Bank of England

Monetary policymaking under uncertainty

Macro Technical Paper No. 5
October 2025

Alex Haberis, Richard Harrison, Kate Reinold and Matt Waldron

A series designed to document models, analysis and conceptual frameworks for monetary policy preparation – they are written by Bank staff to encourage feedback and foster continued model development.

Macro Technical Paper Series

Dr Bernanke's 2024 **review** of the monetary policymaking processes at the Bank of England provided a number of constructive recommendations for reform which we are taking forward.

This included improving our model maintenance and development. Macroeconomic models and frameworks play an important role in the monetary policy process. To maximise the value of macroeconomic models, they must be well documented and continuously improved as time goes on.

This Macro Technical Paper (MTP) series is part of the Bank's response to Dr Bernanke's recommendations. These MTPs are authored by Bank staff, and are intended to document models, analysis, and conceptual frameworks that underpin monetary policy preparation. The models documented in the series will typically be used to assess the current state of the economy, forecast its future, and to simulate alternative paths and policy responses.

Importantly, while each MTP will provide insights about a particular model or modelling framework that is an 'input' to policy, no single model can possibly capture all the relevant features to perform even just one of those roles adequately. Models will inevitably have to be updated and will improve over time, including as they are adapted to different constellations of macroeconomic conditions. This is a natural part of real-time monetary policy making. So, inevitably, no MTP will provide definitive answers.

The Bank seeks to encourage an active and informed debate about its modelling frameworks. Publishing and discussing the analytical work undertaken to support its monetary policy choices is central to this ambition. These MTPs will support a culture of continuous learning in monetary policy making. As time goes on, Bank staff will update and upgrade models, drawing on insights from the frontier of the academic literature. Moreover, this transparency will encourage external engagement from experts to ensure our modelling tools remain fit for purpose.

Clare Lombardelli
Deputy Governor for Monetary Policy

Bank of England

Macro Technical Paper No. 5

Monetary policymaking under uncertainty

Alex Haberis,[1] Richard Harrison,[2] Kate Reinold[3] and Matt Waldron[4]

Abstract

Pervasive and time-varying uncertainty characterises the monetary policy environment. While theory provides insights into setting monetary policy under uncertainty, it is widely recognised that theoretical approaches are too simplified to be directly applicable. This paper discusses three alternative perspectives on the policy problem that can be applied in practice and that acknowledge uncertainty in different ways: a forecast-based, a news-based, and a rules-based perspective. These perspectives are complementary so that insights from all of them may be informative for policymaking. They should also be applied flexibly and with varying emphases depending on judgements about the prevailing level and nature of uncertainty. These observations provide some context for the approach that the Bank of England is taking to adapting monetary policymaking as part of its response to the Bernanke Review.

Key words: Monetary policy, uncertainty, misspecification.

JEL classification: E52, E58.

---

(1) Bank of England and Centre for Macroeconomics. Email: alex.haberis@bankofengland.co.uk
(2) Bank of England and Centre for Macroeconomics. Email: richard.harrison@bankofengland.co.uk
(3) Bank of England. Email: kate.reinold@bankofengland.co.uk
(4) Bank of England. Email: matthew.waldron@bankofengland.co.uk

The authors would like to thank Andy Blake, Bob Hills, Ben Nelson, Huw Pill, James Proudman, Fergal Shortall, Danny Walker and Iain de Weymarn for helpful discussions and comments during the preparation of this paper.

The Bank's macro technical paper series can be found at
www.bankofengland.co.uk/macro-technical-paper/macro-technical-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 2978-3194 (online)

1 Introduction

The presence and implications of uncertainty have been key themes in the development of both the theory of monetary policy and the practice of monetary policymaking.[1] While uncertainty is always a part of the monetary policymaking landscape, the past two decades have been characterised by a marked increase in volatility, as policymakers confronted the consequences of large and unprecedented shocks, against a backdrop of structural change (Broadbent, 2024).

The size and unusual nature of these shocks, their supply-side implications, and the accompanying increase in uncertainty, have all made monetary policy more challenging. In contrast to the fifteen years prior to the Global Financial Crisis (GFC), monetary policymakers have, at various times, been confronted with material trade-offs between deviations of inflation from target and the balance between demand and supply. Indeed, following Russia's invasion of Ukraine, inflation in many countries, including the United Kingdom, rose rapidly to levels not seen since the 1980s (long before the adoption of inflation targeting). While the inflation targeting framework for UK monetary policy has stood up well to these challenges, they have nonetheless stress-tested all aspects of the monetary policymaking process within that framework.

This is why the Bank of England commissioned Dr Ben Bernanke to conduct a review "to consider the appropriate approach to forecasting and analysis in support of decision-making and communications in times of high uncertainty from big shocks and structural change" (Bernanke, 2024). The Bernanke Review made twelve wide-ranging recommendations, covering the support provided by Bank staff to the Monetary Policy Committee (MPC) via modelling, forecasting and broader analysis, the role of forecasts in the MPC's decision-making process, and communication of the outlook and policy.

In its initial response, the Bank welcomed the Review and, given its broad scope and the interconnected nature of the recommendations, recognised that "the consequences could affect many aspects of the monetary policymaking process" (Bank of England, 2024). Accordingly, the Bank is "reforming the whole nose to tail process of monetary policy making and communication" (Lombardelli, 2024).

This paper provides context to these reforms by exploring the implications of uncertainty for monetary policymaking in practice, drawing insights from the analysis for the ongoing changes to monetary policy processes and communication. In so doing, the paper contributes to the literature by setting out an encompassing characterisation of alternative practical perspectives on the monetary policy problem and discussing the role

[1] This is evidenced by a long history of contributions by monetary policymakers. See, for example, Issing (2002), Greenspan (2004), King (2004, 2010), Bean (2005, 2007), Bernanke (2007), Carney (2016), Pill (2022), Lane (2024), Bailey (2025) and Williams (2025).

of judgement, topics that have received relatively little attention in the literature to date. Where insititutional details are important, the paper naturally takes a UK perspective. However, much of the analysis is applicable to monetary policymaking in general.

Assessing the role of uncertainty in monetary policymaking, and how changes in uncertainty might affect its conduct, requires an assessment of the nature of uncertainty confronting policymakers. Uncertainty is pervasive and multi-faceted. Most obviously, policymakers face uncertainty about the future, made relevant by the "long and variable" lags with which policy actions affect economic outcomes (Friedman, 1960, 1961). However, policymakers also face uncertainty about the past and the present. Economic data contain measurement noise and there is even more uncertainty around estimates of inherently unobservable variables relevant for the conduct of policy, such as potential output and the equilibrium real interest rate. More fundamentally, policymakers face uncertainty about the way the economy works, giving rise to uncertainty about how both shocks and policy actions transmit to economic outcomes.

For many aspects of practical policymaking, learning is a crucial and natural response to uncertainty. Learning as data and analyses evolve helps to shed light on the shocks that are affecting the economy, the ways in which those shocks are transmitting, and the appropriateness of the stance of monetary policy. Beyond learning about the current state of the economy, research and analysis helps to develop understanding about how the economy functions. That development in general understanding has helped to underpin the successful monetary policy regimes that have operated across much the world for the past thirty years and continues to deepen understanding of the economy.

Nevertheless, while ongoing learning is essential, both the inherent complexity of modern economies and the structural changes that they are subject to suggest that achieving a complete understanding of how the economy works, or of exactly what forces are driving it, is an impossible task. Some amount of uncertainty is an ever-present and inherent feature of the monetary policymaking landscape.

Both the nature and extent of uncertainty can vary substantially over time. Monetary policy operates in a grey area of uncertainty: policymakers neither have a complete understanding of the state of the economy and its potential evolution; but neither do they lack any basis at all on which to set policy. During some periods, such as the fifteen years of relative stability preceding the GFC, signals from macroeconomic data and models may be relatively more informative for monetary policymaking than at other times, when uncertainty is a much more material consideration, such as the particularly volatile period since 2020.

In light of this, the challenge for policymakers is how best to proceed in practice in an environment of pervasive and time-varying uncertainty. Theory-based approaches to solving the monetary policy problem assume environments that are too simplified to be applied directly to the practical policy problem.Real-world policymaking necessarily relies on judgement and so the challenge can be reframed in terms of how to support judgement-based policymaking most effectively.

One part of the policy problem over which there is more certainty and, therefore, over which less judgement is required relates to the objectives of policy. By specifying a numerical target for inflation, an inflation targeting policy framework establishes a crucial aspect of the objectives of monetary policy with clarity. In this way, the framework provides an environment of "constrained discretion", within which policymakers formulate judgement-based policy (Bernanke and Mishkin, 1997). Importantly, such a regime also reduces private sector uncertainty about policy objectives, helping to ensure that inflation expectations are anchored to the target in the medium and longer term.

Within a framework of constrained discretion, monetary policy should be set in as systematic and, therefore, as predictable, a way as possible (King, 1997). As noted above, quantitative solutions to simplified versions of the monetary policy problem cannot be applied *directly* to practical policymaking. However, insights from these approaches provide valuable ways of thinking about the policy problem – 'perspectives' – that can be used to organise judgement and promote systematic policymaking.

This paper discusses three such perspectives on monetary policymaking, each of which draws on the insights from theory, facilitates the application of judgement, and acknowledges uncertainty in different ways.

A *forecast-based perspective* uses a macroeconomic forecast as the basis for setting policy. It embraces a forward-looking and quantitative approach to policymaking, but relaxes the tight links between the model of the economy, the forecast, and the policy response that is implied by theory, in a way that supports the application of judgement.

A *news-based perspective* involves updating policy over time, as policymakers' understanding of the underlying drivers of the economy evolves, and as new shocks or events occur. Relative to a forecast-based perspective, it embodies greater flexibility over the approach to policy formulation (e.g., more or less forward looking) and does not require (though can be supported by) quantification.

A *rules-based perspective* uses simple rules that relate the policy instrument to a small number of macroeconomic variables as the basis for setting policy. This perspective recognises both the importance of a systematic monetary policy response to the state of the economy and the difficulties in approximating an optimal response when there is substantial uncertainty about both the state and structure of the economy.

In the grey area of uncertainty in which monetary policy operates, these alternative perspectives may be complementary. For example, a rules-based perspective may be a particularly valuable complement to either of the others because, by its nature, it requires less judgement. Similarly, a news-based and forecast-based perspective may be used alongside each other to inform policy in a way that systematically embeds a response to news, but is not completely reliant on a quantitative forecast.

The relevance and usefulness of each of these perspectives depends on the prevailing economic environment and the nature of uncertainty that policymakers face. While all offer some flexibility of application, which allows them to be tailored somewhat to the prevailing environment, they each presume differing degrees of knowledge about the current state of (and outlook for) the economy. For example, a forecast-based perspective may be less reliable in times of heightened uncertainty.

Applying insights from any of these perspectives requires judgement to bridge the gap to the real-world policy problem. In doing so, policymakers need to form a judgement about the nature of the prevailing uncertainties and, conditional on that, the relevance of alternative perspectives for the practical policy problem. Such judgements naturally form part of an assessment of the most appropriate monetary policy strategy to adopt.

Application of any of the perspectives requires a set of supporting 'inputs', including models, data, various forms of intelligence, and a variety of analysis by staff and policymakers. Judgement is also required to synthesise and extract information from this array of supporting inputs. In general, the specific inputs that inform policy decisions will also vary depending on the circumstances, including the degree and nature of the uncertainty, and the policy perspective(s) that policymakers judge to be most appropriate.

While judgement is an essential part of policymaking, there remains a practical question of how individual and collective judgements are formed and, therefore, how judgement-based policymaking can be supported most effectively. There is no single algorithmic way to characterise the formation of judgement. Appropriate 'framing' of policy questions (including via the perspectives described in this paper), the 'trained intuition' of policymakers and staff, and application of 'analogical' (or 'case-based') reasoning may all provide valuable ways to support policy deliberations.

When monetary policy is set by a committee, individual policymakers may naturally form different judgements about all aspects of the policy problem. Indeed, a key strength of committee-based decision-making is the diversity of experiences and approaches that different policymakers bring to the table. This underscores the benefits of supporting judgement-based policymaking by applying different perspectives on the policy problem and using a broad range of different supporting analytical inputs.

In practice, policymakers have indeed applied a variety of approaches in their deliberations, particularly in times of heightened uncertainty. This is evident in the published transcripts from past MPC meetings and the analysis that supported those discussions. The changes to monetary policymaking that the Bank of England is implementing following the Bernanke Review are designed to improve the processes that support the MPC's approach to judgement-based policymaking, consistent with several insights from this analysis.

First, there is a case for broadening the perspectives that are applied to monetary policymaking at the Bank of England, building on the observation that "[a] forecasting framework that performed well during periods in which inflationary shocks were relatively small may be less robust during periods of large and unusual shocks and the associated heightened uncertainty" (Bank of England, 2024).

Second, and related, the emphasis that is placed on different perspectives is likely to vary over time, as the economic environment, and the nature of uncertainty, varies. This may mean placing greater emphasis on forecast-based analysis when the economy is apparently more stable than when uncertainty is a more material consideration. This reasoning is reflected in planned changes to the content of the MPC's quarterly Monetary Policy Report, to reflect a more pluralistic and flexible approach to policymaking (Bank of England, 2025).

Third, consistent with the Review's recommendations, the Bank's modelling framework and toolkit should be improved and extended to provide better support for a more pluralistic approach. One specific area of improvement suggested by this paper is further development of policy analysis tools to facilitate a more thorough treatment of risk management considerations.

This paper connects to several strands of the academic literature. As noted above, the paper contributes to a small literature that considers monetary policymaking in practice, including the role of uncertainty and risk management (e.g., Blinder, 1998; Faust, 2005; Mendes et al., 2017; Garga et al., 2025). The paper also draws on strands of the literature that consider alternative formal treatments of monetary policy under uncertainty (e.g., Hansen and Sargent, 2008; Levin and Williams, 2003; Svensson, 2010) and general discussions and theories of decision-making under uncertainty (e.g., Kay and King, 2020; Spiegelhalter, 2024; Gilboa and Schmeidler, 2001).

The rest of the paper is organised as follows. Section 2 discusses the implications of uncertainty for monetary policy. Section 3 sets out three perspectives on the monetary policy problem, each deriving from different insights from a theory-based benchmark. Section 4 discusses the roles of data and analytical inputs in informing each perspective and the role of judgement in monetary policymaking. Section 5 concludes with a summary of the analysis, highlighting insights for the Bank's response to the Bernanke Review.

2 Monetary policy and uncertainty

> Uncertainty is not just an important feature of the monetary policy landscape; it is the defining characteristic of that landscape. Greenspan (2003).

The problem of how to set monetary policy in the face of uncertainty has been extensively investigated by both academic researchers and practitioners. In part, this reflects the fact that empirical evidence has consistently confirmed the view of Friedman (1961) that "monetary actions affect economic conditions only after a lag that is both long and variable."[2] This section draws on academic research to provide a conceptual framework for understanding the nature of uncertainty characterising the policy landscape, including how it changes over time. It then discusses the specific ways in which the literature has formalised the problem of monetary policy under uncertainty, how these may be located within that conceptual framework, and how insights from theory can be used to inform real-world policymaking.

To provide context for the discussion, it is useful to consider how uncertainty can affect the 'monetary policy problem'. At a high level, the 'problem' for a monetary policymaker is to determine the setting for its instrument(s) that best achieves its objectives. In approaching this problem, a monetary policymaker would naturally seek to address a number of questions. What forces might be pushing the objective variables away from their target values and how might they evolve? How do changes in the instrument(s) affect the objective variables? What is the most appropriate adjustment to the instrument(s) to bring objective variables back to their targets? Monetary policymakers face uncertainty about all the information that informs consideration of these questions:[3]

- Data and measurement uncertainty. Economic data are typically subject to sampling error, are often published with a lag, can be subject to considerable revision over time, and do not always map directly to relevant economic concepts.[4] As a result, there is always uncertainty about the state of the economy and that uncertainty is more pronounced in real time.

[2]Empirical evidence tends to suggest that monetary policy actions have a peak effect on activity and prices of anywhere between 1 and 3 years (see, e.g., Burr and Willems, 2024, for a discussion of the monetary transmission mechanism in the United Kingdom). This is relevant for an assessment of uncertainty in monetary policymaking because it implies that: (a) uncertainty about the future can be important; (b) uncertainty about the monetary transmission mechanism is harder to mitigate because learning about the effects of monetary policy is inevitably a relatively slow process.

[3]There are various ways of classifying sources of uncertainty in monetary policymaking. See also, for example, Mendes et al. (2017) and Bernanke (2007).

[4]The indirect mapping from some data to relevant economic concepts muddies the distinction between data and model uncertainty. This is most apparent when it comes to estimates of fundamentally unobservable variables, like the level of potential supply or the natural rate of interest, because these variables are underpinned by concepts that do not have meaning outside of models.

- Shock or event uncertainty. At any point in time, it is hard to discern precisely all the shocks that are affecting the economy. Furthermore, the future is inherently uncertain. It is impossible to specify all potential future shocks. Even when potential future shocks or events are known (e.g. an event like an election), their outcomes are not.

- Transmission uncertainty. Both the effect of any shock and the effect of any monetary policy response depend on the structure of the economy, including the beliefs and behaviours of households, businesses, and other policymakers. Understanding of this is necessarily incomplete given the factors underpinning decision-making are unobservable and given the complexity of economic relationships, thereby leading to uncertainty about how the economy works.

Further light can be shed on the nature of the uncertainty facing monetary policymakers by distinguishing, conceptually, between resolvable and irresolvable, or 'epistemic', uncertainty.[5] In a hypothetical world in which all uncertainty is resolvable, it would be possible to characterise – in an objectively correct and quantitative way – all relevant uncertainties. In principle, this would allow a policymaker to characterise the joint probability distribution of the current state of the economy and to project that into the future. In this hypothetical world, discussed further in Section 3, a policymaker is able to set policy in an optimal manner – i.e. to achieve the best possible outcomes given the policy objectives.

By contrast, in the real world, in which epistemic uncertainty is a material consideration, policymakers do not know enough about the way the economy works or about the characteristics of the uncertainty they face to characterise the joint probability distribution of outcomes in this way. An immediate implication – that is well-known to monetary policymakers – is that there is no quantitative, optimal policy solution to the real-world monetary policy problem.[6] As noted by Blinder (2007): "the implicit optimisation problem facing a monetary policy committee is far too hard to be solved explicitly. It may not even be well defined."

[5]See, for example, Kay and King (2020) and Spiegelhalter (2024) for more extensive and general discussions. In the context of economics, the distinction between resolvable and irresolvable uncertainty is often attributed to Knight (1921), who defined 'risk proper' as uncertainties that could be objectively characterised (like the outcome of rolling fair dice) and who reserved the term 'uncertainty' for uncertainties that cannot be objectively characterised in this way. The present paper does not use these terms for two reasons. The first is that they create the potential for confusion given that common language usage of the words 'risk' and 'uncertainty' does not conform to Knight's precise distinction. Second, as argued by Spiegelhalter (2024, Chapter 13), "The unfortunate phrase, 'Knightian uncertainty', has come to be used for situations when people "don't know the probability distribution", but this inappropriately implies that probability is an objective property of the world which we happen to not know."

[6]Another implication is that attaching probabilities to eventualities of relevance to a monetary policymaker, such as the distribution of future inflation outturns given some policy plan, is a subjective exercise, regardless of whether or not a model is used to quantify those probabilities. This issue is briefly discussed in Section 4.

One obvious and desirable response to recognising uncertainty is to try to reduce it. As discussed briefly below, in the context of shocks realised over recent economic history, learning about how shocks are transmitting and about how the state of the economy is evolving is a critical part of policymaking. As King (2005) observes "monetary policy in practice is characterised by a continuous process of learning embedded, in the case of the Bank of England, in the rounds of meetings and forecasts that are the daily life of the Monetary Policy Committee."

Beyond day-to-day learning, improvements in data collection and measurement, and research and analysis to improve understanding of how the economy works, are valuable endeavours that have had a tangible impact on the capacity of policymakers to set policy in support of their objectives.

The benefits of this research are reflected in the monetary policy regimes in operation across much of the world. Friedman's (1960) proposal that monetary policy should follow a simple money growth rule was a contribution to a decades-long debate about the appropriate objectives and instruments of monetary policy. The consensus that emerged during the 1990s that monetary policy should aim to stabilise inflation over the medium to long run, using the short-term interest rate as the primary instrument, has proved to be durable in part because of its demonstrable success compared to what went before. This was underpinned by research setting out the conditions under which there is no long-run trade-off between inflation and unemployment for monetary policy to exploit (Phelps, 1967, 1968; Friedman, 1968) and research demonstrating the instability of demand for money due to variation in the (unobservable) velocity of money (see, e.g., Goodhart, 1989, for a discussion of the UK experience and evidence). Subsequent research has deepened understanding of how both the macroeconomy and monetary policy work in many ways, including the joint determination of prices and quantities in the labour market, the nature and consequences of rigidities in price and wage setting, the role of financial frictions in the transmission of shocks and policy, and the importance of accounting for household heterogeneity in the monetary transmission mechanism.[7]

Nevertheless, it remains the case that policymakers' understanding of how the economy works is very far from complete. While ongoing research is essential to continue to enhance understanding over time, realising a complete understanding of how the economy works is an impossible task.[8] Both the economy and the behaviours of agents within

[7] See Taylor and Uhlig (2016) for a large set of surveys that together chart developments in, and provide a summary of, the current state of macroeconomic study. See also Reis (2018) for a contemporary discussion of the state of macroeconomic research.

[8] As has been the case in the past, progress is neither linear nor guaranteed at all: "Our understanding of the economy is [...] constantly evolving, sometimes in small steps, sometimes in big leaps" (King, 2005). It is also possible that what looks like progress at the time may be judged not to be by some observers after the fact: "Econometrics and macroeconomics were active research areas during the 1970s, 1980s and 1990s, and one might therefore have hoped that there would be clear progress [...]. But if

it are far too complex to be fully understood. Furthermore, structural and technological changes, such as globalisation and Artificial Intelligence, effectively create a moving target. Some amount of epistemic uncertainty is, therefore, an inherent feature of the environment in which monetary policymakers operate. Indeed, given that both data and model uncertainty are always present to some degree, uncertainty in the monetary policymaking environment has been described by policymakers as "pervasive" (Greenspan, 2004; King, 2010).

As well as being pervasive, the extent and nature of uncertainty is also state-contingent and, hence, time-varying. For example, the first decade of monetary policy independence for the Bank of England's MPC was characterised by unusually low volatility. Fluctuations in activity were generally driven by relatively small shocks to aggregate demand against a backdrop of steadily and consistently expanding supply. If that period was 'NICE' ('Non-Inflationary and Consistently Expansionary'), the period that has followed has been described as 'NAsTY' ('Not-As-Tranquil Years', Broadbent, 2024).[9] This period included the GFC, the most material change to UK trading arrangements for many decades, the largest pandemic in a century, and the largest war in Europe since 1945. These shocks have been large, mainly global, and have had material consequences for the supply-side of the economy.[10]

The shift from the relatively small 'demand-like' shocks of the NICE period to the relatively large global 'supply-like' shocks of the NAsTY period that followed has had a demonstrable impact on UK monetary policymaking (Broadbent, 2024). Intuitively, if most shocks are 'demand-like', fluctuations in GDP growth, and indicators thereof, can be presumed to contain a strong signal for future inflation and, therefore, monetary policy (Broadbent, 2023). By contrast, supply shocks create short-run trade-offs for monetary policy, whereby GDP and inflation move in opposite directions. So, if there is more uncertainty about whether shocks are 'demand-like' or 'supply-like', it is appropriate to take a stronger signal from a broader range of data (e.g., including labour market data) to try to disentangle the types of shock that are hitting the economy. This may go hand-in-hand with waiting for longer before acting, to the extent that lags in the transmission

there has been progress, it certainly has not been clear, and my own view is that, by and large, the changes in these models over time have been more regress than progress" (Sims, 2002).

[9]A distinction was made in King (2007) between the 1990s, described as 'NICE', and the pre-GFC part of the 2000s, described as 'NOT-SO-BAD' ('Not Of The Same Order But Also Desirable'). Given what has transpired subsequently, it seems reasonable to take the view that the entirety of that period can be described as 'NICE'.

[10]Recommendation 4.e. of the Bernanke Review states that a revamped forecasting framework should include "greater attention to, and ongoing review of, supply-side elements and their role in the determination of inflation and growth. Important supply-side factors include changes in productivity, labour supply, the efficiency of job-worker matching, supply-chain disruptions, and trade policy. Notably, analyses of inflation should consider supply-side factors as well as the state of aggregate demand" (Bernanke, 2024). See also Greene (2025) for a discussion of the importance of supply-side developments in recent years and of the desirability of understanding the supply-side of the economy better.

mechanism generally imply that GDP tends to react to shocks more quickly than labour market quantities and inflation.

As well as generally higher levels of background and supply-side uncertainty, the nature of the specific shocks encountered over recent years implied that there was naturally less understanding of their potential effects. The Covid pandemic is a particularly stark example. Given the unprecedented nature of the shock, there was relatively little historical experience from which to assess its potential effects. Moreover, the range of potential outcomes was very large. In recognition of this extreme uncertainty, the MPC did not publish its usual forecasts for inflation and GDP growth in the May 2020 *Monetary Policy Report*. Instead, it published an illustrative scenario, alongside a discussion of some of the factors likely to be relevant for determining how the economy would evolve. As time progressed, the MPC and other policymakers were able to learn about the economic effects of Covid, and the public health measures taken to contain it. As a result, the level of uncertainty, while still considerable, decreased over time. To a less extreme degree, uncertainty arising from the GFC, Brexit, and Russia's invasion of Ukraine can be characterised in a similar way.

These ideas are illustrated in a subjective manner in Figure 1, adapted from Spiegelhalter (2024). The figure loosely describes uncertainty associated with particular periods or events from recent economic history along two dimensions: the degree of understanding of the range of potential outcomes as relevant for monetary policy formulation; and the strength of the basis for attaching probabilities to those outcomes. At one extreme, discussed above and denoted by (A) in the top left corner, there is no epistemic uncertainty and it is possible to characterise correctly the complete joint probability distribution of all current and future outcomes relevant for monetary policymaking. In this corner, therefore, there is an objectively optimal solution to the monetary policy problem. At the other extreme, denoted by (C) in the bottom right corner, the policymaker is confronted with extreme, or very 'deep', uncertainty. In this case, policymakers have very little understanding of the range of potential outcomes relevant for monetary policymaking, and very little basis on which to attach probabilities to those outcomes.

In general, monetary policy operates in the grey area between the extremes of perfect knowledge and very deep uncertainty. This is depicted by the shaded ovoid in Figure 1. As discussed above, corner (A) is not applicable to real-world monetary policymaking. The existence of epistemic uncertainty implies that the policymaker does not possess the knowledge required to implement an optimal policy solution to the real-world monetary policy problem. While there are circumstances in which policymakers know less about the state of the economy and how it might evolve than they ordinarily do, they are also generally not in a position of very deep uncertainty. Official statistics and other forms of data typically reveal something about what is going on in the economy, and accumulated

Figure 1: The changing nature of uncertainty in recent economic history

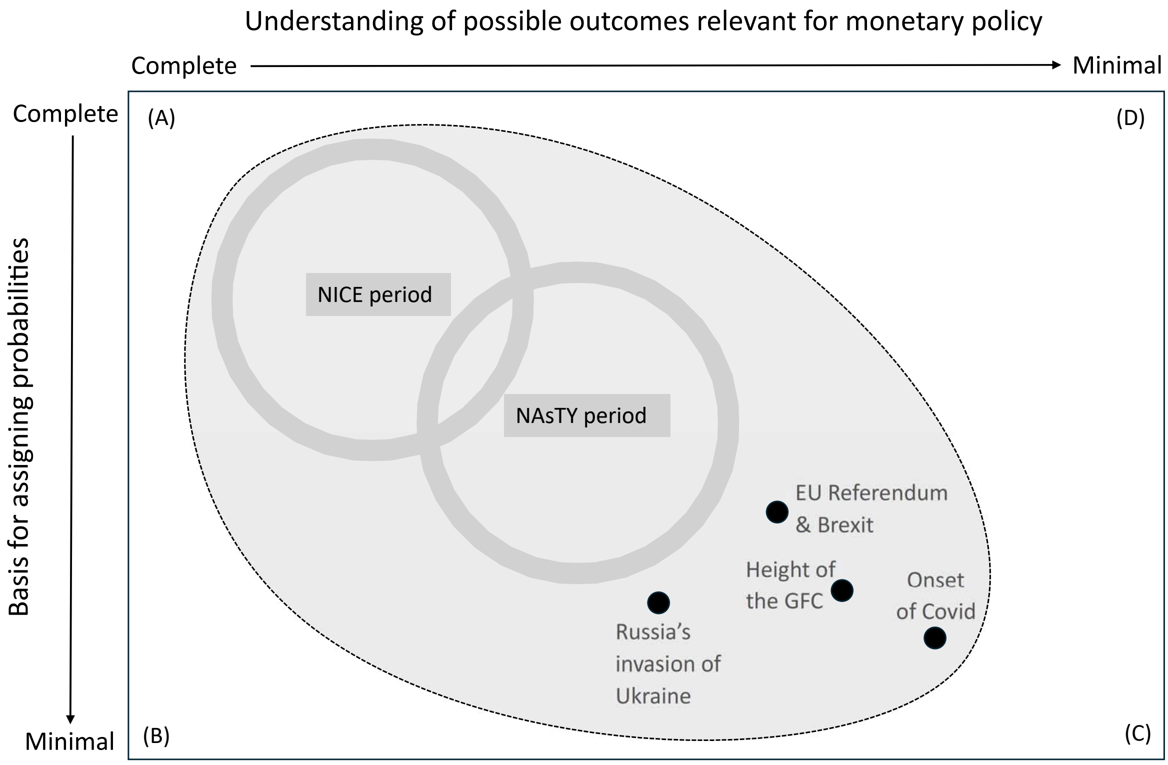

Notes: This uncertainty space is adapted from Spiegelhalter (2024, Chapter 13), who in turn credits Stirling (2010). The axes of the diagram measure the degree of understanding of the range of potential outcomes as relevant for monetary policy decisions (horizontally) and the strength of the basis for attaching probabilities to those outcomes (vertically). The uncertainty space can be used to depict – in an entirely subjective manner – the prevailing level and nature of uncertainty relevant for monetary policy at particular points or during particular periods of recent economic history. As discussed in the main text, monetary policy generally operates in the grey area between the extremes. This area is depicted by the shaded ovoid. See Appendix A for further discussion of this diagram.

knowledge about monetary policy transmission provides at least some basis on which to assess potential policy responses. Therefore, the bottom right corner of very deep uncertainty, denoted by (C), is not relevant in practice either.

Within this grey area of uncertainty, both the quantum and nature of uncertainty varies over time. As in the NICE period, where shocks were generally small and demand-like, there are times when monetary policymakers have more certainty about their assessments of the state of the economy and the outlook. But there are also times, especially after large and novel shocks, where policymakers face considerably more uncertainty.

The challenge for policymakers is how best to proceed in practice, taking uncertainty as a pervasive feature of the monetary policymaking landscape, but also recognising that uncertainty is not constant.

As discussed briefly in Appendix A, there is no unified framework for decision-making under uncertainty. This is mirrored in the monetary policy literature, which has explored three types of formal approach to factoring uncertainty into monetary policy decision-making, each of which embodies a particular form of risk management:[11]

- Extensions of textbook optimal policy (discussed in Appendix B.2), whereby various specific forms of uncertainty can affect optimal policy formulation – i.e. 'certainty equivalence' does not apply. This includes extensions that incorporate relevant non-linearities (e.g. in firms' pricing decisions, Karadi et al., 2024) that affect optimal policy via their interaction with uncertainty about realisations of future shocks. It also includes extensions in which some of or all the model's parameters are treated as random variables (e.g., Söderström, 2002), extensions in which the model is subject to stochastic regime switching via changes in the model parameters (e.g., Blake and Zampolli, 2011), and, under certain circumstances, extensions in which a subset of the model's state variables are observed with measurement error (e.g., Svensson and Woodford, 2003). These types of approach map into corner (A) in Figure 1 because the policymaker knows the true structure of the economy and so is presumed to be able to compute an objectively correct joint probability distribution of outcomes given a state-contingent policy plan.

- Bayesian optimal policy (discussed in Appendix B.3), whereby a policymaker sets policy optimally in a 'hypermodel' constructed by applying prior probabilities to a set of alternative reference models of the economy, which are updated over time according to Bayes' rule. Bayesian optimal policy recognises that the true model is not currently known but, provided that evolving data are informative about the data generating process, also embeds an assumption that the true model is in the set of alternatives under consideration and would be revealed in time (i.e. the hypermodel would converge on the one true model). As such, Bayesian optimal policy could reasonably be described as applying in the region of corner (A) (and then converging into corner (A) asymptotically).

- Robust control (discussed in Appendix B.4), whereby a policymaker is concerned that their reference model of the economy is misspecified, but is unable to attach probabilities to the alternative ways in which that can be true. Under some assumptions about the precise nature of the concern, the policymaker sets policy

[11] Risk management is an umbrella term that describes any situation in which policy decisions deviate from attempting to deliver the best possible outcomes while ignoring uncertainty (i.e. a deviation from certainty equivalence). As the discussion in the bulleted list and in Section 3 makes clear, there are several different ways in which a policymaker could approach risk management in theory, each of which has somewhat different implications for how policy is set. In practical policymaking, the term risk management is synonymous with Alan Greenspan. See, for example, Greenspan (2004) for a discussion of how he viewed risk management and Blinder and Reis (2005) for an analysis of the ways in which risk management appeared to influence US monetary policy during Greenspan's tenure as Chairman of the Board of Governors of the Federal Reserve System.

optimally under the assumption that the model turns out to be the worst (from the perspective of their objectives) in the set under consideration (Hansen and Sargent, 2008).[12] Conceptually speaking, this approach can be associated with corner (B) in Figure 1.

None of these approaches can be applied directly to real-world policymaking. Both the extensions to textbook optimal policy and Bayesian optimal policy are applicable to an area close to or at corner (A) of Figure 1, and robust control is applicable to corner (B). These approaches, therefore, are not directly applicable in the uncertainty grey area.

The gap between these theory-driven approaches and real-world policymaking also reflects that the quantitative prescriptions that arise from their application are, to varying degrees, dependent on the precise model(s) used.[13] Beyond epistemic uncertainty, which implies that models are misspecified in unknown ways, models are also typically subject to known misspecifications.[14] For example, it would be a formidable task to incorporate all of the extensions to the textbook optimal control problem discussed in Appendix B.2 simultaneously. In practice, both of these issues are present and both could underpin Faust (2005)'s view that "all of our models are grossly deficient relative to the ideal, and this cannot be corrected in the medium term."[15]

That formal, quantitative, solutions to simplified versions of the monetary policy problem cannot be applied directly to real-world monetary policy does not imply that they cannot be useful. Indeed, they have provided *qualitative* insights that have informed real-world monetary policymaking.[16] For example, it is well-known to policymakers that there are circumstances in which it may be appropriate to adjust policy cautiously if they are uncertain about the strength of the monetary transmission mechanism (Brainard, 1967), that they may wish to downweight estimates of the output gap if they are more uncertain about those estimates (Orphanides and Williams, 2007b), and that they may wish to adopt a looser policy stance than they otherwise would if the policy rate is at

[12] This approach has recently been extended to incorporate multiple reference models, each of which may be misspecified (Cerreia-Vioglio et al., 2025). This extension has the potential to overcome an important practical drawback of the Hansen and Sargent (2008) framework, highlighted by Sims (2001), that the potential models of the economy being entertained by the policymaker tend not to deviate too far from the reference model, reflecting the precise way in which potential misspecification is modelled and practical constraints on the size of the space of exploration away from the reference model. See also Appendix B.4 for further discussion.

[13] This is even true of the robust control approach, which may seem surprising given its intent. See footnote 12 for a brief discussion.

[14] See Spiegelhalter and Riesch (2011) for a useful discussion of classifying different types of uncertainty around model-based analyses.

[15] Consistent with this and the discussion in this paper, and drawing on concepts from the theoretical computer science literature on complexity, Faust concludes that "monetary policy is hard". This view is closely related to a strand of the social sciences literature that discusses so-called 'wicked' policy problems. The defining characteristic of wicked problems is that they do not have a definitive resolution. See Pill (2023) for a brief discussion of the parallels between monetary policy and wicked policy problems.

[16] This can be interpreted as a form of 'analogical' reasoning. See Section 4 for further discussion.

or near its effective lower bound and unconventional policy is viewed as an imperfect substitute (Evans et al., 2016).

The absence of a formal solution to the practical policy problem implies that subjective judgement is an essential ingredient of real-world policymaking. By its nature, the application of judgement involves policymakers using their discretion to respond to their evolving understanding of the state of the economy and the effects of shocks that arrive While judgement is an issue that will be discussed in more detail in Section 4, two points on the use of discretion and uncertainty are worth highlighting.

First, an inflation targeting framework helps to reduce uncertainty by setting out the objectives of monetary policy as clearly as possible, in an environment of "constrained discretion" (Bernanke and Mishkin, 1997). Such a framework appropriately constrains discretion regarding a critical part of the policy problem – the policymaker's objectives – and helps to reduce associated uncertainty in the private sector about inflation in the medium to long run.[17]

Second, the framework helps to ensure that monetary policy is set in a systematic and, therefore, predictable manner (King, 1997). While quantitative solutions to simplified versions of the policy problem cannot be applied directly to real world policymaking, insights from these approaches can be used to take different perspectives on the policy problem in a way that can help to frame policymaker judgement and promote systematic policymaking.

3 Alternative perspectives on the policy problem

This section explores alternative ways that a theory-based approach to optimal monetary policy can form the basis of different 'perspectives' on the real-world monetary policy problem. The reasoning follows Blinder (1998), who argues that even though formal solutions to simplified versions of the policy problem are not *directly* applicable to real-world policymaking, they may provide high-level insights that can be applied less formally:[18]

> In my view, we must use the [theory-based] framework – with as many complications as we can handle – even in quite an informal way. (Blinder, 1998).

In developing the reasoning for this approach, the discussion starts by outlining a simple textbook approach to the optimal monetary policy problem under the assump-

[17] See Section 3 for further discussion.

[18] The specific framework that Blinder refers to is one he describes as the 'Tinbergen-Theil framework'. His exposition of that approach is very similar (with some simplifications) to the 'textbook' case we consider below and in Appendix B.

tion that the policymaker knows the true model, consistent with corner (A) in Figure 1 and summarised in Figure 2. The subsequent discussion considers how insights from this theory-based approach can be used to inform alternative perspectives on the monetary policy problem, in a manner that relaxes the strict assumptions and supports the application of judgement, including in ways that acknowledge pervasive uncertainty.

Figure 2: Optimal policy in textbook theory

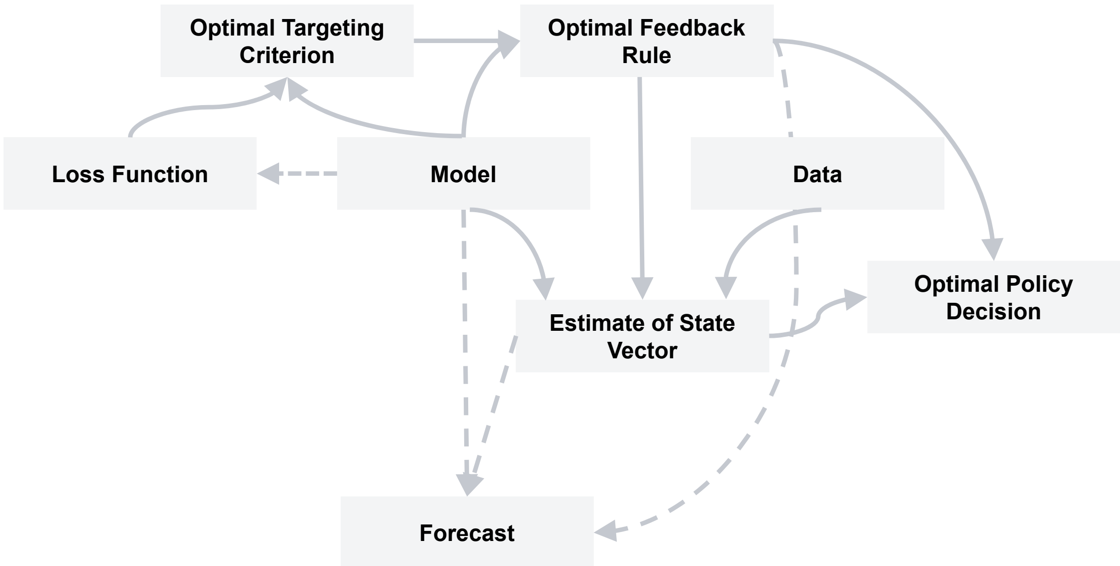

Notes: This figure presents a stylised representation of the relationships between key components of optimal monetary policy formulation according to textbook theory. Solid lines represent connections between necessary components (e.g., computation of the optimal policy decision requires the optimal feedback rule and an estimate of the state vector). Dashed lines denote 'optional' connections (e.g., in some cases, as discussed in footnote 19, the loss function is derived from welfare criteria such as household utility within the model and therefore requires the model). Appendix B.1 presents a formal treatment of the textbook optimal policy problem.

The textbook theory-based approach to optimal monetary policy assumes that the monetary policymaker has: (i) a precisely specified 'objective function' describing their policy goals; and (ii) an accurate model of the economy in which they operate. As discussed in Section 2, an implication of (ii) is that the policymaker can calculate (exactly) the effects of both exogenous shocks and endogenous monetary policy reactions on macroeconomic variables. It is common to specify the policymaker's objective function as a 'loss function', and, without loss of generality, this term will be used in the subsequent discussion.[19]

[19] A quadratic loss function including a weighted average of squared deviations of inflation from target and the output gap is commonly used to represent the monetary policymaker's objectives in the literature on optimal monetary policy. In textbook New Keynesian monetary policy models (see e.g., Woodford, 2003; Galí, 2008), a loss function of this form can be derived as a second-order approximation to the representative agent's utility function. A term involving squared changes in the instrument may also be included, to capture inertia in policy (Woodford, 2003). Such formulations of the loss function are often argued to capture the "flexible inflation targeting" mandates of many central banks (Svensson, 1999, 2000), including the instruction in the MPC's remit to avoid "undesirable volatility in output" (Bean, 1998, 2003; Svensson, 2003a; Carney, 2017). Yellen (2012), and Carney (2017) discuss the consideration

In this stylised case, setting policy optimally amounts to an exercise in calculus: the policy instrument is set to minimise the loss function subject to the constraint imposed by the structure of the economy, as represented by the model. As such, calculating how monetary policy should respond in all circumstances requires just two inputs: the loss function and the model. Combining the solution to the optimisation problem with the data determines the optimal policy decision, as illustrated in Figure 2 and set out more formally in Appendix B. While it is straightforward to compute a forecast by combining the solution to the model and the estimate of the state vector, a forecast is not necessary to compute optimal policy decisions. All the information necessary to set policy in an appropriately forward-looking manner is encapsulated in the model of the economy.

A key observation from this stylised case is that there are very tight links between the (correctly specified) model and optimal policy behaviour. The model determines: the interpretation of incoming data (identifying the underlying shocks that determine them and producing consistent estimates of unobserved state variables, such as potential output); the projection of the current state estimate into the future; and the optimal control of the system. These estimates, optimal policy responses and projections are all jointly determined in a mutually consistent way. This approach can therefore be regarded as 'analytically closed' insofar as the statement of the problem and its solution all exist within the same internally consistent framework; no further external input is required.

Appendix B.2 considers several of Blinder's "additional complications" that have been incorporated into the baseline treatment of optimal monetary policy. These include incorporating parameter uncertainty, non-linearities, imperfect and asymmetric information, structural change, and non-rational expectations. These extensions result from (often long-running) research agendas focused on "formally solving a sequence of progressively more realistic problems", thereby bringing the results closer to the practical monetary policy problem (Faust, 2005).

As noted in Section 2, within the conceptual space of epistemic uncertainty in Figure 1, the textbook theory briefly described above and the extensions summarised in Appendix B.2 are typically applied to environments consistent with the top left corner, (A), where the true model is assumed to be known. Therefore, these theory-based approaches are not *directly* applicable to the grey area in which monetary policy typically operates (indicated by the shaded ovoid in Figure 1).

However, the insights from the theory-based approach can be applied more informally in several ways, providing alternative 'perspectives' on the monetary policy problem that acknowledge uncertainty in different ways. In each case, these insights are informed by

of this type of loss function for practical policy analysis. Alternative forms of policymaker objectives may be appropriate to approximate preferences around skewed risks (Al-Nowaihi and Stracca, 2003) and low-probability extreme events (Svensson, 2003b), for example.

how the theory-based approach "teaches us to think about" the monetary policy problem, in an analogous way to Blinder's example of the lessons from dynamic programming:

> I do not believe it is important for central bankers to acquire any deep understanding of Bellman's principle, still less of the computational techniques used to implement it. *What really matters for sound decision making is the way dynamic programming teaches us to think about intertemporal optimization problems* – and the discipline it imposes. It is essential, in my view, for central bankers to realize that, in a dynamic economy with long lags in monetary policy, today's monetary policy decision must be thought of as the first step along a path. Blinder (1998, emphasis added).

In a similar way, alternative perspectives on the monetary policy problem are based on insights into the general properties of optimal policy behaviour from the theory-based approach. Each perspective therefore follows Blinder's advice and applies the insights "in quite an informal way". This allows greater flexibility compared with the tight quantitative solutions to simplified versions of the monetary policy problem studied in academic research and the simple textbook variant described above.

The discussion of the perspectives below takes as given that there is an over-arching inflation targeting framework to provide an environment of "constrained discretion" for policy formulation. As discussed in Section 2, one benefit of such a framework is that it substantially reduces uncertainty about the objectives of monetary policy.[20] Within this framework, the alternative perspectives on the policy problem draw on different insights from theory in a way that helps to ensure that monetary policy is systematic.[21]

[20] As in many other jurisdictions, the Monetary Policy Committee of the Bank of England's objectives make clear that the inflation target is the ultimate, long-run objective of policy. The remit also recognises, however, that there may be trade-offs between output and inflation in the short to medium run that the MPC should actively manage in pursuing its inflation targeting objective. Since the nature of any trade-off will depend on the circumstances, the remit is not prescriptive about exactly how such trade-offs should be managed. As noted by Carney (2017), "In other words, the monetary policy problem cannot be fully contracted ex-ante." In this way, the formulation of policy objectives could also be regarded as an application of the implications of theory-based insights "in quite an informal way": that low and stable inflation is beneficial is a robust result; however, while the existence of costly short-run tradeoffs between inflation and economic activity is also well established, there is less certainty over the underlying causal relationships and their stability over time. This is consistent with the observation that shorter-run stabilisation objectives of many inflation targeting central banks are typically defined in qualitative rather than quantitative terms.

[21] As noted in Section 2, systematic monetary policy helps to ensure both that inflation expectations are anchored in the medium-to-long term and that policy responses to developments in the economy are as predictable as possible (King, 1997). One way in which the inflation targeting framework itself helps to achieve that is by providing mechanisms for accountability and transparency. For example, in the case of the Bank of England's MPC, communication via the Monetary Policy Report, Monetary Policy Statement, minutes, and speeches are important for setting out individual and collective policy narratives. See, for example, Tucker (2004, 2006) for a discussion of evidence that the advent of Inflation Targeting and the creation of the MPC reduced uncertainty around UK monetary policy.

In this context, the remainder of this section considers three 'analytically open' perspectives on the monetary policy problem:

1. A *forecast-based perspective* uses forecasts for key macroeconomic variables as the basis for formulating monetary policy. It can be understood via an equivalence result that reformulates the optimal policy problem in terms of finding the best achievable forecast for the policymaker's goal variables. It embraces a forward-looking and quantitative approach to policymaking, but relaxes the tight links between the model of the economy, the forecast, and the policy response in a way that supports the application of judgement.

2. A *news-based perspective* involves updating policy over time, as understanding of the drivers of macroeconomic outcomes evolves, and as new shocks or events are identified. It builds on the observation that adjustments to optimal policy decisions over time can be represented as the sum of (a) revisions to optimal responses to past shocks using updated estimates of their effects and (b) the optimal responses to newly arrived shocks.

3. A *rules-based perspective* uses simple instrument rules as a guide to the appropriate setting of policy. It is based on the observation that 'optimal' policy is model dependent and, therefore, attempting to follow such an approach in practice would lead to sub-optimal outcomes to varying degrees, depending on the extent to which the policymaker's model is misspecified. It draws on insights from the robust policy rules literature, applied via an understanding of the current state of the economy, to inform the appropriate policy strategy given prevailing circumstances.

Figure 3 presents a stylised representation of zones within the uncertainty space from Figure 1 in which each of these perspectives may be *particularly* relevant.

Several points are worthy of note. First, the relative locations and sizes of the zones are based on subjective judgements (discussed further below) and are, therefore, intended to illustrate the broader discussion somewhat informally. Second, and relatedly, since the zones are intended to capture situations in which different perspectives may be *particularly* useful, Figure 3 is not intended to suggest that these perspectives have no value outside of those zones.[22] Third, that the zones overlap indicates that there are conditions under which multiple perspectives may offer useful practical policy insights. Finally, a

[22]For example, the discussion below notes that both forecast-based and news-based perspectives could be applied in conditions that satisfy the assumptions of corner (A) in ways that replicate the theory-based approach. By this reasoning, the relevant zones in Figure 3 could be extended to encompass corner (A). However, under those conditions, it is not clear that either perspective would be particularly useful as the methods to replicate optimal policy would be inefficient relative to direct application of the textbook theory-based approach. Consistent with that, the authors' judgement is that neither the forecast-based, nor news-based perspectives would be particularly useful in corner (A).

Figure 3: Uncertainty and alternative perspectives on the practical policy problem

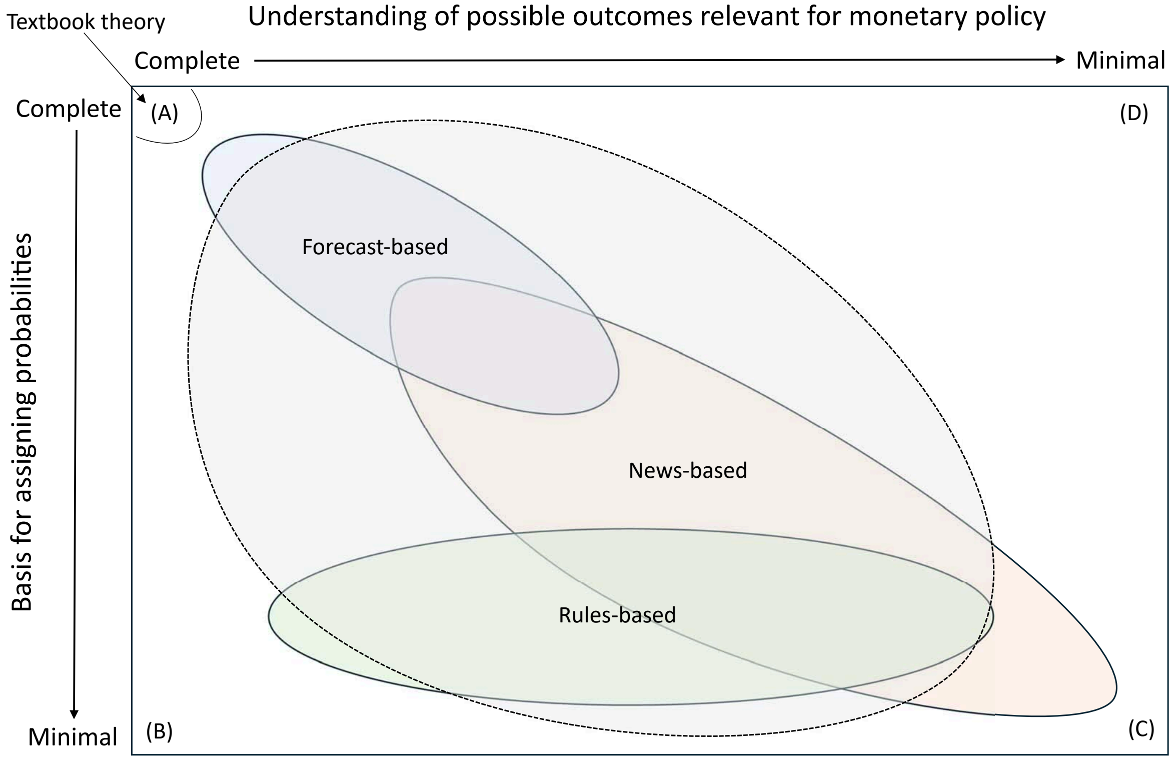

Notes: The uncertainty space is the same as that shown in Figure 1 and is adapted from Spiegelhalter (2024, Chapter 13). The axes of the diagram measure the degree of understanding of the range of potential outcomes as relevant for monetary policy decisions (horizontally) and the strength of the basis for attaching probabilities to those outcomes (vertically). The subset of the space in which monetary policy typically operates is subjectively depicted by the shaded ovoid. See Section 2 for further discussion. The diagram also subjectively depicts the regions of the uncertainty space that may be more applicable to the three alternative perspectives discussed in the main text and the region, very close to corner (A), in which the theory-based approaches discussed in the main text and in Appendix B are directly applicable. See Appendix A for further discussion of this diagram.

judgement that each zone covers a relatively large area suggests that each perspective can be applied in different ways, while remaining consistent with the high-level insights from the theory on which that perspective is based.

3.1 Forecast-based perspective

Forecasts have long played an important role in monetary policymaking. Consistent with the textbook theory, this reflects the evidence that there are significant lags in the monetary transmission mechanism and so monetary policymaking is necessarily forward-looking. A forecast-based perspective retains the same forward-looking and quantitative approach to the policy problem as the textbook theory. But it also acknowledges uncertainty and model misspecification with the implications that: (a) there is an important role for multiple models and judgement in forecasting; (b) textbook optimal policy may deliver poor performance in practice.

The link from the textbook theory to a forecast-based perspective can be more formally understood via an equivalence result that, under certain conditions, it is possible to replicate the results of the theory-based (optimal control) approach by splitting the problem into two parts (Barnichon and Mesters, 2023). Specifically, a baseline forecast (computed under a non-optimal policy path) can be produced separately and then combined with a model of the monetary transmission mechanism to calculate the adjustments to the policy path that minimise the policymaker's loss function. Applying these adjustments to the baseline policy path produces a forecast for the optimal path of the policy instrument(s), together with consistent forecasts of the 'goal variables' in the loss function. The solution to this problem coincides with the model-based optimal control solution provided that the transmission mechanism used to calculate the loss-minimising adjustments to policy is identical to that in the model used for the optimal control solution.[23]

The potential to use this equivalent formulation of the policy problem to inform policymaking has been recognised for some time. Svensson (2003a) presents an informal version of the logic underpinning the Barnichon and Mesters (2023) result, noting that, in principle, a policymaker who considers forecasts constructed under different assumptions about the path for monetary policy, and chooses the one that "looks best" from the perspective of achieving its policy objectives, will uncover the optimal path, thereby solving the optimal policy problem.[24]

Figure 4 depicts how the equivalence result can be applied to permit the combination of multiple models and judgement, informed by a broader range of analysis, to provide a forecast-based perspective on optimal policy. In this example, the forecast is an 'input' to a process that generates an optimal policy path that informs the policy decision.[25] The approach embodies more flexibility than textbook optimal control by allowing different models to be used for the production of the baseline forecast and optimal policy path, and by allowing for the application of judgement. Indeed, models can be separated out from the forecasting process (as indicated by dashed lines) so that, in principle, a model can be used to compute an optimal policy projection using a baseline forecast that has not been produced by, or necessarily even been informed by, that model (e.g., an entirely judgemental forecast, or one generated by a different set of models).

Central banks have developed tools to implement practical variants of the Barnichon

[23]Barnichon and Mesters (2023) demonstrates these results for linear-quadratic models within the set considered in Appendix B.1.

[24]See Svensson (2006) for another early contribution.

[25]See Section 4.1 for a discussion of the role of analytical inputs in informing different perspectives on the practical policy problem.

Figure 4: Example of tools and analysis to support a forecast-based perspective

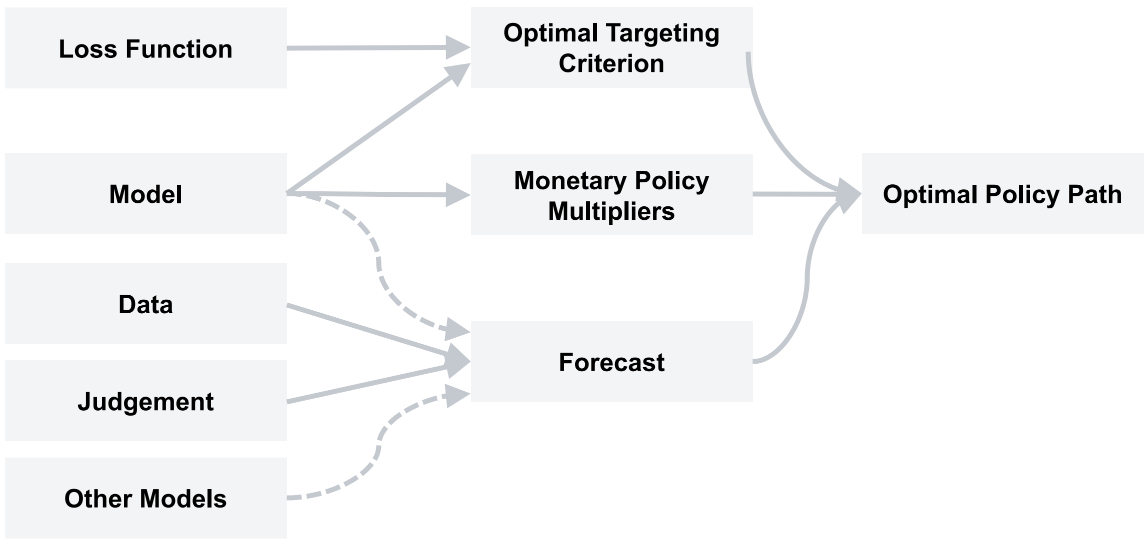

Notes: The presents a stylised diagrammatic example of how a variant of the 'sufficient statistics' approach discussed in the text could be implemented. Solid lines show required components of the approach (e.g., computation of the optimal policy path requires the optimal targeting criterion, a baseline forecast and a set of monetary policy 'multipliers'). Dashed lines denote 'optional' components. For example, the model from which the monetary policy multipliers are derived need not be used to produce the forecast. In that case, another model (or set of models) could be used for that purpose.

and Mesters (2023) approach (the components at the top of the diagram).[26] For example, Alati et al. (2025) explains how this 'optimal policy projection' approach is deployed at the Bank of England and discusses applications to several past Monetary Policy Report forecasts.[27]

One reason why it is typically desirable to apply judgement to a model-based projection is that the model may be known to be misspecified. Existing research has explored tools both to identify misspecification in a baseline model and potentially 'correct' for it using information from other models.[28] However, such approaches typically focus on relatively 'local' misspecifications of a baseline model and in many cases policymakers

[26]See, for example, Svensson and Tetlow (2005), Harrison and Waldron (2021), de Groot et al. (2021), Hebden and Winkler (2024).

[27]As shown in Figure 4, an important output of such methods is a projection of the optimal path(s) for the policy instrument(s). The importance placed on a 'policy path' view of the policy problem in the quote from Blinder (1998) on page 17 rests on the high-level implications of a dynamic programming perspective and, indeed, optimal policy paths can be computed using dynamic programming methods as shown by Harrison and Waldron (2021, Section 4).

[28] For example, the approaches to identifying misspecification explored by Faust (2005, 2012), Faust and Gupta (2012), and Gupta (2016) use a baseline or reference model as a lens through which the reliability of its predictions (or inferences based on the model) can be assessed (using tools developed for prior predictive analysis that are widely used in Bayesian statistics, as set out by Gelman et al. (1995)). Alternative approaches take a more empirical perspective (for example Del Negro and Schorfheide, 2006, 2009; Filippeli et al., 2020). A practical approach to attempting to correct for misspecification using the information from other models is presented in Burgess et al. (2013, Section 7.1), using insights from Alvarez-Lois et al. (2008), Monti (2010) and Caldara et al. (2014) among others.

may be more confident in applying judgement directly to forecasts for endogenous variables using information from models and other sources that are somewhat distant from a baseline model (as discussed in Burgess et al., 2013, Section 6.2). This situation is likely to apply to forecasting in the aftermath of large and novel shocks that are not well-captured by the baseline model (or any other single model), such as the GFC and the Covid pandemic.

A narrow application of a forecast-based perspective via optimal policy projections is less useful when considering policy responses to some forms of uncertainty. The equivalence results of Barnichon and Mesters (2023) only allow for (additive) shock uncertainty. As a result, the effects of relevant uncertainty are assumed to be fully captured by the mean of the baseline forecast density (i.e., certainty equivalence applies).[29]

In contrast, in formal analyses of risk management, separation of the forecast density from policy analysis is no longer possible. For example, as briefly discussed in Section 2 and Appendix B.2.2, Evans et al. (2016) illustrate why risk is a relevant consideration for monetary policy when there is an effective lower bound constraint on the policy rate. Analogous considerations would also apply in the presence of other types of non-linearity, such as those that can appear in price and wage inflation determination at higher rates of inflation (see, for example, Karadi et al., 2024).

Nonetheless, informal application of risk management reasoning – including using quantitative solutions to simplified versions of the problem as guides – could motivate a "tilt" in monetary policy away from that which would be implied by a baseline forecast (Bernanke, 2024, 2025). Similarly, a tilt in policy for risk management reasons can be aided by the production and publication of density forecasts and scenarios (i.e., plausible alternative conditional projections) that explore the potential effects of risk factors on the economic outlook (see, for example, Adrian et al., 2019, 2025).

While quantitative approaches can be useful guides for policy, tilts in policy for risk management reasons are likely to entail significant judgement:

> In pursuing a risk-management approach to policy, we must confront the fact that only a limited number of risks can be quantified with any confidence. [...] Policymakers often have to act, or choose not to act, even though we may not fully understand the full range of possible outcomes, let alone each possible outcome's likelihood. As a result, risk management often involves significant judgment as we evaluate the risks of different events and the probability that our actions will alter those risks. (Greenspan, 2004).

[29] Svensson (2010) discusses how a variety of complementary tools can be used within a forecast-based perspective, including those that incorporate some types of uncertainty (see also the discussion in Appendix B.2 and B.3).

More generally, a forecast-based perspective need not apply the same degree of formality as an optimal policy projection. For example, Qvigstad (2005) documents a set of criteria used by Norges Bank to "find an instrument-rate path such that projections of inflation and output gap 'look good'" in the spirit of Svensson (2003a). While these criteria include direct analogues to the conditions that characterise the trade-off criterion of a theory-based optimal control solution, others are based on considerations that cannot be captured within the textbook theory-based framework.[30] Carney (2017) offers a similar perspective in discussing trade-off management by the Bank of England's MPC in the period between the GFC and the EU Referendum.

Somewhat 'coarser' versions of a forecast-based perspective include those in which the policy path is not computed optimally, perhaps instead using simple policy rules (e.g. that capture relevant elements of the policymaker's reaction function), alternative paths that policymakers deem plausible (Bernanke, 2024), or a path based on market expectations of the policy rate.[31] In these cases, the monetary policy signal that could be taken from such forecasts would depend on the extent to which the resulting forecast delivers policy objectives in the manner that policymakers would prefer.

Since a forecast-based perspective fundamentally requires quantification of the outcomes and probabilities relevant for monetary policymaking, alongside some understanding of the transmission mechanism of monetary policy (extending to complete knowledge in the case of optimal policy projections), it is likely to be more useful in states of the world where policymakers are more confident in the models being used and their ability to apply quantitative judgements (i.e. closer to the top left quadrant of Figure 3).[32] Conversely, a forecast-based perspective becomes less useful in areas of Figure 3 that are more distant from the top left quadrant. In those regions of the uncertainty space, constructing reliable forecasts is more challenging, confidence in estimates of the transmission mechanism may be lower, and policymakers may wish to take other perspectives on the policy problem.

[30] For example, criteria relating to robustness and interest rate smoothing are based on reasoning that is not internalised within the underlying theory-based perspective. Indeed, the overall approach described in Qvigstad (2005) bears many similarities to Blinder's argument for the application of a range of theory-based results in "quite an informal way", consistent with the forecast-based perspective presented here.

[31] See Alati et al. (2025) for an explanation of how simple rules can be implemented and discussion of an application to past Monetary Policy Report forecasts, including comparisons with optimal policy projections.

[32] While a forecast-based perspective does not necessarily require the construction of a complete density forecast, it does require production of moments of central tendency and an assumption about whether or not that central tendency is a sufficient statistic for policy. As discussed above, if it is not sufficient, then that may motivate a 'tilt' in policy away from that implied by a central projection.

3.2 News-based perspective

A news-based perspective is built around the notion that policy should respond systematically to news and that learning about the evolution of the economy is a critical part of policymaking. Unlike a forecast-based perspective, it does not require the construction of a forecast and so may be particularly useful in areas of the uncertainty space in which forecasting is more challenging.

A news-based perspective can be related to the textbook theory by considering the implications of that theory for the revisions to the optimal policy setting (i.e., the optimal policy instrument choice relative to the choice expected at the time of the previous decision). From the viewpoint of theory, the optimal policy revision can be decomposed into the optimal responses to newly arrived shocks and the optimal responses to revisions of the estimated effects of past shocks.[33] Indeed, under assumptions consistent with corner (A) of Figure 3, a procedure that summed revisions to optimal policy on a shock-by-shock basis – across both newly arrived shocks and revisions to the estimated effects of past shocks – would deliver the same results as direct application of the textbook optimal policy theory.[34] It is this equivalence that provides the inspiration for the news-based perspective.[35]

Applying insights of the theory-based approach regarding optimal policy revisions "in quite an informal way" makes it possible to relax the underlying assumptions and apply the news-based perspective in a wide range of conditions. Indeed, the zone in which a news-based perspective may be particularly useful extends across a large swathe of the grey area relevant to monetary policymaking, as shown in Figure 3. Moving further from corner (A) implies that applications of a news-based perspective become increasingly informal and the links with the theory-based approach become weaker.

In regions that are relatively close to corner (A), different models could be used to identify different (current and past) shocks and the optimal responses to them. Aggregating these results provides a guide to optimal adjustment of the policy instrument(s).[36] Such an approach naturally raises the question of whether the results are internally consistent,

[33] See equation (B.17) in Appendix B and the related discussion.

[34] This echoes the equivalence result for the forecast-based perspective described in Section 3.1.

[35] Compared with the 'sufficient statistics' approach discussed in Section 3.1, there has been less academic interest in the news-based perspective. One possible reason is that much of the power of a sufficient statistics approach is that it replicates (under appropriate conditions) direct application of the textbook theory while requiring only a minimal amount of information. By contrast, replication of the textbook theory using shock-based summation requires all of the same information (and is obviously less efficient). Another possible reason is that, as discussed below, application of a news-based perspective in regions of the uncertainty space that are distant from corner (A) typically requires an approach that is tailored to the prevailing conditions and is, therefore, less amenable to analysis of the 'general' case.

[36] The analogue to the decomposition in (B.17) in Appendix B would be computed using different coefficient matrices for different shocks.

though this issue may also be encountered from a forecast-based perspective (e.g., when using multiple models to produce a forecast). As in the forecast-based perspective, the use of multiple models represents a pragmatic approach to misspecification for situations in which the 'true' (internally consistent) model is unknown, or unknowable.[37]

Similar pragmatism may also be applied to the number of shocks that are considered when revising the optimal policy choice. An informal application of the news-based perspective would suggest that a good approximation to the optimal policy revision can be obtained by focusing on the policy implications of the most important shocks (in terms of their likely effects on the optimal instrument setting). Misspecification of the models used for policy analysis (and forecasting) implies that there is an imperfect mapping between the shocks in the models and the events that monetary policy responds to in practice. For this reason, the appropriate policy response to real-world shocks and events inevitably requires judgement which, depending on the shock, can be aided by model-based analysis to varying degrees.[38]

The application of a news-based perspective in the presence of "deep uncertainty", toward corner (C) in Figure 3, would have somewhat weaker links to quantified estimates of optimal policy responses to well-identified shocks in a structural macroeconomic model. However, the way that the news-based perspective "teaches us to think about" the policy problem still has value and provides discipline in policy formulation: "Acknowledging radical uncertainty does not mean that anything goes" (Kay and King, 2020).[39]

In their discussion of decision making in environments of radical uncertainty, Kay and King (2020, Chapter 22) suggest an approach with two key elements: development of a guiding narrative account of "what is going on here?"; and mechanisms to "open that explanation to challenge and be ready to change the guiding narrative when new information emerges."[40] Such an approach echoes core insights from analysis of decision-making

[37] Appendix B.6 argues that using multiple models can be motivated as an appropriate response to inevitable model misspecification in two ways. First, the inherent complexities associated with including relevant features to approximate the real-world policy problem implies that it is generally only feasible to consider them in isolation. As a result, different models should be designed to capture different aspects of the policy problem. Second, using multiple models to analyse a particular part of the policy problem facilitates judgements about the robustness of the results.

[38] When considering model-based analysis, this requires policymakers "to probe more deeply behind the 'labels' that the [. . .] model places on the important shocks driving the data in order to uncover more fundamental stories" (Burgess et al., 2013).

[39] The concept of "radical uncertainty" in Kay and King (2020) is essentially consistent with the "deep uncertainty" concept discussed in connection with Figures 1 and 3 (Spiegelhalter, 2024, Chapter 13).

[40] Kay and King (2020) are concerned with environments close to corner (C) of Figure 3 and they consider approaches that are tailored to such environments rather than taking inspiration from a theory-based benchmark. Indeed, they are sceptical of the reliance on 'shocks' as explanations of economic fluctuations, for example. As the discussion of the news-based perspective demonstrates, this is not inconsistent with the use of high-level insights from theory as a way of thinking about the policy problem. However, the links to the theory-based benchmark are necessarily weaker in environments of "deep" or "radical" uncertainty.

in environments of deep uncertainty, including conviction narrative theory (Tuckett and Nikolic, 2017; Johnson et al., 2023) and case-based decision theory (Gilboa and Schmeidler, 1995, 2001).[41]

In the present context, the monetary policy equivalent of "what is going on here?" could be framed as "what are the policy implications of the shocks or events that are currently affecting the economy?" That question could be informed by consideration of the following general questions:

1. How does what we have learned about the transmission of past shocks or events affect the guiding policy narrative?
   (a) How does the evidence from the most recent information influence our assessment?
   (b) How does the evidence from new analysis influence our assessment?
2. What are the policy implications of any newly arrived shocks or events given the backdrop against which they have arrived?
3. How might we be going wrong?

Questions 1 and 2 are informal analogues of the shock-based decomposition implied by the theory-based approach. Assessments of these questions can be used to inform incremental updates to the guiding policy narrative, which summarises monetary policymakers' interpretation of the effects of shocks on the economy and the appropriate policy response.[42]

Considering questions 1 and 2 provides some degree of challenge to the guiding policy narrative because they embed an assessment of the extent to which newly accumulated

[41] Conviction narrative theory describes processes that develop a narrative that best explains observed data "to imagine possible futures given potential choices" (Johnson et al., 2023). Case-based decision theory provides a framework in which 'similarity judgements' are central to guiding decisions. Roughly speaking, payoffs to actions taken in past (and potentially hypothetical) decision problems are weighted by an assessment of their similarity to the current decision problem and the highest scoring action is chosen. In this way actions chosen are 'best' only with respect to the set of comparators, which may expand as new hypothetical decision problems are considered. See Appendix D for further discussion.

[42] For example, in the spring of 2016, there was substantial uncertainty about the outcome of the EU Referendum due to take place in June of that year. This, in turn, created uncertainty about the signals from economic and financial indicators and the extent to which they reflected households and firms modifying their behaviour in advance of the referendum or other underlying trends. The transcripts of past MPC meetings, released in accordance with the changes following the Warsh (2014) review, show how the MPC approached this at the time. For example, transcripts from the MPC's May 2016 meeting indicate that disentangling EU Referendum effects from other trends and factoring this into their overall policy narrative was a focus for many MPC members (Bank of England, 2016d). It is possible to view these discussions as MPC members' assessing questions 1 and 2, potentially consistent with the news-based perspective. However, many of those MPC members also made reference to the forecast, in line with a forecast-based perspective. This is consistent with the observation that the alternative perspectives discussed in this section are not mutually exclusive.

information and analysis are consistent with the prevailing narrative. However, question 3 provides an explicit way to explore the possibility of more fundamental misspecifications of the assumptions underpinning the guiding narrative.[43,44] This is particularly valuable in "deep uncertainty" regions of the uncertainty space.[45]

In these regions, application of a news-based perspective may be more reliant on qualitative analysis than quantified and/or model-based predictions. Nevertheless, model-based analysis can play important roles in guiding more qualitative assessments. For example: "simple analytical models are immensely valuable as a way of generating insights which can be carried across to the policy process" (King, 2024). This idea is consistent with Blinder's view that the value of some types of analysis is in "how it teaches us to think about" the practical policy problem and, as noted above, with insights from conviction narrative theory and case-based decision theory.[46]

3.3 Rules-based perspective

As Figure 2 makes clear, a correctly specified model is central to the textbook theory discussed above. Indeed, an optimal control approach embeds detailed knowledge of the model structure within the optimal policy response function.[47] Two important corollaries of this observation underpin a rules-based perspective. First, uncertainty about the true model of the economy necessarily limits policymakers' ability to achieve optimal outcomes. Second, and relatedly, attempting to implement an optimal policy derived from a misspecified model could deliver poor outcomes (possibly materially so).

In extreme cases, in the region of "deep uncertainty" close to corner (C) in Figure 3, ignorance about the economy may be sufficiently great that policymakers eschew any attempts at cyclical stabilisation and instead set policy with the aim of achieving satis-

[43]An obvious risk associated with updating a reference narrative incrementally is that errors in the reference narrative (i.e., policymakers' understanding of what is driving economic outcomes) may persist if and when they emerge. These risks apply in principle to all incremental approaches as highlighted by Bernanke (2024) with reference to the incremental construction of MPC forecasts. Incremental approaches, however they are implemented, place a premium, therefore, on challenge and the question of "how might we be going wrong?".

[44]MPC members' discussions in May 2016 also reveal consideration of question 3, with a particular focus on their confidence in the prevailing narrative for the recent slowdown in GDP growth (Bank of England, 2016b).

[45]In contrast, in corner (A), use of a well-specified model to compute the (mathematically formalised) answers to questions 1 and 2 is sufficient.

[46]As discussed further in Appendix D.2, Gilboa et al. (2014) argue that one important role of models is to provide results that can be applied analogically to the practical problem at hand. In contrast to models used for forecasting or policy analysis, models used to support case-based reasoning need not be quantitatively 'plausible' (or 'realistic' in other dimensions). Tuckett et al. (2020) discuss the use of models as " "imaginaries" [...] with heuristic use for playing with *possible futures*" (original emphasis).

[47]The coefficients of the optimal feedback rule (B.7) derived in Appendix B are functions of the structural coefficients of the model.

factory outcomes only on average over long periods of time. Such reasoning underpinned Milton Friedman's advocacy of a 'k-percent rule' for the expansion of the base money supply, which was based on a minimum of economic theory (the long-run applicability of the quantity equation) and supporting empirical evidence:

> Steady monetary growth would provide a monetary climate favourable to the operation of those basic forces [...] that are the true springs of economic growth. That is the most that we can ask of monetary policy at our present state of knowledge. (Friedman, 1968).

While more recent investigations of simple policy rules have been motivated in a similar way (seeking policy rules that deliver 'good' performance in the face of uncertainty), they typically apply more insights from the theory-based approach, and, so, presume more about policymakers' understanding of the economy. For example, they often presume that the policymaker is aiming to achieve some degree of cyclical stabilisation, albeit not in a fully optimal manner.

Two important ingredients for a rules-based perspective are a set of models of the economy and a loss function to evaluate the performance of alternative policy rules.[48] These policy rules are typically 'simple', in that they relate the policy instrument to a small number of macroeconomic variables (for example, inflation and the output gap). By considering a set of models, this perspective acknowledges both that the performance of a particular policy rule will depend on the structure of the economy in which it is employed and that there is *a priori* uncertainty about the model(s) that best approximate the behaviour of the economy. Nonetheless, basing the analysis on a set of dynamic models that are considered plausible approximations to the economy moves beyond Friedman's minimalist approach.[49] Using a loss function to assess the performance of alternative policy rules echoes its role in the theory-based approach as a key metric to assess policy performance.[50]

An influential example of this approach is the work of Levin and Williams (2003), who search for simple monetary policy rules that perform well in three alternative models of

[48] The discussion in the text focuses on rule-based perspectives that explicitly incorporate concerns regarding model uncertainty. A large literature explores the performance of alternative rules within a single model using a loss function to compare their performance, often considering models tailored to study particular aspects of economic behaviour or the structure of the economy. For example, Ball (1999) and Batini et al. (2003) explore the properties of simple rules in an open economy context. Taylor and Williams (2010) provide a comprehensive review of the origins and development of the simple policy rules research agenda.

[49] In the context of Figure 3, this implies that the rules-based perspective can be applied in zones of the uncertainty space to the north and west of corner (C), as shown in that diagram.

[50] Using a loss function also allows direct comparison of performance when policy is set using simple rules and the optimal rule, with the latter providing a natural benchmark.

the US economy.[51] The authors use stochastic simulations of the models to assess the performance of alternative rules using a loss function that approximates the Federal Reserve's dual mandate and use a measure of 'fault tolerance' to assess robustness across rules and models. This involves studying the effects of changes in the values of a rule's parameters on the losses from each of the models.[52] The authors also explore alternative approaches to setting parameter values for the rules, again guided by high-level insights from the theory-based approach.[53]

Many other studies follow a similar approach.[54] In some cases, the structures of the candidate models within the set considered are chosen to explore specific monetary policy challenges. For example, Orphanides and Williams (2007b) consider a set of models in which the monetary policymaker is uncertain about the natural rates of interest and unemployment and in which both policymakers and private agents are learning about the structure of the economy. They seek policy rules that are robust to different types and combinations of uncertainty (captured by different combinations of modelling assumptions) and find that simple rules for the change in the policy rate are particularly robust to uncertainty about natural rates.

These examples illustrate an important aspect of how the rules-based perspective can inform practical monetary policymaking. While models play an important role in the *design* of simple policy rules (in particular, in informing the values of their parameters), they are not required for their *use*. Since simple policy rules include a small number of relevant macroeconomic variables, they are straightforward to evaluate: requiring only data for (and/or estimates of) the variables that appear in the rules.[55] As such, a rule-based perspective may be particularly useful in conditions in which forecasting is difficult, particularly those in the lower region of Figure 3.

Another important aspect of a rule-based perspective is the manner in which the rules are intended to inform monetary policymaking. As noted in Taylor and Williams (2010), insights from simple rules "would serve as a rough benchmark for making decisions, not a mechanical formula."[56] Such an interpretation is consistent with the notion of monetary

[51] This analysis builds on earlier studies, such as Bryant et al. (1993) and Weale et al. (1989).

[52] If the loss function is relatively insensitive to the values of the parameters in the rule, the model economy is relatively fault tolerant. See also Williams (2025) for a discussion.

[53] Choosing parameters for the simple rules to minimise expected losses echoes probability weighting across sub-models within the theory-based Bayesian hyper-model discussed in Appendix B.3. The authors also use a min-max criterion to select parameters of the rules, inspired by the robust control approach discussed in Appendix B.4.

[54] See, for example, Brock et al. (2003, 2007), Levin et al. (1999, 2003), Orphanides and Williams (2002, 2006, 2007a, 2008), Cogley et al. (2011) Taylor and Wieland (2012), and Tetlow (2018).

[55] These considerations may interact with the types of rules considered to be potentially useful for policy, as evidenced in the Orphanides and Williams (2007b) investigation of rules that do not require estimates of latent state variables such as the natural rate of interest.

[56] Similarly, Levin (2014) notes that "it would be inadvisable for policymakers to mechanically follow the prescriptions of a rule whose specification has been permanently fixed" for two reasons. First, events

policy strategy "being *rule-based but not rule-bound*: governed by clear principles that ensure decisions are taken to support the price stability objective, but not mechanically determined by some simple automatic algorithm that fails to recognise changes in the economic environment sufficiently" (Pill, 2024, original emphasis).[57]

Indeed, one dimension in which a rules-based approach may be adjusted to changing circumstances is in the subset of rules that policymakers consider to be more useful. For example, following a large supply shock – such as that which occurred when Russia invaded Ukraine as the global economy was still adjusting to the macroeconomic effects of the Covid pandemic – rules that are robust to uncertainty about the response of inflation expectations and/or intrinsic inflation persistence may be particularly informative. Similarly, following suspected structural change, rules that are robust to the potential consequences of such changes, including to the natural rates of interest and unemployment, may be more informative. As with any judgement about the relevance of each of the alternative perspectives for the monetary policy problem, given the prevailing circumstances, an assessment about which rule(s) may be more informative is a strategic judgement that policymakers must make within a framework of constrained discretion. This issue is discussed further in Section 4.

3.4 Insights and implications

Several insights emerge from the preceding discussion.

The first is that the perspectives are not mutually exclusive and can be viewed as complementary. In some regions of the uncertainty space, away from corner (A) in Figure 3, two or more perspectives may provide valuable insights for monetary policymaking. For example, a news-based and forecast-based perspective may be used alongside each other to inform an overall formulation of policy and accompanying narrative that is forward-looking, but is not completely reliant on a quantitative forecast. Use of more than one perspective is most obvious when the shaded zones in Figure 3 overlap. However, even outside these overlapping areas, insights from all perspectives may be useful, albeit with different emphases. For example, a rules-based perspective may be particular useful as a 'cross check' on either, or both, of the forecast-based and news-based perspectives.[58]

and economic conditions that are poorly captured in any of the models used to formulate the rule may warrant a policy response that deviates from the rule. Second, the set of models used to formulate simple rules is likely to evolve over time in light of structural change in the economy.

[57] Again, this marks a departure from the reasoning in Friedman (1968) which embodies the view that ignorance of how the economy works is sufficiently severe that no additional information could be useful in reliably guiding deviations from the rule.

[58] The use of rules as a cross-check on a forecast-based perspective is discussed in Qvigstad (2005). The intended status of such cross-checks is clear: "These cross-checks will not help the MPC in the difficult deliberations about whether to move the interest rate by 25 basis points, or whether to move

The potential use of a rule-based perspective as a complement to forecast and news-based perspectives can be contextualised in the long-standing 'rules versus discretion' debate. Both a forecast-based and news-based perspective embody more flexibility in the ways that they can be applied via the application of policymaker judgement. This has the advantage that they can be informed by continuous learning about how the economy is evolving. Even within a framework of constrained discretion, a potential concern, however, is that these perspectives afford policymakers too much discretion. Application of a rules-based perspective alongside forecast-based and news-based perspectives may, therefore, provide one way of helping to ensure that policymaking is appropriately consistent and systematic.[59]

Second, as the nature of the economic environment changes, the nature of uncertainty moves around within the grey area in Figure 3, so that different perspectives on the monetary policy problem may become more or less pertinent for practical policymaking. For example, in times of heightened uncertainty, when forecasting may be more difficult, news-based and rules-based perspectives may be particularly useful ways of thinking about the monetary policy problem. In times of lower uncertainty, however, a forecast-based perspective may be an appropriate focal point for policymaking, with rules-based and news-based perspectives playing a supporting role.

A third insight is that assessments of the relevance and usefulness of different perspectives are subjective and, therefore, require judgement. In particular, as discussed in Section 4, this entails forming judgements about both the nature of the current economic environment, including the extent and nature of the uncertainty, and, conditional on that, about which combination of perspectives might be the most appropriate.

An important implication of the analysis is that there is no mechanical link between the insights from any of these perspectives and policy decisions.[60] This is well known

at the next meeting or the following. However, they may warn the MPC if the committee is 250 basis points off!" (Qvigstad, 2005). This is consistent with the idea that the zone in which a rules-based perspective has *particularly* useful insights for policy formulation is somewhat distant from the zone in which a forecast-based perspective would be most useful.

[59]The 'rules versus discretion' debate has a long history. In its original form, of which Friedman's 'k-percent' money growth rule is an exemplar, it concerned the issue of whether or not a discretionary policymaker could deliver better outcomes than a rule, given their state of knowledge and ability. An alternative variant of this debate – that emerged during the 1970s – is that discretionary policy comes with a risk of an inflationary bias, as policymakers seek to exploit the short-run Phillips curve and stimulate aggregate demand (see, e.g., Barro and Gordon, 1983). As noted by Blinder (1998) and Bean (1998), this form of discretionary bias does not seem applicable to monetary policy conducted by independent central banks. The existence of the bias requires the objectives of the central bank to be misaligned with societal objectives or for the central bank to seek to deviate from the objectives they have, both of which are effectively ruled-out by the type of inflation targeting framework in operation in the United Kingdom. In commenting on the original form of the rules versus discretion debate, Blinder (1998) observes that "...while I find the Friedmanite arguments for rules less than persuasive, they cannot be summarily dismissed."

[60]As discussed in Section 2, there is no solution to the real-world monetary policy problem and poli-

to monetary policymakers. For example, in the context of a rule-of-thumb that had at times been used to interpret how monetary policy was related to application of a forecast-based perspective, Bean and Jenkinson (2001) noted that "there is no mechanical link between the central projection for inflation at the two-year horizon and monetary policy." Similarly, monetary policymakers have never mechanically applied prescriptions from policy rules, given the simplifying assumptions on which they are derived.

Indeed, adopting any of these perspectives requires moving away from a fully internally consistent treatment of the monetary policy problem. Moving from this 'analytically closed' environment of theory-based approaches to 'analytically open' perspectives therefore requires the use of judgement to "bridge the gap" between insights derived from these perspectives and the practical problem (Faust, 2005).

4 Bridging the gap: analytical inputs and judgement

If judgement is an essential part of monetary policymaking, helping to bridge the gap from theory to practice, how is it formed, and how can processes support it? By definition, judgement is "the act or process of forming an opinion or making a decision after careful thought." As such, it is naturally subjective. Subjectivity of this kind is an important part of decision making under uncertainty in a wide range of contexts.[61] This section therefore considers the deliberative processes that underpin subjective judgement-based monetary policymaking.

While judgement is subjective, it sits within an overall framework for monetary policy of "constrained discretion." As discussed in Sections 2 and 3, this ensures that the ultimate objective of monetary policy is well specified and establishes explicit mechanisms to ensure the accountability and transparency of decision making. These processes are also informed by an appropriately designed monetary policy strategy, encompassing an articulation of what monetary policy is seeking to achieve and an articulation of how policy will act in order to achieve those objectives.[62]

cymaking necessarily requires judgement. In the context of the alternative perspectives on the monetary policy problem, the forecast-based and news-based perspectives only provide a definitive solution to the monetary policy problem when the assumptions underpinning them coincide with that on which the textbook theory is based (i.e., corner (A) in Figure 3 and therefore away from the environment in which monetary policy typically operates).

[61] Spiegelhalter (2024) discusses the role of subjective judgement in fields including epidemiology, metrology, military intelligence and the law. The application of subjective probabilities also underpins expected utility theory and associated Bayesian approaches (Savage, 1972). Furthermore, Spiegelhalter (2024) argues that the experience of uncertainty itself is a personal perception and, hence, necessarily subjective.

[62] As discussed in footnote 20 on page 17, the MPC's remit is not prescriptive about how the MPC should manage any short to medium-run trade-offs between returning inflation to target and stabilising activity (or about how lags in the monetary transmission mechanism might bear on the speed with

In the context of uncertainty and the alternative perspectives on the policy problem, these processes operate against a backdrop of two higher-level strategic judgements that would form part of the design of an appropriate monetary strategy. First, a subjective assessment of the nature of prevailing uncertainties and, hence, where the practical policy problem lies within the uncertainty space. Such an assessment will be informed by an interpretation of the nature of the monetary policy problem *in general*, including assessments of: the generic nature of the uncertainty; the average 'intensity' of that uncertainty; and the extent to which it varies with changes in the economic environment.[63] Second, and conditional on the first, an assessment of the relevance of insights from alternative perspectives for the practical policy problem.

Revisions to these strategic judgements can be prompted both by novel or large shocks, and by considering evidence gathered following a period of learning.[64] To provide an illustration of the former, Figure 5 depicts a stylised example of the effects accompanying a large and novel economic shock on the nature of uncertainty facing monetary policymakers. Starting from point X, the realisation of the unusual shock leads to an increase in the policymaker's uncertainty as represented by an immediate shift to point Y. This is associated with an immediate reduction in policymakers' understanding of the possible policy-relevant outcomes alongside a decline in the ability to assign probabilities to events. Over time, the accumulation of evidence, research, and analysis results in a more gradual shift to point Z, as indicated by the sequence of smaller arrows.[65]

As the process in Figure 5 unfolds, the first strategic judgement informs assessments of the effect of the shock on the location of the practical policy problem (i.e., the position of point Y). The second strategic judgement is informed by an assessment of the location of the zones in which each of the three perspectives discussed in Section 3 is *particularly* relevant (as illustrated in Figure 3). These assessments ultimately determine the mix of perspectives that are applied to the policy problem as the dynamic learning process

which inflation can be returned to target). Therefore, a well-specified monetary strategy should include an articulation of the objectives, consistent with the remit, but considered in light of the prevailing circumstances. See Garga et al. (2025) for a general discussion of desirable elements of a monetary policy strategy, including risk management considerations.

[63]Judging the nature of uncertainty need not require introspection into one's own ability to assign probabilities because it may be possible to infer subjective probabilities from 'revealed preference' exercises (Savage, 1972, Chapter 3). Moreover, attempting to apply a (subjective) probabilistic approach may be a reliable approach, even in circumstances in which that might appear challenging (Spiegelhalter, 2024, Chapter 6).

[64]Since learning is an ongoing process, there may be benefits from regular low-frequency reassessments of strategic judgements based on the accumulation of evidence and research. For example, such assessments sometimes form part of the monetary policy framework reviews periodically undertaken by some central banks.

[65]While the dynamics in this stylised example are illustrative the qualitative pattern is similar to recent experiences (discussed in Section 2) triggered by the GFC, the Covid pandemic and Russia's invasion of Ukraine. In these cases, there was a very sharp initial shift in the level and nature of uncertainty at the onset of the shock, followed by a gradual and persistent learning process.

Figure 5: Stylised example of effects of a large novel shock and subsequent learning process on subjective uncertainty facing a policymaker

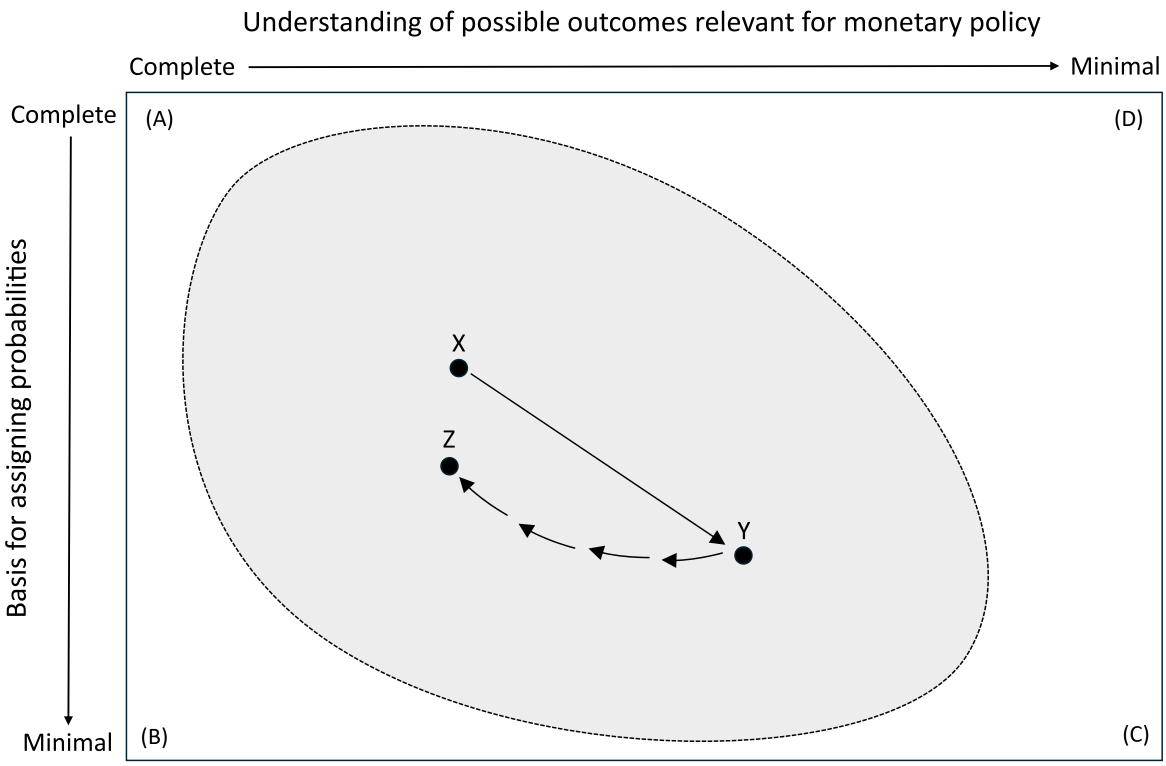

Notes: The axes of the diagram measure the degree of understanding of the range of potential outcomes as relevant for monetary policy decisions (horizontally) and the strength of the basis for attaching probabilities to those outcomes (vertically). The subset of the space in which monetary policy typically operates is subjectively depicted by the shaded ovoid (see Section 2 for further discussion). The diagram also presents a stylised example of the effects of a shock on the nature of uncertainty facing monetary policymakers. Starting from point X, a large and novel shock increases uncertainty as represented by an immediate shift to point Y. This is associated with an immediate reduction in policymakers' understanding of the possible policy-relevant outcomes alongside a decline in the ability to assign probabilities to events. Over time, the accumulation of evidence, research and analysis results in a gradual shift to point Z. See Appendix A for further discussion of the uncertainty space.

continues. As noted above, this judgement is part of a consideration of the appropriate monetary strategy to adopt given the prevailing circumstances.

The preceding discussion implies that, absent the arrival of a large shock or a periodic reassessment, the mix of perspectives applied to the policy problem should be expected to be somewhat persistent over time. Nevertheless, the ability to adjust the mix of perspectives flexibly and rapidly is particularly beneficial when large or novel shocks arrive.

The remainder of the section explores two aspects of judgement-based monetary policymaking conditional on the high-level strategic judgements (i.e., abstracting from the arrival of large shocks and focusing on periods in which the position in the uncertainty space is relatively stable). As such, the mix of perspectives applied to the policy problem and the monetary strategy that is adopted is taken as given. The discussion characterises monetary policy formulation as a process in which analytical inputs inform judgement-based policy decisions (the 'outputs'). Section 4.1 discusses the range of analytical inputs that may be considered by policymakers and how they may be organised and synthesised. Section 4.2 considers how judgements are made and the approaches and processes that can support them.

4.1 Analytical inputs to monetary policymaking

As noted above, monetary policy formulation can be characterised as a process. Decisions about the policy instrument(s) and associated explanation and communication are the 'outputs', which are based on an assessment of a set of 'inputs'. This section discusses examples of possible inputs and how they could be organised.

The alternative perspectives on the policy problem discussed in Section 3 provide a high-level systematic 'framing' of the policy problem (discussed further in Section 4.2). Consistent with this, the alternative policy perspectives can also act as frameworks for organising underlying analysis. That analysis takes the form of 'supporting inputs', some of which are transformed into 'intermediate inputs' to inform judgement-based policy via the lens of a particular perspective. Alternative perspectives on the policy problem may draw on supporting inputs in different ways (as discussed below with reference to some examples). This includes the specific types of input of particular relevance to the perspective, the relative emphasis placed on different types of input, and how they may be combined to inform policy discussions.

An example of this type of approach is presented diagrammatically in Figure 6 and Appendix C provides a summary of different types of inputs that may be used in the formulation of policy.[66] The ovals in the diagram represent the way that different perspectives draw on supporting inputs and, therefore, provide intermediate inputs to judgement-based policymaking.

For example, for a forecast-based perspective, judgement-augmented forecasts could be constructed using data, model(s), and potentially different types of other 'bespoke' supporting inputs, which typically draw on other inputs (such as data or models). Examples of bespoke supporting inputs include scenarios that articulate how the economy, including monetary policy, may respond to a particular constellation of events, or under a specific set of assumptions.[67] These might be considered alongside forecast-based pol-

[66] The diagram is stylised for expositional purposes and does not aim to include an exhaustive set of the inputs.

[67] For example, Bank of England (2025) discusses the use of scenarios that explore the interaction of key assumptions, different shocks, and policy as an input to monetary policy at the Bank of England.

Figure 6: Inputs, alternative perspectives on the policy problem, and judgement

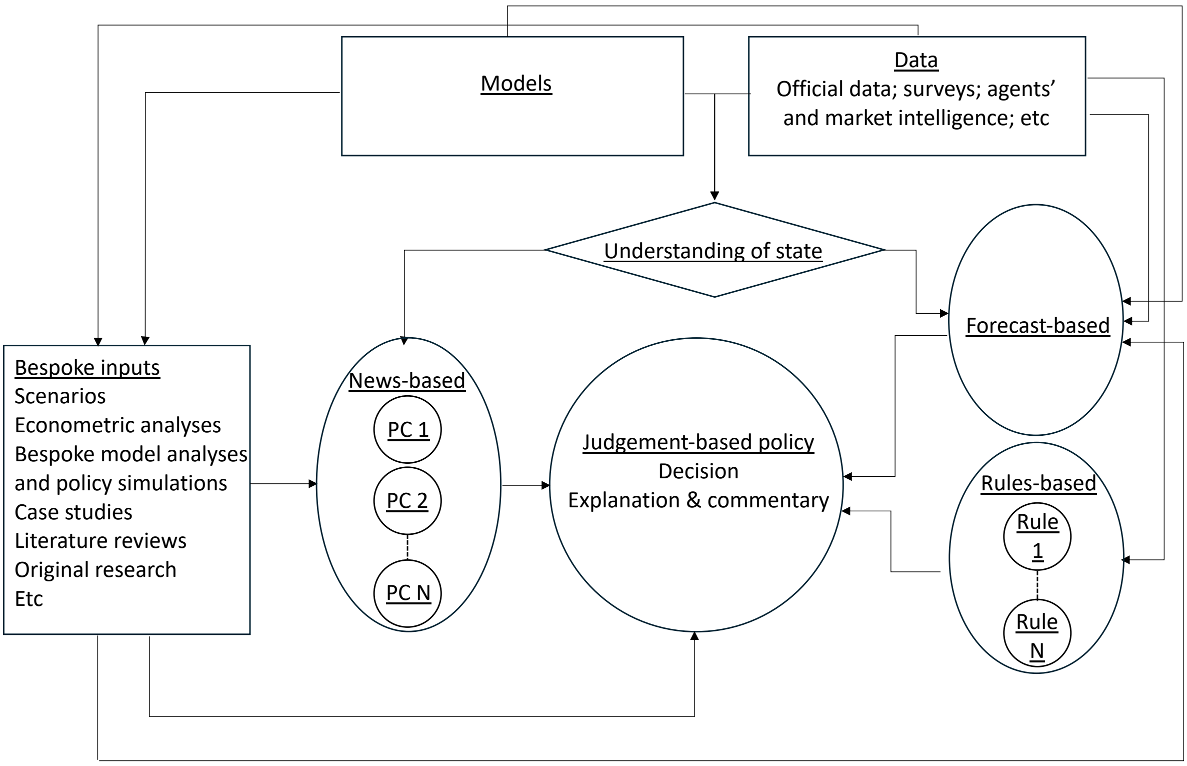

Notes: This figure gives a stylised representation of the flow of inputs to an ultimately judgement-based policy decision, using the different perspectives discussed in Section 3. The inputs listed in the diagram are not intended to be exhaustive. The rectangular boxes correspond to 'supporting' inputs and the ovals correspond to the perspectives. The figure highlights that there may be connections between supporting inputs before they are drawn on by the alternative perspectives.

icy analysis including, for example, the optimal policy projections approach described in Section 3.1.[68] Discussions of this type of forecast and policy analysis could then serve as intermediate inputs to the ultimately judgement-based policy decision.

A news-based perspective could draw on supporting inputs to produce intermediate inputs that address specific 'policy considerations' that apply the general questions set out in the discussion in Section 3.2 to the current policy problem (denoted as 'PC' in Figure 6). As in a forecast-based perspective, these inputs would draw on a process for updating the latest understanding of the state of the economy (either quantitatively or qualitatively), shown in the diamond-shaped box.

A rules-based perspective could use data for (and/or estimates of) the variables that appear in the simple rules that may be most relevant. Additional bespoke inputs may include assessments of the likely performance of the simple rules in the prevailing economic circumstances (e.g., some rules may perform differently in the face of different types of shocks). Such analyses can be combined to form an intermediate input to policy judge-

[68]The model-based policy analysis discussed in Alati et al. (2025) is an example of this kind of input.

ments. Therefore, as discussed in Section 3.3, the rules themselves serve as a benchmark rather than a mechanical prescription.

As noted, there are likely to be many ‘supporting’ inputs to the process of policy formulation (shown in the rectangular boxes in Figure 6). Some of these inputs can be regarded as the building blocks of further analysis, and can be combined in different ways. Key examples include official data, surveys, and other forms of intelligence about the state of the economy and its potential evolution. Economic models, which may have been purpose-built or adapted from academic research, are another form of supporting input. Bespoke supporting inputs would build further on those inputs: for example, econometric and/or model-based analysis. Some, though not necessarily all, inputs in this category may be based on the theory-based approaches located in corner (A) of Figure 3. Other examples include the examination of historical episodes as case studies, reviews of academic research relevant to the current policy problem, and original research by staff and policymakers.

In general, the specific inputs that inform policy decisions will vary depending on the circumstances, including the degree and nature of the uncertainty and the (combination of) policy perspective(s) that policymakers judge to be most appropriate. For example, case studies may be more informative (and bespoke econometric or theory-based analysis less so) when particularly unusual shocks arrive. Moreover, the inputs to policy judgements that are derived from a given perspective might also vary. For example, in the context of a forecast-based perspective, the relative relevance of policy analysis derived from simple rules and from optimal policy projections may change with the circumstances that policymakers face. These observations highlight the key role of judgement in the use of inputs in monetary policy formulation.

4.2 Judgement

As noted above, judgement is an integral part of the process by which policymakers make monetary policy decisions using a wide range of analytical inputs. By its nature, judgement is not algorithmic. Moreover, individual judgements are necessarily subjective. Accordingly, this section does not aim to describe (and certainly not prescribe) processes through which policymakers could or should apply judgement. Nevertheless, considering the context and general nature of the decision problem provides some insights on how to support judgement-based policymaking.

Kay and King (2020) argue that “successful decision-making under uncertainty is a collaborative process.” Indeed, group decision making is favoured in many situations in which both judgement and technical expertise are required (King, 2002; Blinder, 2007).[69] These are among the most important reasons why monetary policy decisions are taken by committees in many jurisdictions (Blinder, 2007).

In the United Kingdom, monetary policy decisions are the result of a vote by members of the Bank of England's MPC and each member is individually accountable to Parliament. Processes to ensure accountability tend not to focus on an individual's votes *per se*, which are revealed simultaneously with each policy decision, but rather on the reasoning and the associated judgements that motivated them.[70]

As highlighted by King (2002), within this institutional structure it is important that the processes supporting decision-making provide space for open discussions and exchanges of view. One purpose of such exchanges is to pool information, views, and judgements, thereby allowing policymakers to learn from each other. For example, evidence from experiments in simplified environments designed to mimic monetary policymaking suggests that groups do, indeed, make better decisions than individuals, with decision makers able to learn from each other to improve their performance (Blinder and Morgan, 2005; Lombardelli et al., 2005).[71]

Nevertheless, even following extensive discussion, differences of view between individuals are to be expected, given the nature of uncertainty and the necessary role of subjective judgement. Indeed, Blinder (2007) observes that individual policymakers "often reach different policy conclusions even though all of them see the same data", noting that "different people bring to the table different decision making methods" and that different judgements "surely stem from different models and different ways of processing the same information."[72] Viewed through the lens of the discussion in this paper, these differences can be interpreted in terms of different judgements about both the relative importance of alternative perspectives on the policy problem and the analysis brought to bear on each of them.[73]

[69]For example, Faust (2005) and King (2002) discuss parallels with medicine and law.

[70]For example, part of the process for ensuring accountability is through regular hearings of the Treasury Select Committee (TSC) on Monetary Policy, in which TSC members question MPC members about their judgements. These hearings may include, for example, discussions of an MPC member's rationale for their votes as set out in speeches and other public commentaries.

[71]There is a large academic literature studying monetary policymaking committees and it is beyond the scope of the paper to provide a comprehensive summary. However, notable contributions include Blinder (2007), Blinder and Morgan (2008), Charness and Sutter (2012), Hansen et al. (2014), Reis (2013), Sibert (2006), and Warsh (2016).

[72]This conjecture echoes King's (2002) observation of the MPC: "Differences of view on our Committee are an honest reflection of the uncertainty about both the data and the structure of the economy."

[73]While individual policymakers may make different judgements, there may also be strong commonalities between individual views. In general, such commonalities are likely to relate to *qualitative* aspects of monetary policy formulation. So, for example, several policymakers may agree on the most important factors that are relevant for the prospects for monetary policy, their relative importance (i.e, their ordinal ranking) and the directional implications of each factor for the policy stance. These policymakers may or may not cast the same votes in a policy meeting, since those votes are informed by a quantitative cali-

Against this backdrop, there remains a practical question of how to support an environment in which individual and collective policy judgements can be developed and shared most effectively. While it is beyond the scope of the present paper to set out detailed prescriptions, some guiding principles emerge from the analysis in Section 3.

In particular, the alternative perspectives on the monetary policy problem described in Section 3 each provide a systematic high-level 'framing' of the policy problem. They thereby provide a 'way to think about' the policy problem, which, in turn, structures the discussions of specific analytical inputs. Kay and King (2020) argue that problem framing is a key aspect of decision-making under uncertainty to which expert advisers may add particular value.[74]

A productive discussion may be usefully framed using the "trained intuition" of policymakers and staff, which describes the way in which expertise, acquired through specialist training and experience, provides consistent ways to understand and interpret evidence (Solow, 1984).[75]

Staff expertise can also be brought to bear in the preparation and synthesis of inputs. A key role is to assist policymakers when interpreting results and forming judgements on their relevance to the practical policy problem. Important requirements in playing this role are a deep understanding of the strengths and deficiencies of the underlying tools and techniques, and an ability to produce and explain an "informed analysis of how the deficiencies might or might not temper the appeal of the [policy] advice" (Faust, 2005).[76]

The application of trained intuition and assessments of policy relevance both have

bration of specific individual judgements (perhaps relating to assessments of the relevance of a particular factor, the quantitative implications of analytical inputs attempting to capture each factor, or both). This observation highlights that the distribution of votes is not a sufficient statistic for the distribution of *qualitative* judgements among a policymaking committee. That in turn implies that both individual and collective judgements are important components of the policy outlook.

[74] Kay and King (2020, Chapter 21) note that "economists cannot tell policy-makers what decisions to make [. . . but] they can help them think about their problems and provide relevant information", defining framing as "identifying critical factors and assembling relevant data [and] applying experience of how these factors have interacted in the past, and making an assessment of how they might interact in the future". The general insight that staff analysis may play an important role in framing discussions also applies to many of the broad range of inputs described in Section 4.1.

[75] Solow's exposition is consistent with Blinder's 'quite informal' application of theory-based reasoning discussed in Section 3: "Suppose all [economics] can do is help us to organize our necessarily incomplete perceptions about the economy, to see connections the untutored eye would miss, to tell plausible stories with the help of a few central principles. [. . .] In that case what we want a piece of economic theory to do is precisely to train our intuition, to give us a handle on the facts" (Solow, 1984). Note that Solow did not intend 'intuition' to supplant analytical rigour: "I hope that no one will fall into the error of thinking that this low-key view of the nature of economics is a license for loose thinking. Logical rigour is exactly as important in this scheme of things as it is in the more self-consciously scientific one" (Solow, 1984).

[76] Faust (2005) focuses primarily on policy analysis using structural macroeconomic models (some of the approaches that staff may use in that context are discussed in footnote 28). However, the point applies more broadly to a wide range of inputs.

some links to analogical reasoning and case-based decision theory (explored further in Appendix D). Such approaches consider the extent of similarity between a piece of analysis (or synthesis) and the practical policy problem in judging the relevance of the former to the latter. This often involves 'arguing by analogy' to establish a link between the analysis and the practical problem.[77] This reasoning implies that even stylised analysis, including in the 'textbook theory' region close to corner (A) of Figure 3, can be brought to bear on the practical policy problem. Such analysis would not be regarded as a direct approximation to the solution of the policy problem (as discussed in Section 2), but rather form part of a broader set of supporting inputs. In this context, expertise is partly reflected in an ability to support (and make) judgements on the relevance of a diverse array of inputs.[78]

The most appropriate approaches to supporting policy judgements will vary on a case-by-case basis. However, some general insights can be gleaned by considering past examples of approaches and analysis that supported policymaking. Those considered below are from transcripts of 2016 MPC meetings and the supporting analysis.[79]

The first example is the February 2016 policy meeting, in preparation for which Bank staff produced a "Quarterly Monetary Policy Analysis Note" (Bank of England, 2016e). The analysis adopts a forecast-based perspective to frame a discussion of the appropriate policy strategy using optimal policy projections similar to those described in Section 3.1. It includes sensitivity analysis to explore the possibilities of model misspecification and a section exploring the mapping of the analysis to the practical policy problem, arguing that "the optimal policy simulations cannot be applied directly to the real-world issues facing the MPC."[80] Indeed, discussion by MPC members in their policy meeting suggests that factors beyond the scope of the model-based analysis would indeed be more relevant to their policy outlook.[81]

The second example is the May 2016 policy meeting (held shortly before the referen-

[77]For example, Kay and King (2020, Chapter 14) distinguish between real-world "mysteries" with no knowable answer and "puzzles" that can be characterised in stylised models (that Savage (1972) called "small worlds") arguing that: "While not providing comprehensive or quantitative answers to economic problems, [stylised models] help us frame arguments to understand better the nature of the mystery, by drawing analogies with a small world in which a puzzle has a determinate answer".

[78]As discussed in Appendix D.3, this view is consistent with the implications of case-based decision theory since expertise partly reflects the ability to learn the 'similarity function' that informs decision making in that framework.

[79]These materials have been released in accordance with changes following the Warsh (2014) review.

[80]Section 2.3 of Bank of England (2016e) is entitled "What do these policy simulations have to say about the real world?". It discusses three features of the practical policy problem that may not be adequately captured in the simulations (questions around the credibility of the implied commitment in the simulation, the extent to which policy should look through the near-term cost shocks driving inflation, and the absence of risks management considerations) and the likely implications for policy in each case. The briefing material therefore provides an "informed analysis of how the deficiencies might or might not temper the appeal of the [policy] advice" as advocated by Faust (2005).

[81]See, for example, the remarks by Nemat Shafik (Bank of England, 2016c, p14).

dum on EU membership) and the supporting analysis. Unsurprisingly, the uncertainty surrounding the referendum result and the potential policy implications were discussed at length.[82] Indeed, with reference to some of the factors that could determine the appropriate monetary policy response to the referendum result, Andrew Haldane (Bank of England, 2016d, p4) assessed them to be "not just unknown, but unknowable", suggestive of a high-level strategic judgement that the policy problem was characterised by deep uncertainty.

Supporting staff analysis explored the potential channels through which a vote to leave the EU could affect the economy.[83] Andrew Haldane (Bank of England, 2016d, p4) regarded the analysis as providing a helpful "framework for us thinking through the monetary policy consequences", indicative of the role of intermediate inputs in framing policy discussions. Indeed, the approach adopted by that analysis provided the basis for the MPC's framing of its external communications (in terms of the "combination of influences on demand, supply and the exchange rate") in the May 2016 Monetary Policy Statement and subsequent communication.[84]

Other staff analysis used stylised simulations to draw out qualitative policy implications of different types of supply shocks under alternative assumptions about private sector learning.[85] The analysis includes highly stylised simulations under alternative assumptions about: the nature of the supply shock; the extent to which agents (including the monetary policymaker in the model) must learn about the nature of the shock; and the conduct of monetary policy (comparing a simple rule with optimal policy conditional on an assumption about the loss function). Despite the highly stylised nature of the model-based simulations, Mark Carney (Bank of England, 2016d, p14) noted their qualitative implications for the appropriate stance of monetary policy.

This discussion demonstrates that policymakers do indeed approach the practical policy problem in a pluralistic way, particularly in times of heightened uncertainty. This involves considering the nature of the uncertainty characterising the environment and assessing different types of inputs in terms of their relevance to the real-world problem, including those based on approaches situated in corner (A) of Figure 3. Moreover, the discussion also suggests that there are benefits to designing processes that support flexibility over the perspectives that policymakers apply to the policy problem, and the inputs they draw upon to facilitate their decision making, within an overall framework of constrained discretion.

[82] See also the discussion in Section 3.2, footnote 42.

[83] "Leave vote channels and updated illustrative quantification" (Bank of England, 2016a).

[84] See Bank of England (2016b).

[85] The investigation was motivated by uncertainty over the potential supply-side effects of a leave vote in the referendum. See "Topical Issues in Monetary Strategy: Monetary Policy and Supply Shocks" (Bank of England, 2016f).

5 Conclusion

Pervasive and time-varying uncertainty characterises the monetary policy environment. As a result, policymakers must pursue their objectives without a complete understanding of the current state of the economy, the economic outlook, or the precise way in which shocks and monetary policy transmit. While theoretical treatments of the monetary policy problem provide many valuable insights, they are too simplified to be applied directly to the real world. Therefore, judgement is an essential part of policymaking in practice.

Nevertheless, theoretical treatments of monetary policy can play an important role in supporting judgement-based policymaking insofar as they can be used to provide alternative ways of thinking about the real-world monetary policy problem. Indeed, as discussed in this paper, theory can be used to provide different perspectives on the monetary policy problem that can be applied in practice to support real-world policymaking. Crucially, such perspectives recognise the inherent uncertainty of the policy environment. There is no unique way to draw on theory to inform real-world policymaking and so policymakers may deploy a variety of different perspectives in practice.

This paper has discussed three alternative perspectives, each of which draws on theory in different ways. First, a forecast-based perspective, which uses forecasts for key macroeconomic variables as the basis for formulating monetary policy. Second, a news-based perspective, which involves updating policy over time, as policymakers' understanding of the drivers of macroeconomic outcomes evolves, and as new shocks or events are identified. Third, a rules-based perspective, which uses simple instrument rules as a guide to the appropriate setting of policy, motivated by a concern for policy to be robust to epistemic uncertainty.

The relevance of the different perspectives will depend on the policymaker's subjective assessment of uncertainty, specifically with regard to their understanding of the possible outcomes relevant for monetary policy, and their assessment of the strength of the basis for assigning probabilities to those outcomes. In the grey area of uncertainty between complete knowledge and very deep uncertainty in which monetary policy operates, each of these perspectives may be informative to some degree. But the extent to which any perspective may be more informative than another depends on the prevailing level and nature of uncertainty.

This analysis has three key implications, consistent with the Bank's response to the Bernanke Review. First, in more uncertain environments, such as the one that has prevailed in the United Kingdom over recent years, a forecast-based perspective on the monetary policy problem is likely to be less useful, and other perspectives more useful.

Broadening the perspectives that are taken, and the analytical inputs in support of that, can better support judgement-based policymaking in recognition of uncertainty. Second, as uncertainty varies over time, the perspectives that are taken to policymaking should vary accordingly. Processes and analysis in support of policymaking should support that flexibility. Third, consistent with many of the recommendations in the Bernanke Review, the Bank's modelling and toolkit should be improved and extended in order to better support a more pluralistic approach.

References

Adam, K. and R. Billi (2006). Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking 38*(7), 1,877–905.

Adam, K. and R. M. Billi (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics 54*(3), 728–752.

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review 109*(4), 1263–1289.

Adrian, T., D. Giannone, M. Luciani, and M. West (2025). Scenario synthesis and macroeconomic risk. Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Al-Nowaihi, A. and L. Stracca (2003). Behavioural central bank loss functions, skewed risks and certainty equivalence. *The Manchester School 71*, 21–38.

Alati, A., M. Arazi, J. Barrdear, A. Gimber, A. Haberis, E. Hughes, L. Laureys, S. Lloyd, O. Ozturk, J. Page, K. Reinold, E. Tong, M. Tong, and M. Waldron (2025). Tools for endogenous monetary policy analysis: Optimal projections and instrument rules. *Bank of England Macro Technical Paper No. 4* (4).

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 32*(8), 2428–2452.

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2000). Robustness, detection and the price of risk. *mimeo*.

Bailey, A. (2025). Monetary policy in uncertain times. Speech given at the Reykjavík Economic Conference.

Ball, L. M. (1999). Policy rules for open economies. In J. Taylor (Ed.), *Monetary policy rules*, pp. 127–156. University of Chicago Press.

Bank of England (2016a). Leave vote channels and updated illustrative quantification. Internal staff briefing note, available at https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy/mpc-documentation/2016/mpc-may-2016.zip.

Bank of England (2016b). Minutes of the Monetary Policy Committee meeting ending on 11 May 2016.

Bank of England (2016c). MPC February 2016 Transcript. Available at https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy/mpc-documentation/2016/mpc-transcript-february-2016.pdf.

Bank of England (2016d). MPC May 2016 Transcript. Available at https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy/mpc-documentation/2016/mpc-transcript-may-2016.pdf.

Bank of England (2016e). Quarterly Monetary Policy Analysis Note. Internal staff briefing note, available at https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy/mpc-documentation/2016/mpc-february-2016.zip.

Bank of England (2016f). Topical Issues in Monetary Strategy: Monetary Policy and Supply Shocks. Internal staff briefing note, available at https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy/mpc-documentation/2016/mpc-may-2016.zip.

Bank of England (2024). Response of the Bank of England to the Bernanke review of forecasting for monetary policy making and communication at the Bank of England.

Bank of England (2025). Monetary policymaking at the Bank of England in uncertain times. *Quarterly Bulletin*. Bank of England.

Barnichon, R. and G. Mesters (2023). A sufficient statistics approach for macro policy. *American Economic Review 113*(11), 2809–2845.

Barro, R. J. and D. B. Gordon (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics 12*(1), 101–121.

Batini, N., R. Harrison, and S. P. Millard (2003). Monetary policy rules for an open economy. *Journal of Economic Dynamics and Control 27*(11-12), 2059–2094.

Bean, C. (1998). The new UK monetary arrangements: a view from the literature. *The Economic Journal 108*(451), 1795–1809.

Bean, C. (2005). Monetary policy in an uncertain world. Speech to the Oxonia Distinguished Speaker Seminar, Oxford.

Bean, C. (2007). Risk, uncertainty and monetary policy. Speech to Dow Jones, at City Club, Old Broad Street.

Bean, C. R. (2003). Inflation targeting: the UK experience. *Bank of England Quarterly Bulletin*.

Bean, C. R. and N. Jenkinson (2001). The Formulation of Monetary Policy at the Bank of England. *Bank of England Quarterly Bulletin*.

Bernanke, B. (2024). Forecasting for monetary policy making and communication at the Bank of England: a review. Review commissioned by the Court of the Bank of England.

Bernanke, B. (2025). Improving Fed communications: A proposal. *Hutchins Center on Fiscal and Monetary Policy, The Brookings Institution*.

Bernanke, B. S. (2007). Monetary policy under uncertainty. Remarks to the Annual Policy Conference of the Federal Reserve Bank of St. Louis.

Bernanke, B. S. and F. S. Mishkin (1997). Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives 11*(2), 97–116.

Blake, A. P. and F. Zampolli (2011). Optimal policy in Markov-switching rational expectations models. *Journal of Economic Dynamics and Control 35*(10), 1626–1651.

Blinder, A. (1998). *Central banking in theory and practice.* MIT press.

Blinder, A. S. (2007). Monetary policy by committee: Why and how? *European Journal of Political Economy 23*(1), 106–123.

Blinder, A. S. and J. Morgan (2005). Are two heads better than one? Monetary policy by committee. *Journal of Money, Credit and Banking 37*(5), 789–811.

Blinder, A. S. and J. Morgan (2008). Do monetary policy committees need leaders? A report on an experiment. *American Economic Review 98*(2), 224–29.

Blinder, A. S. and R. Reis (2005). Understanding the Greenspan standard. *Proceedings of the Jackson Hole Economic Policy Symposium*, 11–96.

Box, G. E. and N. R. Draper (1987). *Empirical model-building and response surfaces.* John Wiley & Sons.

Brainard, W. C. (1967). Uncertainty and the effectiveness of policy. *The American Economic Review 57*(2), 411–425.

Broadbent, B. (2023). Signal versus noise. Speech at the London Business School, London. Bank of England.

Broadbent, B. (2024). From NICE... to not so nice. Speech at the Bank of England.

Brock, W. A., S. N. Durlauf, and K. D. West (2003). Policy evaluation in uncertain economic environments. *Brookings Papers on Economic Activity*.

Brock, W. A., S. N. Durlauf, and K. D. West (2007). Model uncertainty and policy evaluation: Some theory and empirics. *Journal of Econometrics 136*(2), 629–664.

Brumm, J. and S. Scheidegger (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica 85*(5), 1575–1612.

Bryant, R. C., P. Hooper, and C. L. Mann (1993). *Evaluating Policy Regimes: New Research in Empirical Macroeconomics*. Brookings Institution Press.

Burgess, S., E. Fernandez-Corugedo, C. Groth, R. Harrison, F. Monti, K. Theodoridis, and M. Waldron (2013). The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models. *Bank of England Staff Working Paper No. 471* (471).

Burr, N. and T. Willems (2024). About a rate of (general) interest: how monetary policy transmits. *Bank of England Quarterly Bulletin 64*(1), 3–3.

Caldara, D., R. Harrison, and A. Lipińska (2014). Practical tools for policy analysis in DSGE models with missing shocks. *Journal of Applied Econometrics 29*(7), 1145–1163.

Carney, M. (2016). Uncertainty, the economy and policy. Speech given at the Court Room, Bank of England.

Carney, M. (2017). Lambda. Speech given at the London School of Economics.

Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2025). Making decisions under model misspecification. *The Review of Economic Studies*.

Charness, G. and M. Sutter (2012). Groups make better self-interested decisions. *The Journal of Economic Perspectives 26*(3), 157–176.

Cogley, T., B. De Paoli, C. Matthes, K. Nikolov, and T. Yates (2011). A Bayesian approach to optimal monetary policy with parameter and model uncertainty. *Journal of Economic Dynamics and Control 35*(12), 2186–2212.

Cox, D. (1995). Discussion of "Model uncertainty, data mining and statistical inference". *Journal of the Royal Statistical Society. Series A 158*, 455–456.

Davig, T. and E. M. Leeper (2007). Generalizing the Taylor principle. *American Economic Review 97*(3), 607–635.

de Groot, O., F. Mazelis, R. Motto, and A. Ristiniemi (2021). A toolkit for computing Constrained Optimal Policy Projections (COPPs). Working Paper Series 2555, European Central Bank.

Del Negro, M. and F. Schorfheide (2006). How good is what you've got? DGSE-VAR as a toolkit for evaluating DSGE models. *Federal Reserve Bank of Atlanta Economic Review 91*(2), 21.

Del Negro, M. and F. Schorfheide (2009). Monetary policy analysis with potentially misspecified models. *American Economic Review 99*(4), 1415–1450.

Dennis, R. (2007). Optimal policy in rational expectations models: new solution algorithms. *Macroeconomic Dynamics 11*(1), 31–55.

Dennis, R., K. Leitemo, and U. Söderström (2009). Methods for robust control. *Journal of Economic Dynamics and Control 33*(8), 1604–1616.

Druedahl, J. and T. H. Jørgensen (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control 74*, 87–107.

Eichengreen, B. (2011). Economic history and economic policy. *mimeo*.

Esponda, I. and D. Pouzo (2016). Berk–Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica 84*(3), 1093–1130.

Eusepi, S. and B. Preston (2018). The science of monetary policy: An imperfect knowledge perspective. *Journal of Economic Literature 56*(1), 3–59.

Evans, C., J. Fisher, F. Gourio, and S. Krane (2016). Risk management for monetary policy near the zero lower bound. *Brookings Papers on Economic Activity*, 141–219.

Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

Farmer, R. E., D. F. Waggoner, and T. Zha (2009). Understanding Markov-switching rational expectations models. *Journal of Economic Theory 144*(5), 1849–1867.

Farmer, R. E., D. F. Waggoner, and T. Zha (2011). Minimal state variable solutions to Markov-switching rational expectations models. *Journal of Economic Dynamics and Control 35*(12), 2150–2166.

Faust, J. (2005). Is applied monetary policy analysis hard? *mimeo*.

Faust, J. (2012). DSGE models: I smell a rat (and it smells good). *International Journal of Central Banking 8*(1), 53–64.

Faust, J. and A. Gupta (2012). Posterior predictive analysis for evaluating DSGE models. *NBER Working Paper*.

Filippeli, T., R. Harrison, and K. Theodoridis (2020). DSGE-based priors for BVARs and quasi-Bayesian DSGE estimation. *Econometrics and Statistics 16*, 1–27.

Friedman, M. (1960). *A Program for Monetary Stability.* New York: Fordham University Press.

Friedman, M. (1961). The lag in effect of monetary policy. *Journal of Political Economy 69*(5), 447–447.

Friedman, M. (1968). The role of monetary policy. *American Economic Review 58*(1), 1–17.

Gabaix, X. (2020). A behavioral New Keynesian model. *American Economic Review 110*(8), 2271–2327.

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle.* Princeton University Press.

Garga, V., E. Herbst, A. McKay, G. Nicolo, and M. Paustian (2025). Monetary policy, uncertainty, and communications. Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System.

Gaspar, V., F. Smets, and D. Vestin (2006). Adaptive learning, persistence, and optimal monetary policy. *Journal of the European Economic Association 4*(2-3), 376–385.

Gaspar, V., F. Smets, and D. Vestin (2010). Inflation expectations, adaptive learning and optimal monetary policy. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, pp. 1055–1095. Elsevier.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (1995). *Bayesian data analysis.* Chapman and Hall/CRC.

Geweke, J. (2010). *Complete and incomplete econometric models.* Princeton University Press.

Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics 164*(1), 130–141.

Geweke, J. and G. Amisano (2012). Prediction with misspecified models. *American Economic Review 102*(3), 482–486.

Ghosh, A. (2024). Robust comparative statics with misspecified Bayesian learning. *mimeo*.

Gilboa, I., A. Postlewaite, L. Samuelson, and D. Schmeidler (2014). Economic models as analogies. *The Economic Journal 124*(578), F513–F533.

Gilboa, I. and D. Schmeidler (1995). Case-based decision theory. *The Quarterly Journal of Economics 110*(3), 605–639.

Gilboa, I. and D. Schmeidler (2001). *A theory of case-based decisions.* Cambridge University Press.

Goodhart, C. (1989). The conduct of monetary policy. *The Economic Journal 99*(396), 293–346.

Granger, C. W. (1989). Combining forecasts – twenty years later. *Journal of forecasting 8*(3), 167–173.

Greene, M. (2025). The supply side demands more attention. Speech given at the Adam Smith Business School, University of Glasgow.

Greenspan, A. (2003). Monetary policy under uncertainty. *Remarks at a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming*.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. *American Economic Review 94*(2), 33–40.

Grünwald, P. and T. van Ommen (2017). Inconsistency of Bayesian Inference for Misspecified Linear Models, and a Proposal for Repairing It. *Bayesian Analysis 12*(4), 1069–1103.

Guerrieri, L. and M. Iacoviello (2015). Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics 70*, 22–38.

Gupta, A. (2016). A forecasting metric for evaluating DSGE models for policy analysis. *International Journal of Central Banking*.

Hansen, L. P. and T. J. Sargent (2008). *Robustness.* Princeton university press.

Hansen, L. P. and T. J. Sargent (2013). *Recursive models of dynamic linear economies.* Princeton University Press.

Hansen, S., M. McMahon, and C. Velasco Rivera (2014). Preferences or private assessments on a monetary policy committee? *Journal of Monetary Economics 67*, 16–32.

Harrison, R. and M. Waldron (2021). Optimal policy with occasionally binding constraints: piecewise linear solution methods. *Bank of England Staff Working Paper 911*.

Hebden, J. and F. Winkler (2024). Computation of policy counterfactuals in sequence space. *FEDS Discussion Paper, Federal Reserve Board of Governors*.

Issing, O. (2002). Monetary policy in a world of uncertainty. Speech at the Economic Policy Forum in Paris. European Central Bank.

Johnson, S. G., A. Bilovich, and D. Tuckett (2023). Conviction narrative theory: A theory of choice under radical uncertainty. *Behavioral and Brain Sciences 46*.

Karadi, P., A. Nakov, G. N. Barrau, E. Pasten, and D. Thaler (2024). Strike while the iron is hot: optimal monetary policy with a nonlinear Phillips curve. BIS Working Paper, Bank for International Settlements.

Kay, J. and M. King (2020). *Radical uncertainty: Decision-making for an unknowable future*. Hachette UK.

King, M. (1997). The inflation target five years on. Bank of England Quarterly Bulletin.

King, M. (2002). The Monetary Policy Committee: Five Years On. *Bank of England Quarterly Bulletin*.

King, M. (2004). What fates impose: Facing up to uncertainty. Eighth British Academy Annual Lecture, London.

King, M. (2005). Monetary policy: Practice ahead of theory. Mais Lecture, Cass Business School, London.

King, M. (2007). The MPC ten years on. Lecture delivered to the Society of Business Economists.

King, M. (2010). Uncertainty in macroeconomic policy making: art or science? Speech at the Royal Society conference on 'Handling Uncertainty in Science', London.

King, M. A. (2024). Inflation targets: Practice ahead of theory. *NBER Working Paper 32594*.

Knight, F. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.

Kulish, M. and A. Pagan (2017). Estimation and solution of models with expectations and structural changes. *Journal of Applied Econometrics 32*(2), 255–274.

Lane, P. R. (2024). Monetary policy under uncertainty. Keynote speech at the Bank of England Watchers' Conference 2024, King's College London.

Lanzani, G. (2024). Dynamic concern for misspecification. *mimeo*.

Levin, A., V. Wieland, and J. C. Williams (2003). The performance of forecast-based monetary policy rules under model uncertainty. *American Economic Review 93*(3), 622–645.

Levin, A. T. (2014). The design and communication of systematic monetary policy strategies. *Journal of Economic Dynamics and Control 49*, 52–69.

Levin, A. T., V. Wieland, and J. Williams (1999). Robustness of simple monetary policy rules under model uncertainty. In *Monetary policy rules*, pp. 263–318. University of Chicago Press.

Levin, A. T. and J. C. Williams (2003). Robust monetary policy with competing reference models. *Journal of Monetary Economics 50*(5), 945–975.

Lombardelli, C. (2024). Managing the present, shaping the future. Speech at the Bank of England Watchers Conference.

Lombardelli, C., J. Proudman, and J. Talbot (2005). Committees versus individuals: An experimental analysis of monetary policy decision-making. *International Journal of Central Banking 1*(1), 181–205.

Lucas, R. E. (2011). What economists do. *Journal of Applied Economics 14*(1), 1–4.

Maliar, L., S. Maliar, and P. Winant (2019). Will artificial intelligence replace computational economists any time soon? *CEPR Discussion Paper 14024*.

Matsui, A. (2000). Expected utility and case-based reasoning. *Mathematical Social Sciences 39*(1), 1–12.

Mendes, R., S. Murchison, and C. A. Wilkins (2017). Monetary policy under uncertainty: Practice versus theory. *Bank of Canada Staff Discussion Paper*.

Molnár, K. and S. Santoro (2014). Optimal monetary policy when agents are learning. *European Economic Review 66*, 39–62.

Monti, F. (2010). Combining judgment and models. *Journal of Money, Credit and Banking 42*(8), 1641–1662.

Morgan, M. S. (2012). *The world in the model: How economists work and think.* Cambridge University Press.

Nakov, A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking 4*(2), 73–127.

Orphanides, A. and J. Williams (2007a). Inflation targeting under imperfect knowledge. In F. Mishkin and K. Schmidt-Hebbel (Eds.), *Monetary policy under inflation targeting.* Central Bank of Chile, Santiago, Chile.

Orphanides, A. and J. C. Williams (2002). Robust monetary policy rules with unknown natural rates. *Brookings Papers on Economic Activity 2002*(2), 63–145.

Orphanides, A. and J. C. Williams (2006). Monetary policy with imperfect knowledge. *Journal of the European Economic Association 4*(2/3), 366–375.

Orphanides, A. and J. C. Williams (2007b). Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics 54*(5), 1406–1435.

Orphanides, A. and J. C. Williams (2008). Learning, expectations formation, and the pitfalls of optimal control monetary policy. *Journal of Monetary Economics 55*, S80–S96.

Phelps, E. S. (1967). Phillips curves, expectations of inflation, and optimal unemployment over time. *Economica 34*, 254–281.

Phelps, E. S. (1968). Money-wage dynamics and labor market equilibrium. *Journal of Political Economy 76*, 678–711.

Pill, H. (2022). Monetary policy with a steady hand. Speech given at the Society of Professional Economists conference.

Pill, H. (2023). Inflation is a 'wicked problem'. Speech at a fireside chat with Warwick Think Tank on the UK Economy, Bank of England.

Pill, H. (2024). Monetary policy strategy. Speech at Cardiff University Business School.

Qvigstad, J. F. (2005). When does an interest rate path "look good"? criteria for an appropriate future interest rate path – a practitioner's guide. *Norges Bank Staff Memo 05/6*.

Reis, R. (2013). Central bank design. *Journal of Economic Perspectives 27*(4), 17–44.

Reis, R. (2018). Is something really wrong with macroeconomics? *Oxford Review of Economic Policy 34*(1-2), 132–155.

Sala, L., U. Söderström, and A. Trigari (2008). Monetary policy under uncertainty in an estimated model with labor market frictions. *Journal of Monetary Economics 55*(5), 983–1006.

Sargent, T. J. (1999). *The conquest of American inflation.* Princeton University Press.

Savage, L. J. (1972). *The foundations of statistics.* Courier Corporation.

Seneca, M. (2020). Risk shocks and monetary policy in the new normal. *International Journal of Central Banking*.

Sibert, A. (2006). Central banking by committee. *International Finance 9*(2), 145–168.

Sims, C. A. (2001). Pitfalls of a minimax approach to model uncertainty. *American Economic Review 91*(2), 51–54.

Sims, C. A. (2002). The role of models and probabilities in the monetary policy process. *Brookings Papers on Economic Activity 33*(2), 1–62.

Söderström, U. (2002). Monetary policy with uncertain parameters. *Scandinavian Journal of Economics 104*(1), 125–145.

Solow, R. M. (1984). Mr Hicks and the classics. *Oxford Economic Papers 36*, 13–25.

Spiegelhalter, D. (2024). *The art of uncertainty: how to navigate chance, ignorance, risk and luck*. Random House.

Spiegelhalter, D. J. and H. Riesch (2011). Don't know, can't know: embracing deeper uncertainties when analysing risks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 369*(1956), 4730–4750.

Stirling, A. (2010). Keep it complex. *Nature 468*(7327), 1029–1031.

Svensson, L. E. (1999). Inflation targeting: some extensions. *Scandinavian journal of Economics 101*(3), 337–361.

Svensson, L. E. (2000). Open-economy inflation targeting. *Journal of International Economics 50*(1), 155–183.

Svensson, L. E. (2003a). The inflation forecast and the loss function. In P. Mizen (Ed.), *Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart*. Edward Elgar Publishing.

Svensson, L. E. (2003b). Optimal policy with low-probability extreme events. *NBER Working Paper No. 10196*.

Svensson, L. E. (2006). The instrument-rate projection under inflation targeting: the Norwegian example. *CEPS Working Paper 127*.

Svensson, L. E. (2010). Inflation targeting. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, pp. 1237–1302. Elsevier.

Svensson, L. E. and R. J. Tetlow (2005). Optimal policy projections. *International Journal of Central Banking 1*(3).

Svensson, L. E. and N. Williams (2007). Bayesian and adaptive optimal policy under model uncertainty. *NBER Working Paper No. 13414*.

Svensson, L. E. and N. Williams (2008). Optimal monetary policy under uncertainty in DSGE models: a Markov jump-linear-quadratic approach. *NBER Working Paper 13892*.

Svensson, L. E. and M. Woodford (2003). Indicator variables for optimal policy. *Journal of Monetary Economics 50*(3), 691–720.

Svensson, L. E. and M. Woodford (2004). Indicator variables for optimal policy under asymmetric information. *Journal of Economic Dynamics and Control 28*(4), 661–690.

Taylor, J. B. and H. Uhlig (Eds.) (2016). *Handbook of Macroeconomics* (1 ed.), Volume 2 of *Handbooks in Economics*. Amsterdam: Elsevier.

Taylor, J. B. and V. Wieland (2012). Surprising comparative properties of monetary models: results from a new model database. *Review of Economics and Statistics 94*(3), 800–816.

Taylor, J. B. and J. C. Williams (2010). Simple and robust rules for monetary policy. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3 of *Handbook of Monetary Economics*, pp. 829–859. Elsevier.

Tetlow, R. J. (2018). Real-time model uncertainty in the United States: “robust” policies put to the test. *International Journal of Central Banking*.

Tucker, P. (2004). Risk, uncertainty and monetary policy regimes. Speech given at the UK Asset and Liability Management Association.

Tucker, P. (2006). Uncertainty, the implementation of monetary policy and the management of risk. Speech given at the Association of Corporate Treasurers.

Tuckett, D., D. Holmes, A. Pearson, and G. Chaplin (2020). Monetary policy and the management of uncertainty: a narrative approach. *Bank of England Staff Working Paper No. 870*.

Tuckett, D. and M. Nikolic (2017). The role of conviction and narrative in decision-making under radical uncertainty. *Theory & psychology 27*(4), 501–523.

Warsh, K. (2014). Transparency and the Bank of England’s Monetary Policy Committee. Review commissioned by the Bank of England.

Warsh, K. M. (2016). Institutional design: Deliberations, decision, and committee dynamics. In J. Cochrane and J. Taylor (Eds.), *Central Bank Governance & Oversight Reform*. Hoover Institution, Stanford University.

Watson, J. and C. Holmes (2016). Approximate models and robust decisions. *Statistical Science*, 465–489.

Weale, M., A. Blake, N. Christodoulakis, J. E. Meade, and D. Vines (1989). *Macroeconomic Policy: Inflation, Wealth and the Exchange Rate.* Unwin-Hyman.

Williams, J. C. (2013). A defense of moderation in monetary policy. *Journal of Macroeconomics 38*, 137–150.

Williams, J. C. (2025). Uncertainty and robust monetary policy. Remarks at the Reykjavík Economic Conference, Reykjavík, Iceland.

Woodford, M. (2003). *Interest and prices: foundations of a theory of monetary policy.* Princeton University Press.

Yellen, J. L. (2012). Perspectives on monetary policy. *speech at the Boston Economic Club Dinner, Boston, MA, June 6*, 2014–27.

A The two-dimensional epistemic uncertainty space

Figure 1 from Section 2, reproduced below without the historical reference points, depicts a conceptual epistemic uncertainty space that describes uncertainty along two dimensions: (i) the extent of understanding about the range of potential outcomes as relevant for monetary policy decisions (horizontally); (ii) the strength of the basis for attaching probabilities to those outcomes (vertically). This appendix discusses the diagram in more detail and relates the discussion of monetary policy under uncertainty in the main text to a broader discussion of uncertainty.

Figure A.1: Two-dimensional epistemic uncertainty space

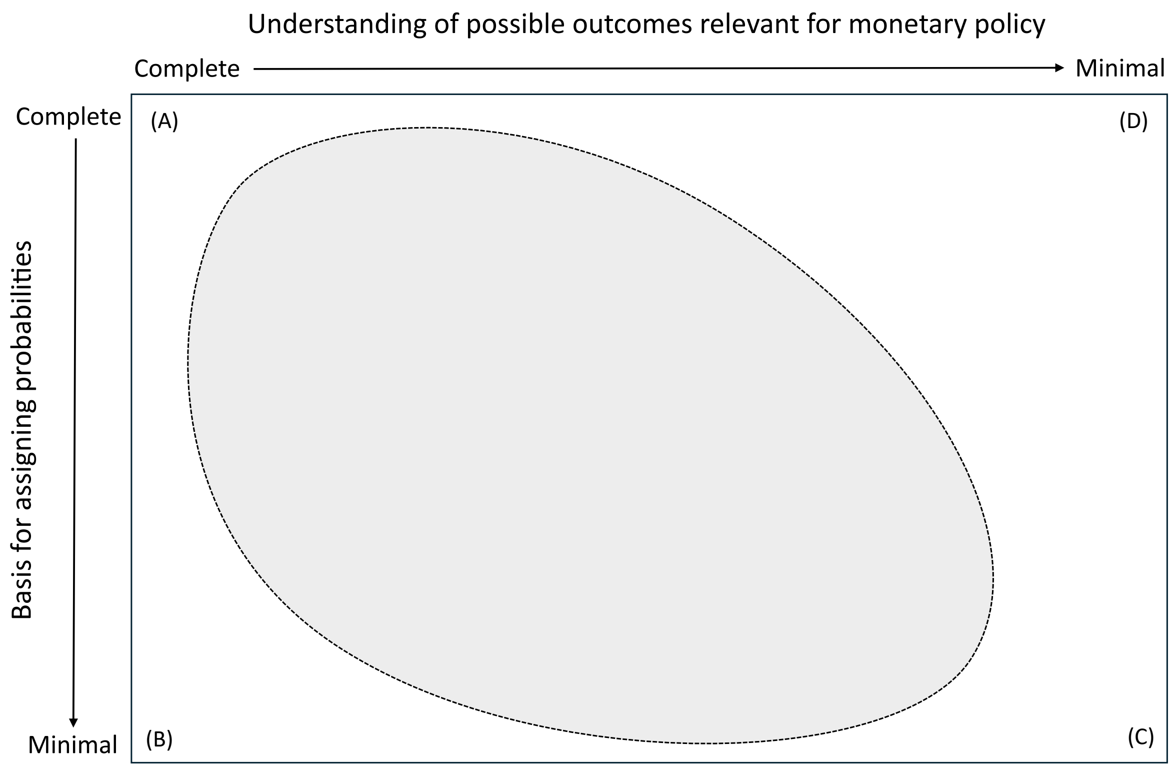

Notes: This uncertainty space is adapted from Spiegelhalter (2024, Chapter 13), who in turn credits Stirling (2010).

The uncertainty space covers a range of possibilities from complete understanding of the joint probability distribution of outcomes in corner (A) to a position of extreme ignorance, sometimes known as very "deep uncertainty", in corner (C). Corners (B) and (D) are intermediate cases, in which the policymaker has perfect knowledge of outcomes but a very weak basis on which to attach probabilities to those outcomes or vice versa.

As subjectively depicted by the shaded ovoid, monetary policy typically operates in the grey area between these extremes. That is, monetary policymakers neither have perfect knowledge, nor do they operate in an environment of very deep uncertainty. As discussed in the main text, the extent and nature of uncertainty faced by monetary policymaker varies over time, as the economic environment varies. There are periods where policymakers may be more confident in specifying potential outcomes and attaching probabilities to those, but there are also periods where that is likely to be more difficult.[86]

The shape of the ovoid reflects the observation that strength in understanding of possible outcomes tends to go hand-in-hand with strength in the basis for attaching probabilities to those outcomes. For example, to the extent that structural or econometric models are informative, they would contain some information about both outcomes and probabilities. Indeed, the limiting case of corner (A) is the hypothetical case in which the model of the economy is known. Corner (C), on the other hand, could be described as a position in which models are essentially completely uninformative.

The shape of the ovoid also reflects that it is arguably easier to imagine situations in which more is known about the range of possible outcomes than the probabilities of those outcomes than the obverse of that. Put differently, it is easier to imagine situations closer to corner (B) than corner (D). In particular, there are situations in which economic theory or history provides some basis on which to understand the range of potential outcomes, but econometric evidence is not informative about the likelihood of those outcomes. For example, following Russia's invasion of Ukraine at a time when the global economy was still adjusting to the economic effects of the Covid pandemic, rapid and large increases in global commodity and goods prices pushed inflation rates across the world to levels not seen since the 1970s and 1980s, periods with very different macroeconomic and monetary policy regimes. In those circumstances, it was reasonable to suppose that price and wage setting could adapt to materially higher inflation in a way that made high inflation more persistent. However, since inflation had remained well below these levels during the entire period over which present monetary policy regimes have operated, econometric evidence was relatively uninformative about the potential for this to occur.

Spiegelhalter (2024, Chapter 13) associates the quadrant nearest corner (D) with situations in which it is only possible to specify a subset of the possible outcomes. In such cases, it is always possible to create a 'catch-all' category, like 'other', and assign it a probability. As an example, Spiegelhalter (2024) cites the Bank of England's fan charts – visual representations of density projections for inflation, GDP growth and unemployment – in which 10% of the probability mass is reserved for possible outcomes in the blank space above the top and below the bottom of the shaded fans. In this way, the probability of those outcomes is specified, but the outcomes themselves are not. While this is a useful example of where it might be helpful to attach a subjective probability to a potentially large range of possible outcomes without specifying what those outcomes are,

[86] As noted by Spiegelhalter (2024, Chapter 13), "There is no sudden jump into deep uncertainty; we wade in along a continuum in which there are increasing difficulties in both specifying possible outcomes and assessing possibilities through judgement and models."

it can reasonably be argued that this situation is not very close to corner (D), given that outcomes covering 90% of the distribution are well-specified. Perhaps more importantly from the perspective of Figure A.1 and the discussion in this paper, it is not clear how much relevance, if any, potential outcomes in the blank space outside the fan charts have for monetary policymaking in practice.

This discussion can also be related to the discussion in Faust (2005), who frames his analysis of applied monetary policy via an assessment of the potential for macroeconomic models to quantify joint probability distributions in a way that would mean that little judgement is required to bridge from the model-based policy solution to real-world policymaking. In the framework depicted in Figure A.1, this would be equivalent to an assessment that the environment is close to corner (A) *and* that what is known about the world and the way the economy works can be encapsulated in a model. Faust's (2005) conclusion is that this is not an accurate description of the state of affairs:

> In stark form, all our policy analysis models are grossly deficient relative to the ideal and will be so for the indefinite future [...] By ideal here I mean roughly that the characterization of the business cycle and policy problem embedded in the models is sufficiently good that little judgment is required in getting from the model to an arguably optimal outcome in practice. (Faust, 2005).

That assessment could be supported by a view that the underlying state of knowledge is not sufficiently close to the ideal and/or that there are material practical constraints on the extent to which what is known about the economy can be captured in any model.

Section 2 includes a brief discussion of different approaches to forming quantitative solutions to simplified versions of the monetary policy problem from the literature that might be associated with particular quadrants of Figure A.1:

- Near corner (A): Bayesian optimal policy with textbook optimal policy as a special case and with or without risk management arising from particular forms of known uncertainties. See Appendix B.1, B.2 and B.3 for details and further discussion.
- Near corner (B): Robust control, robust policy rules assessed from known models. See Appendix B.4
- Near corner (C): Friedman-esque rules, narratives.
- Near corner (D): There are no methods that are obviously applicable, which probably reflects the fact that uncertainty in the real world – as applicable to monetary policymaking – is less likely to be well-described by this zone. There are variations

of Bayesian and robust control analysis that seek to acknowledge that the true model is not in the set under consideration in more or less informal ways. As a formal example, the Cerreia-Vioglio et al. (2025) approach described in Appendix B.4 does not presume that the policymaker's reference model set includes the true model. As an informal example, Cromwell's rule applied to Bayesian statistics states that it is often sensible to retain a small probability for events outside of the specified set, often generically labelled as 'other' (Spiegelhalter, 2024).

These methods developed for monetary policymaking under uncertainty in different contexts can be related to more general discussions of approaches to dealing with uncertainty, such as those in Stirling (2010) and Spiegelhalter (2024, Chapter 13). There is, however, an important distinction between general analytical approaches to dealing with uncertainty and monetary policy decision-making under uncertainty. The dynamic control problem of monetary policymakers is fundamentally different to the problem of how to quantify or approach uncertainty in other contexts in which the actor has no ability to affect outcomes. For example, monetary policymakers have the potential capacity to alter economic outcomes in the future across the entire distribution. This means that forecasting cannot be separated from monetary policy decision-making.[87] The same problem does not generally arise for other economic forecasters.

There is also an important distinction between monetary policymaking under uncertainty and more general theories of decision-making under uncertainty. In particular, a careful in-principle distinction should be drawn between the ways in which people make decisions as part of their daily lives and the ways in which decisions are made by a group of experts employed to make those decisions and who are held accountable for them. For example, given limited time, the costs of cognition or the costs of acquiring information, it may be optimal for people to use a subset of the available information and heuristics in their daily decision-making. While there is no single framework that monetary policymaker can use, it is reasonable to expect them to be more thorough and deliberate in their approach.[88]

[87] Under some assumptions (discussed in Section 3.1), it is possible to separate the forecast problem from the decision problem. But these assumptions only hold in corner (A) of Figure A.1 and so are not directly applicable in the grey area of uncertainty in which monetary policy operates. See Section 3.1 for further discussion.

[88] For example, Sims (2001) argues that "The criteria for acceptable shortcuts in decision-making by a central bank should generally be much stricter than those applying to, say, a consumer buying a new washing machine."

B Model-based optimal control

This Appendix reviews approaches from the literature based on model-based optimal control. These approaches are developed for environments with properties that are close to corner (A) in Figure 1 and therefore represent simpler and 'easier' environments than those typically facing monetary policymakers.

The baseline case is arguably a 'Bayesian hyper-model' which resides close to the boundary of the 'textbook theory' area in Figure 3. That is because the true model is not known with certainty during the Bayesian learning process, though will in general be 'learnable' (so that the learning process converges on point (A) eventually). However, for simplicity discussion of the Bayesian case is deferred to Appendix B.3 and the initial focus is an exposition of the (simpler) nested case of a single model that is known with certainty.

B.1 A basic linear quadratic framework

The intention is to provide a brief, but relatively self-contained, description of a standard approach that has been used for over four decades. It draws very heavily on Harrison and Waldron (2021), who in turn start from the structure set out in Burgess et al. (2013).

Our starting point is a discrete time, infinite horizon linear(ized) rational expectations model. The discrete 'period' or 'date' is indexed by $t = 1, \dots, \infty$ and the model is written in the following form:

$$H^F \mathbb{E}_t x_{t+1} + H^C x_t + H^B x_{t-1} = \Psi z_t \tag{B.1}$$

where x_t is a $n_x \times 1$ vector of endogenous variables, z_t is a $n_z \times 1$ vector of exogenous shocks, and $\mathbb{E}_t$ represents the mathematical expectation conditional on period-t information. The $n_x \times n_x$ matrices H^F, H^C, H^B and the $n_x \times n_z$ matrix Ψ are coefficient matrices.

The exogenous shocks, z, are draws from a standard normal distribution:

$$z_t \sim N(0, I_{n_z}) \tag{B.2}$$

where I_{n_z} is the $n_z \times n_z$ identity matrix, implying identically and independently distributed shocks.

Optimal policy

The starting point for considering optimal policy is a version of the model in which the policy rules describing the behaviour of the policy instruments are removed from the model (B.1). Doing so gives the following representation:

$$\widetilde{H}^F_{\tilde{x}}\mathbb{E}_t\widetilde{x}_{t+1} + \widetilde{H}^C_{\tilde{x}}\widetilde{x}_t + \widetilde{H}^B_{\tilde{x}}\widetilde{x}_{t-1} + \widetilde{H}^F_r\mathbb{E}_t r_{t+1} + \widetilde{H}^C_r r_t = \widetilde{\Psi}_{\tilde{z}}\widetilde{z}_t \tag{B.3}$$

where $\widetilde{x}$ denotes the $n_{\tilde{x}} \times 1$ vector of non-policy endogenous variables, r is the $n_r \times 1$ vector of policy instruments and $\widetilde{z}$ denotes the $n_{\tilde{z}} \times 1$ vector of non-policy shocks.[89] Without loss of generality, lags of the policy instruments, r, are excluded from (B.3).[90] This ensures that the instruments are not state variables in the system, thereby simplifying the form that the solution takes.

For completeness, the equation(s) for the policy instrument(s) that are removed from (B.1) are:

$$r_t = \widehat{H}^F\mathbb{E}_t\widetilde{x}_{t+1} + \widehat{H}^C\widetilde{x}_t + \widehat{H}^B\widetilde{x}_{t-1} + \widehat{\Psi}z_{r,t} \tag{B.4}$$

where $z_{r,t}$ are the exogenous shocks to the policy equations.

A linear-quadratic optimal policy problem implies that the instruments r are chosen to minimize a quadratic loss function given by:

$$\mathcal{L}_t \equiv \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left\{ (\widetilde{x}_{t+i})' W (\widetilde{x}_{t+i}) + (r_{t+i})' Q (r_{t+i}) \right\} \tag{B.5}$$

where $0 < \beta < 1$ is the discount factor and W and Q are $n_{\tilde{x}} \times n_{\tilde{x}}$ and $n_r \times n_r$ positive semi-definite weighting matrices.

The optimal policy problem is to minimize (B.5) subject to (B.3).

A further (possible) constraint is whether or not policy must be set in a time-consistent manner: if so, the optimal policy is a 'discretionary' equilibrium; if not, the solution represents the 'commitment' equilibrium. Without loss of generality the remainder of the exposition assumes time-consistent policy.[91]

[89]To move from (B.1) to (B.3), the n_x $(= n_{\tilde{x}} + n_r)$ variables, x, are partitioned into policy instruments, r and non-policy variables $\widetilde{x}$. The n_r equations that describe the behavior of the policy instruments are removed, leaving a system of $n_{\tilde{x}}$ equations. The matrices $\widetilde{H}^F_{\tilde{x}}, \widetilde{H}^C_{\tilde{x}}, \widetilde{H}^B_{\tilde{x}}$ are $n_{\tilde{x}} \times n_{\tilde{x}}$ coefficient matrices formed by extracting the relevant rows and columns from H^F, H^C, H^B. The $n_{\tilde{x}} \times n_r$ coefficient matrices $\widetilde{H}^F_r$ and $\widetilde{H}^C_r$ are constructed analogously. The $n_{\tilde{x}} \times n_{\tilde{z}}$ matrix $\widetilde{\Psi}_{\tilde{z}}$ is found by removing the rows of Ψ corresponding to the policy equations and the columns corresponding to any policy shocks (that appear solely in the policy equations).

[90]This is not a restrictive assumption since any model that does contain lags of the instrument(s) can be rewritten (by introducing appropriate identities if necessary) so that no instrument lags appear.

[91]Dennis (2007) and Harrison and Waldron (2021) also consider optimal commitment policies. The analogues of the equilbrium conditions shown in the main text have a similar form, but respond to a vector that includes a vector of 'co-states' in addition to the state vector $\widetilde{x}_{t-1}$. The co-states represent the Lagrange multipliers on the optimal commitment problem and encode the effects of past policy commitments on current behaviour.

For optimal time-consistent policy, the equilibrium satisfies:

$$\widetilde{x}_t = B_{\tilde{x}\tilde{x}}\widetilde{x}_{t-1} + \Phi_{\tilde{x}\tilde{z}}\widetilde{z}_t \tag{B.6}$$

$$r_t = B_{r\tilde{x}}\widetilde{x}_{t-1} + \Phi_{r\tilde{z}}\widetilde{z}_t \tag{B.7}$$

These equations can be combined (by stacking them) to give a solution for the complete state vector, x:

$$x_t = Bx_{t-1} + \Phi\widetilde{z}_t \tag{B.8}$$

The derivation of the equilibrium dynamics of the system makes use of the first order condition of the optimisation problem:

$$Qr_t - \left(\Theta^{-1}\widetilde{H}_r^C\right)' (W + \beta V_{\tilde{x}\tilde{x}})\, \widetilde{x}_t \quad = \quad 0 \tag{B.9}$$

where

$$\Theta = \widetilde{H}_{\tilde{x}}^C + \widetilde{H}_{\tilde{x}}^F B_{\tilde{x}\tilde{x}} + \widetilde{H}_r^F B_{r\tilde{x}} \tag{B.10}$$

and

$$V_{\tilde{x}\tilde{x}} = (B_{\tilde{x}\tilde{x}})' \, W B_{\tilde{x}\tilde{x}} + (B_{r\tilde{x}})' \, Q B_{r\tilde{x}} + \beta\,(B_{\tilde{x}\tilde{x}})' \, V_{\tilde{x}\tilde{x}} B_{\tilde{x}\tilde{x}} \tag{B.11}$$

Equation (B.9) relates the optimal choice for the instruments to the optimal choice for the endogenous variables.[92] This representation of the first order condition for policy – sometimes called the 'targeting criterion' – can often provide more intuition about the nature of optimal policy behaviour.[93] In particular, it expresses monetary policy behaviour in terms of the endogenous variables within the model and only features the policy instrument directly if the loss function specifies that Q is non-zero.

Inspection of equations (B.9)–(B.11) reveals that the equilibrium satisfies a fixed point in the matrices $\widetilde{H}_{\tilde{x}}^F, \widetilde{H}_{\tilde{x}}^C, \widetilde{H}_r^F, \widetilde{H}_r^C, B_{\tilde{x}\tilde{x}}, B_{r\tilde{x}}, W, Q, V_{\tilde{x}\tilde{x}}$.[94] Since this system does not include $\widetilde{\Psi}_{\tilde{z}}$, the solution is invariant to the variance of the shocks.[95] This *certainty equivalence* result implies that optimal responses to shocks are unaffected by the degree of uncertainty.

[92] The optimal choice for the instruments depends on three effects: (i) a direct effect of the instruments on the contemporaneous period loss, which depends on Q; (ii) an indirect effect of the instruments on the contemporaneous period loss via its effect on the endogenous variables, which depends on W; (iii) an indirect effect of the instruments on the discounted sum of expected future losses, which depends on $\beta V_{\tilde{x}\tilde{x}}$. This final term captures the weights on expected future state variables in determining the loss in period $t+1$.

[93] When anticipated disturbances are non-zero, the targeting criterion becomes dependent on the values of those disturbances, as shown by Harrison and Waldron (2021).

[94] Computation of the equilibrium typically involves an iterative scheme (Dennis, 2007; Harrison and Waldron, 2021).

[95] Recall that the model is written under the assumption that the elements of $\widetilde{z}_t$ have unit variance, so the standard deviation of shocks is encoded withing the $\widetilde{\Psi}_{\tilde{z}}$ matrix.

Model-based filtering and projection

The (true) state vector is related to a set of observable data, y, according to a measurement equation:

$$y_t = D + Gx_t \tag{B.12}$$

where D is a vector of constants (measurement error is considered in Appendix B.2).

The assumption that the state is not directly observable means that the policymaker must form an estimate of the state ($x_{t|t}$), conditional on the history of observed data $\{y_s\}_{s=1}^{t}$. The Kalman filter exploits information about the structure of the economy (i.e., the model) and the measurement equations to construct the best estimate of the state, x, given the observable data, y. So the assumptions about the structure of the economy – the model – will affect the estimates of the states.

Forecasts can be computed by projecting forward the state and shock estimates using (B.8):

$$\mathbb{E}_t x_{t+h} = B^h x_{t|t} = B^h \left(Bx_{t-1|t} + \Phi\widetilde{z}_{t|t}\right) \tag{B.13}$$

Forecasts for the instruments can be found by extracting the relevant rows of the forecasts for the entire state vector or by applying the same logic as in equation (B.13) using the partitioned solution matrices, $B_{r\tilde{x}}$ and $\Phi_{r\tilde{z}}$. These results can be used to consider the revision to the expected policy instrument:

$$r_{t|t} - \mathbb{E}_{t-1} r_t = \underbrace{B_{r\tilde{x}}\left(\widetilde{x}_{t|t-1} - \widetilde{x}_{t-1|t-1}\right)}_{\text{Effects of revision to state estimate}} + \underbrace{\Phi_{r\tilde{z}}\widetilde{z}_{t|t}}_{\text{Effect of new shocks}} \tag{B.14}$$

The innovation in the optimal policy setting (relative to that expected in the previous period) can thus be decomposed into a term that records the effect of the revision to the state estimate given updated information and a term that records the effect of newly arrived shocks.

The second term can be further decomposed:

$$\Phi_{r\tilde{z}}\widetilde{z}_{t|t} = \sum_{j=1}^{n_{\tilde{z}}} \Phi_{j,r\tilde{z}}\widetilde{z}_{j,t|t} \tag{B.15}$$

where $\Phi_{j,r\tilde{z}}$ and $\widetilde{z}_{j,t|t}$ denote the j-th column and j-th row of $\Phi_{r\tilde{z}}$ and $\widetilde{z}_{t|t}$ respectively. This decomposition expresses the effect of newly arrived shocks on the updated optimal policy setting in terms of the contributions of each individual shock.

The same approach can be applied to the update to the estimated state vector by iterating (B.8) backward in time:[96]

$$B_{r\tilde{x}}\left(\widetilde{x}_{t|t-1}-\widetilde{x}_{t-1|t-1}\right)\approx B_{r\tilde{x}}\sum_{j=1}^{n_{\tilde{z}}}\sum_{s=1}^{t-1}B_{\tilde{x}\tilde{x}}^{s-1}\Phi_{j,\tilde{x}\tilde{z}}\left(\widetilde{z}_{j,t-s|t}-\widetilde{z}_{j,t-s|t-1}\right) \tag{B.16}$$

Plugging (B.15) and (B.16) into (B.14) gives the following decomposition, which demonstrates the equivalence of the news-based perspective with textbook optimal policy under conditions of corner (A) in Figure 3, as discussed in Section 3.2:

$$r_{t|t}-r_{t|t-1}\approx\underbrace{B_{r\tilde{x}}\sum_{j=1}^{n_{\tilde{z}}}\sum_{s=1}^{t-1}B_{\tilde{x}\tilde{x}}^{s-1}\Phi_{j,\tilde{x}\tilde{z}}\left(\widetilde{z}_{j,t-s|t}-\widetilde{z}_{j,t-s|t-1}\right)}_{\text{Revised estimate of effects of past shocks}}+\underbrace{\sum_{j=1}^{n_{\tilde{z}}}\Phi_{j,r\tilde{z}}\widetilde{z}_{j,t|t}}_{\text{Effect of new shocks}} \tag{B.17}$$

B.2 Extensions

This section provides a (far from exhaustive) summary of extensions to the simple baseline case considered in Appendix B.1.

B.2.1 Parameter uncertainty

The baseline case assumes that the parameters of the model (B.3) – that is, the matrices $\widetilde{H}_{\tilde{x}}^{F}, \widetilde{H}_{\tilde{x}}^{C}, \widetilde{H}_{\tilde{x}}^{B}, \widetilde{H}_{r}^{F}, \widetilde{H}_{r}^{C}, \widetilde{\Psi}_{\tilde{z}}$ – are known with certainty. If instead the parameters are considered as random variables, then the certainty equivalence property of the baseline linear quadratic case no longer holds.

Brainard (1967) considers a simple (static) linear quadratic example in which the coefficient linking the policy instrument and the (single) goal variable (the 'policy multiplier') is a random variable. Accounting for this 'multiplier uncertainty' reveals that the policy instrument should be adjusted by less in response to shocks, compared to a case in which the best estimate of the multiplier is treated as if it is known with certainty. The result reflects Jensen's inequality given the convexity of the quadratic loss function with respect to the uncertain policy multiplier.

In richer settings, the loss function may or may not be convex with respect to the uncertain parameter(s) and so the 'Brainard conservatism' result may or may not hold in those cases (Söderström, 2002; Sala et al., 2008; Williams, 2013).[97] Appendix B.3

[96] Equation (B.16) is an approximation because it omits the contribution of the revision to the estimate of the initial state vector $\widetilde{x}_{0|t}-\widetilde{x}_{0|t-1}$. Since that revision is pre-multiplied by $B_{\tilde{x}\tilde{x}}^{t}$, and the eigenvalues of $B_{\tilde{x}\tilde{x}}$ lie within the unit disk, its contribution to the state estimate revision in period t is typically very small for a sufficiently large sample of data, t.

[97] Even under certainty equivalence, as in the baseline case, the mapping from the structural matrices

considers parameter uncertainty from a Bayesian perspective.

B.2.2 Non-linearities and discontinuities

Linear models such as that considered in Appendix B.1 are popular because they are very tractable and allow for relatively easy analysis of probabilities and risks (the expectations operator is a linear operator). This is a simple implication of the certainty equivalence result.

In comparison, non-linear models are typically less tractable which implies that the 'curse of dimensionality' emerges quite rapidly: only relatively small non-linear rational expectations models can be reliably solved. That said, progress is being made on non-linear solution methods that can be applied to larger-scale models typically used in central banks (see Brumm and Scheidegger, 2017; Druedahl and Jørgensen, 2017; Maliar et al., 2019, for example)

In some important cases, piecewise linear approximations may be sufficiently accurate as demonstrated by Guerrieri and Iacoviello (2015). Harrison and Waldron (2021) develop piecewise linear techniques to study optimal policy (under both commitment and discretion) in the presence of occasionally binding constraints on both the policy instrument(s) and other endogenous variables.

However, the above methods embed a variant of certainty equivalence (often termed 'perfect foresight') and therefore omit the implications of the *risk* of encountering non-linearities or discontinuities for expectations and hence decisions. Fully non-linear solutions of simple monetary policy models indicate that these effects may be substantial in the vicinity of non-linearities and discontinuities (Adam and Billi, 2006, 2007; Nakov, 2008).[98] Evans et al. (2016) interpret the effects of uncertainty in the vicinity of the zero lower bound through the lens of a 'risk management' strategy, though the effects of non-linearities and discontinuities on policy behaviour are likely to be specific to both the model and context under consideration.

In the presence of non-linearities, one approach to studying optimal monetary policy is Ramsey policy. In this case, the policymaker maximises household utility subject to a set of non-linear implementation constraints introduced by the equilibrium conditions

$\widetilde{H}^F_{\tilde{x}}, \widetilde{H}^C_{\tilde{x}}, \widetilde{H}^B_{\tilde{x}}, \widetilde{H}^F_r, \widetilde{H}^C_r, \widetilde{\Psi}_{\tilde{z}}$, to the optimal response coefficients is non-linear and generally does not have a closed form solution. For DSGE models, the structural matrices themselves are typically non-linear functions of a vector of 'deep' parameters, θ, and so the effects of the non-linearity of the mapping for the implications of parameter uncertainty for the loss function, and hence optimal policy behaviour, are further compounded. Hence the implications of parameter uncertainty for optimal policy are typically heavily dependent on the precise details of the model.

[98] Seneca (2020) explores the effects of changes in uncertainty (risk shocks) on optimal time-consistent policy within the same framework.

of the model under consideration. In this set up, the policymaker is an expected utility maximiser, and would exhibit risk aversion consistent with societal preferences. Therefore, given the uncertainty associated with the realisation of future shocks, this would imply some form of precautionary (i.e. risk management-like) behaviour. For example, Karadi et al. (2024) adopt this approach in the context of non-linearities in the Phillips curve.

In general non-linear treatments of the policy problem, it typically remains the case that optimal policy is characterised by a function mapping the state to the policy instrument. However, the optimal feedback rule analogous to (B.7) is (a) non-linear and (b) uncertainty dependent (i.e., the function is not invariant to the stochastic properties of the exogenous shocks).

B.2.3 Imperfect and asymmetric information

Svensson and Woodford (2003, 2004) consider optimal policy when agents in the economy have limited information about the true state vector. In particular, they consider cases in which a subset of state variables are observed with measurement error:

$$y_t = D + Gx_t + V\omega_t \tag{B.18}$$

which generalises (B.12) by adding iid standard normal measurement errors, ω.

If the optimal policy is time consistent, a certainty equivalence principle applies. This means that the optimal feedback coefficients for the instruments in equation (B.7) and hence the equilibrium dynamics of the state vector (B.6) are unaffected by the measurement equation (B.18). However, optimal outcomes themselves are affected by the presence of data uncertainty because the optimal estimates of the state vector depend on the set of variables that are observed and the stochastic properties of the measurement errors attached to them. Under commitment, the filtering and control problems are no longer independent. So the structure of the observable information set affects the coefficients of the optimal feedback rule.

These general results apply both in cases the policymaker and private agents share the same limited information (Svensson and Woodford, 2003) and in cases in which private agents are fully informed (Svensson and Woodford, 2004). However, there are no robust general results for how data uncertainty affects optimal policy behaviour. In other words, the optimal policy response to particular shocks could be stronger, weaker, more or less gradual depending on the specific form of data uncertainty and the precise structure of the underlying model.

B.2.4 Structural change

There is a relatively rich literature on DSGE models that incorporate Markov switching parameters, often built on the Farmer et al. (2009, 2011) solution method. Much of that literature considers Markov-switching behaviour on the part of policymakers following simple rules, inspired by Davig and Leeper (2007). Those applications capture switches in policy regime and hence a related form of policy uncertainty. However, Kulish and Pagan (2017) develop a solution algorithm for linear rational expectations models with structural change, potentially in the wider economy (rather than policy behaviour).

Blake and Zampolli (2011) study optimal time-consistent policy when model parameters follow a Markov process. They show that the equilibrium dynamics have a similar form to (B.6), with coefficient matrices dependent on the parameter 'regime'. Svensson and Williams (2008) study (timeless perspective) optimal commitment policies in a similar setting. In both cases, certainty equivalence no longer holds in the sense that uncertainty about regime switching affects optimal policy behaviour.

B.2.5 Non-rational expectations

In forward-looking macroeconomic models, it is standard to assume 'rational' (or model-consistent) expectations. Indeed, the baseline approach sketched out in Appendix B.1 employs the same assumption.

While rational expectations are the default assumption, there are extensive research agendas exploring alternatives. This partly reflects the sheer number of ways to plausibly deviate from rational expectations – indeed Sargent (1999) refers to the "wilderness of bounded rationality".

Some alternative assumptions can be incorporated within the baseline framework set out in Appendix B.1 since the resulting models can be written in the same form, as in (B.1). For example, adaptive expectations in which the expectations of (some) agents are determined by time-invariant linear functions of lagged endogenous variables can typically be encoded into (B.1). Similarly, deviations from rational expectations that involve multiplicative distortions to rational expectation operators, such as that developed by Gabaix (2020), can also be directly incorporated.

However, a popular deviation from rational expectations is to assume that agents (possibly including the monetary policymaker) engage in a dynamic learning process. Such a learning process captures the idea that agents may have less than complete knowledge of the economic environment, but seek to adopt expectation formation approaches that are consistent with the data that they observe (Evans and Honkapohja, 2001).

Even within a linear framework, dynamic learning typically implies that the model coefficients (i.e., $\widetilde{H}_{\tilde{x}}^{F}, \widetilde{H}_{\tilde{x}}^{C}, \widetilde{H}_{\tilde{x}}^{B}, \widetilde{H}_{r}^{F}, \widetilde{H}_{r}^{C}, \widetilde{\Psi}_{\tilde{z}}$ in (B.3)) are not fixed, but instead vary over time. As a result, the baseline approach cannot be applied in these cases.

Indeed, time-variation in the model structure complicates the optimal control problem considerably and typically only relatively simple models have been studied.[99] For example, Molnár and Santoro (2014) and Eusepi and Preston (2018) study optimal policy in the textbook New Keynesian model, finding that private sector learning creates an intertemporal tradeoff between stabilising current and future economic conditions. Gaspar et al. (2006, 2010) consider learning about other aspects of the economy, including inflation persistence and the central bank's inflation target.

Other specific deviations from full information rational expectations may be motivated by particular features of the economic environment being modelled. For example, in the regime-switching context discussed in Appendix B.2.4, Blake and Zampolli (2011, Section 7) examine optimal time-consistent policy in cases in which the policymaker and private agents have different beliefs about the regime transition probabilities.

B.3 Bayesian hypermodels

This Appendix formalises a Bayesian approach to model uncertainty using what Hansen and Sargent (2008) refer to as a 'hypermodel', formed by combining a group of distinct models ('sub models'). Combination is by probability weighting the models, with probabilities for each model updated according to Bayes's rule. Implementing such an approach in practice is naturally more complicated than with a single (fixed) model. But, *in principle*, the nuts and bolts of the single model approach are preserved, precisely because the Bayesian approach combines several models to form a single 'hypermodel'.

The implications for optimal control are therefore, in principle, relatively straightforward: "forming a hypermodel would allow the decision maker to proceed with business as usual, albeit with what may be a more complex model and a computationally more demanding control problem" (Hansen and Sargent, 2008, p14).

To fix ideas, consider a finite set of n_M models: $\mathcal{M} = \{M_1, \ldots, M_{n_M}\}$, with $n_M \geq 2$. The policymaker's uncertainty is captured by a $n_M \times 1$ vector, p. The elements of p are the probabilities attached to each model in the set being the true model.

Each sub-model M_i has the same general form as the simple baseline case set out in

[99] Research into the optimal parameter values for *simple* rules is more straightforward and some of the results are discussed in Section 3.3.

Appendix B.1:[100]

$$\widetilde{H}_{\widetilde{x}}^{F,i}(\theta_i)\,\mathbb{E}_t\widetilde{x}_{t+1} + \widetilde{H}_{\widetilde{x}}^{C,i}(\theta_i)\,\widetilde{x}_t + \widetilde{H}_{\widetilde{x}}^{B,i}(\theta_i)\,\widetilde{x}_{t-1} + \widetilde{H}_{r}^{F,i}(\theta_i)\,\mathbb{E}_t r_{t+1} + \widetilde{H}_{r}^{C,i}(\theta_i)\,r_t = \widetilde{\Psi}_{\widetilde{z}}^{i}(\theta_i)\,\widetilde{z}_t \tag{B.19}$$

There are two differences compared with the simple case considered in Appendix B.1. First, the matrices of the model are indexed by i, reflecting the fact that different models have different economic structures. Second, the coefficient matrices depend on a vector of 'deep' parameters, θ.[101] For each model $M_i \in \mathcal{M}$, the vector of parameters is $\theta_i \in \Theta_i$.

This setup implies that there is a hierarchy of uncertainty facing the policymaker. At the highest level, they are uncertain about the true model of the economy. Expectations are computed as a probability-weighted sum over the implications from each model in $\mathcal{M}$. Conditional on a belief about the structure of the economy (i.e., given a specific $M_i \in \mathcal{M}$) there is uncertainty about the value of the parameter vector $\theta_i \in \Theta_i$. Finally, conditional on a specific model M_i and associated parameter vector θ_i there is uncertainty about the shocks $\widetilde{z}$ that generate the observed data.

This hierarchical structure implies that expectations are formed with respect to the relevant conditional densities. For the loss function, for example:

$$\mathbb{E}_t\mathcal{L}_t = \sum_{i=1}^{n_M} p_{i,t} \int_{\theta\in\Theta_i} \mathcal{L}_t\left(M_i, \theta, \widetilde{x}_t\left(M_i, \theta_i\right)\right)\phi_{i,t}\left(\theta\right)d\theta \tag{B.20}$$

where $\phi_{i,t}$ is the policymaker's posterior density of θ_i at date t.

The Bayesian approach implies that all probability distributions are updated using Bayes's rule. For example, the posterior density of θ_i can be constructed by conditioning on M_i being 'true' and applying Bayesian estimation to that sub-model using standard techniques (see, for example, Burgess et al., 2013, Section 6.2.2). The probabilities attached to each model are similarly updated according to:

$$p_{i,t} \equiv p\left(M_i|Y^t\right) = \frac{\ell\left(Y^t|M_i\right)p_{i,0}}{\sum_i \ell\left(Y^t|M_i\right)p_{i,0}} \tag{B.21}$$

where $Y^t \equiv \{y_1, \ldots, y_t\}$ is the sample of observable data, ℓ is the marginal data density

[100] Writing the models in this form requires that $\widetilde{x}$ represents the vector formed by the union of the vectors of 'non-policy' endogenous variables across the n_M models (and analogously for $\widetilde{z}$). Thus if a particular variable does not appear in a particular model the loadings in the columns of the matrices will be zero. To ensure all matrices are conformable, additional equations can be added to set 'missing' variables in a given model to zero. For the policy problem discussed below to be well defined requires that: (a) all models in the set include the variables that appear in the policymaker's loss function; (b) all models in the set include all of the variables with non-zero loadings on the measurement equations that described the observed data.

[101] Referring to θ as a vector of deep parameters is without loss of generality as it could simply reflect a vector of all elements of the relevant matrices (e.g., for a semi-structural or VAR model).

(conditional on a specific model) and $p_{i,0}$ is the policymaker's initial prior probability on M_i.[102]

Characterising the loss function itself $\mathcal{L}_t\left(M_i, \theta, \widetilde{x}_t\left(M_i, \theta_i\right)\right)$ is challenging even when the objective function is quadratic, because of the the Bayesian updating processes (e.g., equation (B.21)). As discussed in Svensson and Williams (2007), this implies that the loss function is no longer quadratic in the state vector (since the terms used to compute expectations on the right-hand side of (B.20) themselves depend on non-linear functions of the state). Convexity or otherwise of the loss function (or the value function derived by Svensson and Williams (2007) using their dynamic programming approach) determines whether or not a Bayesian policymaker has the incentive to experiment to accelerate the learning process.[103]

The non-linearity of the problem leads Svensson and Williams (2007) to solve the optimal policy problem using global solution methods.[104] They also explore a variant of the policy problem in which the policymaker does not internalise the effects of their actions on their future learning (hence eliminating the incentive to experiment) that they call 'Adaptive Optimal Policy'. That approach simplifies the problem by ensuring that value functions are conditionally quadratic, thereby enabling the use of linear quadratic techniques.[105] Under both Bayesian and adaptive optimal policy, beliefs (as encoded in p_t) form part of the state vector in the analogous formulation of equations (B.6) and (B.7).[106]

B.4 Robust control

This Appendix considers robust control following a textbook introductory treatment Hansen and Sargent (2008, Chapter 2), applied to a 'structural form' representation of the model.

Following Dennis et al. (2009), the *approximating* model structure is given by (B.3)

[102] The marginal data density integrates over the posterior density for θ_i, though this dependence is suppressed for notational convenience, instead using the conditioning on M_i to indicate it.

[103] Prospective future losses may be much smaller if the (future) policymaker is certain about the model of the economy. This means that a policymaker today may have an incentive to create near-term losses (e.g., by generating a recession) if doing so accelerates the learning process (e.g., if the models in the set differ markedly in terms of the response of certain variables to a recession).

[104] The cases they consider are simpler in some respects than the approach laid out here since there is no uncertainty about θ_i conditional on M_i and more complicated in other dimensions as they allow the true model to change over time according to a (hidden) Markov process. The variant sketched in this Appendix maps into their framework under the assumption that the Markov transition matrix is an identity matrix (see Svensson and Williams, 2007, Section 5)

[105] This result relies on no uncertainty regarding θ_i conditional on M_i.

[106] Since Svensson and Williams (2007) consider optimal commitment problems, the state vector includes Lagrange multipliers on the policy problem as co-state variables (see the discussion in footnote 91).

and the policymaker entertains a *distorted* variant perturbed by additional disturbances, ξ_t:

$$\widetilde{H}^F_{\tilde{x}}\mathbb{E}_t\widetilde{x}_{t+1} + \widetilde{H}^C_{\tilde{x}}\widetilde{x}_t + \widetilde{H}^B_{\tilde{x}}\widetilde{x}_{t-1} + \widetilde{H}^F_r\mathbb{E}_t r_{t+1} + \widetilde{H}^C_r r_t = \widetilde{\Psi}_{\tilde{z}}\left(\widetilde{z}_t + \xi_t\right) \tag{B.22}$$

The policy problem then becomes

$$\min_{\{r_{t+i}\}_{i=0}^{\infty}} \max_{\{\xi_{t+i}\}_{i=0}^{\infty}} \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left\{ \left(\widetilde{x}_{t+i}\right)' W \left(\widetilde{x}_{t+i}\right) + \left(r_{t+i}\right)' Q \left(r_{t+i}\right) \right\} \tag{B.23}$$

subject to (B.22) and

$$\mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \xi'_{t+i}\xi_{t+i} \leq \eta \tag{B.24}$$

Solving the robust control problem simply amounts to adding an additional set of constraints and first order conditions to the policy problem (i.e., those associated with the maximisation) and solving them alongside the policymaker's first order conditions.[107]

Despite initial appearances, modelling misspecification using a set of additive disturbances is not particularly restrictive in terms of the richness of misspecification that can be entertained (Hansen and Sargent, 2008). That is because the 'misspecification constraint' (B.24) places no direct restriction on the autocorrelation structure of the shocks.

A common narrative description of the problem is to imagine a fictitious 'evil agent' (Hansen and Sargent, 2008) that chooses shocks $\{\xi_{t+1}\}_{t=0}^{\infty}$ to *maximise* policymaker losses, subject to the misspecification constraint, (B.24). An important benefit of the evil agent analogy is that it emphasises that the policymaker is not in control of the 'worst case' misspecification. That is done by the (maximising) evil agent: their role is to make the stabilisation problem as difficult as possible. This observation also suggests that robust control techniques could potentially be used as a model diagnostic tool: to help understand the aspects of approximating models in which misspecification would be particularly damaging (Hansen and Sargent, 2008, Chapter 1.6).

Importantly, the constraint (B.24) implies that the extent of misspecification is limited. One reason for this is that there typically exists some upper bound for η beyond which the evil agent can render the model "uncontrollable" (Hansen and Sargent, 2013), thereby inflicting infinitely large losses on the policymaker: a "breakdown" of the policy problem.[108]

[107] A 'dual' formulation of the problem adds a term in $\xi'_{t+i}\xi_{t+i}$ (with a negative weighting coefficient) to (B.23) (i.e., to form a Lagrangean that should be minimised). The size of the weighting coefficient can then be set to deliver a desired η in equilibrium.

[108] Since the disturbances ξ_t can be functions of the state vector $\widetilde{x}_{t-1}$, it is possible to chose the disturbances so that the reduced form solution is explosive. Therefore, in most models there is a maximal amount of misspecification that can be entertained before the loss function becomes unbounded.

However, in many applications, the extent of misspecification is determined by the idea that a policymaker should not be able to easily detect that their approximating model (B.3) did not generate the data they observe. In that case, η is calibrated to match a desired 'detection error probability' (Anderson et al., 2000): specifically, the probability that a policymaker would correctly realise that the data they observe was generated by the 'worst case' model rather than the approximating model. Indeed, Hansen and Sargent (2008, p.27) use this argument as justification for the assumption that the policymaker does not engage in a learning process to uncover the true model.

B.5 'Open set' approaches

As discussed in Appendix B.3, a Bayesian hyper model approach allows for model uncertainty by positing a set of sub-models. A policymaker's uncertainty is expressed in the form of a probability distribution across sub-models. These probabilities are updated over time according to Bayes's rule, starting from a set of (potentially subjective) initial prior probabilities. Over a (potentially very long) period of time, Bayesian updating will reveal the true model.[109]

While the Bayesian approach allows for uncertainty over sub-models, a key assumption is that the set of models is *complete* in the sense that it contains the true model of the economy: an assumption that defines "M-closed" problems in the context of forecasting. An alternative, "M-open", assumption is that the true model does not lie within the set. Watson and Holmes (2016) use "D-open" to describe decision problems (rather than forecasting problems) in which the decision-maker's model set is incomplete. This 'open set' alternative perspective is, in principle, more applicable to areas of the uncertainty space most relevant for real-world monetary policymaking, illustrated in Figure 1.[110]

A useful starting point is the analysis of M-open forecasting problems. If a Bayesian approach is applied in an M-open setting, in many cases, the probability placed on the best-fitting model will (often quickly) converge to 1, even though that model is not the 'true model' (Geweke, 2010).[111] However, forecasts from the 'best' model (identified in

[109] This implies that individual policymakers with initially different (non-degenerate) prior distributions across models will eventually converge on the same beliefs (i.e., a unit probability on the true model). Learnability of the true model depends on the specification of the measurement equation (B.12), which could in principle be specified in a way that made it impossible to learn the true model, even in the long run.

[110] The Bayesian hyper-model approach arguably addresses three of the five 'layers' of the hierarchy of uncertainty set out by Spiegelhalter and Riesch (2011), but abstracts from more fundamental uncertainties regarding 'acknowledged' and 'unknown' inadequacies of the sub models. This perspective shares some similarities with Faust's (2005) advice to study the implications of "gross deficiencies" of monetary policy analysis (i.e., "acknowledged inadequacies") when considering the implications for practical policy.

[111] Geweke (2010) provides examples in which the model with the smallest Kullback-Leibler distance to the (unknown) true model will asymptotically receive 100% weight. However, this is not a general result

this way) are often outperformed by those from a weighted combination of the forecasts of some or all sub-models (see, for example, Granger, 1989). That is, when all sub-models are misspecified, better forecasts may be obtained by placing some weight on some/many/all of them, rather than choosing the forecasts of the 'best' model.

M-open approaches offer methods to assign and update weights applied to forecasting models within an incomplete set. For example, Geweke (2010) and Geweke and Amisano (2011, 2012) develop methods with desirable properties.[112] However, the assumption that the true model may not (or, more realistically, will not) be within the set of sub-models typically requires moving away from Bayes's rule as a method of forecast combination.[113] There is no unique way to weight models and research continues to develop new (often context specific) weighting approaches. Importantly, the weights attached to each sub-model by such approaches are *not* interpreted as probabilities, since every sub-model is potentially misspecified. Nonetheless, in forecasting problems it may be optimal to place non-zero weights on sub-models that are certainly misspecified.

In comparison with forecasting problems, formal analysis of decision theory in an environment with multiple misspecified models is somewhat less developed. The decision problems analysed by this literature are more stylised than the structural models used for optimal policy analysis with a known model (or hyper-model) such as those set out in earlier appendices. In particular, the models considered typically have very limited dynamics and rarely consider behavioural equations that incorporate expectations.[114]

Nevertheless, some papers in this developing literature suggest that optimal policy behaviour may exhibit some plausible qualities.

Lanzani (2024) suggests that the optimal reaction function is a weighted combination of robust control reaction functions across models. The weights depend on an updating equation that looks very similar to Bayes's rule. The parameters governing the degree of misspecification in the robust control problem applied to each model depend on their congruence to the data.[115] So the policymaker becomes more concerned about the misspecification of a given model if it starts to fit the data less well than usual. These properties of equilibrium behaviour imply that the policymaker's concern for misspecification varies over time, consistent with the notion set out in Section 2 that the nature of

(Grünwald and van Ommen, 2017), so it is possible in some cases that misspecified Bayesian learning does not even converge to the 'best' model within a set.

[112] One is that the true model will be correctly identified, in the (unlikely) event that it *is* contained in the set of models.

[113] Some approaches seek to 'adapt' Bayesian updating (e.g., Grünwald and van Ommen, 2017).

[114] Incorporating rational forward-looking behaviour presumably requires assumptions about the relative information sets of private sector decision makers and the policymaker in an environment in which none of the models are 'true'.

[115] The parameter referred to here is η in Appendix B.4. There is a (typically) different η for each model, which varies over time according to how well each model fits the data.

uncertainty facing monetary policymakers is state-contingent.

Conversely, Ghosh (2024) studies a similar setting and finds that the policymaker gravitates towards an optimal feedback rule tailored to the 'least bad' model. This equilibrium (a Berk-Nash equilibrium, Esponda and Pouzo, 2016) also arises in Lanzani (2024) for certain characterisations of policymaker preferences.

Cerreia-Vioglio et al. (2025) also study a situation in which the policymaker uses a set of models that are regarded as simplified representations of the true data generating process. Their simplicity implies that they are misspecified with respect to the observed empirical properties of the data, but are *structured* in the sense that that they can be used to study the implications of policy actions. The extent of misspecification of the structured models is measured relative to a set of unstructured (empirical) models. This setting allows uncertainty across models to be combined with concerns about misspecification, and both Bayesian and robust control methods are used.

B.6 Implications for monetary policy formulation

This appendix extracts key insights from the preceding discussion for practical monetary policymaking.

As noted in Sections 2 and 3, the baseline theory-based approach outlined in Appendix B.1 represents an uncertainty environment close to corner (A) in Figure A.1 which is somewhat distant from the typical 'grey area' in which monetary policy operates. Therefore, the theory-based approach cannot be *directly* applied to the practical monetary policy problem.

The extensions discussed in Appendix B.2 summarise research agendas that relax some key assumptions of the baseline case. These extensions allow for some aspects of model and data/shock uncertainty and structural departures from certainty equivalence. These research agendas have tended to evolve somewhat separately as each strand of literature typically focuses on extending the baseline case in one 'direction'. That is standard practice for two reasons. First, relaxing one set of assumptions at a time makes it easier to isolate their implications (relative to a well understood stylised benchmark case). Second, computational complexity increases (sometimes dramatically) when these assumptions are relaxed making it substantially harder to consider more general cases in which several assumptions are relaxed.

Of course, in practice *all* of the issues tackled in Appendix B.2 are, at least to some extent, *ever present* in the practical monetary policy problem. Given the complexities involved in incorporating each extension individually, the prospect of developing a "grand unified model" addressing all relevant features of the practical policy problem is at best remote and most likely impossible (Faust, 2005). These considerations apply with even more force to possible extensions of the Bayesian hyper-model discussed in Appendix B.3 given the substantial additional complexity attached to a 'baseline' hyper-model (Hansen and Sargent, 2008). Furthermore, these approaches all embody assumptions that make them applicable near or at corner (A) in Figure A.1 and so they do not acknowledge the epistemic uncertainty that policymakers face.

The practical relevance of the robust control approach discussed in Appendix B.4 depends in part on implementation choices. One implication of the 'detection error probability' approach to calibration is that, almost by construction, the range of model misspecification that can be entertained by the policymaker is relatively 'small'. In some sense, this does not matter much if the approximating model fits the data very well – since there are likely to be many other models that fit well, but have somewhat different implications for the conditional mean and variance of a projection and hence present non-trivial stabilisation challenges.

However, this observation presents some challenges in (directly) applying the robust control approach to cases that lie within the uncertainty 'grey area' in which monetary policy typically operates. Importantly, if the approximating model is to believed be objectively 'bad', it is unclear whether robust control methods can incorporate enough misspecification to be practically relevant.[116]

Some of the extensions considered in Appendix B.5 allow for the arguably more plausible case in which the policymaker entertains a set of specific and distinct candidate approximating models, none of which is believed to be 'true'. These approaches (and subsequent developments to incorporate suitable dynamics) may provide a richer baseline framework for formal analysis of monetary policy within the 'grey area'. At present, however, there is no consensus baseline approach to decision-making when the policymaker relies on multiple misspecified models.

As noted in Section 2, this line of reasoning is consistent with Faust's (2005) view that "all our policy analysis models are grossly deficient relative to the ideal". However, the recognition that none of the available models can be used to directly solve the monetary policy problem should not lead inexorably to the abandonment of model-based analysis – a retreat to "econometric nihilism" (Blinder, 1998).

A more productive approach is to embrace the fact that all models are misspecified and build that into their use in practical monetary policymaking. One such view is set out by Cox (1995):

[116] In the context of Figure A.1, this amounts to an assumption that the practical policy problem lies some way to the right of the left edge of the uncertainty space.

> [I]t does not seem helpful just to say that all models are wrong. The very word "model" implies simplification and idealization. The idea that complex physical, biological or sociological systems can be exactly described by a few formulae is patently absurd. The construction of idealized representations that capture important stable aspects of such systems is, however, a vital part of general scientific analysis and statistical models, especially substantive ones [...] (Cox, 1995).

In other words, "[...] all models are wrong, *but some are useful*" (Box and Draper, 1987, emphasis added). The practical challenge is how to make the best use of the misspecified models at policymakers' disposal.

An obvious response to inevitable misspecification is to use many models. Two lines of reasoning suggest such an approach.

First, given the inherent complexities of including Blinder's "additional complications", there is a case for developing different models that are specialised for particular purposes. Indeed, this is the typical route of progress in most strands of the structural modelling literature in macroeconomics (including, as discussed above, those outlined in Appendices B.2–B.5): modelling effort and complexity is focused on the parts of the model that matter most for the specific question at hand; other aspects are kept as simple (and standard) as possible.

Second, for any particular purpose that models are used, using many models provides a way to gather evidence about the robustness of results across them. As discussed in Section 3.3, this view is embedded within the approach to simple rules in a rules-based perspective on the policy problem. However, it applies more broadly. For example, a forecast-based perspective could include cross-checks using forecasts from a variety of alternative models.

This reasoning tends to lead towards a pluralistic approach to the use of models in the policy process. A wide range of model-based analysis – using different models to address different questions – would be combined with other analysis (including those of the types described in Section 4.2).

The fact that different models are optimised to support different aspects of monetary policy formulation may give rise to apparently conflicting results:

> Given the current deficiencies in our knowledge, the best model from the perspective of any one role taken in isolation is unlikely to be the best from the narrow perspective of the others. Indeed, optimizing the models for these separate roles independently would almost surely give rise to models that are, at least in some respects, mutually inconsistent. (Faust, 2005).

As discussed in Appendix D, case-based (or 'analogical') reasoning can be useful in assessing the implications from alternative model-based analyses for the practical policy problem even if the former is very stylised. In other words, the 'gap' between the model-based analysis and the practical policy problem may be large. In an echo of Cox's (1995) view of models as simplifications of reality, Lucas (2011) describes structural macroeconomic models as "make-believe economic systems", arguing that such models provide "the only way we have found to think seriously about reality", in large part because macroeconomics is not an experimental science. See Appendices D.2 and D.3 for further discussion.

For these reasons, consistent with the examples given in Section 2, formal model-based policy analysis using the methods described in this Appendix can provide valuable, albeit necessarily incomplete, insights for practical policymaking.

C Analytical inputs to support policymaking

The table below summarises a range of analytical inputs that may be used to support policymaker, corresponding to the diagrammatic representation in Figure 6.

Type	Input	Description
Supporting inputs	Official data	Data published by national statistical authorities, other official data providers, financial market data, etc.
	Surveys and other intelligence	Evidence regarding private sector behaviour, views, and expectations gathered from a variety of sources, such as business surveys, central banks' own contacts (e.g. the Bank's Agents' contacts), and so on.
	Scenarios	Economic assessments that relax certain assumptions relative to a base case. They may be sensitivity analysis or fuller articulations of the possible effects on the macroeconomy of some salient risk.
	Empirical / econometric analyses	Data-led, including more formal econometric, analyses of key questions relevant for policy formulation.
	Bespoke model analyses and policy simulations	Model analyses designed to inform a particular policy consideration. These may be more complex than regular policy analysis, may include different modelling approaches to capture key non-linearities and risks, may be forecast-based, but need not be, and could use small models.
	Case studies	Historical and international case studies to extract lessons from situations in which some analogies can be drawn to present circumstances. These may include quantitative and qualitative analysis.
	Literature reviews	Discursive analysis that summarises and applies largely qualitative insights from the academic literature to current policy questions.
	Original research	Staff and policymakers undertake original research, which may ultimately yield insights for policy both while the research is underway and after it has been published.
Intermediate inputs	Nowcasting and nearcasting	Estimates for state of the economy (including both observable and unobservable variables) and short-term forecasts for key variables.
	Forecasts (point and density)	Forecasts for relevant macroeconomic variables. These may be produced with significant amount of expert (policymaker and/or staff) judgement, or may be relatively 'hands-free' model projections.
	Regular policy analysis	Implications of optimal policy projects, simple/standard policy rules (including that used in baseline forecast construction), and robust policy rules. These may be based on a baseline forecast and/or particular scenarios.
	Analysis of key policy considerations	Staff-led synthesis setting out analysis of key considerations for policy formulation.

It is important to note that the list of inputs in the table is not intended to be exhaustive. In addition, as discussed in Section 4, different perspectives will draw on a potentially different set of supporting inputs. In turn, the perspectives will tranform these inputs into intermediate inputs to the ultimate judgement-based policy decision in different ways.

In the context of the Bank of England and the processes supporting the MPC, including ongoing changes to processes and communication in response to the Bernanke Review, it is useful to highlight a number of the inputs in the table. For example, the Bank of England *Quarterly Bulletin* article "Monetary policymaking at the Bank of England in uncertain times" (Bank of England, 2025) provides details on the practical use of scenarios as an input to monetary policy, describing how scenarios can be used to work through the policy implications of different outcomes, explore the diversity of views across the MPC, and communicate a reaction function. In addition, Alati et al. (2025) discuss how the Bank has approached standard policy analysis, conditional on a (judgemental) baseline forecast as well as scenarios. This includes optimal policy projections (OPPs) around the forecast and scenarios, consistent with the sufficient-statistics approach discussed in Section 3.1. Further to the OPPs, Alati et al. (2025) also discuss how it is possible to assess the policy prescriptions from a set of simple instrument rules, conditional on the forecast and/or scenarios under consideration.

D Case-based decision theory and monetary policy formulation

This appendix sets out some potential implications of case-based decision theory (CBDT) for monetary policy formulation.

At the outset, it is useful to reiterate that there is no single or algorithmic way to describe how judgement is formed and policy decisions are made. Consistent with that, the intention is not to suggest that CBDT is a normative prescription for decision making under uncertainty. Instead, the intention is to highlight some potential similarities between the framing of decision problems suggested by the theory and real-world policy making. This in turn suggests that the implications of the theory could be used "in quite an informal way" to structure decision-making and the processes that support it, in an analogous way to the analysis of alternative perspectives on the monetary policy problem presented in Section 3.

Gilboa and Schmeidler (2001) describe the intended scope of their theory with reference to the dominant approach in economics, namely expected utility theory:

> [...] our goal is not to fine-tune expected utility theory as a descriptive theory of decision making in situations described by probabilities or states of the world. Rather, we wish to suggest a framework within which one can analyze choice in situations that do not fit existing formal models very naturally. Our theory is just as idealized as existing theories. We only claim that in many situations it is a more natural conceptualization of reality than are these other theories. (Gilboa and Schmeidler, 2001, Chapter 1).

The relevance of CBDT is therefore specific to the context in which decisions are taken. In particular, the assumptions underpinning expected utility theory are equivalent to those that apply to corner (A) in the uncertainty space shown in Figure 1.[117] While those assumptions never hold literally, they may provide a useful starting point for application "in quite an informal way" to the practical policy problem, in analogous fashion to the role of textbook theory in shaping the alternative perspectives discussed in Section 3.

However, the assumptions underpinning CBDT are more relevant to situations further from corner (A):

> We may thus refine Knight's dichotomous distinction between risk and uncertainty (Knight, 1921) by introducing a third category of structural ignorance:

[117] That is, a decision-maker must be able to conceive of all feasible states, evaluate utility in those states and apply (subjective) probabilities to them.

> “risk” refers to situations where probabilities are given; “uncertainty” – to situations in which states are naturally defined, or can be simply constructed, but probabilities are not. Finally, decision under “structural ignorance” refers to decision problems for which states are neither (i) naturally given in the problem; nor (ii) can they be easily constructed by the decision maker. EUT [i.e., expected utility theory] is appropriate for decision making under risk. [. . .] However, in cases of structural ignorance CBDT is a viable alternative to the EUT paradigm. (Gilboa and Schmeidler, 2001, Chapter 2).

In terms of Figure 1, Gilboa and Schmeidler’s concept of “structural ignorance” is consistent with the “deep uncertainty” and “radical uncertainty” environments discussed by Spiegelhalter (2024) and Kay and King (2020) in the vicinity of corner (C) (with Knight’s (1921) concepts of “risk” and “uncertainty” corresponding to corners (A) and (B) respectively). Indeed, Gilboa and Schmeidler (2001, Chapter 2) argue that “case-based decision theory is probably the most natural description of decision making when the decision problem is amorphous and there are insufficient data to analyze it properly.”

This reasoning suggests that CBDT and the benchmark expected utility theory are complementary approaches that may differ in the circumstances in which they may be most useful:

> To a large extent, EUT and CBDT are not competing theories; they are different conceptual frameworks, in which specific theories are formulated. Rather than asking which one of them is more accurate, we should ask which one is more convenient. (Gilboa and Schmeidler, 2001, Chapter 4).

Indeed, under the same informational assumptions (i.e., in corner (A) of Figure 1) CBDT and expected utility theory can be shown to be equivalent (Matsui, 2000). However, when close to corner (A), expected utility offers a more direct approach to decision-making and CBDT may be inefficient and over-engineered in such environments.[118] Hence, CBDT is arguably more relevant in situations further from corner (A).

Appendix D.1 provides a brief review of the underlying theory. Applications of case-based reasoning in macroeconomics are discussed in Appendix D.2 and potential applications to monetary policy formulation are discussed in Appendix D.3.

[118] This mirrors the discussion of the relative efficiency of forecast-based and news-based approaches when close to corner (A) in Section 3.2.

D.1 Case-based decision theory

This Appendix sets out a brief summary of the theory of case-based decision-making developed by Gilboa and Schmeidler (2001), using a variant that extends Gilboa and Schmeidler (1995).

Following Gilboa and Schmeidler (2001, Chapter 2), a case is defined as a triple (p, a, r), where p is a problem, a is an act, and r is a result. A "problem" refers to a description of, or, equivalently, story about, a choice situation. An "act" is a choice that is made in a particular problem. The "result" is the outcome of that choice. M denotes the memory of such cases – that is, problems that an agent and others have encountered in the past. A decision-making agent has a utility function, $u(r)$, which assigns a numerical value to results and a similarity function, $s(\cdot)$, which maps pairs of cases to the unit interval.[119] Facing a new problem, p, the agent chooses an act, a, to maximise:

$$U(a) = \sum_{r \in R} \sum_{(q,b,t) \in M} s\left((p, a, r), (q, b, t)\right) u(r) \tag{D.1}$$

Equation D.1 says that, for an agent facing a new problem p, the utility from an act is the result of a double summation. The inner summation, evaluated over all cases in the decision-maker's memory, sums the product of the similarity of each case in memory with the current case and the resulting utility, for each combination of action and result applied to the current case. The outer summation aggregates the inner summation over all results and therefore requires the decision maker to be able to imagine the utility of every result (even if it has not been experienced within the memory of cases).[120] The generalised case in which the decision-maker may consider hypothetical results seems best suited to the nature of cases that may be considered relevant by real-world monetary policymakers as discussed in Appendix D.3.

Several important similarities and differences with expected utility theory are worthy of note. First, the similarity function $s(\cdot)$ is subjective and therefore specific to the decision-maker. The standard formulation of expected utility theory also assumes that the probabilities used to evaluate utility from outcomes are subjective (Savage, 1972).[121]

[119] The similarity function is, therefore, is defined over two triples. In equation (D.1), the second triple is denoted (q, b, t).

[120] Gilboa and Schmeidler (2001) use the example of a decision-maker who has experience of a soft drink vending machine and has experienced the case ('soft drink machine','push button A','receive drink A'). The formulation in (D.1) allows the decision-maker, on encountering a sandwich vending machine for the first time, to hypothesise that ('sandwich machine','push button B','get sandwich B') is similar to the drinks machine case despite never having encountered a sandwich machine (and therefore never having experienced the result 'get sandwich B').

[121] Indeed, Gilboa and Schmeidler (2001, Chapter 4) argue that similarity judgements underpin subjective probabilities: "Our ability to discuss counterfactuals in an intelligent manner, as well as our ability

Second, the maximand in (D.1) is cumulative in u, in contrast to expected utility theory in which expected utility is a probability weighted average of utility in each state. This implies that increasing memory is a key mechanism through which learning occurs, in contrast with expected utility theory in which the space of feasible states is known from the outset.[122] Finally, in the extended version of CBDT considered here, 'memory' is interpreted more broadly than an individual's recollections of their own experiences of the results of their own actions. Thus memory may include observations or accounts of the actions of others, the results of hypothetical experiments that may be similar to the case under consideration, and so on. This interpretation of memory is particularly important when considering the types of analysis and evidence that supports monetary policy formulation, as discussed in Appendix D.3.

D.2 Case-based reasoning in macroeconomics

While the case-based decision theory presented in Appendix D.1 provides a formal theory of how decisions might be made (particularly in environments of "deep uncertainty"), it also provides a useful framework for interpreting some aspects of macroeconomic research and how policy implications are derived from such research. In particular, case-based *reasoning* (or "analogical reasoning") is used (explicitly or otherwise) in many areas of economics. In the context of the present discussion, a "case" is any piece of analysis that may have relevant insights for the decision problem at hand.

Perhaps the best-known description of analogical reasoning in macroeconomics is Lucas (2011), which explores the potential implications of changes in the US money supply on US inflation via a thought experiment involving a change in the price of tickets in a hypothetical amusement park:[123]

> To apply the knowledge we have gained about [the hypothetical amusement park], we must be willing to argue by analogy from what we know about one situation to what we would like to know about another, quite different situation. And, as we all know, the analogy that one person finds persuasive, his neighbor may well find ridiculous. (Lucas, 2011).

to assign probabilities to events, relies on our subjective similarity judgments. We therefore conclude that the notion of similarity is primitive. It lies at the heart of probability assignments, as well as at the heart of induction."

[122] In expected utility theory, learning is possible only by changes to subjective probabilities (e.g., by Bayesian updating).

[123] Another well-known example is the Nobel prize-winning paper on adverse selection, which is framed using an analogy of a hypothetical second-hand car market (Akerlof, 1970). Indeed, this type of approach has a long history in economics (Morgan, 2012, Chapter 5). See also Eichengreen (2011) for a discussion of the role of analogical reasoning in applying lessons from economic history.

Viewed from the perspective of the theory in Appendix D.1, Lucas describes both the role of the similarity function in assessing the relevance of the model-based result for real-world policymaking and the fact that the similarity function is subjective. Lucas also acknowledges that the model-based environment is "quite different" to the real world, potentially allowing for a wider range of similarity judgements across individuals. However, with respect to the predictions of the model-based experiment, he also argues that "the clarity with which these effects can be seen is the key advantage of operating in simplified, fictional worlds" (Lucas, 2011). So, conversely, the fact that the simplified thought experiment generates sharp and easily understood implications may make it *easier* for individuals to make a similarity judgement.

Gilboa et al. (2014) argue that these perspectives pervade many aspects of economic research. For example, the continued use of models built on evidently false assumptions makes sense if those models are valued primarily for their capacity to provide useful cases, rather than the accuracy with which they describe reality. Another example is the value placed on axiomatic foundations for key model components. Axiomatic foundations are not useful because they extend the results or predictions of an extant model, but rather because they provide additional possible bases for similarity judgements on the same model.[124] In macroeconomics, this perspective is likely to apply to the appeal of frictions that can be 'microfounded' for similar reasons. Finally, echoing Lucas, the high value placed on results that are intuitive may reflect the fact that they provide a stronger basis for judging similarity and relevance.

These observations suggest that many aspects of cases in a decision-maker's 'memory' may be important for supporting similarity judgements, consistent with the use of a generalised version of case-based decision theory in Appendix D.1.

D.3 Implications for monetary policy formulation

Many of the observations regarding the potential value of case-based reasoning in macroeconomics carry across more-or-less directly to the analysis that supports monetary policy formulation. However, some specific implications are worthy of discussion.

The fact that decision-making using case-based reasoning requires a holistic judgement mirrors the use of multiple 'inputs' in policy formulation. Therefore, inputs (or combinations thereof) can be regarded as cases and "All cases, real, experimental and theoretical, are aggregated, weighing their similarity and relevance, to generate predictions for the case in hand" (Gilboa et al., 2014). This interpretation makes clear, as

[124] For example, in forming a similarity judgement, as well as considering "is it plausible that this person behaves roughly consistently with maximising an expected utility function?", a decision maker can *also* consider "is it plausible that this person has transitive preferences?"

discussed in Appendix D.1, that 'memory' can include hypothetical results (including, for example, model-based simulations and scenarios). Indeed, in the context of monetary policymaking (compared with day-to-day decision making) such results may form a relatively large portion of 'memory'. Note also that "weighing" is distinct from "weighting", as the latter suggests a linear combination of the insights from each case. Part of the reason is that different cases will have policy-relevant implications in different 'spaces' (e.g., quantitative versus qualitative insights) and so cannot be meaningfully aggregated in a convex combination.[125] Moreover, different aspects of different inputs will influence the similarity judgements that determine their relevance for the practical policy decision.

The importance of 'memory' for decision-making suggests that it is good practice to draw on the insights from a diverse range of inputs. In the context of committee-based monetary policymaking, memory is collective and so the pooled 'case histories' of a diverse group of expert policymakers and staff can expand the evidence base for policymaking considerably.

Furthermore, expertise among policymakers and staff can support learning in two ways:

> [F]irst, by introducing more cases into memory; second, by refining the similarity function based on past experiences. Knowing more is but one aspect of learning. Making better use of known cases is another. Correspondingly, the notion of expertise combines knowledge of cases and the ability to focus on the relevant analogies. (Gilboa and Schmeidler, 2001, Chapter 7).

Since "knowledge of the similarity function is inherently subjective" (Gilboa and Schmeidler, 2001), it is possible that "honest people can disagree" on the relevance of a given input to the practical policy problem, (Lucas, 2011). Moreover, different inputs may have conflicting implications for policy formulation. Therefore, it is to be expected that individual policymakers may form different judgements from the same set of inputs.

[125] For example, Gilboa et al. (2014) argue that "An economist who is asked to make a prediction will then use case-based reasoning to learn from empirical data, experiments, theoretical models [...], historical examples, casual observations and computer simulations."