8.2 shows a shock-based decomposition for COMPASS that has been augmented to include VAT (and energy) shocks in a ‘hybrid decomposition’. Hybrid decompositions combine shock-based decompositions with off-model information and judgement. Such decompositions can be written as follows (where x^j represents one or more judgements about off-model effects (which could itself be decomposed) and z^{-j} represents all shocks that do not form part of an explanation for the judgement):¹⁰⁸

$$x_t = B^t x_0 + x_t^j + \sum_{s=1}^t B^{t-s} \Phi z_s^{-j} \quad (48)$$

Shock-based decompositions and hybrid variants are designed to provide a story in terms of the fundamental drivers of the data, the exogenous disturbances. At times, it is also useful to be able to produce accounting decompositions using the structural equations of a model. For example, consumption in COMPASS is comprised of the consumption of rule-of-thumb households and optimising households (see equation (A.288) in Appendix A). The left panel of Figure 11 shows a decomposition using that identity: it shows the contribution of the consumption of each of those types of household to aggregate consumption. These could be interpreted as the relative contributions of current and permanent income in determining aggregate consumption in COMPASS over the sample period.

Figure 11: Example decompositions using COMPASS



Notes: The left panel shows a decomposition of consumption into the contributions of optimising and rule-of-thumb households. The right panel shows a ‘flexible’ decomposition in which an additional equation has been used to decompose annual consumer price inflation into domestic costs, imported costs and margins. Both decompositions use the estimated COMPASS model described in Section 4.3.

Finally, MAPS also includes functionality for ‘flexible’ decompositions. These can be used to produce alternative equation-based decompositions by substituting one or

¹⁰⁸And it is straightforward to derive an analogous expression as the basis for hybrid decompositions of forecasts using equation (47).

more equations through other equations, by combining shock-based and equation-based decompositions together, or by defining new equations consistent with the structure of the model being used.¹⁰⁹ For example, consumer price inflation in COMPASS is defined by a Phillips curve (see equation (10) on page 16). But, by identity, inflation can also be defined as the sum of costs and margins. Flexible decomposition functionality in MAPS makes it possible for Bank staff to define additional equations (which do not explicitly form part of the set of equations in the model) and use them to decompose inflation in the back data and/or forecast. The right panel of Figure 11 shows an example using the ‘costs and margins’ decomposition of inflation described above.¹¹⁰

6.2.6 Expectations & policy analysis

As Section 6.2.4 highlighted, the default assumption for ‘fixes’ is that they are applied using anticipated shocks. One implication of that assumption is that agents’ expectations and the resulting forecast paths are the same. This makes it straightforward for Bank staff to trace through the implications of a particular judgement for agents’ expectations and for the response of monetary policy.

One routine exception to this approach is the application of policy conditioning paths.¹¹¹ As discussed in Section 8.4, the published *Inflation Report* GDP and inflation projections are conditioned on a measure of market expectations of the term structure of interest rates, a path for the exchange rate,¹¹² and the latest announced level of asset purchases. These paths are imposed in exactly the same way as other conditioning information using the toolkit described in Section 6.2.4 with the exception that, unlike most other judgements, these paths are applied using unanticipated shocks.

Although standard and a useful benchmark, the assumption of rational, model-consistent expectations is a strong one. In exploring scenarios in the forecast, it is sometimes useful to explore deviations from rational expectations. To that end, MAPS includes a toolkit whereby expectations for the endogenous variables can be determined by arbitrary backward-looking rules.¹¹³

In light of the Stockton Review and the Bank’s subsequent response (as briefly discussed in Section 2.3), future development of MAPS is likely to include enhancements to the expectations toolkit to incorporate learning and development of a monetary policy toolkit to include automation for the creation of models with monetary policy rules that minimise a desired loss function and models with fully optimal monetary policy.¹¹⁴

6.2.7 Non-Linear Backward-Looking models in MAPS

As described above, the main class of model permitted in MAPS is linear state space (LSS) models. But in order to support the suite of models described in Section 5, MAPS also permits non-linear backward-looking (NLBL) models. These can be defined as follows,

¹⁰⁹This technology also allows for automated grouping of terms in a decomposition. See Figure 10 for an example.

¹¹⁰The same decomposition is also used as part of the analysis of credit frictions in Section 8.3.

¹¹¹Although the application of judgement using anticipated shocks is the most common case, there are exceptions to that as the economics of a particular judgement dictates.

¹¹²Applying 50% weight to an off-model UIP condition and 50% weight to a random walk

¹¹³See, for example, Harrison and Taylor (2012) for a discussion of alternative approaches to and implications of this type of expectations formation.

¹¹⁴These kinds of issues have been explored by other central banks. See, for example, Adolfson et al. (2011).

where x is a vector of endogenous variables, z a vector of exogenous variables and res a vector of residuals.¹¹⁵

$$f(x_t, x_{t-1}, z_t, z_{t-1}, res_t) = 0 \quad (49)$$

More precisely, MAPS supports NLBL models that are recursive in nature. The above non-linear system of equations can be described as recursive if it can be analytically solved to give the following:

$$x_t = g(x_{t-1}, z_t, z_{t-1}, res_t) \quad (50)$$

The COMPASS ‘Post-Transformation Model’ and the Balance Sheet Model described in Section 5.3 are both good examples of large, recursive NLBL models. In both cases, many of the variables that form part of the exogenous set are from COMPASS. For example, the Post-Transformation Model can be used to produce a projection for unemployment given a forecast for total hours from COMPASS (and other inputs).

Relative to the LSS modelling framework, the functionality in MAPS that supports NLBL models is more limited. Specifically, MAPS contains a set of tools for building judgemental projections using NLBL models (similar to those described above for LSS models).¹¹⁶ In addition, all of the tools in EASE for visualisation and transformation of time series data can be used for NLBL models in exactly the same way as for LSS models.

¹¹⁵As with LSS models, this does not restrict the lag order of the system because lag identities can be added.

¹¹⁶The possibility of over-identifying fixes to endogenous variables in NLBL models is not catered for reflecting that NLBL models tend to have much less rich structures (which restricts the feasibility of using many of the residuals to fix a particular variable) and that there is no obvious identification scheme given that in many cases these models are not estimated and so there is no metric for the relative standard deviations of the residuals.

7 Forecasting with misspecified models and the role of judgement

As described in Section 2.3, our approach takes as its starting point the view that “all models are wrong, some are useful”. This motivates use of significant judgement in the construction of the forecast, with this judgement informed by a range of models. In this section, we describe how this approach is applied to a forecasting platform based around a central organising model that we know to be misspecified. The approach outlined in this section will be applied to realistic examples in Section 8.

The design of the forecasting platform is based on the view that a structural central organising model is a useful device for helping to identify the economic shocks driving the data and hence construct a baseline narrative for the data and the forecast. But of course we need to be careful when interpreting the results of this approach. Specifically, just because our model identifies, say, a significant ‘wage mark-up’ shock as important in particular episodes does not mean that we should believe that sharp quarter-to-quarter changes in the competitiveness of the labour market are the key drivers of those episodes. As Chari et al. (2008) point out, standard New Keynesian DSGE models often imply that this type of shock plays an important role in business cycle fluctuations. However, if interpreted literally, the implied quarter-to-quarter changes in labour market competitiveness seem implausible. So it would be unwise to interpret the ‘wage mark-up shock’ identified by the model as changes in the elasticity of labour demand. Indeed, Chari et al. (2008) show that it is impossible to disentangle the effects of this shock from competing labour market assumptions.¹¹⁷ This means that we need to probe more deeply behind the ‘labels’ that the central organising model places on the important shocks driving the data in order to uncover more fundamental stories and to reflect that in the forecast accordingly. Moreover, as is discussed below, COMPASS abstracts from a number of important economic channels that are relevant for the outlook. Capturing the effects of these missing channels in a coherent way is an important part of dealing with misspecification.

In Section 7.1, we describe our general approach for probing and adjusting the provided by the central organising model, COMPASS. In Section 7.2, we consider some challenges to this approach and the implications for the design of the forecasting platform.

7.1 The ‘misspecification algorithm’

Our approach can be summarised in terms of a generic ‘misspecification algorithm’, which staff work through when trying to incorporate the effects of misspecification into a forecast organised using COMPASS. The generic algorithm we follow is as follows:

1. Understand the economics of the misspecification and how it relates to the economic behaviour incorporated in COMPASS.
2. Quantify the effects of the misspecification.

¹¹⁷Chari et al. (2008) also argue that such concerns make it difficult to believe that these models can be used reliably for welfare analysis. While we accept the sentiment of these arguments, we do not agree with the implication that DSGE models are inherently useless for forecast or policy analysis. Instead, a pragmatic approach suggests that we can use these models as a useful starting point for a deeper economic inquiry.

3. Incorporate the quantitative effects into the forecast built using COMPASS.

Of course, implementing this algorithm is not always straightforward and necessitates the use of judgement. In the following sections, we outline some of the techniques we use in each step, highlighting how the design of the forecast platform facilitates their use.

7.1.1 Understanding the economics of the misspecification

We can group COMPASS misspecifications into two broad categories: the behavioural assumptions giving rise to the model structure may be incorrect; and there may be important shocks and transmission channels that are simply missing from COMPASS.¹¹⁸

The first step in investigating these misspecifications is to understand the assumptions underlying the central organising model and how those assumptions affect the model's interpretation of the past and the forecast. Several features of the forecasting platform support this step. As noted in Section 2.3, COMPASS has deliberately been designed to be relatively small and simple. The theoretical foundations underpinning the behaviour in the model have been chosen to be as standard, well used and well understood as possible. These design choices make it easier for staff to understand the underlying assumptions about economic behaviour and hence where these assumptions may be misspecified. MAPS has been designed to analyse model behaviour efficiently and to make the results easily interpretable. MAPS includes tools to extract, understand and interrogate the model's explanation of the data and the forecast. For example, the ability to specify 'flexible' decompositions of model equations allows model users to examine model-based explanations of the data and forecast efficiently using narratives that are particularly relevant to the model (see Section 6.2.5). Taken together, these design features help staff to understand quickly 'COMPASS-based' explanations of the past data, the forecast and simulations such as impulse responses.

Once we know how COMPASS may be wrong, understanding the potential implications of a given misspecification typically requires us to use models that incorporate the alternative assumptions about behaviour and/or shocks and transmission channels that are missing from COMPASS. For example:

- Although COMPASS includes a share of households that can be viewed as credit constrained (in the sense that they spend their current labour income), that share is assumed to be constant. In reality, the share of households facing credit constraints is likely to depend on the state of the economy (for example, the strength of wages, the levels of interest rates and asset prices). This means that the effects of interest rates and labour income on consumption could be time-varying in a way that depends on variables that are not included in COMPASS. One implication is that the contributions of interest rates and labour income in COMPASS-based accounts of consumption are likely to be understated during episodes in which aggregate activity and asset prices are weak and interest rates are high. So staff may 'aim off' COMPASS-based explanations of the data and forecast appropriately in this instance.
- COMPASS abstracts from financing frictions on investment expenditure, so there is no role for a 'financial accelerator' mechanism such as that modelled by Bernanke et al. (1999). The presence of this type of mechanism would imply that investment

¹¹⁸Since all models are wrong, we know that both of these misspecifications are present.

decisions are constrained by a firm's ability to borrow which in turn depends on an external finance premium based on the firm's net worth. This means that equity prices would affect the real return on capital and hence investment decisions over and above a measure of Tobin's Q based solely on the risk free real interest rate, which is the treatment in COMPASS.¹¹⁹ One implication is that COMPASS-based accounts of investment in terms of a risk free real interest rate measure of Tobin's Q could understate the boost to investment expenditure generated by rising asset prices (to the extent that asset prices increase by more than the increase consistent with the change in real interest rates).

7.1.2 Quantifying the effects of misspecification

The key tools for helping staff to quantify the effects of misspecification in COMPASS are models within the suite. As explained in Section 5, there is a broad range of models in the suite, some of which are intended to capture channels that are misspecified or missing entirely from COMPASS. These models may be empirically focused (for example: an estimated equation for a particular variable in COMPASS; a VAR model of a subset of COMPASS variables) or derived from explicit behavioural assumptions (for example, a calibrated DSGE model containing sectors or behaviour excluded from COMPASS). The key requirement for these particular suite models to be useful for making adjustments to the COMPASS narrative in a systematic way is that they can provide simulations or forecasts that illustrate how a particular misspecification in COMPASS affects the macroeconomic variables in the model. For example, as discussed in Section 5.2.1, the suite contains a version of COMPASS extended to include an explicit treatment of the role of energy. This model can be used to quantify the effects of exogenous shocks to energy prices on the variables that are included in COMPASS. Since, as discussed in Sections 2.1 and 8, the forecast is built in incremental steps by incorporating judgements as marginal updates to the forecast, the marginal effects of exogenous shocks to energy prices can be 'layered on' to a forecast constructed using COMPASS.

Several features of the forecasting platform support this process. The fact that COMPASS is relatively small and simple, built using well-used theoretical foundations, facilitates the process of creating expanded versions of the model to analyse missing shocks, sectors and transmission channels. The IT infrastructure allows us to use the same set of tools to analyse all of the models in the suite and the EASE user interface allows staff to view outputs from these models alongside those from COMPASS.

7.1.3 Incorporating the quantitative effects of misspecification

Once we have a quantification of the effects of a particular misspecification from the suite of models, we need to use that information to adjust the forecast being constructed using COMPASS. Mechanically, the only way to do this is to use shocks to move the endogenous variables in COMPASS in a way that captures the quantitative effect of the misspecification. MAPS provides a very powerful and flexible set of tools for doing this. But judgement is required to select the best method of applying the shocks and, in particular, the subset of shocks that should be used. To guide our judgement we rely on insights from two distinct literatures on model misspecification.

¹¹⁹See Section 4.2.3 for a discussion of investment determination in COMPASS. The absence of financial frictions means that equity prices are given by the shadow price of capital (Tobin's Q) in COMPASS. See footnote 35.

The ‘Business Cycle Accounting’ (BCA) literature pioneered by Chari et al. (2007) contains two important results. First, a number of papers have demonstrated that it is possible to represent a wide variety of models in terms of a very simple ‘prototype’ Real Business Cycle (RBC) model that contains stochastic disturbances to the equilibrium conditions (known as ‘wedges’).¹²⁰ For example, Chari et al. (2007) show that ‘detailed economy’ models with input-financing, financial accelerators, sticky prices and sticky wages can all be mapped into a single prototype RBC model with time-varying wedges. In particular, it is possible to find sequences of wedges in the simple prototype model such that allocations in that model are identical to those in the ‘detailed’ model. In doing so, the sequences of wedges in the simple model required to match the allocations in the detailed model are typically correlated. So the wedges in the BCA approach are not interpreted as ‘structural shocks’.¹²¹

The second strand of literature is the analysis of the ‘robust control’ approach to model misspecification exemplified by Hansen and Sargent (2007). In this approach it is assumed that the decision maker has access to a ‘reference model’ which describes their best estimate of how the economy works. However, the decision maker suspects that their reference model may be misspecified and expresses their uncertainty around the true model in terms of a set of alternative specifications for the shocks driving the model. Hansen and Sargent (2007, p18–19) describe how a range of deeper misspecifications of the structure of the model can manifest themselves as specifications for the shocks which depend on the variables in the model in complex, possibly non-linear ways.

Both the BCA and robust control approaches imply that misspecification of a baseline model can manifest itself in terms of correlation among the shocks to that model. So an important part of our approach is to try to capture the effects of misspecification using a constellation of shocks to COMPASS. In particular, it is likely to be appropriate to use several shocks to try to correct for misspecification of just one assumption underlying the model.

For example, consider the case in which firms can vary their utilisation of labour in the production process (‘labour hoarding’).¹²² In this case, the value added production function in COMPASS is misspecified because it assumes that value added depends on total hours worked but not the fraction of hours devoted to production activity.¹²³ One way to incorporate the effects of a reduction in labour utilisation on value added would be to use a negative TFP shock to reduce the value added produced given the observed level of *total* hours worked. But such a modification would have effects elsewhere in the model. For example, the TFP shock would increase the marginal cost of value added producers, putting upward pressure on value added prices, which is unlikely to be consistent with the underlying story (reduced labour utilisation). In this case, shocks to value added markups can be used to offset the effects on marginal cost of the TFP shock. Depending on the precise way in which variable labour utilisation is thought to affect production, we may also expect the equations governing wage setting and employment to be affected by such a misspecification. This may require additional shocks to be used to ensure that the overall response of the model captures the effects of the misspecification of the production

¹²⁰The ‘wedges’ can be thought of in terms of imaginary distortionary taxes that push equilibrium prices and allocations away from their undistorted values.

¹²¹See Sustek (2011) for a discussion of the relationship between DSGE and BCA approaches. Our view is that the methodological differences between proponents of the BCA and DSGE approaches reflect some fundamental differences in view over how models should be used.

¹²²Models incorporating this type of assumption include Wen (2004) and Kim and Lee (2007).

¹²³See equation (1) on page 14.

function.

The example above shows that the process of selecting which shocks to use to incorporate the effects of misspecification is very similar to the method used to demonstrate the ‘equivalence’ of detailed and prototype models in the BCA literature. This approach relies on familiarity with COMPASS in order to trace through the potential effects of a particular misspecification. The fact that COMPASS is relatively small and simple and based on well-understood theories helps staff to implement this process efficiently. In cases in which we are able to select an appropriate constellation of shocks to use to incorporate the quantitative effects of misspecification, the MAPS toolkit allows us to select the most likely sequences of those shocks that generate the desired quantitative effects on the endogenous variables.¹²⁴ The example above also shows that understanding the effects of a misspecification requires a good understanding of the economics of the particular case in hand. In many cases, suite models can be used as a useful aid in developing that understanding.

In addition to the tools within MAPS, other approaches can be used to incorporate the quantitative effects of misspecification into a forecast constructed using COMPASS. Caldara et al. (2012) explore the propagation of shocks originating in sectors that are not present in a central organising model. The authors show how a structural model can be augmented to create a semi-structural version of the model in which the shock processes are generalised to depend on a small set of factors, driven by innovations to variables that capture the effects of missing shocks. The coefficients of the factor structure can be estimated by matching impulse responses of the augmented central organising model to those generated by a suite model or by Bayesian estimation using time series techniques. Monti (2010) shows how the Kalman filter can be used to combine forecasts from external sources (for our purposes, suite models) with those from a central organising model. This approach has the benefit of taking proper account of the fact that the suite model information represents a forecast for the variable(s) of interest. Therefore, it is particularly useful if the information from the suite model takes the form of a forecast for a particular set of variables rather than a marginal response to a particular shock (often known as a ‘multiplier’).

Some methods imply that all of the shocks in COMPASS should be used to correct for misspecification. For example, some suite models, such as extended versions of COMPASS, provide information about how all of the variables in COMPASS would respond to a particular misspecification. In these cases, it is possible to adjust the data for the COMPASS raw observables over the past and the forecast for the effects of a particular missing shock or transmission mechanism. This ‘cleaned’ data can then be interpreted through COMPASS using the Kalman filter and smoother as explained in Section 6.2.1. Examples of this approach are illustrated in Section 8.2.

7.2 Alternative approaches

Our strategy is to attempt to capture the effects of misspecification by applying constellations of correlated shocks to COMPASS. This approach is motivated by our desire to use a structural model as the central organising framework for building the forecast.

¹²⁴As discussed in Section 6.2.4, it is possible to specify whether the shocks are anticipated or unanticipated by agents in the model. The default treatment is to assume that shocks are anticipated, so that agents’ expectations for the variables directly affected by the misspecification are consistent with the quantification of the misspecification applied to the model.

Naturally, our approach places a considerable burden on the central organising model, which needs to be sufficiently credible in order to provide a baseline narrative for the data and the forecast. As the number of off model interventions to incorporate misspecification increases, it becomes more difficult to maintain coherence and consistency with that baseline narrative. There are at least two alternatives to our approach.

One approach would be to expand COMPASS to include more channels and shocks. This seems feasible given that our suite includes extended versions of COMPASS. Of course, there are also some difficulties associated with this approach: if the model is to be estimated, then computational considerations place a (practical) upper bound on the number of observable variables; and larger models are inherently harder to understand and explain to busy policymakers. But even if this strategy is a desirable long-term objective, in our experience it is likely that the issues relevant for policy discussions develop more quickly than the operational forecast models used to support those discussions can be feasibly expanded. In these situations, a small and simple central model is likely to offer more flexibility in dealing with such challenges.

Another approach is to consider a more formal model averaging exercise in which the forecasts from COMPASS and other suite models are weighted together using statistical criteria.¹²⁵ While such an approach may improve forecasting performance, it becomes more difficult to construct a consistent narrative for the resulting forecast: a tradeoff identified in the discussion of the design of the forecasting platform in Section 2. Indeed, the ‘core/non-core’ (CNC) approach used in the staff’s previous central organising model, BEQM can be regarded as a form of model averaging: a ‘core’ DSGE model is combined with a set of VECM-style ‘non-core’ equations. The non-core equations can be thought of as measurement equations, that also include a role for variables that are not modelled inside the core, as explained by Alvarez-Lois et al. (2008). As noted by Alvarez-Lois et al. (2008), one drawback of the CNC approach is that the projections of the non-core variables do not form part of the information set used by agents in the core model. This makes it difficult to produce fully model-consistent projections for scenarios that involve an alternative path for a non-core variable.¹²⁶ The tools available in the current forecasting platform allows us to incorporate such information into COMPASS much more coherently.

¹²⁵A variant on this theme is the ‘DSGE-VAR’ approach of Del Negro and Schorfheide (2004). Del Negro and Schorfheide (2009) compare the DSGE-VAR approach with various approaches to generalising the shock processes of the model to include additional correlation as approaches for dealing with model misspecification. Although COMPASS is relatively small and simple, applying the DSGE-VAR approach to a model of this scale would be challenging.

¹²⁶Indeed, this observation leads Alvarez-Lois et al. (2008) to suggest adjustments to the forcing processes driving the central organising model, similar in spirit to the approach of Caldara et al. (2012).

8 The forecasting platform in action

The aim of this section is to demonstrate how the models and tools described earlier in the paper can be used in practice. We present four case studies, showing how COMPASS, the suite of models and the associated toolkit can be applied to particular economic problems, designed to be representative of those typically encountered during production of the MPC's forecast.

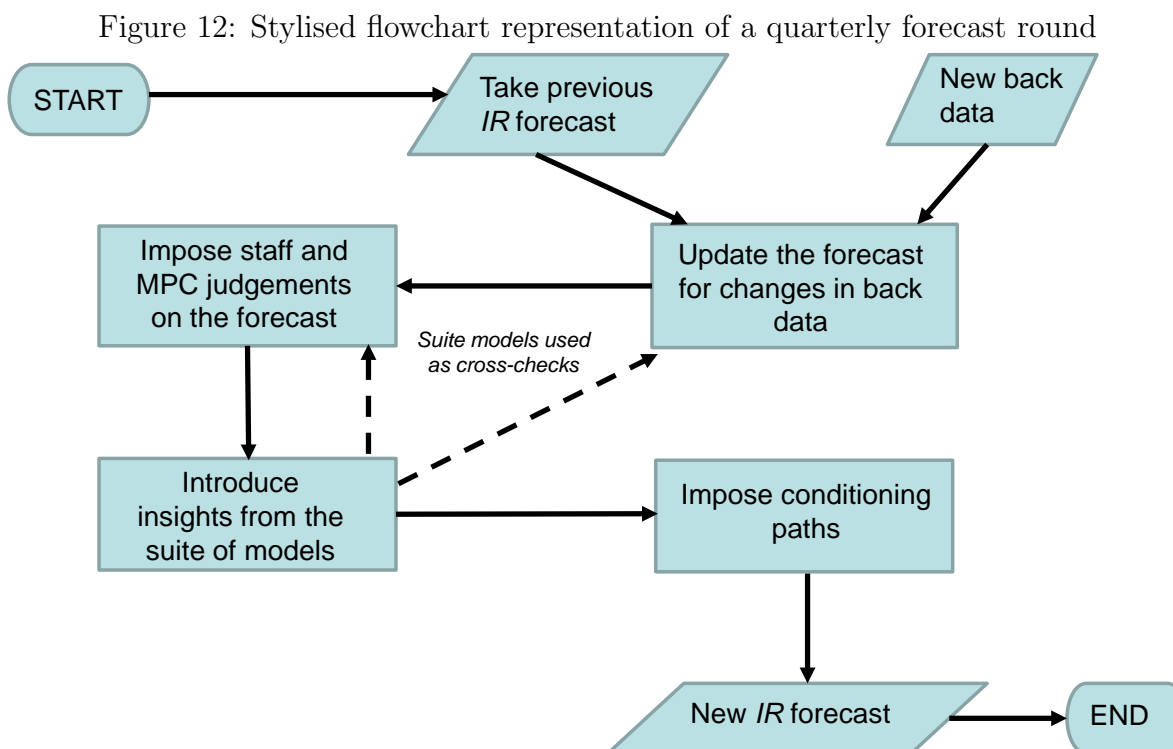


Figure 12 gives a stylised depiction of the key steps in the process of generating each *Inflation Report* forecast.¹²⁷ There are three points in particular worth emphasising about the forecast process:

- First, the use of judgement, by both Bank staff and the MPC, is central to the process. As explained in Section 2, models play a supporting, rather than a starring role. But using appropriate models can provide important insights, and models can also act as a valuable disciplining device. The production of one UK forecast would typically make use of at least 40 suite models alongside COMPASS. Depending on the context, suite models may be used to process news, motivate particular judgements, or to cross-check profiles for particular variables.
- In practice the forecast round involves several iterations of most of the steps shown in Figure 12. For example, official data are released and updated a number of times each quarter, so the back data are usually analysed more than once. There are also at least five meetings between the MPC and Bank staff at which the forecast is discussed with a view to making judgemental adjustments. Reflecting these considerations, a single forecast round normally involves several hundred runs of

¹²⁷Readers who are interested in a more detailed description of the Bank's forecast process, including the interaction between the MPC and Bank staff, are referred to Bean and Jenkinson (2001).

COMPASS, producing different projections at each step, as the forecast is gradually refined and alternative scenarios are explored.

- The diagram shows that the starting point of each forecast round is the previous *Inflation Report* forecast, so each new forecast is built incrementally. This means that one common task during the round is to “process news”: that is, to establish the extent to which historic data or the conditioning paths have changed since the previous forecast was published, and to update the forecast in light of that. It also means that previous staff and MPC judgements are incorporated into subsequent forecasts automatically, though of course they are routinely analysed and reviewed.

The rest of this section provides some concrete examples of those processes. Section 8.1 illustrates the first process in Figure 12, namely the updating of the forecast for changes in published data. Section 8.2 uses the example of VAT to show how a suite model with a channel omitted from COMPASS can be used to incorporate the economics of that channel into the forecast. Section 8.3, which addresses financial frictions, shows how a variety of different suite models can be used in parallel to explore the economics of an important policy issue and then inform MPC judgements. Finally, Section 8.4 illustrates how conditioning paths can be applied to the forecast with a particular focus on the expected path for policy.

All of the examples in this section are illustrative and have been simplified somewhat to aid their exposition. They are representative of the type of analysis and judgement carried out during forecast rounds, but they do not correspond to actual judgements which have ever been applied by Bank staff or the MPC. The precise implementation of such judgements would, in any case, depend on the context in which they were applied.

8.1 Introducing data news

As Figure 12 shows, the production of any new forecast involves a step where the forecast is updated for new information about the past. This section provides an illustration of how COMPASS and the range of tools described in this paper can be used to make those incremental changes. We use a specific example, updating the November 2011 forecast to take account of the economic data that were released between the November 2011 *Inflation Report* and the closure of the following forecast in February 2012.

Throughout the section, the focus will be on the fifteen variables that COMPASS treats as observable. In each case, the paths of those fifteen variables are draws from a probability distribution, both over the past (where published data are subject to revision, and hence uncertain) and over the forecast (where the MPC’s judgements are all about the balance of risks). However, for simplicity, we focus on the most likely (modal) paths throughout this discussion.

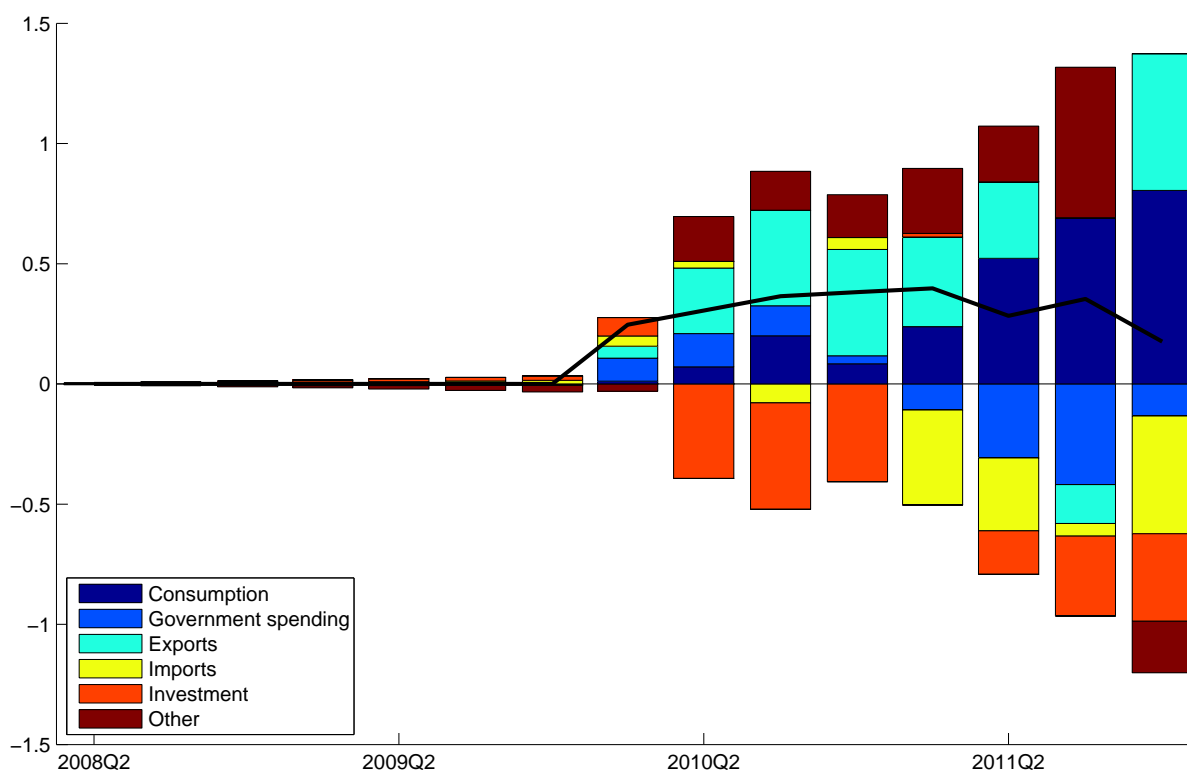
When the November 2011 forecast was published, there would have been published data available for all those variables up until at least 2011Q2. For 2011Q3, a subset of the data would have been available, though in some cases some “nowcasting” would have been needed to smooth the ‘ragged edge’ associated with data being released at different times.¹²⁸ However, the values for the variables in 2011Q4 would have been almost wholly forecast.

¹²⁸Some examples may be helpful here. For the interest rate and exchange rate, data are available more or less immediately. For monthly series like inflation and wages, typically there will be one or two monthly observations available for the latest quarter when the forecast is closed, and Bank analysts

However, by the time the February 2012 forecast was being finalised, there would have been data published for all variables up to 2011Q3, and partial information available for 2011Q4. Much of this would have represented “news” relative to the previous forecast. Further, there would have been revisions to the data for time periods prior to 2011Q3. In some data releases these can go a long way back into the past, though in practice more recent observations tend to be the ones which are most subject to revision.

In the example here, there was some news in all fifteen variables, with the exception of Bank Rate, where the November 2011 forecast had been conditioned on the assumption that Bank Rate would stay at 0.5% for some time, and that had proved to be correct. There were revisions back to 2005 for some variables, though the vast majority of the revisions affected data for 2010 and 2011. The most significant revisions were in the expenditure breakdown of GDP. As Figure 13 shows, the level of GDP was higher in February 2012 than had been anticipated in November 2011, though most of that related to upward revisions to 2010 data: the profile for growth in 2011 was actually a little weaker. The level of household consumption was revised up, business investment was revised down, and there were also noticeable revisions to government spending, exports and imports.

Figure 13: Revision to GDP data between November 2011 and February 2012 forecasts



Notes: Revisions to the level of GDP are measured in per cent. Contributions from the main expenditure components are shown. Dwellings investment and stockbuilding form part of the ‘other’ category.

Faced with this new information, Bank staff would need to take a view on whether to

would use other available evidence, such as more timely survey information, to impute the remaining months. In the case of GDP, there is a preliminary estimate available when the *Inflation Report* is published, but the expenditure breakdown is not available, as the lag on publication for that is around seven weeks from the end of the relevant quarter. For the November 2011 forecast, for example, the 2011Q3 values for the expenditure breakdown would have been based purely on staff nowcasts.

change the forecast beyond 2011Q4 and, if so, how. This is first and foremost a question of judgement, though the forecast platform offers a range of tools which can help inform that. As one example, we show here how we can use the MAPS toolkit to decompose the data news into the contributions from the different shocks within COMPASS.

The theory behind this is explained in Section 6.2.4, where we describe how the Kalman filter can be used to uncover estimates for the model variables given the data and the structure of COMPASS.¹²⁹ Because of the general equilibrium nature of the model, a given endogenous variable can be affected by many different shocks. It is therefore important to consider the news in all of the endogenous variables together, since it is their joint behaviour which will determine the estimates of the shocks that best explain them, which will in turn determine how the forecast changes.


Figure 14 shows one example of this type of exercise. It shows the *change* in the level of GDP relative to the November 2011 forecast decomposed into the shocks in COMPASS following the introduction of news in the back data. Since there are eighteen of those shocks, they have been collected into five groups for clarity of exposition. The chart also shows how those shocks would contribute to changes in the forecast for GDP in the absence of any additional judgement by Bank staff.

The positive blue bars show the contribution to the GDP news of the labour-augmenting productivity shock, η^{LAP} . This particular shock is assumed to have permanent effects in COMPASS, so in the absence of judgement, it contributes to higher GDP throughout the forecast period. The dark blue bars show the contribution of risk premium shocks, η^B , and the cyan bars the contribution of investment adjustment cost shocks, η^I . Both those shocks are temporary ones, though the contribution of investment cost shocks is estimated to be more persistent. The orange and yellow bars show the contributions from import preference and export preference shocks, respectively, and finally the maroon bars show the contribution of the other thirteen shocks. Most of those have a limited effect beyond the first few quarters of the forecast.

Given this information, there are usually two key judgements that Bank staff have to make. First, does the identification of shocks look sensible, given the revisions to the data that were observed? Second, is the resulting change to the forecast sensible? We tackle those questions in turn, continuing to use the same case study.

¹²⁹This can be thought of as the “best” (minimum variance) explanation of the news given the linearity of COMPASS and the assumption of normality of the shock and measurement error distributions.

Figure 14: A shock-based interpretation of the news in the level of GDP



Notes: A COMPASS shock decomposition of the news in the level of GDP (measured in per cent) using the Kalman filter. LAP denotes the ‘labour augmenting productivity’ shock.

8.1.1 Does the identification of shocks look sensible given the data?

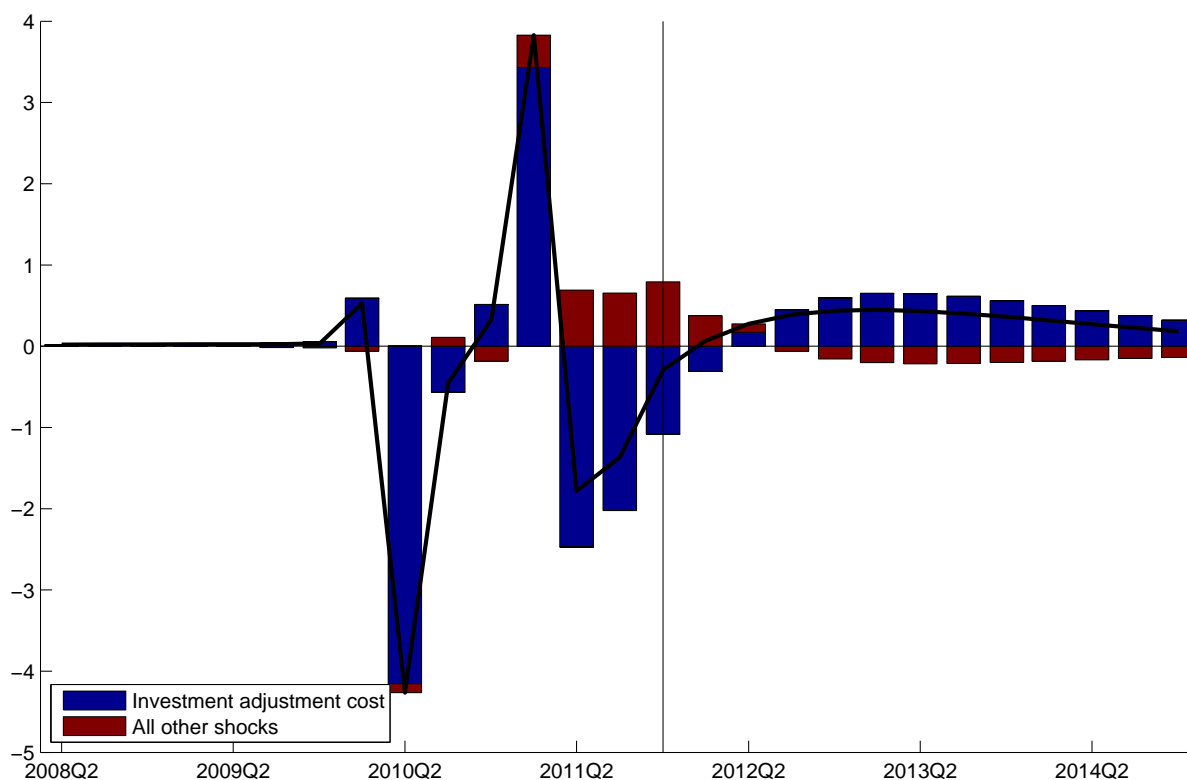
For simplicity we focus here solely on the explanation for the news in GDP given in Figure 14. Analogous decompositions can be produced for any variable in COMPASS, since they are all functions of the same set of eighteen shocks.

In this case it is fairly straightforward to understand why COMPASS has chosen this particular combination of shocks. To begin with, comparing Figure 14 with Figure 13 shows a close link between the news in particular expenditure components and the expenditure-specific shocks in COMPASS. For example, much of the upside news in consumption is explained as a risk premium shock. Similarly, much of the downside news in business investment is explained in terms of an investment adjustment cost shock. This is made clear in Figure 15, which decomposes the news in investment into the investment adjustment cost shock (blue bars), and all of the other shocks (red bars).

The other important contribution in Figure 14 comes from the shock to labour-augmenting productivity. This is also easy to rationalise given the data news: productivity was revised up over the past, due to upside news in GDP (Figure 13) and downside news in hours worked (employment), and this was consistent with small upside news in the level of wages.

More formally, for any given shock it is possible to use the MAPS toolkit to identify which endogenous variables were instrumental in driving the estimate of that shock when news was being introduced. Figure 16 demonstrates this for the investment adjustment cost shock, η^I , and shows that the downside news in business investment was the dominant factor in leading to this particular shock being selected. This panel of charts can be viewed

Figure 15: Shock-based interpretation of the news in quarterly business investment growth

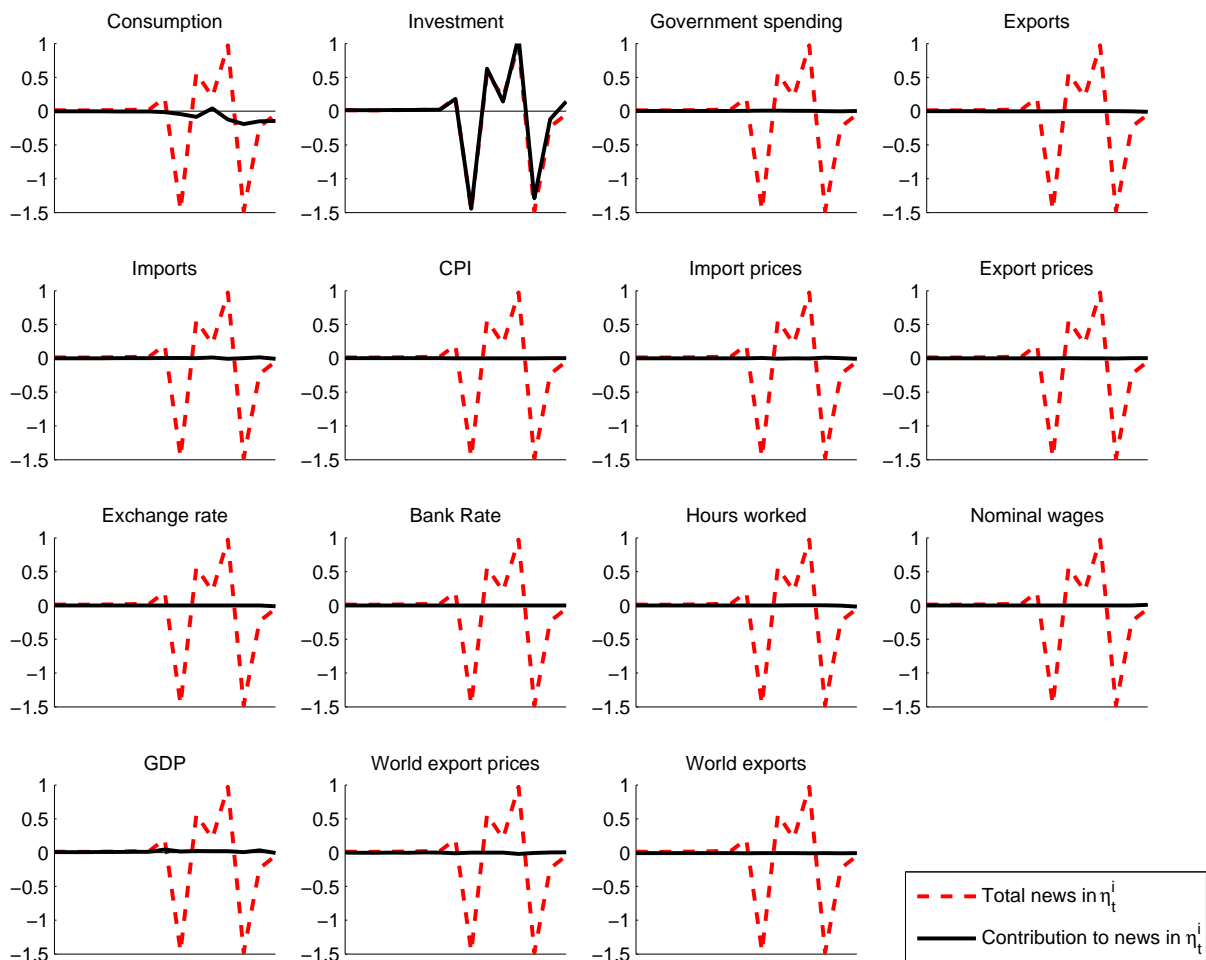


Notes: A COMPASS shock decomposition of the news in quarterly growth of business investment (measured in percentage points) using the Kalman filter.

in some sense as the ‘inverse’ of Figure 15, which shows the contributions of shocks to the news in an endogenous variable, rather than vice versa.

Clearly one would not want to place too literal an interpretation on COMPASS’s explanation of the news in the data: a user of COMPASS would not claim to be able to measure firms’ investment adjustment costs or the risk premium applying to optimising households at a quarterly frequency and be able to detect small changes in them! However, in most cases, Bank staff find the way in which the Kalman smoother identifies shocks over the past is a reasonable starting point. And, indeed, there are often occasions in which Bank staff may prefer to adopt a different explanation, for example if they had prior information about the factors driving historic data. A recent example where this was useful was during the Queen’s Diamond Jubilee, where it would not have made sense to attempt to explain all of the variation in output and employment in 2012Q2 using structural shocks. In that situation, one possible strategy is to pre-filter the data before introducing news.

Figure 16: Effects of data revisions on estimated investment adjustment cost shock



Notes: Each panel shows the effect of backdata revisions in a particular observable variable to the change in the Kalman smoothed estimate of the investment cost shock, η^I (black solid line). The red dashed line shows the change in the Kalman smoothed estimate in response to backdata revisions in *all* observable variables.

8.1.2 How should the forecast be changed in light of the data news?

Figure 14 shows how, in the absence of further judgement, COMPASS would unwind these newly estimated shocks and thereby update the GDP forecast. The changes to the forecast depend on: first, the size of the shocks identified over the past; and second, the persistence of those shocks, which is determined by the model estimation. In this example, most of the changes to GDP *growth* were made to the first year of the forecast, with the level being about 0.1-0.2% higher in the medium term on account of the permanent productivity shock identified by the model, which by definition does not unwind.

Bank staff need to form a view on whether this baseline treatment is plausible. In this particular example, the treatment of business investment offers a case study of why and how one might want to use a different treatment to that suggested by the model. Note that business investment growth over the forecast is actually stronger after the introduction of news, because COMPASS primarily attributes the new lower level of investment to a temporary shock, which then unwinds. This is most apparent in Figure 15, where quarterly growth rates average around 0.5pp more during the first half of the forecast.

In this example, one could make a strong case that the shocks which had lowered

investment over the past were likely to prove more persistent than the model treatment suggested. Equivalently, if one were sceptical about the merit of having a stronger investment growth forecast in February 2012 than in November 2011, a better alternative would be to carry through the lower level of investment over the past throughout the forecast period. There are two economic arguments one could appeal to to justify that alternative treatment:

- First, the evidence from the suite of investment models at the time was not supportive of a stronger profile for investment growth. Figure 8 in Section 5 shows the individual model forecasts at the end of the November 2011 round, alongside the MPC's forecast at the time. At that time, four of the seven models were suggesting weaker investment growth than the MPC's forecast. In particular, the money, lending and investment system in the investment suite (described in Section 5.4), which among those models accords most weight to financial channels and spare capacity, was below the MPC's November forecast. This provides an example of how different suite models, emphasising different economic channels, can be employed as a cross-check to inform the main forecast.
- Second, the wider economic environment at the time did not suggest that a sharp recovery in investment was imminent. The euro area crisis had been intensifying in the second half of 2011 and a number of external commentators had been revising their growth forecasts for the UK downwards.¹³⁰ As the February 2012 *Inflation Report* said, "The level of business investment is projected to recover only gradually in the near term, given the degree of slack within businesses, continuing uncertainty about future demand, and restricted access to credit". In those circumstances staff would have been more likely to adopt a cautious approach to the processing of data news, a view that, on balance, would probably have been justified by data outturns since 2011.


Having made a decision not to take COMPASS's treatment of the news as given, Bank staff would then have to decide how to implement the judgement to weaken investment. One approach would be to manipulate the investment forcing process, $\hat{\varepsilon}^I$. This represents the cumulative total of previous shocks to investment, taking account of their decay. Figure 17 shows how this adjustment could be done: rather than assume no impact on the level of investment by the end of the forecast, the extra shocks to investment identified over the past are assumed to persist for longer. The effect of this is to lower growth rates of investment throughout the forecast, but particularly early on.

As explained in Section 6.2.4, the MAPS inversion algorithm makes it straightforward for Bank staff to impose this particular profile for $\hat{\varepsilon}^I$. The desired path for $\hat{\varepsilon}^I$ can be 'fixed' using a suitable shock or combination of shocks to best capture the economics of the judgement in hand. The MAPS inversion technology can then be used to determine the values of the shocks which are needed over the forecast in order to generate the desired path.

In this instance, the economics of the judgement are best captured using the investment adjustment cost shock, η^I . The main reason for that is that the judgement is simply offsetting a previous change to the forecast (made when the data news was first introduced), and that was imposed using the investment adjustment cost shock. The standard

¹³⁰See Page 49 of the November 2011 *Inflation Report*.

Figure 17: Applying judgement to the investment adjustment cost forcing process



Notes: The solid black line shows the change in the forcing process, ε^I , associated with the investment adjustment cost shock generated by the revisions in past observable data. The red dashed line shows a judgemental profile for the change in the forcing process applied to lower the growth rates of business investment in the forecast.

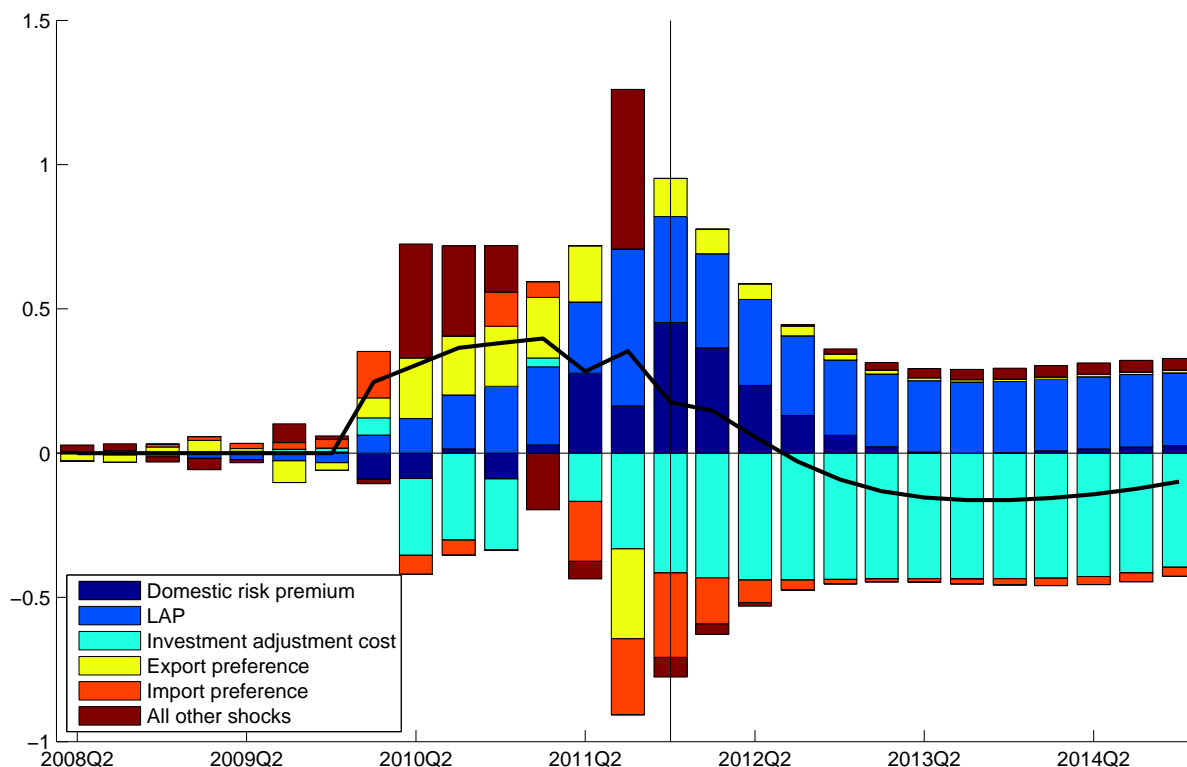
practice would be to assume that all of these forthcoming shocks are anticipated by agents in the model.

When this judgement is applied, all other variables in the model respond to the new profile of investment adjustment cost shocks which have been incorporated into the forecast. In this context, because the balance of shocks is negative and the shock selected is a demand shock, the effect will be to reduce the level of GDP and to lower inflation over the early part of the forecast.

The GDP impact is evident from Figure 18, which shows the shock decomposition of the news in GDP following the extra adjustment. It is identical to Figure 14, except for the fact that investment shocks continue to drag on the level of GDP over the forecast. The net effect of “introducing news” and imposing the additional judgement is that rates of GDP growth were slightly weaker at the end of this process than they had been in November, whereas under a “vanilla” COMPASS treatment with no additional intervention, they would have been slightly stronger. Note that in practice, Bank staff would repeat this treatment for each one of the shocks in COMPASS, verifying that they were content with the changes to the forecast caused by the process of introducing news.

This section has demonstrated how COMPASS and its associated tools can be used to interrogate a set of new back data and to identify the shocks that drive it. It also showed how the forecast can be updated in the light of those shocks. Most importantly, however, it emphasised the importance of judgement, including the use of the suite of models as a cross-check, in updating the forecast.

Figure 18: Shock decomposition of the news in the level of GDP, following staff judgement



Notes: A COMPASS shock decomposition of the news in the level of GDP (measured in per cent) using the Kalman smoother, following the judgemental adjustment of the forcing process for in the investment adjustment cost shock shown in Figure 17. LAP denotes the 'labour augmenting productivity' shock.

8.2 Incorporating effects of VAT changes

An important part of the forecast process is applying judgement to incorporate the implications of shocks or events that are not captured in COMPASS. As outlined in Sections 2 and 7, the forecast platform has been explicitly designed with this issue in mind. This would be a challenge regardless of the choice of central forecast model, but it is closer to centre stage given that COMPASS is designed to be relatively small and so excludes more economic mechanisms than it otherwise might.

This sub-section considers changes in the standard rate of Value Added Tax (VAT) as a case study for how the suite of models and IT infrastructure can be used to incorporate judgement about economic mechanisms that do not form part of COMPASS into the central projection. It starts by illustrating how the forecast platform could have been used to adjust the COMPASS judgemental forecast following the announcement by the government that it would temporarily cut the rate of VAT in December 2008 and throughout 2009.¹³¹ It then illustrates how the same set of tools can be used to incorporate the story for changes in VAT into a COMPASS-based narrative for the endogenous variables over the past. Finally, it demonstrates that the same techniques can be used to incorporate the effects of shocks to energy prices into the narrative.

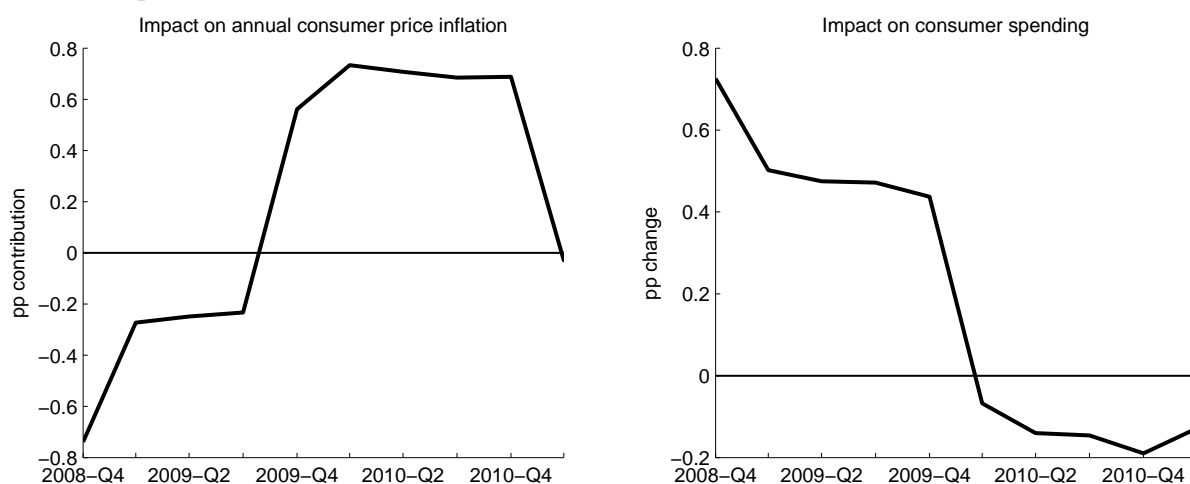
On 24th November 2008, the Chancellor announced that the rate of VAT would be cut from 17.5% to 15% on 1st December 2008 before reverting back to 17.5% on 1st

¹³¹Of course, this is a hypothetical example, given that the forecasting platform described in this paper did not exist at this time.

January 2010. If COMPASS included a consumption tax, then the forecast could have been updated for this news in a similar way to that discussed in Section 8.1. But, given that COMPASS does not include a tax on consumption, the challenge would have been to incorporate the implications of this news into the forecast that would have been published in the February 2009 *Inflation Report*.

The approach taken in this sub-section is a good example of the generic approach taken to tackling similar problems since the COMPASS forecast platform was introduced. Other examples of where this approach has been used in practice include judgements about the role of changes in energy prices (see page 84 of this sub-section) and credit frictions (see Section 8.3). In the case of VAT, the approach combines off-model analysis and information about the effect of VAT changes on inflation with a structural model and the MAPS toolkit. Specifically, the Bank has constructed in-house estimates of the effect on inflation of the temporary reduction in the rate of VAT outlined above. These estimates are formed by combining an empirical estimate of the effect of VAT changes on inflation¹³² overlaid with judgement using intelligence gathered by the Agents. The left panel of Figure 19 shows the latest in-house estimate of the contribution to annual CPI inflation of the temporary VAT reduction during December 2008 and 2009.¹³³

Figure 19: Estimated impact of the temporary reduction in VAT during 2009 on inflation & consumption



Notes: The left panel shows the latest in-house estimate of the effect of the temporary reduction in 2009 on annual CPI inflation. The right panel shows the implications of that for consumer spending from a simulation using the model described in the text.

The wider macroeconomic implications of these VAT-induced changes to inflation can be traced out by combining the off-model estimates shown in the left panel of Figure 19 with a structural model. We use an extended version of COMPASS modified to incorporate VAT, combined with the MAPS judgement toolkit described in Section 6.2.4 to assess the consequences of the changes in VAT under the assumption that the cut in VAT to 15% was unanticipated, but that the reversal of that cut was anticipated.¹³⁴ As

¹³²This is an empirical estimation using time dummies to capture the impact of VAT changes while controlling for other determinants.

¹³³The estimates used in this section are the latest internal estimates and so benefit from the effect of hindsight. The Bank's real-time estimates of the effect the announced changes in VAT used the same approach of combining empirical estimates with judgement.

¹³⁴More specifically still, the extended version of COMPASS includes a variable that measures the

described below, using the MAPS toolkit to account correctly for the anticipated reversal of the temporary VAT rate cut is an important part of the story for the wider economic effects.

In this extended version of COMPASS, a cut in VAT reduces the wedge between final output prices and consumer prices. This wedge has implications for the measure of labour income that is relevant for consumer spending with direct implications for the spending of rule-of-thumb households whose spending changes following an increase in VAT in proportion to the change in their real incomes.¹³⁵ The change in VAT also has an effect on the spending of optimising households despite it having no impact on their lifetime wealth.¹³⁶ That reflects changes in the ex-ante real interest rate induced by the temporary change in VAT. In line with the way the MPC responded to the VAT change,¹³⁷ the policy maker in the model is assumed not to respond to the direct effect of changes in VAT on inflation. This means that the changes in inflation expectations that occur in anticipation of VAT going back up to 17.5% lead to (close to) proportional changes in the real interest rate, inducing unconstrained (optimising) households to substitute their spending over time. Simply put, unconstrained households take advantage of what they know to be temporarily lower prices to spend more while VAT is lower with a corresponding reduction in their spending when VAT goes up.¹³⁸ The overall effect on consumption - the weighted sum of the effect of the temporary cut on rule of thumb and optimising households - is shown in the right panel of Figure 19. The model suggests the cut in VAT could have temporarily boosted consumer spending by around 0.5-0.7pp. Of course, as discussed on page 69 at the beginning of this section, the precise impact that any future changes in VAT might have on the published *IR* projections is a matter of MPC judgement, which would take into account the circumstances at that time.¹³⁹

This extended version of COMPASS also traces out the effect of the change in VAT on all the other observable variables in COMPASS and so could have been used to update the February 2009 *Inflation Report* forecast for changes in VAT.¹⁴⁰

A related use of this approach is in incorporating the effect of changes in VAT into the direct effect of VAT changes on inflation. In the experiment, this variable is fixed to the contributions shown in the left panel of Figure 19 using the consumption tax rate shock under the assumption that that shock was unanticipated in 2008Q4 and anticipated thereafter (i.e. that agents in the model did not anticipate the initial shock, but anticipated the inflationary consequences of it thereafter).

¹³⁵See Section 4.2.3 for a description of the modelling of the household sector in COMPASS.

¹³⁶That is because in this extended version of COMPASS a reduction in VAT reduces government revenue, leading to a proportional increase in the lump-sum tax levied on optimising households in order to maintain the level of government spending.

¹³⁷From the minutes of the December 2008 MPC meeting: “Although the temporary reduction in VAT would lead to some volatility in inflation over the next two years, the new fiscal plans were unlikely to have a significant effect on inflation beyond that period. The Committee noted that under the terms of its remit, it was required to look through short-run movements in inflation in order to avoid undesirable volatility in output”.

¹³⁸This intertemporal substitution effect was discussed in a box on page 31 of the February 2009 *Inflation Report*: “Lower VAT is likely to increase demand in 2009 because it will encourage households to bring forward spending while the lower rate is in force”.

¹³⁹For example, a judgement on the overall impact on GDP would likely require an analysis of the effects of VAT changes on relative prices and substitution effects across different expenditure components. A range of models would typically be required to support such analysis.

¹⁴⁰That is not to say that the BEQM-based February 2009 *IR* forecast did not contain the effect of the VAT announcement. Indeed, VAT was included in BEQM itself. Rather, the purpose of the discussion is to illustrate that a smaller, simpler central model imposes no impediment to updating the forecast for economic channels that are not directly modelled within COMPASS.

COMPASS narrative for the data. As described in Section 8.1, one way of decomposing news in the data is into the contributions of the structural shocks that drive COMPASS. It is also possible to use the same approach (as described in Section 6.2.5) to compute the contributions of the structural shocks to the data itself, rather than the news in the data. This can be useful as a starting point for thinking about what has been driving the economy. Figure 20 shows the contributions of the shocks in COMPASS (grouped using the technology described in Section 6.2.5) to deviations of annual inflation from the target between 2005 and 2011.¹⁴¹ The most striking feature of the chart is that it suggests that a large part of the rise, fall and subsequent rise in inflation in 2008, 2009 and 2010 can be attributed to domestic price markup and monetary policy shocks.¹⁴² The analysis of the temporary VAT cut in 2008-9 described above suggests that this interpretation is questionable.

But the narrative embodied in Figure 20 can be adjusted to account for the role of VAT in explaining the evolution of inflation (and other endogenous variables in COMPASS) over the past. That can be done in three steps, which follow the logic of the misspecification algorithm discussed in Section 7.1. First, the simulation described above can be augmented (with suitable modifications) to include the effect of the increase in the rate of VAT from 17.5% to 20% announced in June 2010 and applicable from 4th January 2011. The result is a set of time series for the observable variables in COMPASS that record the simulated effect of all the VAT changes observed since 2008. Second, these simulated paths can be stripped out of the data to produce a synthetic dataset that excludes the estimated effect of VAT changes. That synthetic dataset can then be decomposed into the structural shocks in the same way as above to produce a decomposition of the data that excludes the estimated VAT effect. Third, the results of the simulation can be added on to incorporate the simulated effect of VAT changes. The resulting hybrid decomposition is shown in Figure 21. In line with the discussion above, it shows that changes in VAT have been important in explaining movements in inflation over this period and that once changes in VAT have been accounted for the role of domestic price mark-up and monetary policy shocks is reduced.

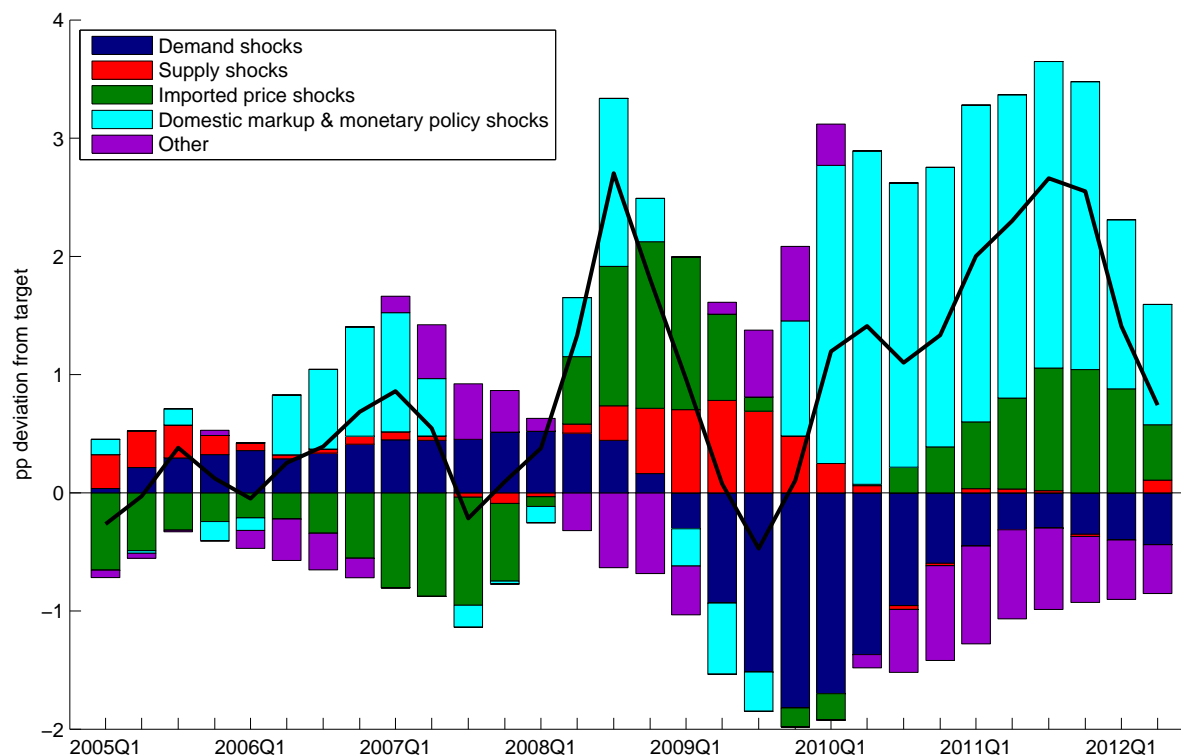
The same algorithm can be applied to augment the narrative for an account of the role of energy price shocks. More specifically, we can use the energy suite model described in Section 5.2.1 to simulate the macroeconomic effects of changes in the price of energy since 2004.¹⁴³ As the discussion in Section 5.2.1 highlighted, the assumed behaviour

¹⁴¹Note that as in all the applications described in this section of the paper, the data used here is the latest vintage of data available as at February 2013, which is the vintage of data used to estimate the version of COMPASS in this paper (see Section 4.3).

¹⁴²Ignoring the influence of the output gap (which in any case has a smaller weight), the Taylor rule in COMPASS prescribes that interest rates should go up when inflation is above target and down when inflation is below target. Any deviation of interest rates from that prescription is explained in COMPASS as being due to monetary policy shocks. As discussed in Section 7, misspecification implies that this may not always be the correct, 'structural' interpretation. Indeed, the rest of this sub-section demonstrates that a failure to account for the way policy responds to VAT distorts COMPASS's interpretation of the data (and the behaviour of interest rates in particular) over this period.

¹⁴³More specifically, we use an approach that combines off-model empirical information with the suite model in a very similar way to that underpinning the VAT simulations described above. We use the energy price shock in the model to 'fix' a variable measuring the direct contribution of energy price changes to annual inflation to an in-house estimate of the direct impact of changes in energy prices on consumer price inflation based on observed changes in utility and petrol prices (and their time-varying weights in the consumer basket). The term shock is used quite loosely here. Movements in energy prices reflect underlying shocks to the supply of and demand for energy, which can emanate from a variety of


Figure 20: Naive grouped shock-based decomposition of annual inflation using COMPASS



Notes: The decomposition was constructed using the tools described in Section 6.2.5. The groups are defined in the following way: ‘Demand shocks’ includes the contributions of all domestic demand shocks and the world output shock; ‘Supply shocks’ includes the two productivity shocks and the labour supply shock; ‘Imported cost shocks’ includes the exchange rate risk premium shock, the world export price shock and the import price mark-up shock; ‘Domestic price mark-up and monetary policy shocks’ includes what it says on the tin (including the wage mark-up); ‘Other’ includes the import preference shock, the UK export price mark-up shock and the contribution of the initial conditions which is negligible.

of monetary policy is crucial in determining the macroeconomic effects of energy price shocks. For the purposes of this simulation we use an assumption analogous to that used in the VAT suite model: we assume that the monetary policy maker ‘looks through’ the direct effect of energy price changes on inflation. That is, the policymaker does not respond to changes in inflation driven by changes in utility and petrol prices, arising from their inclusion in the basket of consumer goods used to compute the CPI. The policymaker may respond to the indirect effects, but they turn out to be small, reflecting two partly offsetting effects. On the one hand, energy is a factor of production, so increases in energy prices raise inflationary pressure by increasing firms’ marginal costs. However, on the other hand, energy price increases reduce consumer spending (primarily through rule-of-thumb households), reducing the demand for labour and hence wages, which pushes down on marginal costs. In general equilibrium, these two effects more or less offset, which means that the overall impact of the simulated changes in energy prices on inflation is similar to the direct ‘consumer basket’ effect. Figure 22 adjusts the COMPASS shock-based inflation narrative to incorporate the results of the energy price simulation using the algorithm discussed above. Energy prices rose sharply during 2008, fell back during 2009 and then rose sharply again during 2010. This pattern is reflected in the contribution of energy price shocks in the chart. Figure 22 also demonstrates that the underlying sources with a range of different consequences.

Figure 21: Grouped VAT hybrid decomposition of annual inflation using COMPASS and a model from the suite (before accounting for energy)




Notes: This hybrid decomposition was constructed in the same way as described in Section 6.2.5. The groups are as defined in Figure 20.

role of domestic price mark-up and monetary policy shocks in explaining inflation over recent years is even further reduced once energy price changes are taken into account.¹⁴⁴

This sub-section has illustrated how the judgemental forecast and narrative in COMPASS can be adjusted in a coherent and systematic way using the suite and the IT infrastructure for shocks or events that are not captured by COMPASS. This is a key part of the overall design of the forecast platform and one that is explored further in the next sub-section, which describes how the suite can be used to adjust the forecast for the implications of financial frictions.

¹⁴⁴The remaining shocks may be a reflection of the sharp depreciation of sterling that took place in 2008.

Figure 22: Grouped VAT & energy hybrid decomposition of annual inflation using COMPASS and models from the suite



Notes: See notes for Figure 21.

8.3 Incorporating effects of financial frictions

Because there are no financial frictions within COMPASS, we use suite models to help us think through the implications of shocks caused by or transmitted through such frictions. In this section, we illustrate how we might have used the new forecasting platform to examine the implications of the large increase in credit spreads associated with the financial turmoil in late 2008. The suite of models, by its nature, offers a multiplicity of approaches, and the discussion here does not imply that this method is the only one which can be applied to the problem. Further, credit spreads are just one of many channels through which the effects of financial frictions may be observed, but we focus on them here because they have been viewed by some as important summary indicators of financial conditions.¹⁴⁵

We proceed in four steps. First, we outline a process for measuring and forecasting credit spreads that can provide a broad-brush assessment of the current state of (and prospects for) financial conditions. We examine the change in the projection for credit spreads between the August and November 2008 forecast rounds – the ‘news’ in credit spreads – a period of intense financial stress surrounding the collapse of Lehman Brothers. Next, in Section 8.3.2, we discuss two suite models that can help us to quantify the effects of financial shocks (associated with movements in credit spreads) on macroeconomic variables. One model includes explicit behavioural assumptions about the banking sector; the other is an empirical model augmented with assumptions that help to identify shocks to the demand for and supply of credit.

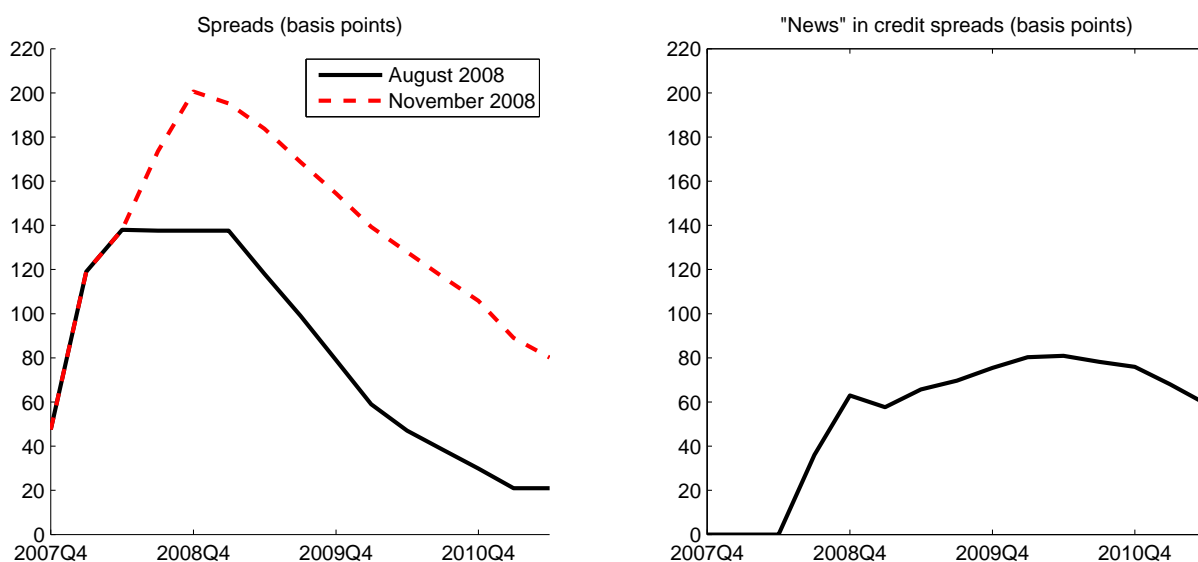
¹⁴⁵See, for example, Gilchrist and Zakrajsek (2012).

In Section 8.3.3 we use each suite model to provide a quantification of the effects of financial shocks that generate a change in credit spreads equal to the ‘news’ recorded between the August and November 2008 forecast rounds. In each case, the credit spreads news is generated as the endogenous response to different types of financial shock, using the MAPS toolkit described in Section 6.2.4.

Finally, in Section 8.3.4, we show how the macroeconomic effects of financial shocks implied by the suite models can be incorporated into a forecast built using COMPASS. Again the MAPS toolkit is used to mimic the suite model responses of the macroeconomic variables in COMPASS, as endogenous responses to an appropriate choice of shocks. The choice of which shocks to use for this step is an essential part of the process. We also briefly discuss how we can use other suite models as cross-checks on the alternative quantifications of the effects of financial shocks produced using COMPASS.¹⁴⁶

8.3.1 Financial frictions and credit spreads

Figure 23: Illustrative credit spreads forecasts and credit spreads ‘news’



Notes: The left panel shows forecasts for credit spreads based on information available in the August and November 2008 forecast rounds. The right panel shows the “news” in the credit spreads forecast, computed as the difference between the August and November forecasts. See the main text for a discussion of the construction of these forecasts.

Figure 23 depicts data for and forecasts of the economy-wide spread of borrowing (and deposit) rates over Bank Rate (the ‘credit spread’) based on information available in August 2008 and November 2008. Forecasts are plotted over a three year horizon, corresponding to the typical focus of *Inflation Report* forecasts. Staff construct this credit spread by measuring the marginal cost of funds to different groups of households and corporates and aggregating them using the appropriate population shares.¹⁴⁷ For example, among households four groups are identified: depositors; low loan to value (LTV) secured borrowers; high LTV secured borrowers; and unsecured borrowers. Data

¹⁴⁶As explained in Section 5, another role of the suite is to provide such cross-checks, which may ultimately guide judgements made by the staff and MPC when producing forecasts.

¹⁴⁷The credit spread is measured relative to the level of the data in 2007Q3, before the dramatic tightening of credit conditions in light of the financial crisis. This is simply a normalisation.

on the marginal rates at which each group can borrow are weighted by the share of each group in the total population of households.¹⁴⁸ Some judgement can be applied to this approach. For example, credit rationing effects can be proxied by reallocating the weights between different groups, capturing the notion that borrowers may be forced to access credit on less favourable terms (for example, some borrowers may no longer be able to access secured credit and hence become forced to borrow at unsecured rates). The basic approach for corporates is analogous.¹⁴⁹ To produce the forecast, staff make a judgement on the likely steady state of the banking sector and the level of bank funding costs implied by that. The forecast is built up by considering a range of individual loan and deposit products and using error correction models to estimate the rate at which changes in funding costs might pass through to interest rates quoted by the banks. This process also relies heavily on staff judgement: for example, information from the Bank's Credit Conditions Survey or conversations with staff from the major banks should give an insight into how those banks are planning to price their products in the near term, and hence whether the rate of adjustment implied by those models is plausible. Likewise, the Bank's regional Agents provide information on how credit conditions faced by households and businesses are changing.

The aggregate credit spread forecast is a weighted average of the forecasts for household and corporate spreads, using weights based on the stocks of outstanding household and corporate bank debt. In principle, information from household and corporate credit conditions can be applied to the forecast separately. Indeed, the effect of a given change in the credit conditions faced by households may well be expected to be rather different to the same change in credit conditions faced by firms. However, for simplicity of exposition, we will use the aggregate credit spread in the analysis that follows.

In the subsequent sections, we assess the possible effects of the 'news' in the credit spread forecasts using two suite models. We focus on the 'news' in credit spreads because, as explained in Section 8, the forecast is constructed in an incremental fashion, starting from the previous *Inflation Report* forecast. So, we assume that the macroeconomic effect of the expected path of credit spreads in August 2008 (the solid blue line in the left panel of Figure 23) was accounted for in the August 2008 *Inflation Report* forecasts. The task is to update the forecast in light of the new information contained in the updated forecast for credit spreads (the red dashed line in the left panel of Figure 23). Consequently, here we will focus on the marginal responses of suite models and COMPASS to financial shocks that give rise to the news in credit spreads, plotted in the right panel of Figure 23.

Of course credit spreads respond endogenously to changes in general economic conditions, so not all of the movement in credit spread forecasts between forecast rounds will reflect financial factors. However, here we simplify the exercise by treating the credit spreads news as if it was driven by a single exogenous financial shock. This seems like a reasonable assumption, given the events of the period we focus on.¹⁵⁰

¹⁴⁸The information on borrowing rates is derived from quoted rates data. The population shares are calibrated using data from the British Household Panel Survey and FSA's Product Sales Database.

¹⁴⁹The marginal borrowing rates are based on corporate bond yields.

¹⁵⁰We limit the scope of the exercise to include updating the forecast only. From the perspective of the November 2008 forecast round, some of the news in credit spreads was contained in new data on credit spreads for 2008Q3, which would necessitate an exercise to capture the effects of this new information on our interpretation of the past data – as implemented for the case of VAT in Section 8.2. However, for simplicity we analyse the news in the credit spreads as if it applied entirely to the forecast.

8.3.2 The suite models

In this section we briefly describe two suite models in which credit spreads are modelled explicitly.¹⁵¹

The first model was developed by Gertler and Karadi (2011) (henceforth, the ‘GK’ model). Like COMPASS, the GK model features optimising, forward-looking households and firms with rational expectations.¹⁵² But the GK model also includes a banking sector, which intermediates household savings to firms. Firms borrow from banks to finance investment in physical capital. It is assumed that households are unable to monitor the behaviour of banks perfectly, so the banks are able to divert a fraction of the funds for their own use. As a result of this so-called agency problem, the optimal loan contract implies that the value of assets that banks may borrow from households is limited by banks’ net worth, with bank leverage determined endogenously. This contract ensures that, in equilibrium, banks do not have an incentive to divert funds and gives rise to a ‘financial accelerator’ in which banks’ balance sheet conditions play an important role in the propagation of shocks. We set the parameter values of the model equal to the posterior mean parameter values estimated in a version of the GK model by Villa and Yang (2011) on UK data for the period 1979–2010. We replace the monetary policy reaction function in the GK model with the reaction function from COMPASS (see equation (22)). This modification is useful in ensuring that differences in the behaviour of the GK model (relative to COMPASS) stem from different assumptions about the structure of the economy rather than the monetary policy reaction function.

The second suite model is the Structural Vector Autoregression (SVAR) model developed by Barnett and Thomas (forthcoming) (henceforth, the ‘BT’ model). The SVAR is a predominantly empirical model, with the equation describing the behaviour of each variable being a linear function of lags of that variable and lags of the other variables in the model.¹⁵³ The BT model is estimated for the UK economy for the period 1966 to 2012 using data for GDP growth, inflation, corporate bond spreads, equity prices, bank lending growth and a long-term gilt yield. The gilt yield is used to identify monetary policy shocks, given that the estimation sample includes a period during which asset purchases have been used as the primary monetary policy instrument.¹⁵⁴ Using a mixture of timing and sign restrictions, the authors identify six shocks, one of which is a shock to the supply of credit by banks. Barnett and Thomas (forthcoming) demonstrate that this shock contributed significantly to the rise in credit spreads and in particular to the decline in the stock of lending relative to trend since the onset of the financial crisis in

¹⁵¹Models that include endogenous determination of credit spreads include Bernanke et al. (1999), Carlstrom and Fuerst (1997) Cúrdia and Woodford (2011), Christiano et al. (2010) and Gilchrist and Zakrajsek (2012). Unsurprisingly, further development of models of credit spreads and their effects on macroeconomic variables has been an active area of research over recent years. However, we focus on two models for brevity.

¹⁵²In particular, the GK model includes similar assumptions about the ‘real rigidities’ facing agents, such as habit formation and investment adjustment costs, and prices are assumed to be sticky. However, the GK model abstracts from a number of features included in COMPASS, including: nominal wage rigidities; international trade in imports and exports; and the presence of ‘rule of thumb’ households.

¹⁵³VAR models have formed part of the Bank’s set of economic models for many years: see Bank of England (1999, Chapter 5).

¹⁵⁴This is because one of the key mechanisms through which the MPC’s asset purchases have affected the macroeconomy is through their effect on longer term gilt yields (see Joyce et al. (2011a), Joyce et al. (2011b) and the discussion in Section 8.4 below). A drawback for our purposes is that the BT model is silent on the implications of shocks for Bank Rate, which is the primary monetary policy instrument in both the GK model and COMPASS.

2008.¹⁵⁵ As the BT model can be represented as a linear state space (LSS) model, we implement it in MAPS, setting the parameter values to the posterior means estimated by Barnett and Thomas (forthcoming).

8.3.3 Quantifying the effects of financial shocks

In this section, we use the suite models described in Section 8.3.2 to quantify the impact of the credit spreads news described in Section 8.3.1 on key macroeconomic variables. In the suite models, credit spreads are determined endogenously, by a range of different shocks. To assess the effects of a particular change in the forecast for credit spreads, we use the MAPS ‘impose judgement’ tools, which allows us to uncover the most likely sequences of shocks that give rise to the desired behaviour of particular endogenous variables over the forecast.¹⁵⁶ The choice of which shocks to use to implement the judgement is crucial.

We use shocks to banks’ net worth and ‘capital quality’ as the candidate instruments to apply the news in credit spreads.¹⁵⁷ Gertler and Karadi (2011) use capital quality shocks to analyse the effects of the US financial crisis, arguing that “in this rough way, we capture the broad dynamics of the sub-prime crises” (p27).¹⁵⁸ Villa and Yang (2011) show that, when the GK model is estimated on UK data, capital quality and banks’ net worth shocks are important determinants of the dynamics of credit spreads over the financial crisis.¹⁵⁹

Figure 24 shows the implications for lending, GDP and CPI inflation using the GK model under two assumptions about the shocks that generated the credit spreads news depicted in Figure 23.¹⁶⁰ In the left column, the news in credit spreads is implemented using shocks to banks’ net worth. In the right column shocks to both banks’ net worth and capital quality are used. In both cases, the shocks are assumed to be fully anticipated. This means that, in both cases, agents revise their forecast of credit spreads precisely in line with the credit spreads news shown in the right panel of Figure 23. Of course, the effects on the variables shown in the two columns of Figure 24 differ because different shocks have been used to implement the experiments.

Figure 24 shows that the broad pattern of responses is similar. The financial shocks push up credit spreads and hence the cost of capital. Banks reduce lending and firms reduce investment, giving rise to a fall in GDP and inflation. Monetary policy responds to

¹⁵⁵See Charts 4 and 10 of Barnett and Thomas (forthcoming), respectively.

¹⁵⁶For more details, see Section 6.2.4.

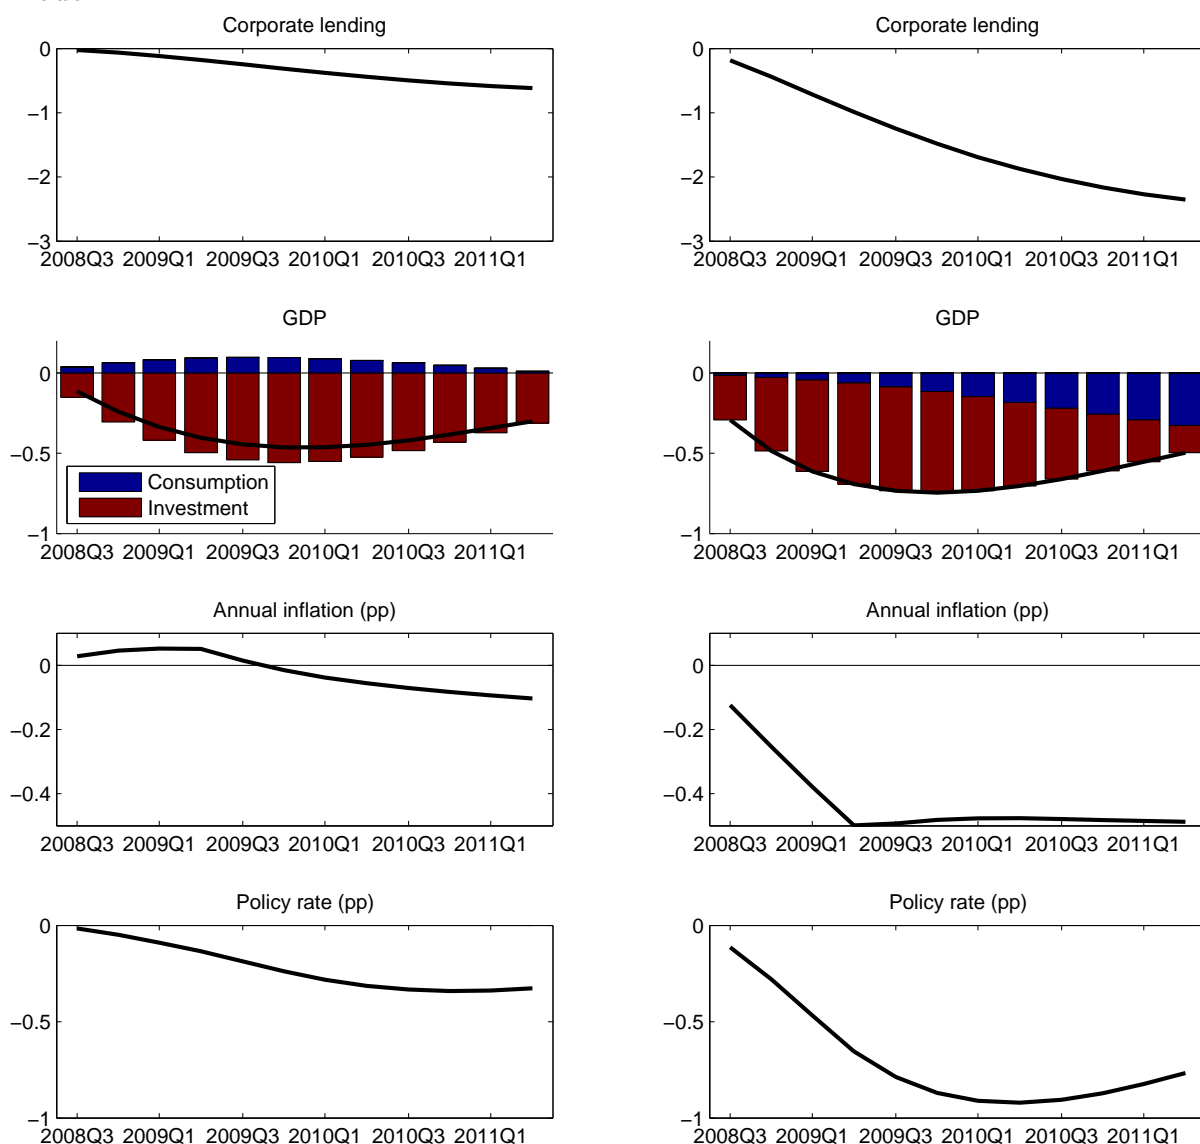
¹⁵⁷Shocks to capital quality directly affect the value of the physical capital that firms invest in and hence the value of loans made by banks. Gertler and Karadi (2011) show that the effects of an anticipated reduction in future capital quality that does not actually transpire (a “news shock”) generates similar effects to a shock to capital quality. Negative shocks to banks’ net worth represent an exogenous transfer of wealth from banks to households (see Gertler and Karadi (2011, p27)) and proxy shocks that inhibit banks’ ability to lend, through a reduction in the availability of retained earnings for this purpose.

¹⁵⁸Gertler and Karadi (2011) argue that the capital quality shock represents economic obsolescence rather than reductions in the quantity of physical capital. This shock is designed to capture important elements of the financial crisis, stemming from a change in the underlying quality of intermediary assets (or indeed beliefs about the future value of those assets).

¹⁵⁹See Figure 4 of Villa and Yang (2011).

¹⁶⁰At first sight, the size of the effects may seem rather small, given the extent of the financial crisis. However, recall that we are examining the estimated effects of the change in credit spreads shown in the right panel of Figure 23, which represents only a small fraction of the overall increase in credit spreads during the financial crisis. Villa and Yang (2011) demonstrate that the model’s estimate of the *total* effect of shocks that increased credit spreads over the financial crisis is sizeable, reducing the level of GDP by around 5% relative to trend.

Figure 24: Quantifying the effects of financial shocks in the Gertler and Karadi (2011) model



Notes: Each column shows the effects of using anticipated shock sequences to generate the profile of credit spreads ‘news’ in the right panel of Figure 23. In the left column shocks to banks’ net worth are used. In the right column, shocks to banks’ net worth and shocks to the quality of capital are used. All responses are plotted in percentage deviations from the baseline forecast or percentage point deviations (pp) where stated.

weaker output and inflation by reducing the short-term nominal interest rate. However, effects on lending, GDP and inflation are larger when the shocks are applied using capital quality shocks as well as bankers’ net worth shocks (right column). Although, by construction, the path for credit spreads in the two simulations is identical, the cost of capital increases more sharply in the simulation using both shocks (the right column). This is evident based on the larger decline in the policy rate (relative to the fall in inflation) in the charts in the right column, consistent with an initial increase in the short-term real interest rate. In contrast, the short term real interest rate *falls* persistently in the left column (using only shocks to banks’ net worth). In that case, the persistent reduction in the real interest rate is sufficiently persistent to induce an increase in household con-

sumption, as evidenced by the decomposition of GDP in the left column.¹⁶¹ So when shocks to both capital quality and banks' net worth are used, there is a sharper increase in the real cost of capital and hence a larger decline in investment, capital formation and corporate lending. The weakness in activity in this case reduces wage costs by enough to outweigh the effects of the rise in the cost of capital on production costs and inflation falls more substantially.

Part of the reason for the increase in consumption when only shocks to banks' net worth are used is that the shock represents a redistribution of resources from banks to households. Using capital quality shocks in addition to net worth shocks allows us to proxy an event in which the underlying value of assets on banks' balance sheets declines, which seems more in line with the narrative of the financial turbulence in the US provided by Gertler and Karadi (2011). Together with Villa and Yang's finding that capital quality shocks were important determinants of UK corporate bond spreads over the period we are studying, we will focus on the results using shocks to both capital quality and banks' net worth to apply the credit spreads news in the analysis that follows.

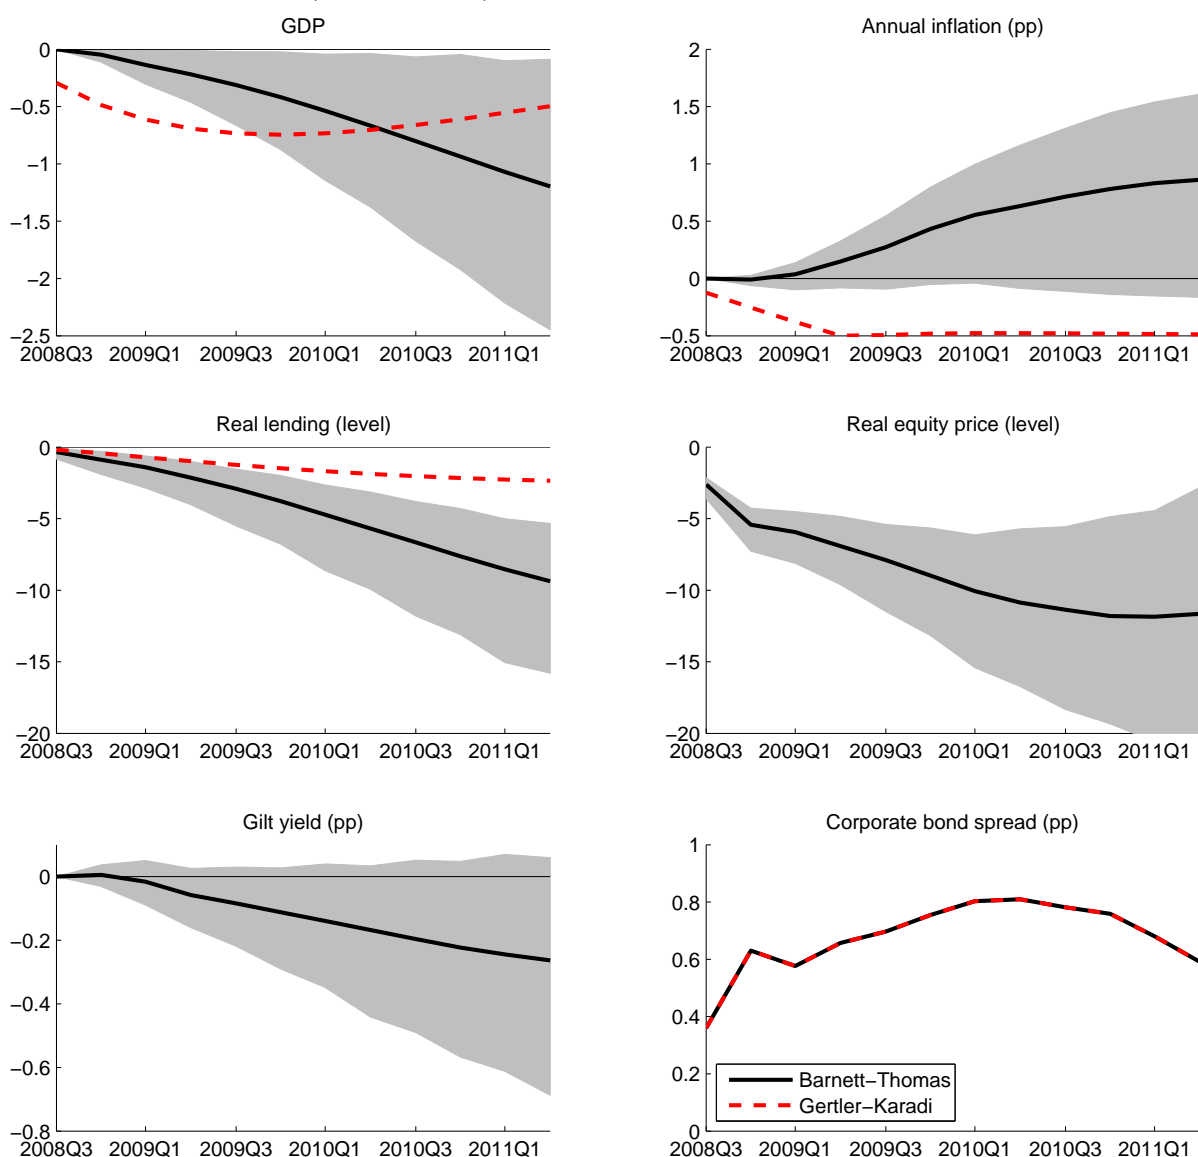
Turning to the BT model, we choose to implement the news in credit spreads using the credit supply shock, as Barnett and Thomas (forthcoming) find that this shock plays an important role in explaining movements in UK credit spreads during the financial crisis. The BT model is backward looking, in the sense that the effects of expected future shocks cannot be explicitly identified. Therefore, the credit supply shocks that generate the change in credit spreads are, by definition, unanticipated. Figure 25 plots the behaviour of the variables in the BT model (solid black lines) in response to a sequence of credit supply shocks that generate the required increase in credit spreads. The grey swathes shows the range of responses that lie between the 16th and 84th percentiles of the distribution generated using draws from the posterior distribution of the SVAR parameters. We also plot the responses of the GK model (using dashed red lines), where broadly comparable variables are available.

Figure 25 shows that the responses of GDP to financial shocks in our two suite models have quite different magnitudes and dynamics. The GDP response in the BT model builds more gradually over time, whereas GDP responds more quickly in the GK model, with the peak impact occurring after about one year. The peak impact is also somewhat larger than in the GK model. In part these differences can be attributed to the fact that the shocks used to implement the experiment are fully anticipated by households and firms in the GK model. Forward-looking households and firms realise that financial shocks will increase credit spreads sharply and keep them elevated for some time. This realisation means that households and firms react immediately to a large change in the outlook for credit spreads over the next three years (the response is relatively fast). It also means that households and firms are able to mitigate the overall effects of the shocks to some extent: the optimal response is to front-load the reduction in investment to help postpone the required fall in consumption.¹⁶² As noted above, the shocks used to implement the

¹⁶¹This effect is very small in the parameterisation of the model used by Gertler and Karadi (2011), but is more noticeable for the parameterisation we use here, based on the estimation results of Villa and Yang (2011).

¹⁶²As explained in Section 6.2.4, the MAPS toolkit allows judgements to endogenous variables to be imposed using either anticipated shocks, unanticipated shocks or a mixture of both. Using unanticipated shocks to banks' net worth and capital quality to deliver the credit spreads news does indeed result in much more gradual responses, as households and firms gradually recognise that the increase in credit spreads will be very persistent. But the magnitude of these responses is also smaller, because the full extent of the persistent increase in credit spreads only becomes apparent over time.

Figure 25: Comparing effects of financial shocks in the Gertler and Karadi (2011) and Barnett and Thomas (forthcoming) models



Notes: Responses of the Barnett and Thomas (forthcoming) and Gertler and Karadi (2011) models to shocks that replicates the profile of credit spreads “news” depicted in the right panel of Figure 23. All responses are plotted in percentage deviations from the baseline forecast or percentage point deviations (pp) where stated. The grey swatches show the 16th and 84th percentiles of the distribution of responses generated using draws from the posterior distribution of the parameters of the Barnett and Thomas (forthcoming) model.

experiment in the BT model are unanticipated. In each period of the simulation, an additional credit supply shock is applied to the model, so that the model responses represent the cumulated responses to a sequence of negative credit supply shocks.

Figure 25 shows that the responses of inflation in the two suite models are somewhat different: inflation rises in the BT model.¹⁶³ Based on their posterior parameter estimates,

¹⁶³As noted earlier, inflation falls in the GK model. The reduction in inflation is persistent because of the protracted output dynamics which depress marginal production costs and hence inflationary pressure for a relatively long time. The simulation results in Gertler and Karadi (2011, Figure 2) demonstrate that the financial accelerator mechanism increases the persistence of the model’s responses to shocks,

Barnett and Thomas (forthcoming) find that, on average, inflation rises in response to a (negative) credit supply shock that increases credit spreads. However, the posterior parameter uncertainty in this case is sufficiently high that there is a non-negligible probability that the inflation response is negative. We see a similar result in the top right panel of Figure 25 since the posterior uncertainty interval around the inflation response includes a region in which the inflation response is mildly negative. The response of lending is larger in the BT model, though caution is required in interpreting this result as the concepts of lending differ between the two models.¹⁶⁴

8.3.4 Mimicking the effects of financial shocks using COMPASS

In this section we show how the forecasting platform can be used to apply quantitative insights from our suite models to a forecast that is built up using COMPASS. Once again, we will use the MAPS inversion toolkit to impose profiles for variables in COMPASS using shocks designed to mimic the economic effects of financial shocks as quantified by our suite models. And again shock selection will be crucial when assessing the implications for other variables. This means that we need to choose the quantitative effects from the suite models that we wish to incorporate and also a selection of COMPASS shocks to impose them.

In Section 8.3.3, we saw that both the GK and BT models predicted that GDP falls in response to financial shocks that generate a rise in credit spreads, although the dynamics are somewhat different. However, while financial shocks leading to a rise in credit spreads reduce inflation in the GK model, inflation rises in the BT model. Given the ambiguity of the suite model predictions for inflation, we therefore focus on the effects of credit spreads on GDP and its components.¹⁶⁵ While the BT model does not include expenditure components, the GK model implies that investment responds more quickly than consumption. We choose to impose this expenditure split when applying the quantitative effects of financial shocks to COMPASS.¹⁶⁶

In choosing shocks to use to implement the consumption and investment profiles, we aim to select shocks that mimic the economic mechanisms included in the suite models, drawing on insights from previous studies. This analysis leads us to choose domestic risk premium, investment adjustment cost and total factor productivity shocks to impose the suite model quantifications on COMPASS.¹⁶⁷ We assume that these shocks are anticipated by agents in the model, mirroring the assumptions used in the GK model analysis in Section 8.3.3. The use of the domestic risk premium shock is motivated by the fact that inspection of the equation describing aggregate consumption behaviour reveals that these shocks affect consumption in a similar way to the risk free real interest rate.¹⁶⁸ Moreover,

including that of inflation.

¹⁶⁴In the GK model, only firms borrow, so the lending variable in this model is best interpreted as corporate borrowing. In contrast, the BT model is estimated using data on total bank lending.

¹⁶⁵A key factor in determining the likely inflationary effects of a financial shock that increases credit spreads is the effect of that shock on costs and potential supply. In practice, the MPC have used a wide range of models and analysis to consider that effect.

¹⁶⁶Specifically, we use the ratios of the consumption and investment contributions to the GK response in the left panel of Figure 24 to compute paths of consumption and investment consistent with the shares of these expenditure components in COMPASS. This delivers profiles for consumption and investment compatible with the total responses for GDP from the GK and BT models that we then impose on COMPASS.

¹⁶⁷That is, we use $\hat{\varepsilon}^B$, $\hat{\varepsilon}^I$ and $\hat{\varepsilon}^{TFP}$.

¹⁶⁸See (15) in Section 4.2.3.

Cúrdia and Woodford (2010) find that in a model with explicit financial frictions, a term capturing credit spreads enters the model in the same way as a risk premium shock. This type of reasoning has led a number of authors to use this shock to mimic the effects of rises in the effective real interest rates facing households arising from tightening credit conditions.¹⁶⁹ The investment adjustment cost shock is chosen to ensure that investment decisions are directly affected by the financial shock we are mimicking. Justiniano et al. (2011) find that this shock explains a significant fraction of US business cycle fluctuations. Moreover, they find that the time series of that shock implied by their estimated model is highly correlated with a measure of corporate bond spreads.¹⁷⁰ We choose the TFP shock because the capital quality shock used to implement the experiments using the GK model in Section 8.3.3 has a direct impact on the production function, analogous to a shock to total factor productivity.

Figure 26 shows the results of applying the suite model quantifications for consumption and investment consistent with their GDP responses in COMPASS, using the selection of shocks discussed above. The left column of charts shows the results using the quantification from the GK model and the right column shows the results from the BT model quantification. We discuss each column in turn.

By construction, the contributions of consumption and investment to the GDP response in the left column are identical to the contributions in the right column of Figure 24 using the GK model. However, the GDP response itself is smaller because of offsetting effects from net trade: effects that are absent by construction in the GK model since it lacks an endogenous determination of net trade. These offsetting effects are driven by a small, but persistent depreciation in the real exchange rate (a fall in the real exchange rate represents a *depreciation* of the domestic currency). The exchange rate depreciation is prompted by a reduction in the policy rate, brought about by weaker inflation and activity. Inflation falls as weaker activity reduces domestic cost pressures, though initially there is a some partially offsetting effect from higher import price inflation as a result of the exchange rate depreciation.¹⁷¹

The story in the right hand column, based on the BT suite model quantification, is qualitatively similar, although GDP falls more slowly, inducing slower falls in the policy rate and inflation. In this case, the sum of the consumption and investment contributions is equal to the GDP response for the BT model plotted in Figure 25.¹⁷² As with the results based on the quantification from the GK model, inflation falls, though there is a partial offset from higher import price inflation. Barnett and Thomas (forthcoming) present regression-based evidence suggesting that the the exchange rate depreciates in response to credit supply shocks in the BT model and that the resulting increase in import price

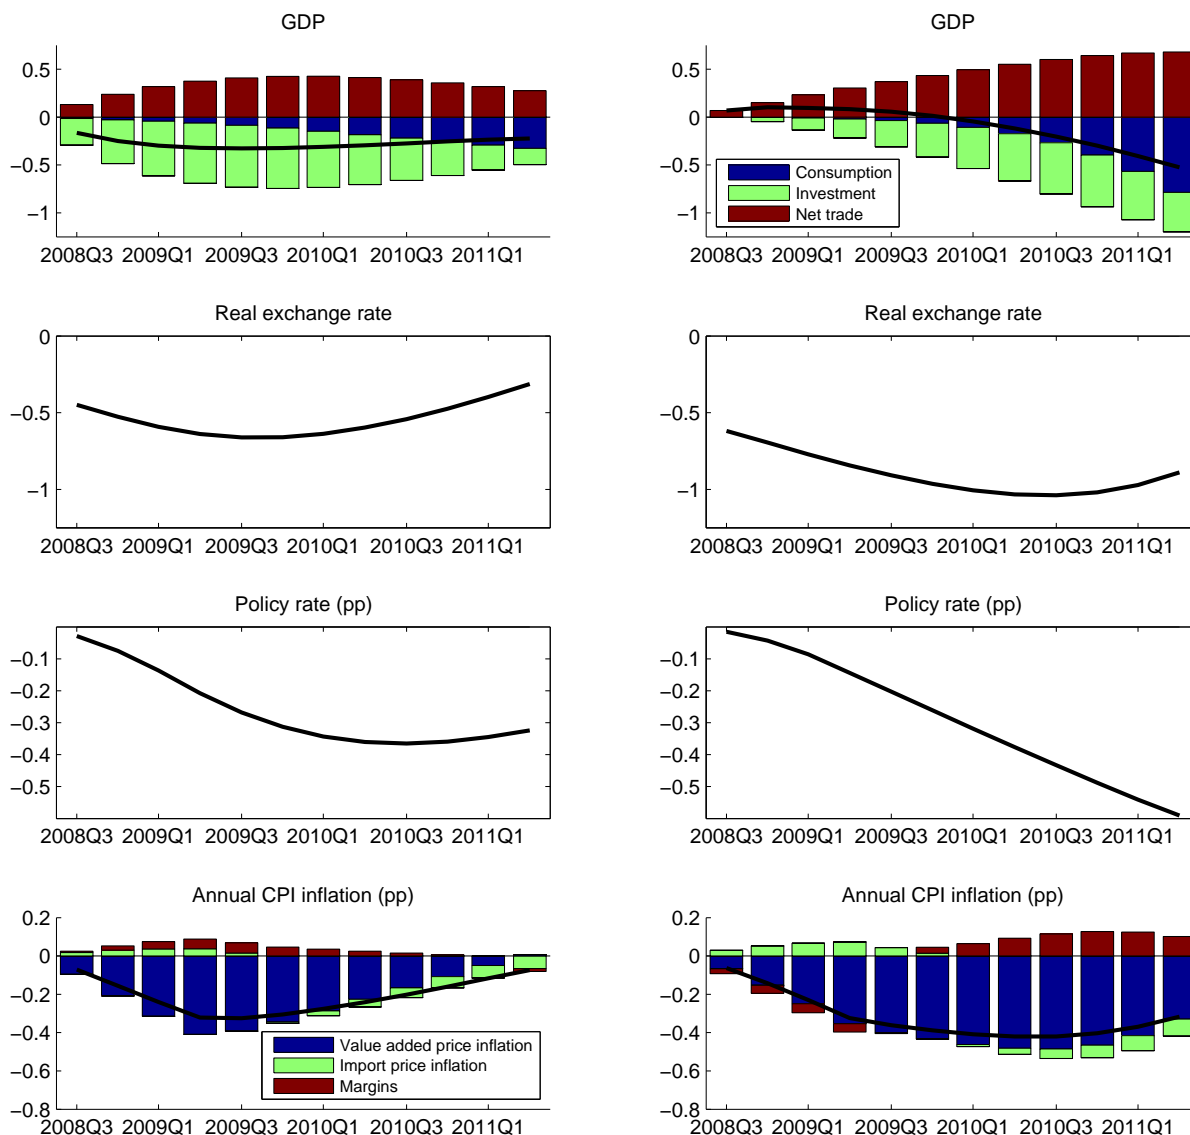
¹⁶⁹See, for example, Eggertsson and Krugman (2012).

¹⁷⁰Justiniano et al. (2011, p115) also discuss how inspection of the structure of their model helps to interpret the result: “In our model, there is no explicit role for financial intermediation. [...] However, the transformation of foregone consumption (real saving) into future productive capital depends on its relative price, which in equilibrium is affected by μ [the investment adjustment cost shock]. [...] Thus, one possible interpretation of the random term μ is as a proxy for the effectiveness in which the financial sector channels the flow of household savings into new productive capital”. Given the similarities between the two models, these arguments can also be applied to COMPASS.

¹⁷¹The decomposition of inflation was produced using a ‘flexible’ decomposition using an additional equation that defines CPI inflation as a markup over value added inflation and import price inflation. See Section 6.2.5 for a brief description of the MAPS toolkit that produces this type of decomposition.

¹⁷²As explained above, the relative importance of consumption and investment is determined by the relative importance of these expenditure components in results from the GK model in the right column of Figure 24.

Figure 26: Mimicking the effects of financial shocks in COMPASS



Notes: The charts show responses of selected variables in COMPASS when the effects of financial shocks on GDP implied by the Barnett and Thomas (forthcoming) and Gertler and Karadi (2011) models are applied. All responses are plotted in percentage deviations from the baseline forecast or percentage point deviations (pp) where stated.

inflation may drive the initial increase in inflation in the BT model shown in Figure 25.

The key differences between the experiments using quantitative information from the BT and GK models are the nature of the investment response and the overall size of the effects on activity and inflation. The GK quantification implies that investment should fall more substantially over the first year or so. To examine the relative plausibility of the two quantifications for investment, we could use the investment suite (as applied in Section 8.1). In terms of the overall size of the effects, both candidate quantifications depicted in Figure 26 have the feature that the weakness in domestic demand induced by the financial shock is partially offset by an improvement in net trade. This observation illustrates the general issues with using quantitative responses from models with a simplified treatment of the expenditure composition of GDP (such as the GK model) or an absence of any information about expenditure components (such as the BT model). In terms of the

specific implications for our experiments, the financial shocks underlying the changes in credit spreads over the period we are analysing are perhaps most naturally regarded as global shocks which also impacted the UK's major trading partners. So a fuller analysis of the implications of this shock for the UK would require an assessment of the impacts of the shock on world demand. Moreover, we note that the significant contribution of financial services to UK exports may suggest a weaker outlook for exports and the exchange rate, which could be incorporated in the COMPASS simulations using export preference shocks.¹⁷³

These considerations demonstrate that the use of suite models can never be mechanical. Significant judgement is typically required to incorporate all of the factors relevant to a particular shock or event.

Our two suite models have given us two alternative sets of adjustments to a COMPASS-based forecast that may help account for the news in credit spreads. One way to assess the alternative adjustments would be to examine the implications for a broader range of economic variables, using models in the suite designed to produce forecasts of additional variables. One approach would be to assess the results of the simulations in Figure 26 using the balance sheet model described in Section 5.3.2).¹⁷⁴ In this case, such an exercise implies that the alternative quantifications of the effect of credit spreads from the GK and BT models have very similar implications for the key balance sheet variables.

8.4 Incorporating policy changes

Inflation Report forecasts have been traditionally based on the assumption that a small set of variables follow trajectories determined by particular conventions, often using information from financial markets or other external sources.¹⁷⁵ As shown on page 69, staff apply the 'conditioning paths' for the relevant variables towards the end of the forecast process. The conditioning path for Bank Rate is derived from market expectations of the policy rate.¹⁷⁶

Of course, market expectations of Bank Rate are unlikely to coincide exactly with the path for Bank Rate implied by the COMPASS-based forecast.¹⁷⁷ To impose the conditioning path for Bank Rate, the MAPS toolkit is used to apply a sequence of unanticipated shocks to the monetary policy reaction function, so that the path for Bank Rate coincides with the path derived from market expectations. Producing a forecast conditioned in this way can be justified under the assumption that the shocks are sufficiently small relative to the statistical distribution of monetary policy shocks that they are unlikely to alter agents' beliefs about the monetary policy reaction function: they are

¹⁷³The importance of financial services in UK exports is discussed in Kamath and Paul (2011).

¹⁷⁴To implement this, we would feed the results from the COMPASS simulations in Figure 26 into the post-transformation model (see Section 5.3.2) and the results from the post-transformation model into the balance sheet model.

¹⁷⁵In particular, it is assumed that asset prices (Bank Rate and the sterling effective exchange rate) and fiscal policy (spending and taxation rates) follow paths derived from external sources. For more details of the assumptions used, see, for example, Monetary Policy Committee (2013).

¹⁷⁶Projections based on the alternative assumption that Bank Rate remains constant over the forecast horizon are also routinely published. These conditioning paths are implemented in the same way as the market curve.

¹⁷⁷The COMPASS-based forecast of Bank Rate will be determined by the monetary policy reaction in the model, based on the profiles for inflation and activity, which in turn will be influenced by judgements applied during the production of the forecast.

“modest interventions” in the terminology of Leeper and Zha (2003).¹⁷⁸

An alternative to unanticipated monetary policy shocks is to assume that the shocks to the monetary policy reaction function are fully anticipated by agents in the model. This assumption is also a strong one since, taken literally, it corresponds to the assumption that the policy deviates from the monetary policy reaction function that stabilises inflation in the model. Perhaps unsurprisingly, for prolonged deviations, forward looking models can generate quite striking results in response to these simulations.¹⁷⁹

As shown in Section 6.2.4, the MAPS toolkit allows us to use both anticipated and unanticipated shocks to impose the profiles for endogenous variables. So, in principle, it is possible to impose the market curve for the policy rate using, say, anticipated shocks for the first few quarters of the forecast and unanticipated shocks thereafter. However, unanticipated shocks to the monetary policy reaction function are chosen for simplicity and convention, since this has been the method for imposing interest rate conditioning paths with the central organising models that were previously used to support the production of *Inflation Report* forecasts.

Since March 2009, Bank Rate has been held at the historically low level of 0.5% and the MPC has been using asset purchases as the instrument of monetary policy: a policy known as ‘quantitative easing’ (QE). Benford et al. (2009) and Joyce et al. (2011b) elucidate a number of channels through QE may affect the economy. First, purchases of assets (bonds) held by the private sector could increase the prices of those assets: there may be ‘portfolio balance’ effects. As bond prices increase, yields fall and private sector borrowing costs are reduced, stimulating aggregate demand. Second, because asset purchases are financed by the creation of central bank money, they lead to an increase in reserve balances held by banks at the central banks.¹⁸⁰ The increase in reserve balances may facilitate an expansion in bank lending. Third, asset purchases may improve market functioning by increasing liquidity through actively encouraging trading. Such an effect would be expected to reduce illiquidity premia and increase asset prices. Fourth, asset purchases may provide a useful signal about the future course of monetary policy by demonstrating policymakers’ resolve to prevent inflation significantly undershooting the target in the medium term. Fifth, to the extent that asset purchases lead to higher asset prices, they may help to support consumer confidence and hence households’ willingness to spend and firms’ willingness to invest.

Given the highly stylised treatment of asset markets in COMPASS, there is no way to incorporate quantitative easing directly into the forecast: again we rely on the suite of models.¹⁸¹ To think through the economics of the transmission mechanism of QE we can

¹⁷⁸Adolfson et al. (2005) investigate whether forecasts conditioned on the assumption of a constant policy rate satisfy the modesty criterion of Leeper and Zha (2003) using an estimated DSGE model for the euro area.

¹⁷⁹Del Negro et al. (2012) document very large effects of such policy experiments and observe that they appear to be generated by implausibly large equilibrium movements of long-term interest rates. Laséen and Svensson (2011) show that prolonged anticipated positive deviations of the policy rate from the reaction function in RAMSES (the DSGE model developed for forecasting and policy analysis at the Riksbank) can generate very large falls in inflation and in some cases, a *rise* in inflation. Carlstrom et al. (2012) show that the latter result also appears in simple New Keynesian models that exhibit inertia in the Phillips curve.

¹⁸⁰As explained by Benford et al. (2009), when the central bank purchases an asset from a non-bank asset holder, the central bank credits the seller’s bank’s reserve account at the central bank and the seller’s bank credits the asset seller with a deposit.

¹⁸¹Note that it is not the absence of government debt issuance in COMPASS that is crucial here, but rather the absence of suitable frictions (for example, in asset markets or expectations formation) that

use models that feature explicit behavioural assumptions giving rise to a role for QE. For the most part, these models focus on the portfolio rebalancing transmission mechanism by incorporating some form of imperfect substitutability among assets.¹⁸² Many of these models are based on the approach introduced by Andrés et al. (2004): see, for example, Chen et al. (2012), Dorich et al. (2011) and Harrison (2012).

To provide a quantitative estimate of the effects of QE policies, we can use empirical models estimated to take account of the potential links between asset purchase policies, asset prices and macroeconomic variables. Joyce et al. (2011b) review recent research on this issue. Most approaches adopt a two step approach. The first step is to estimate the effects of QE on asset prices or another intermediate variable such as the money supply. The second step is to estimate the effects of movements in the intermediate variable on the macroeconomy. A benefit of this two step approach is that it is possible to use longer samples of data to estimate the macroeconomic effects (since QE has only been in operation since 2009), which should produce more precise estimates. But a drawback is that the second step relies on relationships between asset prices, money and macroeconomic variables over a period during which QE was not in operation. The approaches documented by Joyce et al. (2011b) are:

- A ‘bottom up’ approach mapping from estimates of the effects of QE on asset prices to estimates of the effects of changes in asset prices on demand. Estimates of the effects of QE on asset prices (see for example Joyce et al. (2011a)) suggest that the MPC’s asset purchases up to and including February 2010 had a cumulative effect that reduced long-term gilt yields by around 100 basis points. Estimates of the wealth elasticity of consumption and investment to asset price changes can then be used to quantify the effects on aggregate demand. A Phillips curve type relationship can then be used to map the effects of the change in aggregate demand to inflation, given an assumption about the impact of QE on aggregate supply.
- A structural vector autoregression (SVAR) approach based on quarterly data for real GDP growth, CPI inflation, long-term government bond yields and Bank Rate. This model can be used to simulate the effects of a reduction in long-term bond yields on real GDP and CPI inflation, under the assumption that Bank Rate does not respond (since it is constrained by its lower bound).
- The multiple time series model approach of Kapetanios et al. (2012) which uses a range of empirical models to perform counterfactual policy simulations in which gilt yields are reduced by 100bps and the policy rate is unchanged.
- The monetary approach of Bridges and Thomas (2012) which first estimates the effects of QE on the money supply and then uses two alternative approaches (an SVAR and a set of money demand equations) to estimate the resulting effects on activity and inflation.

The quantitative results from this set of models are summarised by Joyce et al. (2011b, Table C, p210), which indicates that there is some uncertainty over the effects on both GDP and inflation. This type of analysis has informed judgements by the MPC on the effects of quantitative easing on the economy. As with all judgements of this nature,

could give QE traction.

¹⁸²Joyce et al. (2011b) argue that much of the evidence on the estimated impact of QE on gilt prices is consistent with the portfolio balance channel.

it is likely to be refined and adjusted as further evidence is gathered. The effects of quantitative easing are applied to COMPASS across a wide range of variables (using the MAPS toolkit to select the most likely sequence of shocks to deliver them). This allows the MPC to consider the effects of changes in asset purchases on the forecast. Staff typically produce such simulations towards the end of the forecast round, in the forecast meetings close to the MPC's policy meeting.

Looking ahead, the recommendations of the Stockton review¹⁸³ call for more routine and wide-ranging policy analysis to be included in the forecast process. The Bank's response to the review indicates that there are plans to move in this direction.¹⁸⁴ COMPASS should be well-suited to support these developments, given the explicit behavioural underpinnings of the model.

¹⁸³ Available at <http://www.bankofengland.co.uk/about/Pages/courtreviews/default.aspx>.

¹⁸⁴ Available at <http://www.bankofengland.co.uk/publications/Pages/news/2013/051.aspx>.

9 Conclusions

This paper has documented the components of the Bank of England's new forecasting platform, and given concrete examples of how models and tools that form part of it can be applied to judgement-based forecasting. At the time of writing, the new platform has been in operation for around a year and a half, and has played a crucial role in supporting the Bank's forecast process over that period. The central model, COMPASS, has provided the organising framework for the MPC's analysis and has also been used in sensitivity analysis. The suite of models has also continued to be used intensively over that period, offering both cross-checks on the forecast and a means of motivating additional judgements which can be applied to COMPASS. MAPS and EASE, the new IT tools, have been integral to the success of the new platform since it was introduced in autumn 2011.

The new platform has also supported internal processes well. The high level of engagement from the MPC in producing each quarterly forecast means that staff have to be able to produce iterations of the forecast to tight deadlines. They must also produce scenario analysis to support MPC discussions at a sequence of meetings in the weeks leading up to the publication of each *Inflation Report*. The improvements in the IT infrastructure and the smaller central model have both helped to increase the efficiency with which Bank staff can update the forecast, and freed up time in which they can think more deeply about the underlying economic issues.

An important aim of this paper has been to demonstrate that the approach of a smaller central organising model, surrounded by a rich suite of other models, can be effective in practice. We have explained our approach to dealing with known misspecifications in COMPASS, shown how those can be quantified and illustrated how the effects of the "missing" economic channels can be incorporated back into COMPASS using suitable shocks. We took the example of frictions in the financial sector, and showed how careful analysis of different suite models can be used to motivate judgements to the forecast which seek to capture the effect of those frictions.

There are still many avenues for future work, and Bank staff expect to make ongoing improvements to all four components of the forecasting platform.

COMPASS, like most macroeconomic models in regular use at policy institutions, is likely to evolve over time for at least two reasons. First, the parameters of the model will be re-estimated on a regular basis (probably annually). Second, Bank staff expect to make alterations to the structure of the model as they learn more over time about its performance. Staff also have the option of adding economic channels to COMPASS, where the benefits of including them are judged to outweigh the costs relative to the alternative of modelling them in the suite.

The suite of models, by its nature, will also evolve over time. As the attention of policymakers switches from one economic question to another, and advances are made in economic modelling, some models will be discarded from the suite, while others will be built and added to it. There are some areas where forthcoming work is planned. For example, the recent Stockton Review concluded that there was some scope for enhancing the role for analysis of monetary policy strategy within the forecast process. This could lead to the construction of more suite models, to incorporate different processes for agents' formation of expectations, and different policy rules. This is also likely to require some incremental improvements in the MAPS toolkit, to support that analysis.

Finally, a continual review is in place to assess the performance of the new forecasting

platform and, in particular, to monitor whether the forecasting process might be supported even more effectively by an alternative type of central model. One advantage of the flexible IT infrastructure described in this paper is that it can be used to support a variety of models. Hence, any decision to use a different central organising model would be straightforward to effect, without there being any need to make potentially costly changes to MAPS or EASE. Bank staff will communicate any changes to the forecasting platform, including to COMPASS itself, on a regular basis.

Bibliography

- ADJEMIAN, S., H. BASTANI, F. KARAM, M. JUILLARD, J. MAIH, F. MIHOUBI, G. PERENDIA, M. RATTO, AND S. VILLEMOT (2011): “Dynare: Reference Manual Version 4,” Dynare Working Papers 1, CEPREMAP.
- ADOLFSON, M., S. LASÉEN, J. LINDÉ, AND L. E. SVENSSON (2011): “Optimal monetary policy in an operational medium-sized DSGE model,” International Finance Discussion Paper 1023, Board of Governors of the Federal Reserve System.
- ADOLFSON, M., S. LASÉEN, J. LINDÉ, AND M. VILLANI (2005): “Are Constant Interest Rate Forecasts Modest Interventions? Evidence from an Estimated Open Economy DSGE Model of the Euro Area,” Sveriges Riksbank Working Paper 180, Sveriges Riksbank.
- (2007): “Bayesian Estimation of an Open Economy DSGE Model with Incomplete Pass-Through,” *Journal of International Economics*, 72, 481–511.
- ALVAREZ-LOIS, P., R. HARRISON, L. PISCITELLI, AND A. SCOTT (2008): “On the application and use of DSGE models,” *Journal of Economic Dynamics and Control*, 2428–2452.
- AN, S. AND F. SCHORFHEIDE (2007): “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26, 113–172.
- ANDERSON, G. AND G. MOORE (1985): “A linear algebraic procedure for solving linear perfect foresight models,” *Economics Letters*, 17, 247–252.
- ANDREASEN, M. M. (2010): “How to maximize the likelihood function of a DSGE model,” *Computational Economics*, 35, 127–154.
- ANDRÉS, J., J. LÓPEZ-SALIDO, AND E. NELSON (2004): “Tobin’s imperfect asset substitution in optimizing general equilibrium,” *Journal of Money, Credit and Banking*, 36, 665–91.
- ANDRLE, M., T. HLEDIK, O. KAMENIK, AND J. VLCEK (2009): “Implementing the New Structural Model of the Czech National Bank,” Working Paper 2009/2, Czech National Bank, Research Department.
- BANERJEE, R. AND N. BATINI (2003): “UK Consumers Habits,” Discussion Papers 13, Monetary Policy Committee Unit, Bank of England.
- BANK OF ENGLAND (1999): *Economic models at the Bank of England*, London: Bank of England.
- (2000): *Economic models at the Bank of England: September 2000 update*, London: Bank of England.
- BARNETT, A. AND R. THOMAS (forthcoming): “Has weak lending in the UK been driven by demand or supply shocks?” Bank of England Working Paper, Bank of England.
- BASU, S. AND J. G. FERNALD (2002): “Aggregate productivity and aggregate technology,” *European Economic Review*, 46, 963–991.

- BATINI, N. AND E. NELSON (2001): “Optimal horizons for inflation targeting,” *Journal of Economic Dynamics and Control*, 25, 891–910.
- BEAN, C. (2007): “Is there a new consensus in monetary policy?” in *Is there a New Consensus in Macroeconomics?*, ed. by P. Arestis, chap. 9.
- BEAN, C. AND N. JENKINSON (2001): “The formulation of monetary policy at the Bank of England,” *Bank of England Quarterly Bulletin*, 41, 434–441.
- BENEŠ, J., A. BINNING, M. FUKAČ, K. LEES, AND T. MATHESON (2009): *K.I.T.T.: Kiwi Inflation Targeting Technology*, New Zealand: Reserve Bank of New Zealand.
- BENFORD, J., S. BERRY, K. NIKOLOV, C. YOUNG, AND M. ROBSON (2009): “Quantitative Easing,” *Bank of England Quarterly Bulletin*, 49, 90–100.
- BENITO, A. AND H. MUMTAZ (2006): “Consumption excess sensitivity, liquidity constraints and the collateral role of housing,” Bank of England working papers 306, Bank of England.
- BENITO, A., J. WHITLEY, AND G. YOUNG (2001): “Analysing corporate and household sector balance sheets,” *Bank of England Financial Stability Report*, 160–174.
- BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): “The financial accelerator in a quantitative business cycle framework,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 21, 1341–1393.
- BLANCHARD, O. J. AND C. M. KAHN (1980): “The Solution of Linear Difference Models under Rational Expectations,” *Econometrica*, 48, 1305–11.
- BOX, G. AND N. DRAPER (1987): *Empirical Model-Building and Response Surfaces*, London: John Wiley and Sons, New York, NY.
- BOX, G., G. JENKINS, AND G. REINSEL (1970): *Time series analysis: forecasting and control*, Wiley.
- BRIDGES, J. AND R. THOMAS (2012): “The impact of QE on the UK economy - some supportive monetarist arithmetic,” Bank of England Working Paper 442, Bank of England.
- BRIGDEN, A. AND P. MIZEN (1999): “Money, credit and investment in the UK corporate sector,” Bank of England Working Paper 100, Bank of England.
- BRUBAKK, L., T. A. HUSEBØ, J. MAIH, K. OLSEN, AND M. ØSTNOR (2006): “Finding NEMO: Documentation of the Norwegian economy model,” Staff Memo 2006/6, Norges Bank.
- BUNN, P. AND G. YOUNG (2004): “Corporate capital structure in the United Kingdom: determinants and adjustment,” Bank of England Working Paper 226, Bank of England.
- BURRIEL, P., J. FERNÁNDEZ-VILLAVERDE, AND J. F. RUBIO-RAMÍREZ (2010): “MEDEA: a DSGE Model for the Spanish Economy,” *SERIEs*, 1, 175–243.

- CALDARA, D., R. HARRISON, AND A. LIPÍŃSKA (2012): “Practical tools for policy analysis in DSGE models with missing channels,” Finance and Economics Discussion Series 2012-72, Board of Governors of the Federal Reserve System (U.S.).
- CANOVA, F. AND L. SALA (2009): “Back to square one: Identification issues in DSGE models,” *Journal of Monetary Economics*, 56, 431–449.
- CARLSTROM, C., T. FUERST, AND M. PAUSTIAN (2012): “Inflation and output in New Keynesian models with a transient interest rate peg,” Bank of England Working Paper 459, Bank of England.
- CARLSTROM, C. T. AND T. S. FUERST (1997): “Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis,” *American Economic Review*, 87, 893–910.
- CARROLL, C., M. SOMMER, AND J. SLACALEK (2012): “Dissecting Saving Dynamics: Measuring Wealth, Precautionary, and Credit Effects,” IMF Working Papers 12/219, International Monetary Fund.
- CARROLL, C. D., J. SLACALEK, AND M. SOMMER (2011): “International Evidence on Sticky Consumption Growth,” *The Review of Economics and Statistics*, 93, 1135–1145.
- CHARI, V. V., P. KEHOE, AND E. MCGRATTAN (2007): “Business cycle accounting,” *Econometrica*, 75, 781–836.
- (2008): “New Keynesian models: not yet useful for policy analysis,” *Federal Reserve Bank of Minneapolis Research Department Staff Report*.
- CHEN, H., V. CÚRDIA, AND A. FERRERO (2012): “The Macroeconomic Effects of Large-scale Asset Purchase Programmes,” *Economic Journal*, 122, F289–F315.
- CHRISTIANO, L., R. MOTTO, AND M. ROSTAGNO (2010): “Financial factors in economic fluctuations,” Working Paper Series 1192, European Central Bank.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- CHRISTOFFEL, K., G. COENEN, AND A. WARNE (2008): “The new area-wide model of the euro area - a micro-founded open-economy model for forecasting and policy analysis,” Working Paper 944, European Central Bank.
- CHUNG, H. T., M. T. KILEY, AND J.-P. LAFORTE (2010): “Documentation of the Estimated, Dynamic, Optimization-based (EDO) model of the U.S. economy: 2010 version,” Finance and Economics Discussion Series Working Paper 2010-29, Board of Governors of the Federal Reserve System.
- CÚRDIA, V. AND D. FINOCCHIARO (2013): “Monetary regime change and business cycles,” *Journal of Economic Dynamics and Control*, 37, 756–773.
- CÚRDIA, V. AND M. WOODFORD (2010): “Credit Spreads and Monetary Policy,” *Journal of Money, Credit and Banking*, 42, 3–35.

- (2011): “The central-bank balance sheet as an instrument of monetary policy,” *Journal of Monetary Economics*, 58, 54–79.
- DAVIDSON, J. E. H., D. F. HENDRY, F. SRBA, AND J. S. YEO (1978): “Econometric Modelling of the Aggregate Time-Series Relationship between Consumers’ Expenditure and Income in the United Kingdom,” *Economic Journal*, 88, 661–92.
- DEL NEGRO, M., M. GIANNONI, AND C. PATTERSON (2012): “The forward guidance puzzle,” Staff Reports 574, Federal Reserve Bank of New York.
- DEL NEGRO, M. AND F. SCHORFHEIDE (2004): “Priors from General Equilibrium Models for VARS,” *International Economic Review*, 45, 643–673.
- (2009): “Monetary Policy Analysis with Potentially Misspecified Models,” *American Economic Review*, 99, 1415–50.
- (forthcoming): “DSGE model-based forecasting,” in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmerman, vol. 2A, chap. 2.
- DORICH, J., R. R. MENDES, AND Y. ZHANG (2011): “Introducing Multiple Interest rates in ToTEM,” *Bank of Canada Review*, 2011, 3–10.
- EDGE, R. M., M. T. KILEY, AND J.-P. LAFORTE (2007): “Documentation of the Research and Statistics Division’s estimated DSGE model of the U.S. economy: 2006 version,” .
- (2010): “A comparison of forecast performance between Federal Reserve staff forecasts, simple reduced-form models, and a DSGE model,” *Journal of Applied Econometrics*, 25, 720–754.
- EGGERTSSON, G. B. AND P. KRUGMAN (2012): “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *The Quarterly Journal of Economics*, 127, 1469–1513.
- ELDER, R. AND J. TSOUKALAS (2006): “Investing in inventories,” *Bank of England Quarterly Bulletin*, 46, 155–160.
- ELLIS, C. AND S. PRICE (2004): “UK Business Investment and the User Cost of Capital,” *Manchester School*, 72, 72–93.
- FACCINI, R. AND C. HACKWORTH (2010): “Changes in output, employment and wages during recessions in the United Kingdom,” *Bank of England Quarterly Bulletin*, 50, 43–50.
- FERNANDEZ-CORUGEDO, E., M. MCMAHON, S. MILLARD, AND L. RACHEL (2011): “Understanding the macroeconomic effects of working capital in the United Kingdom,” Bank of England working papers 422, Bank of England.
- GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 5, 227–270.

- GERTLER, M. AND P. KARADI (2011): “A model of unconventional monetary policy,” *Journal of Monetary Economics*, 58, 17–34.
- GERTLER, M. AND N. KIYOTAKI (2010): “Financial Intermediation and Credit Policy in Business Cycle Analysis,” in *Handbook of Monetary Economics*, ed. by B. M. Friedman and M. Woodford, Elsevier, vol. 3 of *Handbook of Monetary Economics*, chap. 11, 547–599.
- GEWEKE, J. (2005): *Contemporary Bayesian Econometrics and Statistics*, New Jersey: Wiley.
- GILCHRIST, S. AND E. ZAKRAJSEK (2012): “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102, 1692–1720.
- GÓMEZ, A. G., L. MAHADEVA, D. RODRÍGUEZ, AND L. E. ROJAS (2009): “Monetary Policy Forecasting in a DSGE Model with Data that is Uncertain, Unbalanced and About the Future,” Borradores de Economía 559, Banco de la Republica de Colombia.
- GREENSLADE, J. V. AND M. PARKER (2012): “New Insights into PriceSetting Behaviour in the UK: Introduction and Survey Results,” *Economic Journal*, 122, F1–F15.
- HAMILTON, J. D. (1986): “State-space models,” in *Handbook of Econometrics*, ed. by R. F. Engle and D. McFadden, Elsevier, vol. 4.
- (1994): *Time Series Analysis*, London: Princeton University Press.
- HANSEN, L. P. AND T. SARGENT (2007): *Robustness*, Princeton.
- HANSEN, N. AND S. KERN (2004): “Evaluating the CMA Evolution Strategy on Multimodal Test Functions,” Tech. rep.
- HARRISON, R. (2012): “Asset purchase policy at the effective lower bound for interest rates,” Bank of England Working Paper 444, Bank of England.
- HARRISON, R., K. NIKOLOV, M. QUINN, G. RAMSAY, A. SCOTT, AND R. THOMAS (2005): *The Bank of England Quarterly Model*, London: Bank of England.
- HARRISON, R. AND O. OOMEN (2010): “Evaluating and estimating a DSGE model for the United Kingdom,” Bank of England Working Paper 380, Bank of England.
- HARRISON, R. AND T. TAYLOR (2012): “Non-rational expectations and the transmission mechanism,” Bank of England Working Paper 448, Bank of England.
- HARRISON, R., R. THOMAS, AND I. DE WEYMARN (2011): “The impact of permanent energy price shocks on the UK economy,” Bank of England Working Paper 433, Bank of England.
- HUGHES, A. AND J. SALEHEEN (2012): “UK labour productivity since the onset of the crisis – an international and historical perspective,” *Bank of England Quarterly Bulletin*, 52, 138–146.
- ISKREV, N. (2010): “Local identification in DSGE models,” *Journal of Monetary Economics*, 57, 189–202.

- JOYCE, M., A. LASAOSA, I. STEVENS, AND M. TONG (2011a): “The Financial Market Impact of Quantitative Easing in the United Kingdom,” *International Journal of Central Banking*, 7, 113–161.
- JOYCE, M., M. TONG, AND R. WOODS (2011b): “The United Kingdoms quantitative easing policy: design, operation and impact,” *Bank of England Quarterly Bulletin*, 51, 200–212.
- JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): “Investment Shocks and the Relative Price of Investment,” *Review of Economic Dynamics*, 14, 101–121.
- KAMATH, K. AND V. PAUL (2011): “Understanding recent developments in UK external trade,” *Bank of England Quarterly Bulletin*, 51, 294–304.
- KAPETANIOS, G., V. LABHARD, AND S. PRICE (2006): “Forecasting using predictive likelihood model averaging,” *Economics Letters*, 91, 373–379.
- (2007): “Forecasting using Bayesian and information theoretic model averaging: an application to UK inflation,” Working Paper 07/15, Department of Economics, City University London.
- (2008): “Forecast combination and the Bank of England’s suite of statistical forecasting models,” *Economic Modelling*, 25, 772–792.
- KAPETANIOS, G., H. MUMTAZ, I. STEVENS, AND K. THEODORIDIS (2012): “Assessing the Economywide Effects of Quantitative Easing,” *Economic Journal*, 122, F316–F347.
- KIM, D. AND C.-I. LEE (2007): “On-the-Job Human Capital Accumulation in a Real Business Cycle Model: Implications for Intertemporal Substitution Elasticity and Labor Hoarding,” *Review of Economic Dynamics*, 10, 494–518.
- KLEIN, P. (2000): “Using the generalized Schur form to solve a multivariate linear rational expectations model,” *Journal of Economic Dynamics and Control*, 24, 1405–1423.
- LASÉEN, S. AND L. E. SVENSSON (2011): “Anticipated Alternative policy Rate Paths in Policy Simulations,” *International Journal of Central Banking*, 7, 1–35.
- LAXTON, D., D. ROSE, AND A. SCOTT (2009): “Developing a Structured Forecasting and Policy Analysis System to Support Inflation-Forecast Targeting (IFT),” IMF Working Papers 09/65, International Monetary Fund.
- LAYARD, R., S. NICKELL, AND R. JACKMAN (1991): *Unemployment: Macroeconomic Performance and the Labour Market*, Oxford University Press.
- LEEPER, E. M. AND C. A. SIMS (1994): “Toward a modern macroeconomic model usable for policy analysis,” Working Paper 94-5, Federal Reserve Bank of Atlanta.
- LEEPER, E. M. AND T. ZHA (2003): “Modest policy interventions,” *Journal of Monetary Economics*, 50, 1673–1700.
- LYNCH, B. AND S. WHITAKER (2004): “The new sterling ERI,” *Bank of England Quarterly Bulletin*, 44, 429–441.

- MAIH, J. (2010): “Conditional forecasts in DSGE models,” Working Paper 2010/07, Norges Bank.
- MILLARD, S. (2011): “An estimated DSGE model of energy, costs and inflation in the United Kingdom,” Bank of England Working Paper 432, Bank of England.
- MONETARY POLICY COMMITTEE (1999): *The transmission mechanism of monetary policy*, London: Bank of England.
- (2013): “Inflation Report,” February 2013, Bank of England.
- MONTI, F. (2010): “Combining Judgment and Models,” *Journal of Money, Credit and Banking*, 42, 1641–1662.
- MURCHISON, S. AND A. RENNISON (2006): “ToTEM: The Bank of Canada’s New Quarterly Projection Model,” Technical Report 97, Bank of Canada.
- PAGAN, A. (2003): “Report on modelling and forecasting at the Bank of England,” Tech. rep., Bank of England.
- SCHMITT-GROHE, S. AND M. URIBE (2003): “Closing small open economy models,” *Journal of International Economics*, 61, 163–185.
- SCHORFHEIDE, F. (2000): “Loss function-based evaluation of DSGE models,” *Journal of Applied Econometrics*, 15, 645–670.
- SIMS, C. A. (2002): “Solving Linear Rational Expectations Models,” *Computational Economics*, 20, 1–20.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97, 586–606.
- SUSTEK, R. (2011): “Monetary Business Cycle Accounting,” *Review of Economic Dynamics*, 14, 592–612.
- TOVAR, C. E. (2008): “DSGE Models and Central Banks,” Economics Discussion Papers 2008-30, Kiel Institute for the World Economy.
- VILLA, S. AND J. YANG (2011): “Financial intermediaries in an estimated DSGE model for the United Kingdom,” Bank of England Working Paper 431, Bank of England.
- WAGGONER, D. F. AND T. ZHA (1999): “Conditional Forecasts In Dynamic Multivariate Models,” *The Review of Economics and Statistics*, 81, 639–651.
- WEN, Y. (2004): “What Does It Take to Explain Procyclical Productivity?” *The B.E. Journal of Macroeconomics: Contributions*, 0, 5.
- WHITLEY, J. (1997): “Economic models and policy-making,” *Bank of England Quarterly Bulletin*, 37, 163–173.