

ESSAYS ON
UNCONVENTIONAL MONETARY POLICY

Adrian Paul

A thesis submitted in partial fulfilment
of the requirements for the
Doctor of Philosophy in Economics.

St. Peter's College,
University of Oxford.
Michaelmas Term, 2017.

ESSAYS ON UNCONVENTIONAL MONETARY POLICY

Adrian Paul*

Abstract

This is a brief summary of the three chapters comprising my D.Phil. thesis.¹

*St. Peter's College, University of Oxford. Michaelmas Term, 2017. A thesis submitted in partial fulfilment of the requirements for the D.Phil. in Economics. I would like to thank Professor Martin Ellison for his guidance.

¹Approximate word count = (Approximate number of words per page) $\times$ (Approximate number of pages) = $400 \times 190 = 76,000$.

1 Summary

This D.Phil. thesis discusses three issues associated with unconventional monetary policy under uncertainty. The first two chapters are largely theoretical; the third chapter is primarily empirical.

1.1 Robust monetary policy with an unconventional instrument

In the first chapter, we investigate robust monetary policy with an unconventional instrument.

Monetary policy since the global financial crisis of 2007/08 has had to contend with the ‘unknown unknowns’ associated with Knightian uncertainty rather than the ‘known unknowns’ associated with risk. By using techniques from the theoretical literature on robust control, we find that when a policymaker knows neither the parameters of his model nor the probability distribution from which those parameters are drawn, more fear of model misspecification calls for more aggressive use of both conventional and unconventional instruments. Moreover, the greater the policymaker’s doubts about the effect of asset purchases relative to the effect of interest rate changes, the greater the relative zeal with which he should pursue the former.

Critics of the U.S. Federal Reserve have argued that “unwarranted pessimism” about the effectiveness of quantitative easing (QE) inhibited the response of policymakers to the 2007/08 financial crisis. We find that this relative passivism during QE2 may instead have been the optimal response to less fear of model misspecification following QE1. Rather than the FOMC’s return to activism during QE3 implying that its passivism during QE2 was undue, it may have been the latter that was warranted and the former that was undue.

1.2 The portfolio balance channel under uncertainty

In the second chapter, we investigate the portfolio balance channel of unconventional monetary policy, in an environment characterised by uncertainty.

In the aftermath of the global financial crisis, aggregate output collapsed and macroeconomic uncertainty spiked. In response, the large-scale asset purchase programmes pursued by major central banks extended well beyond purchases of risk-free government

bonds to include the accumulation of private-sector credit instruments with increasingly uncertain returns. We develop a theoretical model in which two assets are imperfect substitutes and households' preferences over the composition of their portfolios are robust to the prevalence of Knightian uncertainty. Because relative asset quantities matter, quantitative easing has real effects. Moreover, it is households' portfolio preferences that determine the “bang for the buck” associated with central bank asset purchases. We show that a flight to safety induced by the uncertainty accompanying a recession undermines the effectiveness of QE – just when it is needed most.

1.3 Debt duration as a policy instrument

In the third chapter, we investigate the empirical consequences of changes in the government's balance sheet.

Theory suggests that it is the net composition of the government's *consolidated* balance sheet – which incorporates asset purchases by the monetary authority as well as debt issuance by the fiscal authority – that determines the duration of public debt held by the private sector, and therefore influences the real economy. Seen from this perspective, the transmission channel from unconventional monetary policy to output and inflation is not especially unconventional.

We estimate a structural vector autoregression with time-varying parameters based on a sample of U.S. data stretching back to 1975. We find that an exogenous increase in the average maturity of government debt held in private hands reduces inflation and increases unemployment. These effects are three or four times larger in the depths of the Great Recession than in the midst of the Great Moderation – a finding which underscores the importance of imperfect asset substitutability, but also suggests that the joint behaviour of the U.S. Treasury and the U.S. Federal Reserve in the aftermath of the global financial crisis of 2007/08 amounted to quantitative *tightening* rather than quantitative easing.

I. ROBUST MONETARY POLICY WITH AN UNCONVENTIONAL INSTRUMENT

Adrian Paul*

Abstract

Monetary policy since the global financial crisis of 2007/08 has contended with the ‘unknown unknowns’ associated with Knightian uncertainty rather than the ‘known unknowns’ associated with risk. We find that when a policymaker knows neither the parameters of his model nor the probability distribution from which those parameters are drawn, more fear of model misspecification calls for more aggressive use of both conventional and unconventional instruments. Moreover, the greater the policymaker’s doubts about the effect of asset purchases relative to the effect of interest rate changes, the greater the relative zeal with which he should pursue the former. Critics of the U.S. Federal Reserve argue that “unwarranted pessimism” about the effectiveness of quantitative easing (QE) inhibited the post-crisis monetary policy response. We find that this relative passivism during QE2 may instead have been the optimal response to less fear of model misspecification following QE1. Rather than the FOMC’s return to activism during QE3 implying that its passivism during QE2 was undue, it may be the latter that was warranted and the former that was undue.

*St. Peter’s College, University of Oxford. Michaelmas Term, 2017. This is the first chapter of my D.Phil. thesis: *Essays on Unconventional Monetary Policy*.

1 Introduction

“As we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don’t know we don’t know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones.”

Donald H. Rumsfeld,
U.S. Secretary of Defense,
12 February 2002.

In this paper, we adapt the theory of robust control pioneered by Hansen and Sargent (2008) to address three issues associated with policymaking under uncertainty. Each of these issues emerged in the aftermath of the global financial crisis of 2007/08.

First, we introduce an unconventional monetary policy instrument into a simple New Keynesian model subject to Knightian uncertainty. This approach overturns the policy attenuation result derived by Williams (2013). He prescribes that the responsiveness of both the conventional and the unconventional instrument of monetary policy should be more muted the more imperfectly the central banker understands the effects of his actions on the economy. This precautionary passivism, however, assumes that the policymaker faces only *known* unknowns: while he does not know the realisation of the parameters governing the transmission mechanism of monetary policy, he *does* know the distribution from which those parameters are drawn. By contrast, in an environment of *unknown* unknowns, we find that the responsiveness of both the conventional and the unconventional policy instrument should be more aggressive the more the central banker doubts his model of the economy. Robustness implies that greater uncertainty around the effect of short-term interest rate changes and around the effect of long-term asset purchases calls for doing more of both, not less of either.

Second, we incorporate asymmetric fear of model misspecification into the policymaker’s problem, to confront the ‘pretence of knowledge’ syndrome described by Caballero (2010). He argues that the false precision proclaimed by modern macroeconomics has rendered most academic models incapable of dealing with the new and unexpected. Caballero (2010) asserts that good models must instead allow a policymaker to assign different weights to different policy prescriptions, depending on his degree of conviction in the mechanism from which each prescription follows. The robust control approach allows

us to formalise the way in which a policymaker should react when he doubts some aspects of his model more than others. We find that the more the policymaker doubts the effect of unconventional monetary policy relative to the effect of conventional monetary policy, the greater the relative zeal with which he should pursue the former. Robustness suggests that relatively little confidence in quantitative easing should imply relatively more of it.

Third, we rationalise the “unwarranted pessimism” that Romer and Romer (2013) ascribe to the U.S. Federal Reserve during much of 2010 and 2011. They contend that the Federal Open Market Committee (the FOMC) held unduly pessimistic beliefs about the stimulative power of unconventional monetary policy from the end of the recession in late 2009 through much of 2012. Those beliefs, Romer and Romer (2013) argue, “led to a marked passivity in policymaking.” Inspired by a different defence of the FOMC – constructed in a theoretical setting by Ellison and Sargent (2012), and articulated in a practical setting by Kohn (2013) – we examine post-crisis U.S. monetary policy decisions through the lens of a model of endogenous pessimism that emerges from the robust control approach. Rather than the FOMC’s passivism during QE2 being unjustified, robustness suggests that it may have been the optimal response to a decline in policymakers’ fears of model misspecification following QE1. The FOMC’s return to relative activism during QE3 then appears to violate robust optimality, given that fears of misspecification look to have receded further following the experience of QE2.

After the collapse of Lehman Brothers in September 2008, monetary policymakers around the world exhausted the scope of ‘conventional’ policy measures and resorted to new ‘unconventional’ instruments – despite having had little experience with them, and despite having considerable doubts about their likely impact.¹ The U.S. Federal Reserve, for example, cut its target federal funds rate to a range of between 0% and 0.25% in December 2008. In a succession of large-scale asset purchases (LSAPs), the Federal Reserve accumulated some \$2 trillion in agency debt, mortgage-backed securities and Treasury bonds over the course of 2009. This phase of purchases is typically known as ‘QE1.’ It was followed by a second phase of purchases (‘QE2’) during late 2010 and early 2011. From late 2012 through much of 2014, a third phase of LSAPs (known as ‘QE3’) expanded the Federal Reserve’s asset purchase programme to over \$4 trillion (24%

¹We refer interchangeably to these ‘unconventional instruments’ of monetary policy as: ‘central bank purchases of government debt,’ ‘asset purchases,’ ‘central bank balance sheet interventions,’ ‘quantitative easing’ and ‘QE.’

of GDP). The Bank of England, the Bank of Japan and the European Central Bank adopted similar bond purchase programmes.²

Figure 1 illustrates the unprecedented nature of monetary policy intervention during the crisis. The panel on the left shows the co-ordinated reduction in short-term nominal interest rates in the U.S., the Euro area and the U.K.: policy rates were cut from between 4% and 5% in Autumn 2008 to effectively zero by the turn of 2009. The panel on the right indicates the scale of the central bank balance sheet interventions that followed.

Unconventional monetary policy was initially characterised as an emergency measure to stimulate aggregate economic activity when the conventional instrument of monetary policy – the short-term nominal interest rate – had been driven to its perceived lower bound. The introduction of such an unprecedented policy tool, however, exacerbated the potential for pervasive model misspecification (something that had long troubled central bankers, even those who operated within a conventional policy regime).

Back in January 2004, U.S. Federal Reserve Chairman Alan Greenspan asserted that: “uncertainty is not just a pervasive feature of the monetary policy landscape; it is the defining characteristic of that landscape...one is never quite sure what type of uncertainty one is dealing with in real time, and it may be best to think of a continuum ranging from well-defined risks to the truly unknown.”

In December 2013, San Francisco Federal Reserve President John Williams warned: “the claim that the Fed is responding insufficiently to the shocks hitting the economy rests on the assumption that policy is made with complete certainty about the effects of policy on the economy. Nothing could be further from the truth. Policymakers are unsure of the future course of the economy and uncertain about the effects of their policy actions. Uncertainty about the effects of policies is especially acute in the case of unconventional policy instruments such as using the Fed’s balance sheet to influence financial and economic conditions.”³

²The Bank of England, by the end of 2013, had accumulated around £375 billion in gilts via its Asset Purchase Facility (amounting to 26% of GDP). Over the same period, the Bank of Japan’s asset purchase programme had grown to around ¥224 trillion (equivalent to 47% of GDP).

³Financial regulators, too, wrestled with unknown channels of contagion across asset markets after the collapse of Lehman Brothers. The Bank of England’s head of financial stability, Andy Haldane, has noted: “There are natural limits to what you can know. And taking seriously not knowing, taking seriously ignorance, I think is tremendously important for public policy. Humbling, but important.”

How to conduct monetary policy optimally in the context of Knightian uncertainty is a perennial problem that has come to command urgent attention. Monetary policy in the midst of crisis is beset by unknown unknowns. The modern central banker – armed with both a conventional and an unconventional instrument – has some subjective assessment of the impact of his decisions on the economy. The problem, however, is not only that he does not know the potency of his unconventional policy instrument, he does not even know the mechanisms by which it transmits through the economy.

We investigate optimal monetary policy when a central banker suspects that his approximating model of the economy is close to the true model of state transition, but he does not know the probability distribution over plausible alternative models. We derive robustly optimal reaction function coefficients for the short-term interest rate and long-term asset purchases when – cognisant of the limitations of his knowledge, but unable to enumerate them – the central banker fears that unknown unknowns undermine his model of monetary transmission. Of all the unknowns he must guard against, these tend to be the most difficult ones.

This paper proceeds as follows. Section 2 surveys the literature. Section 3 formalises robustness. Section 4 specifies the model economy. Section 5 derives optimal policy. Section 6 presents key results under ‘symmetric’ fear of model misspecification: robust monetary policy calls for doing more, not less. Section 7 explores ‘asymmetric’ fear of misspecification: the greater the relative uncertainty associated with unconventional policy, the more aggressively it should be used. Section 8 examines whether the FOMC has – given the evolution of its economic forecasts – acted more or less aggressively with its unconventional instrument over time: there is evidence that policymakers’ fears of misspecification receded after QE1. Section 9 concludes.

2 Literature

This paper draws on three strands of literature. Its innovation lies in addressing the individual concerns of each within a common framework. Robust control allows us: (i) to challenge policy attenuation, by specifying an environment of model uncertainty rather than parameter risk, (ii) to confront the ‘pretence of knowledge,’ by formalising the decision problem facing a policymaker who has less confidence in some aspects of his model than others, and (iii) to infer the FOMC’s fears of model misspecification in real time, by examining the empirical responsiveness of its unconventional policy instrument to the

evolution of its forecasts for the real economy.

To derive the policy accentuation result, we combine the stylised open-economy model of conventional monetary policy described by Ball (1999) with the approach to formalising unconventional policy taken by Williams (2013), while incorporating a preference for robustness using the techniques pioneered by Hansen and Sargent (2008).

Following Ball (1999), the model we consider consists of three equations: an IS curve, a Phillips Curve, and an asset market equation linking monetary policy with the exchange rate. There are no micro-foundations and there is no forward-looking behaviour. The omission of any channel for expectations sacrifices sophistication (and realism) for tractability (and clarity). The aim is to capture the effects of multiple monetary policy instruments in a parsimonious way that can elucidate the implications of the policymaker's fear of model misspecification. Despite its simplicity, this three-equation system is similar in spirit to the more complicated, dynamic stochastic general equilibrium (DSGE) models currently used by many central banks.⁴

Our model differs from Ball (1999) only in that the policymaker now has two levers with which to control the economy rather than one. He has recourse to both a conventional policy instrument (the short-term interest rate) and an unconventional policy instrument (purchases of long-term government debt). Following Williams (2013) and Gertler and Karadi (2011), we posit that the transmission mechanism of unconventional monetary policy is structurally similar to that of conventional policy, though its potency differs. Interventions on one margin, therefore, do not exactly offset interventions along the other. This assumption of structural similarity is consistent with empirical understanding of the effects of central bank balance sheet policies, among academics (Chung et al. (2012), for example) and among policymakers themselves (Bullard (2013), for example).⁵

As in traditional analyses such as Levin and Williams (2003), the policymaker has no commitment technology, he targets zero inflation, and he faces a textbook New Keynesian model hit by independent, identically distributed shocks. As in Svensson (1997), the op-

⁴See, for example, Smets and Wouters (2003), Christiano et al. (2005) and Smets and Wouters (2007).

⁵Comparing empirical evidence from two particular episodes – the June 2013 FOMC decision (which was more hawkish than financial markets expected) and the September 2013 decision (which was more dovish) – Bullard (2013) finds “striking confirmation that changes in the expected pace of purchases act just like conventional monetary policy.”

timal policy – that is, the reaction function which minimises the sum of the unconditional variances of those state variables about which the policymaker cares – is a version of a Taylor Rule.

Modelling choice under uncertainty has a short history in macroeconomics, but a deep tradition in microeconomics. From Knight (1921) to Gilboa and Schmeidler (1989), microeconomic theorists have come to distinguish between two types of subjectively uncertain belief. An ‘unambiguous’ belief conforms to Knight’s definition of risk: it captures outcomes with known, objective probabilities. An ‘ambiguous’ belief conforms to Knight’s definition of uncertainty: it refers to outcomes that cannot be assigned probabilities. An ambiguity-averse decision-maker chooses only those actions whose consequences are robust to the imprecision of his knowledge of the odds. Ambiguity aversion is inconsistent with subjective expected utility theory, yet it is commonly observed among individuals.⁶

Macroeconomic theorists, however, have largely neglected any formal treatment of model uncertainty in favour of a Bayesian approach to parameter risk.⁷ Brainard (1967) shows that if a policymaker does not know the particular value of a parameter capturing the impact of policy on the economy – but does know the probability distribution from which that parameter’s value is drawn – he should attenuate his policy response to shocks. Blinder (1999) calls such gradualism “extremely wise.” He translates Brainard’s conservatism principle into the key ingredient of a central banker’s cookbook for practical application: “Step 1: Estimate how much you need to tighten or loosen monetary policy to ‘get it right.’ Then do less.”⁸

⁶The smooth ambiguity model of Klibanoff et al. (2005) attempts to rectify this, capturing individual behaviour that *is* influenced by how well the decision-maker knows relevant probabilities, thereby resolving the behavioural paradox observed in Ellsberg (1961) examples.

⁷As John Kay notes though, the issue was the subject of Keynes’ fellowship dissertation at King’s College, Cambridge. Kay writes: “For Keynes, probability was about believability, not frequency. He denied that our thinking could be described by a probability distribution over all possible future events, a statistical distribution that could be teased out by shrewd questioning – or discovered by presenting a menu of...opportunities. In the 1920s he became engaged in an intellectual battle...Keynes and Knight lost...and Ramsey and Savage won...the probabilistic approach has maintained academic primacy ever since.”

⁸Moving beyond Brainard (1967) – though maintaining the restrictive assumption that the distribution of unknown parameters is known – the literature on policymaking under risk has refined attenuation-type results to incorporate ‘active learning.’ Wieland (2000), for example, shows that it may be optimal for a monetary policymaker to experiment over initially unknown parameters. Policy might optimally be chosen in a way that is detrimental to current outcomes, but yields new information that will improve future performance.

The defence of monetary policy moderation set out by Williams (2013) is an extension of this Bayesian approach to the conduct of quantitative easing. The central banker faces parameter risk rather than model uncertainty. He has both a conventional and an unconventional instrument available, he does not observe the realisations of the parameters governing his model, but he does know the probability distributions from which those parameters are drawn. Williams (2013) asserts that the optimal response of both the short-term interest rate and the quantity of long-term assets purchased is more muted than in the case of certainty equivalence.

Since the global financial crisis of 2007/08, the conduct of monetary policy has been better characterised by the *unknown* unknowns associated with Knightian uncertainty than the *known* unknowns associated with risk. It may well be too demanding to expect policymakers to know the probability distribution from which key parameters are drawn.⁹ There are insufficient historical data available – on the effectiveness of QE, for example – to infer the likelihood of certain scenarios. Without knowing these probabilities, it is impossible for the policymaker to identify the uniquely optimal course of action. Confining ourselves to known unknowns seems less important than understanding how better to respond to unknown unknowns.

A model can be understood as a probability distribution over a sequence. The rational expectations hypothesis imposes a communism of models – one in which the agents (determining an outcome), the econometrician (studying an outcome) and Nature (affecting an outcome) all share the *same* probability distribution over the sequence of outcomes. Lucas (1976) and Sargent and Wallace (1975) have shown that this assumption of communism has precise consequences for the design and impact of macroeconomic policy. It is more realistic, however, to think of models as approximations. One model approximates another. The actual data generating process is unknown. The truth may be unknowable. In formulating the techniques of robust control, Hansen and Sargent (2008) show that the implications of macroeconomic theory differ markedly when decision-makers acknowledge model misspecification and attempt to accommodate the approximation errors that it creates.

Robust control advocates that an apprehensive policymaker, who knows neither the

⁹In emphasising the importance of moving beyond the study of risk towards uncertainty, Sargent (2008) questions why the former is useful when, after all, “a Bayesian knows the truth from the beginning.”

specific values of key parameters nor the probability distribution from which they are drawn, should choose the policy under which the largest possible loss he would incur – across all the potential scenarios he is willing to envisage – is smaller than the corresponding ‘worst-case scenario’ under any alternative policy.¹⁰ Robust control is closely related to the risk sensitive preference formulation explored by Whittle (1990) – an observational equivalence we will exploit in Section 5, but violate in Section 7 – and the engineering concept of H^∞ optimal control, developed by Başar and Bernhard (1995).

Relatively few academic papers use robust control to analyse optimal monetary policy; even fewer focus on discretionary decision-making; none explores optimal policy when multiple instruments are available. Dennis (2013) shows that a central bank’s preference for robustness can substitute for credibility when credibility is low, motivating the central bank to adhere to policy announcements that would otherwise be time-inconsistent. Leitemo and Söderström (2004) show that, in a closed economy with only one conventional instrument available, the effects of model misspecification fear are unambiguous: the robust policy always responds relatively more aggressively to shocks. By contrast, Leitemo and Söderström (2008) show that, in a small open economy, a greater preference for robustness may induce a central banker with only one policy instrument to sometimes respond to shocks more aggressively, and sometimes respond more cautiously.

While the robust control approach to optimal monetary policy under uncertainty has received some attention from policymakers – Bernanke (2007) notes, for example, that “concern about worst-case scenarios...may lead to amplification rather than attenuation in the response of the optimal policy to shocks...Indeed, intuition suggests that stronger action by the central bank may be warranted to prevent particularly costly outcomes” – it is the Bayesian approach to optimal monetary policy under risk that has dominated the academic literature. After the experience of 2007/08, however, guarding against model misspecification in monetary policy constitutes a more pervasive (and more pressing) concern than mitigating parameter risk.

The importance of allowing the policymaker’s fear of model misspecification to be asymmetric follows from the ‘pretence of knowledge’ syndrome first postulated by von Hayek (1974). In a post-crisis assessment of the state of modern, academic macroeco-

¹⁰This min-max approach is often associated with Wald (1950), who argued that it delivered “a reasonable solution of the decision problem when an *a priori* distribution in [the state space] does not exist or is unknown.”

nomics, Caballero (2010) argues that the dynamic stochastic general equilibrium (DSGE) approach “has become so mesmerised with its own internal logic that it has begun to confuse the precision it has achieved about its own world with the precision that it has about the real one.” As the global financial crisis of 2007/08 made clear, this approach left macroeconomic policymakers “overly exposed to the new and unexpected” – ignoring as it does Hayek’s warning of the dangers associated with presuming a precision and degree of knowledge that we do not possess.

Caballero (2010) argues that – given the reaction of human beings to the truly unknown is fundamentally different from their reaction to the risks associated with a known environment – “it is not nearly enough...to do Bayesian estimation of the dynamic stochastic general equilibrium model, for the absence of knowledge is far more fundamental than such an approach admits.” Rather, macroeconomists must explicitly acknowledge the ignorance of policymakers: “the solution is not simply to explore a wide[r] range of parameters for a specific mechanism. The problem is that we do not know the *mechanism*, not just that we don’t know its strength.”

Without formalising how to do so, Caballero (2010) asserts that good models must allow policymakers to “assign different weights to those [policy prescriptions] that follow from blocks over which we have true knowledge, and those that follow from very limited knowledge.” Our robust control approach accommodates precisely those concerns. It allows us to reduce the amount – and the type – of knowledge that policymakers and economic agents are assumed to possess. Furthermore, it enables us to formalise the policy implications of such asymmetric fear.

There is, by now, a significant body of literature – both theoretical and empirical – assessing the effectiveness of unconventional policy. The ‘Modigliani-Miller theorem for open market operations’ put forward by Wallace (1981) makes large-scale asset purchases irrelevant.¹¹ To escape this result, the theoretical literature has turned to various forms of imperfect asset substitutability. Assuming a ‘preferred habitat’ channel of transmission, Ellison and Tischbirek (2013) find that – even when the conventional policy instrument is not constrained by the zero lower bound – the unconventional instrument can prove a

¹¹In a baseline New Keynesian model, Eggertsson and Woodford (2003) show that injecting reserves in exchange for long-term assets is a neutral operation. In a similar model incorporating credit frictions, Curdia and Woodford (2011) show that quantitative easing in the strict sense is ineffective. Market participants ensure no arbitrage opportunities persist; asset purchases have no effect on the economy.

valuable addition to the toolkit of a stabilising central bank.

Several empirical papers find that central bank purchases of government debt after the onset of the global financial crisis were indeed effective in reducing long-term interest rates. Krishnamurthy and Vissing-Jorgensen (2011) provide an overview.¹² However, evidence of the stimulative effect of those asset purchases on the wider real economy has been more mixed.

Chen et al. (2012) find that unconventional policy had only a modest macroeconomic effect. They simulate an asset purchase programme in the U.S. and find that GDP growth increases by less than a third of a percentage point, while inflation barely changes relative to the absence of intervention. Without a commitment to keep the short-term nominal interest rate at its lower bound for an extended period of time, they find an even smaller effect from large-scale asset purchases.

By contrast, Chung et al. (2012) find that the Federal Reserve’s second round of large-scale asset purchases, despite reducing the term premium on long-term interest rates only slightly (by 20 basis points), increased the level of GDP and the rate of inflation substantially (by 0.6% and 0.1%, respectively). Baumeister and Benati (2013) find that a 60 basis point reduction in the term premium accompanied a 3% increase in U.S. GDP growth and a 1% increase in inflation.¹³

Despite the variance in empirical estimates of the effectiveness of QE, some authors used their calculations to urge the FOMC to act more aggressively in real time. At the end of QE1, for example – by which time the Federal Reserve had purchased close to \$2 trillion in long-term assets – Gagnon (2009) argued that the appropriate stance of U.S. monetary policy required a further 75 basis point reduction in the 10-year Treasury yield. To achieve this, he argued that “the Fed would need to buy about \$2 trillion in debt securities (with an average maturity of roughly seven years) over and above what the Fed

¹²Estimates of the reduction in the 10-year Treasury yield (per \$100 billion of assets purchased) resulting from the Federal Reserve’s first round of QE range from 3 basis points (Hamilton and Wu (2012); Gagnon et al. (2011)) to 15 basis points (D’Amico and King (2013)). Table 1 of Chen et al. (2012) also provides a useful summary.

¹³In the U.K., Joyce et al. (2011) report that the Bank of England’s asset purchases raised the level of GDP by between 1½% and 2% and increased inflation by between ¾ and 1½ percentage points (commensurate with a cut in the policy rate of between 150 and 300 basis points). They warn, however, that “these estimates are clearly highly uncertain, particularly as none of the methods used to produce them fully capture all the likely transmission channels.”

has already committed to buy. These purchases would be announced now but could be implemented over the course of 2010.”

Against the backdrop of this empirical evidence, Romer and Romer (2013) argue that unduly pessimistic beliefs about the power of monetary policy “led to a marked passivity in policymaking” in the U.S. after the end of the recession in 2009.¹⁴ Citing public comments made by FOMC members, Romer and Romer (2013) assert that in much of 2010 and 2011 monetary policymakers explicitly acknowledged that the perceived costs associated with unconventional instruments had “muted their policy response.” The authors conclude that “beliefs that the benefits of expansion are small and the costs potentially large appear to have led monetary policymakers to eschew more aggressive expansionary policy...In hindsight, these beliefs may be judged too pessimistic.”¹⁵

This conclusion, however, has not been without its critics. Former Federal Reserve Vice Chairman Kohn, in his comments on Romer and Romer (2013), notes that “an important handicap in making such a judgement is that we are in uncharted waters...[we] are highly uncertain about what forces are holding back economic growth. That uncertainty is compounded by the use of new policy tools. We have little or no empirical basis for making judgements about the effects – costs or benefits – of large scale asset purchases.” Kohn (2013) resists the suggestion in Romer and Romer (2013) that the gradual evolution of unconventional policy towards a more aggressive setting in late 2012 is proof that the Federal Reserve should have been more aggressive to begin with. He instead argues that the subsequent shift towards activism is evidence that policymakers updated their understanding of the effects of QE as they accumulated experience – as they began to realise, for example, “just how little inflation and inflation expectations have been affected by unconventional policies.” Stressing that the Federal Reserve “can’t and shouldn’t ignore its calculation of costs and benefits,” but should rather “adjust policy as needed” over

¹⁴Indeed, the authors argue that this excessive humility has been a source of the Federal Reserve’s failures over the past century: in the 1930s, excessive pessimism about the power of monetary policy exacerbated the Great Depression; in the 1970s, unwarranted pessimism about the potential for policy tightening thwarted attempts to rein in the Great Inflation. Central bankers, Romer and Romer (2013) advocate, “should have a balance of humility and hubris. They need a sound knowledge of both the limitations and the powers of monetary policy.”

¹⁵Compare the humility which Romer and Romer (2013) ascribe to Federal Reserve policymakers with the hubris that Cochrane (2011) ascribes to them: “it is dangerous for the Fed to claim immense power, and for us to trust that power, when it is basically helpless.” This contrast is testament to the uncertainty surrounding the impact of unconventional monetary policy.

time, Kohn (2013) concludes that: “I’m not sure it’s justifiable to argue that the FOMC has...been less forceful than warranted in real-time policymaking.”

Our robust control approach to policymaking under uncertainty delivers a model of *endogenous* optimism and pessimism. This approach allows us to discriminate between the views espoused by Romer and Romer (2013) and the defence mounted by Kohn (2013).¹⁶ From a normative perspective, robust optimality implies that – when faced with a negative shock such as the global financial crisis of 2007/08 – more fear of misspecification calls for an expansionary monetary policy that is more aggressive, not less. Greater humility should imply greater activism. From a positive perspective, examining the empirical sensitivity of the Federal Reserve’s monthly asset purchases to changes in FOMC forecasts for key macroeconomic variables allows us to infer how policymakers’ fears of model misspecification evolved during the three phases of QE conducted after 2008.

3 Robustness

In this section we set out the general, theoretical underpinnings of robust control, as formalised by Hansen and Sargent (2008). We apply this methodology to the specific problem faced by a post-crisis policymaker in Section 4. Analytical derivations of solution techniques are relegated to the Appendix.

Robust control seeks to deliver a disciplined response function to a decision-maker who fears that the framework he is using to make decisions may be misspecified.

3.1 Entropy

The decision-maker envisages a linear model describing the transition of a set of state variables x between time t and $t+1$, contingent on his choice of control variables u at time t and the realisation of a vector of white-noise shocks ε at time $t+1$. His approximating model is:

$$x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1} \tag{1}$$

¹⁶This application of robust control is inspired by Ellison and Sargent (2012). In an entirely different context, these authors use similar techniques to assert (far more definitively) that a preference for robustness can help explain why the FOMC seems to ignore the projections made by its staff, despite recognising the merit of those projections as forecasts of future economic conditions.

The decision-maker is concerned that this model is only an imperfect approximation to the ‘true’ model describing the evolution of the state variables. He is prepared to entertain alternative models in a non-denumerable set that ‘closely’ surrounds his approximating model. This lack of conviction can be captured by specifying a suitable distorted model as:

$$x_{t+1} = Ax_t + Bu_t + C(\varepsilon_{t+1} + w_{t+1}) \quad (2)$$

where w_{t+1} represents a perturbation to the conditional mean of the process for x_{t+1} . Under the approximating model, $w_{t+1} = 0$. Under distorted models, w_{t+1} is non-zero. Specification errors w_{t+1} , disguised by innovations ε_{t+1} , feed back arbitrarily onto the history of past realisations of x . The conditional probability distributions for next-period’s vector of state variables under the approximating model and under the distorted model can be written as:

$$f^{app}(x_+|x) \sim N(Ax + Bu, CC') \quad (3)$$

$$f^{dist}(x_+|x) \sim N(Ax + Bu + Cw, CC') \quad (4)$$

where the subscript $_+$ denotes next period’s values.

The extent to which the decision-maker’s approximating and distorted models are ‘close’ is measured by the conditional relative entropy $I(.,.)$ between them:

$$I(f^{app}, f^{dist})(x) = \int \log \left(\frac{f^{dist}(x_{t+1}|x_t)}{f^{app}(x_{t+1}|x_t)} \right) f^{dist}(x_{t+1}|x_t) dx \quad (5)$$

As the two conditional distributions converge, $\frac{f^{dist}(x_{t+1}|x_t)}{f^{app}(x_{t+1}|x_t)}$ tends to unity and $I(f^{app}, f^{dist})(x)$ tends to zero from above.¹⁷

To formalise the decision-maker’s preference for robustness, we impose an exogenous entropy constraint on his actions. In a dynamic setting, this constraint bounds the dis-

¹⁷Equation (5) shows that the entropy of the distortion w is the expectation of the log-likelihood ratio with respect to the distorted distribution. We appeal to this concept in Section 5.2 when restricting our policymaker’s problem to a set of empirically plausible distortions via ‘detection error probabilities.’ In econometrics, conditional relative entropy is referred to as the Kullback-Leibler distance; the same concept is also used in the literature on rational inattention.

counted discrepancy between his approximating and distorted models:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t I(w_{t+1}) \leq \eta \quad (6)$$

where $\beta \in (0, 1]$ is the decision-maker's discount factor, $\mathbb{E}_0$ is the mathematical expectation operator (evaluated at $t = 0$) with respect to the distribution of ε_t , entropy is a function of the distortion w_{t+1} and $\eta \in (0, \bar{\eta}]$. The parameter η regulates the extent of model misspecification against which the decision-maker wishes to guard. The larger is η , the greater are his misspecification fears and the larger is the set of alternative models under consideration.¹⁸

3.2 Max-min

The decision-maker seeks an optimal policy that is robust to distortions in his approximating model, subject to those distortions satisfying the conditional relative entropy constraint. His per-period return function deteriorates quadratically over the set of state and control variables, x and u , each weighted by the positive semi-definite matrices Q and R .

$$r(x, u) = -x'Qx - u'Ru \quad (7)$$

One way of framing the decision-maker's desire for robustness is to imagine that he faces the following max-min constraint problem (I):

¹⁸ η is bounded above by $\bar{\eta}$ in order to ensure that the problem is well-defined: if η were unbounded, the decision-maker would be attempting to guard against model misspecification of ever-increasing severity. In the limit, he would then be unable to avoid returning an infinitely bad state. The case at $\bar{\eta}$ is sometimes known as the point of 'neurotic breakdown.' Engineers devising systems with maximal robustness typically seek a value of η as high as possible (known, in that literature, as working under the ' H^∞ norm').

$$\begin{aligned}
& \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t r(x_t, u_t) \\
& \text{s.t.} \\
& x_{t+1} = Ax_t + Bu_t + C(\varepsilon_{t+1} + w_{t+1}) \\
& \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t I(w_{t+1}) \leq \eta
\end{aligned} \tag{I}$$

He chooses actions $\{u_t\}_{t=0}^{\infty}$ that maximise his discounted return $\sum_{t=0}^{\infty} \beta^t r(x_t, u_t)$, subject to (i) the distorted model of state transition, and (ii) the actions of a hypothetical ‘evil agent,’ who chooses distortions $\{w_{t+1}\}_{t=0}^{\infty}$ to cause him the most harm possible while still respecting the conditional relative entropy constraint. This malevolent *alter ego* personifies the decision-maker’s doubts. A fictitious opponent, his attempt to minimise the same objective function that the policymaker maximises induces a mental, zero-sum game between the two players. The evil agent’s antagonistic pessimism is instrumental in achieving robust optimality: it allows the decision-maker to design a policy that guards against the worst possible scenario that can be inflicted on him.

Substituting the conditional relative entropy constraint directly into the objective function converts the constraint problem into an equivalent multiplier problem.¹⁹ In Appendix A, we show that since the approximating and distorted models are linear and the return function is quadratic, if shocks follow a Normal distribution (a class known as ‘linear-quadratic Gaussian (LQG) robust control problems’) then our measure of conditional relative entropy in equation (5) simplifies to $I(w_{t+1}) = \frac{1}{2} w'_{t+1} w_{t+1}$.²⁰

The decision-maker’s constraint problem then becomes (II):

¹⁹Hansen and Sargent (2008) show that, under certain conditions, the policy function that solves the constraint problem (I) is equivalent to that which solves the multiplier problem (II). These conditions – which essentially allow us to invoke the Lagrange multiplier theorem – hold throughout this paper.

²⁰This simplification is derived in a general specification of an LQG problem in Appendix A.3.

$$\begin{aligned}
& \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \frac{1}{2} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{r(x_t, u_t) + \beta \theta w'_{t+1} w_{t+1}\} \\
& \text{s.t.} \\
& x_{t+1} = Ax_t + Bu_t + C(\varepsilon_{t+1} + w_{t+1})
\end{aligned} \tag{II}$$

The scalar multiplier $\theta \in [\underline{\theta}, +\infty)$ in problem (II) is a direct counterpart of the parameter η in the constraint problem (I). Both capture how tightly the entropy constraint binds. θ can be understood as the Lagrange multiplier on the time-zero discounted relative entropy constraint: the shadow price of robustness. The lower bound $\underline{\theta}$ corresponds to the upper bound $\bar{\eta}$ as a parameterisation of the ‘point of neurotic breakdown.’ The greater the decision-maker’s potential doubts, the larger the set of plausible alternative models he might consider, the higher the upper bound $\bar{\eta}$ on the permissible discrepancy between the approximating and distorted models in the constraint problem. The corollary of this is a lower lower bound $\underline{\theta}$ multiplying the entropy contribution of a given distortion $w'_{t+1}w_{t+1}$ in the multiplier problem, which implies that the malevolent agent is penalised less heavily in the minimand and the decision-maker is therefore penalised more heavily in the maximand.

Most solutions to this problem exploit a principle of modified certainty equivalence: the stochastic problem (II), in which ε_{t+1} is a random variable, can be treated *as if* it were a deterministic problem, in which $\varepsilon_{t+1} = 0$. Even after setting ε_{t+1} to zero, however, the remaining distortions w_{t+1} are pre-multiplied by C and thus feed back onto the state. The variance and covariance of shocks (CC') nevertheless influence robustly optimal decisions, because noise in the system affects the ability of the ‘malevolent agent’ to distort the decision-maker’s approximating model for a given limit on conditional relative entropy.²¹

Modified certainty equivalence implies that the Bellman equation corresponding to (II) is (III):

²¹The dependence of the policy rule on the volatility matrix lends itself naturally to the analogy between robust control and risk sensitivity, as we will see in Section 3.4.

$$\begin{aligned}
-x'\mathbb{V}x &= \max_u \min_w \{-x'Qx - u'Ru + \beta\theta w'w - \beta x'_+\mathbb{V}x_+\} \\
&\text{s.t.} \\
x_+ &= Ax + Bu + Cw
\end{aligned} \tag{III}$$

where $r(x, u)$ is replaced by its quadratic form using equation (7) and the subscript $+$ generalises next-period variables. $\mathbb{V}$ denotes the decision-maker's value function: defined over state variables x , it encodes the impact of current decisions on future outcomes.

3.3 Solution

Appendix A.1 describes how the decision-maker's recursive problem can be solved by iterating to convergence on the algebraic Riccati equations associated with his dynamic program (III). This yields the optimal decision rule:

$$u = -\mathbb{F}x \tag{8}$$

where

$$\mathbb{F} = \beta(R + \beta B'\mathcal{D}(\mathbb{P})B)^{-1}B'\mathcal{D}(\mathbb{P})A \tag{9}$$

in which

$$\mathcal{D}(\mathbb{P}) = \mathbb{P} + \mathbb{P}C(\theta\mathbb{I} - C'\mathbb{P}C)^{-1}C'\mathbb{P} \tag{10}$$

in which $\mathbb{P}$ is an appropriate fixed-point of the value function $\mathbb{V}$ in (III). $\mathbb{I}$ denotes the identity matrix.

The policy function (8) has a simple form: optimal control requires the control vector u to react linearly to the state variables x , although the co-efficient matrix $\mathbb{F}$ is a non-linear function of the structural matrices A , B and R . Because shocks are additive and the policymaker's payoff is quadratic, by modified certainty equivalence the policy reaction function is independent of stochastic shocks ε .

The optimal responsiveness of policy depends *directly* on: (i) the decision-maker's discount factor, (ii) the relative weights he places on deviations in his control variables,

and (iii) the transition law implied by his approximating model, and *indirectly* (via $\mathcal{D}(\mathbb{P})$) on: (i) the relative weights he places on deviations in state variables, (ii) the multiplier capturing how tightly his entropy constraint binds, and (iii) the variance-covariance matrix of shocks to his approximating model. The last of these is simply CC' : that which was originally the volatility exposure matrix now also becomes an impact matrix for misspecification.

As for the decision-maker's malevolent opponent, the optimal choice of his vector of control variables w , given state variables x , is given by:

$$w = \mathbb{K}x \tag{11}$$

where

$$\mathbb{K} = \theta^{-1}(\mathbb{I} - \theta^{-1}C'\mathbb{P}C)^{-1}C'\mathbb{P}(A - B\mathbb{F}) \tag{12}$$

in which $\mathbb{F}$ is determined by (9).

Equation (11) determines the distortions implicit in the ‘worst-case scenario’ as a function of: (i) the decision-maker's robustly optimal policy function, (ii) the law of transition implied by the decision-maker's approximating model, (iii) the variance-covariance matrix of shocks to the system, and (iv) the tightness of the entropy constraint.

Equations (8) and (11) constitute the general solution to the LQG problem. Substituting this solution into equations (1) and (2) implies that, under the approximating model:

$$x_+ = (A - B\mathbb{F})x + C\varepsilon_+ \tag{13}$$

and under the distorted model:

$$x_+ = (A - B\mathbb{F} + C\mathbb{K})x + C\varepsilon_+ \tag{14}$$

Since $\mathbb{F}$ and $\mathbb{K}$ are both functions of θ , the state transition laws under the approximating model and the worst-case model are ‘exponentially twisted’ according to the decision-maker's preference for robustness. The malevolent agent is the instrument by which context-specific caution is introduced into the decision-maker's control law. Pes-

simism is endogenous.

Before we apply this exponential twisting to the specific problem faced by a monetary policymaker with an unconventional instrument, we state one further equivalence result which will be useful.

3.4 Risk sensitivity

Following Jacobson (1973) and Whittle (1990), Appendix A.2 shows that the decision rule which solves the decision-maker’s robust control problem by imposing a constraint on conditional relative entropy (as in (III)) is *observationally equivalent* to the decision rule which solves an alternative stochastic problem in which the decision-maker does *not* fear model misspecification but instead has risk sensitive preferences.²² To capture risk sensitivity, we can simply apply an exponential transform to the continuation value in the decision-maker’s dynamic program. This pessimistically twists the conditional density of the approximating model by weighting more heavily those outcomes that have lower continuation values. The exponential transformation is parameterised by $\sigma \leq 0$. A greater absolute value of σ captures greater risk sensitivity. The key implication of this observational equivalence is that:

$$\theta \equiv -\sigma^{-1} \tag{15}$$

The inverse mapping between θ and σ formalises the close relationship between the macroeconomic literature on min-max robust control (Hansen and Sargent (2008)) and the microeconomic literature on ambiguity aversion (Gilboa and Schmeidler (1989)). Risk sensitive preferences are a special case of Epstein and Zin (1989) preferences.²³ As $|\sigma| \rightarrow 0$, we know that $\theta \rightarrow \infty$ from (15), so $\mathcal{D}(\mathbb{P}) \rightarrow \mathbb{P}$ in (10), and the robustly optimal solution to the decision-maker’s problem collapses to the standard solution under no fear

²²Hansen and Sargent (2008) also recast the decision-maker’s problem in the frequency domain, showing that the entropy criterion expresses fear of model misspecification by imposing additional concavity on the decision-maker’s return function. This curvature induces a preference for smoothness across frequencies (instead of across states of nature) which “looks like risk aversion.”

²³As we show in Appendix A.2, the observational equivalence between robust control and risk sensitivity implies that (a more general) analogue of the policymaker’s Bellman equation (III) can be written as: $\mathbb{V}(x) = \max_u \left\{ r(x, u) + \frac{1}{\sigma} \log \left(\mathbb{E} \left[e^{\sigma \beta \mathbb{V}(x_+)} \right] \right) \right\}$. This underpins the assertion that, to capture risk sensitivity, we can simply add an exponential transform to the continuation value in the decision-maker’s dynamic program.

of misspecification.²⁴ As the decision-maker's risk sensitivity fades, his preference for robustness disappears.

After specifying a simple model of the economy in Section 4, we will investigate the implications of a monetary policymaker's fear of model misspecification in Section 5 by examining robustly optimal reaction functions under various values of σ .²⁵

4 Model

Our model of the economy is heavily simplified. It is taken from Ball (1999), but is augmented to accommodate an unconventional instrument of monetary policy in addition to the conventional short-term interest rate. The model consists of three equations: (i) an IS curve, (ii) a Phillips Curve, and (iii) an asset market equation linking monetary policy to the exchange rate.

We begin by asserting that – through purchases of long-term government debt – the central bank has complete control over the long end of the yield curve. By determining the quantity of long-term government debt he will buy, the monetary policymaker can *perfectly* control its price.²⁶ This implies that, while conventional monetary policy is implemented by setting the short-term interest rate, we can think of unconventional monetary policy as being implemented by choosing the desired long-term interest rate.

All parameters are positive. All variables are measured as deviations from trend. One unit of time is a quarter. Shocks are serially uncorrelated and mutually orthogonal.

$$y_t = -\tau r_{t-1} - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t \quad (16)$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \quad (17)$$

²⁴This is the certainty-equivalent solution to the single-agent decision problem.

²⁵The tightness of the entropy constraint that corresponds to the policymaker's degree of risk sensitivity will be left in the background, though the identity in (15) will ensure that every parameterisation of risk sensitivity (σ) that we consider has a corresponding entropy penalty on misspecification (θ).

²⁶This assumption will be relaxed in Section 7. By modelling unconventional monetary policy in this way, we take as given the stock of long-term government debt issued by the fiscal authority. We therefore ignore any government budget constraint.

$$e_t = \psi r_t + \mu b_t + v_t \quad (18)$$

In the IS curve (16), the log of real output y depends positively on its own lag, negatively on the lagged value of the short-term real interest rate r , negatively on the lagged value of the long-term real interest rate b (which is controlled perfectly by central bank purchases of government debt) and negatively on the log of the lagged real exchange rate e (where a rise in e constitutes an appreciation). Demand shocks are captured by ξ . The fraction $\frac{\phi}{\tau}$ regulates the relative output effect of unconventional versus conventional monetary policy. The long-term interest rate b is normalized such that zero corresponds to a case in which unconventional monetary policy has no impact on the output gap.

In the Phillips Curve (17), the change in the inflation rate π depends positively on lagged output and negatively on the lagged change in the exchange rate. The strength of the currency affects inflation because it is passed directly into import prices. Inflation shocks are captured by η .

In the asset market equation (18), movements in both the short and long end of the yield curve affect the exchange rate. A fall in either short- or long-term interest rates makes domestic assets less attractive relative to foreign assets, leading to a currency depreciation. The fraction $\frac{\mu}{\psi}$ regulates the relative exchange rate effect of unconventional versus conventional monetary policy. Exchange rate shocks – such as disturbances affecting investor confidence, market expectations of future policy or foreign interest rates – are captured by v .

Each instrument of monetary policy (r and b) affects inflation through two channels – one direct and one indirect. The indirect channel takes two periods to pass through: a monetary tightening reduces output after one period (via (16)) and thereby inflation after another period (via (17)). The direct channel takes one period to pass through: a monetary tightening causes a contemporaneous appreciation of the exchange rate (via (18)) that reduces inflation directly after one period (via (17)). This lag structure reflects the view that the direct transmission mechanism of monetary policy (operating via asset markets through the exchange rate) is faster than the indirect transmission mechanism of monetary policy (operating via the real economy through the output gap).

5 Policy

The three equations described in Section 4 constitute the policymaker’s approximating model of the economy. This section explores how he is to set policy optimally.

The policymaker is infinitely-lived. His objective is to minimise the present discounted value of his loss function. His per-period loss function is quadratic. It consists of the weighted sum of expected squared fluctuations (equivalently, the weighted sum of variances) in the output gap (defined as the percentage deviation between the actual level of output and its trend) and the inflation gap (the difference between the rate of inflation and the central bank’s zero target). The policymaker’s loss also includes the weighted sum of the variances of his two policy instruments. To model the policymaker as being averse to excess volatility in the short-term interest rate is commonplace; there has been little research, on the other hand, into the potential costs of long-term purchases of government debt. Without modelling the latter, we impose that the policymaker expects the costs of balance sheet interventions (their contribution to risks to financial stability or their tendency to disanchor inflation expectations, for example) to increase quadratically, just as he does the disutility of policy rate variability.²⁷

The central banker’s two policy instruments are his control variables: (i) a conventional monetary policy instrument (the short-term real interest rate, r) and (ii) an unconventional monetary policy instrument (the long-term real interest rate, b). As Ball (1999) shows, when it comes to conventional policy, controlling the short-term real interest rate in this model is equivalent to controlling the real exchange rate, e .²⁸

The policymaker’s linear-quadratic control problem can therefore be written as (IV):

²⁷In August 2012, Federal Reserve Chairman Bernanke listed four potential costs associated with unconventional policies: they “could impair the functioning of securities markets,” “reduce public confidence in the Fed’s ability to exit smoothly,” create “risks to financial stability,” and create “the possibility that the Federal Reserve could incur financial losses.”

²⁸In reality, the central bank’s conventional policy instrument is typically the short-term *nominal* interest rate rather than the short-term *real* interest rate. This distinction is immaterial in this model.

$$\begin{aligned}
& \min_{\{e_t, b_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{x_t' Q x_t + u_t' R u_t\} \\
& \text{s.t.} \\
& y_t = -\tau r_{t-1} - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t \\
& \pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \\
& e_t = \psi r_t + \mu b_t + v_t
\end{aligned} \tag{IV}$$

in which the non-negative, diagonal coefficient matrices Q and R capture, respectively, the policymaker's relative aversion to inflation versus output deviations, and his relative distaste for variability in short- versus long-term interest rates. $\mathbb{E}_0$ is the mathematical expectation evaluated (at time zero) with respect to the distribution of shocks hitting the economy.

We can express the policymaker's three constraints in (IV) as a special case of the general approximating model (1). Appendix B.1 shows this explicitly. In $x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1}$, we set $x_t \equiv (y_t \ \pi_t \ e_{t-1} \ b_{t-1} \ v_t)'$, we set $u_t \equiv (e_t \ b_t)'$ and we let $\varepsilon_{t+1} \equiv (\xi_{t+1} \ \eta_{t+1} \ v_{t+1})'$, together with structural coefficient matrices A , B and C .²⁹

Given the general distorted model (2), the vector of perturbations specialises in our case to $w_{t+1} \equiv (w_{\xi,t+1} \ w_{\eta,t+1} \ w_{v,t+1})'$. The policymaker's quadratic loss function in (IV) can also be expressed as a special case of the general return function (7), by appropriate choice of Q and R .

Our baseline calibration starts with the small, open economy parameters chosen by Ball (1999). As Table 1 shows, the policymaker discounts future losses at an annual rate of around 4% ($\beta = 0.99$; one period is a quarter), he places equal weight on the volatility of output, inflation and his policy instrument(s) ($Q_{yy} = Q_{\pi\pi} = R_{ee} = R_{bb}$), the effect of

²⁹As Appendix B.1 shows, in state-space form, x is a 5×1 matrix of state variables – including output and inflation; u is a 2×1 matrix of control variables – the conventional instrument (the exchange rate, or equivalently the short-term interest rate) and the unconventional instrument (interest on long-term debt); ε is a 3×1 matrix of shocks to output, inflation and the exchange rate; A is a 5×5 coefficient matrix capturing the autoregressive components of the system of transition equations; B is a 5×2 coefficient matrix capturing monetary policy transmission; and C is a 5×3 coefficient matrix capturing the variance and covariance of shocks. The dimensions of the weighting matrices attached to the state variables (Q) and to the control variables (R) in the policymaker's return function are 5×5 and 2×2 respectively.

the exchange rate on output is the same as the effect of the change in the exchange rate on the change in inflation ($\delta = \gamma = 0.2$), output is quite persistent ($\lambda = 0.8$) and the slope of the Phillips Curve is just under a half ($\alpha = 0.4$). When the conventional policy instrument alone is available (column (i)), the total output loss from a one-point rise in the interest rate is the sum of the direct effect of the interest rate and the indirect effect through the exchange rate ($\tau + \delta\psi = 1.0$). When the unconventional policy instrument is also available (column (ii)), the impact of a given change in the desired long-term interest rate on output is two-thirds that of the same change in the conventional policy instrument ($\frac{\phi}{\tau} = \frac{2}{3}$); its impact on the exchange rate is three-fifths that of conventional policy ($\frac{\mu}{\psi} = \frac{3}{5}$).

Table 1: THE BASELINE CALIBRATION

			(i) Conventional Policy only	(ii) Conventional and Unconventional Policy
Policymaker's Preferences	Discount factor	β	0.99	0.99
	Weight on output and inflation	Q	$\mathbb{I}$	$\mathbb{I}$
	Weight on instrument volatility	R	$\mathbb{I}$	$\mathbb{I}$
Structure of Economy	Effect of exchange rate on output	δ	0.2	0.2
	Effect of exchange rate on inflation	γ	0.2	0.2
	Persistence of output	λ	0.8	0.8
	Effect of output on inflation	α	0.4	0.4
Impact of Conventional Policy	Effect of short-term interest rate on output	τ	0.6	0.6
	Effect of short-term interest rate on exchange rate	ψ	2	2
Impact of Unconventional Policy	Effect of long-term interest rate on output	ϕ	0	0.4
	Effect of long-term interest rate on exchange rate	μ	0	1.2

Notes: Setting $Q = R = \mathbb{I}$ implies that $Q_{yy} = Q_{\pi\pi} = R_{ee} = R_{bb} = 1$.

Given that all parameters are positive, the sign of almost every element of the structural coefficient matrices A , B and C is invariant to the particular calibration chosen. The one exception is the element ($\frac{\tau\mu}{\psi} - \phi$) that governs the effect of unconventional policy on output. Its sign is calibration-specific. To ensure that an increase in the long-term interest rate depresses output, we rule out calibrations of the model in which $\frac{\phi}{\tau} < \frac{\mu}{\psi}$.³⁰

³⁰This follows directly from the sign restriction: $\frac{\tau\mu}{\psi} - \phi < 0 \iff \tau\mu < \phi\psi \iff \frac{\phi}{\tau} > \frac{\mu}{\psi}$. If we let

5.1 Why not ignore robustness?

Just as in Sargent’s criticism of Ball (1999), the specification (IV) simultaneously incorporates too *little* rational expectations and too *much* rational expectations. The private sector uses too little rational expectations (at least relative to the policymaker), and the policymaker uses too much (given how little rational expectations are built into the specification of firms’ and households’ behaviour).

Ignoring the need for robustness would be tantamount to treating the model economy as a time-invariant system of stochastic difference equations – one which presents state and control vectors to the policymaker, and one which the policymaker views as known and fixed with respect to his own decisions. But calibrating parameters rather than estimating them is tantamount to treating the model as an approximation – one that is not to be taken as literal, empirically. Ruling out model misspecification *a priori* therefore implies a conviction that is incompatible with choosing calibration over estimation.

The specification (IV) is also vulnerable to the Lucas critique. One defence is that, while it lacks microfoundations, a stylised model of this kind is nevertheless a good *approximation* to a truer, more complex general equilibrium model of the economy. But while this defence weakens Sargent’s first criticism (maintaining that the reason private agents use so little rational expectations is because they only formulate an approximating model), it simultaneously strengthens Sargent’s second criticism (begging the question as to how, then, the policymaker can use so much rational expectations).

The incongruity associated with calibrating a model that is supposedly known and fixed, and with endowing different people in the same model with different degrees of rationality, can only be resolved once a preference for robustness displaces the rational expectations hypothesis. If the model described in Section 4 is indeed an approximation, the optimal policy function that solves problem (IV) should accommodate the approximation errors that the policymaker will inevitably make. The policymaker might suspect that the model that he envisages in equations (16)-(18) is *close* to reality but, acknowledging that it is not *true*, his objective must be to find a robustly optimal policy rule that performs well across a range of distorted models in the vicinity of the truth.

$\mathcal{U}_y = \frac{\phi}{\tau}$ characterise the relative power of unconventional policy in the calibration of the IS curve and we let $\mathcal{U}_e = \frac{\mu}{\psi}$ characterise the relative power of unconventional policy in the calibration of the asset market equation, this sign restriction amounts to $\frac{\mathcal{U}_y}{\mathcal{U}_e} > 1$. In the baseline calibration of Table 1, $\frac{\mathcal{U}_y}{\mathcal{U}_e} = 1.11$.

To examine the implications of this preference for robustness, we exploit the observational equivalence described in Section 3.4. We allow the policymaker to have various degrees of risk sensitivity (σ) and examine the robustly optimal policy function that solves problem (IV) in each case. Our solutions are numerical; they follow the iteration of algebraic Riccati equations described in Section 3.3.

5.2 How much robustness?

We have introduced a new parameter σ into the analysis. This parameter has – so far – been allowed to vary freely.³¹ To discipline our choice of σ , we calculate the detection error probability $P(\sigma)$ associated with each value of risk sensitivity under consideration. $P(\sigma)$ allows us to bound the size of the non-denumerable set of distorted state transition equations against which it is empirically plausible for our policymaker to seek robustness.³² Following Anderson et al. (2003), this calibration of scepticism is derived from a statistical theory for model selection.

At a given point in time, the policymaker should – to some degree – have already noticed ‘easily’ detectable forms of model misspecification from *past* time series data and consequently eliminated them. Relative to the rate at which *future* data become available, however, the policymaker’s discount factor will make him sufficiently impatient that he cannot wait for those new data to arrive in order to resolve his remaining model misspecification fears for him. The probability of detection error is the likelihood that, given the finite sample of data at his disposal and his specific value of σ , the policymaker’s approximating and distorted models are statistically indistinguishable.³³ Seen another

³¹Given Lucas’ warning that applied economists should “beware of theorists bearing free parameters,” this approach should raise alarm bells. As shown in Section 3, σ in the risk sensitive formulation has an exact analogy in alternative specifications of the policymaker’s problem – namely η in the constraint problem and θ in the multiplier problem.

³²One rationale for the policymaker’s preference for robustness – even in ‘normal times’, with only one conventional instrument – is Friedman’s observation that lags in the effects of monetary policy are “long and variable”. That lags are long is expressed in the specification of equations (16)–(18); that lags are variable is expressed in the non-denumerable set of distortions w that feed back onto the history of state variables x . In this context, we could interpret $P(\sigma)$ as regulating just *how* variable those lags might be.

³³For a given fear of model misspecification, a distinct decision rule is associated with each approximating model. This respects the Lucas critique. However, for a given fear of misspecification and a given approximating model, the policymaker uses the *same* decision rule over a *set* of models surrounding the approximating model. This ostensibly violates the Lucas critique. The policymaker would justify this violation by appealing to parameterisations of the problem that deliver detection error probabilities *large enough* to make members of that set of models indistinguishable from the approximating model, given

way, a detection error probability is the probability that an econometrician makes an incorrect inference about whether observed data are generated by the approximating model or the worst-case distorted model.

The greater the discrepancy between the policymaker’s approximating model and his distorted model, the greater his fear of misspecification, the smaller the malevolent agent’s penalty multiplier θ , the larger the absolute value of σ , the easier it is to detect misspecification and the lower the probability of detection error $P(\sigma)$. If the approximating and worst-case models were identical, the policymaker would have no fear and the probability of detection error would be exactly half. Appendix B.2 specifies the precise methodology we use to simulate detection error probabilities, given the policymaker’s approximating model in problem (IV). Figure 2 plots $P(\sigma)$ based on 500 hypothetical simulations from our model, each postulating a 25-year history of realisations of state variables.

For comparative statics, we calibrate three distinct values of σ . These illustrate how robustly optimal policy differs when the monetary policymaker has ‘high fear,’ ‘moderate fear’ and ‘low fear’ of model misspecification. These three mindsets correspond to probabilities of detection error of 10%, 25% and 40% respectively. No preference for robustness corresponds to a ‘no fear’ detection error probability of 50%. As Figure 2 shows, these four values of $P(\sigma)$ imply values of σ of approximately -0.055 , -0.035 , -0.015 and 0 .

More fear implies more robustness, which equates to greater risk sensitivity and a lower probability of detection error, as the approximating and distorted models become easier to distinguish statistically. While we initially took the constraint on conditional relative entropy as *exogenous* in equation (6), disciplining misspecification via detection error probabilities serves to *endogenise* the policymaker’s pessimism.

6 Policy Accentuation

In this section, we present our main results under symmetric fear of model misspecification.³⁴

the data he has available.

³⁴When fear is symmetric, a specification error in any one of the state-evolution equations in the policymaker’s approximating model is treated the same as any other. The policymaker’s problem (III) imposes this restriction, in that the coefficient θ on distortions w in the entropy penalty $\theta w'w$ is a scalar. In Section 7, we allow for misspecification fear to be asymmetric.

We begin with the policy accentuation result under the baseline calibration, and then consider its sensitivity to changes in key parameters. Our findings contradict Williams (2013). He finds that, when parameters are subject to risk, the optimal response of both conventional and unconventional policy instruments is *more muted* than in the case of certainty equivalence. His defence of moderation conforms to Brainard’s principle of conservatism. By contrast, we find that when the policymaker’s approximating model is subject to uncertainty, the robustly optimal response of both conventional and unconventional instruments is *more aggressive* than in the case of no uncertainty. The graphical results later in this section provide deeper intuition.

6.1 Accentuation

Table 2 illustrates the first key result, under the baseline calibration in Table 1. From equation (8) we know that the policymaker’s robustly optimal decision rule is given by $u = -\mathbb{F}x$, where the 2×5 matrix $\mathbb{F}$ is a function of σ (alternatively, θ) and $\mathbb{F}$ captures the responsiveness of each of the policymaker’s two policy instruments to changes in the five state variables of the system. Table 2 reports the optimal responsiveness of conventional policy to output ($\mathbb{F}_{ey}$), the optimal responsiveness of conventional policy to inflation ($\mathbb{F}_{e\pi}$), the optimal responsiveness of unconventional policy to output ($\mathbb{F}_{by}$) and the optimal responsiveness of unconventional policy to inflation ($\mathbb{F}_{b\pi}$) for different degrees of misspecification fear.

When the policymaker has no fear of model misspecification (Panel (a)) and two instruments are available (column (ii)), he uses both for stabilization, though he is considerably more responsive with conventional policy than he is with unconventional policy to a given deviation in output or inflation. As the policymaker’s fear of model misspecification increases from ‘low’ to ‘high’ (Panels (b)-(d)), he uses both his conventional and his unconventional monetary policy instrument more aggressively.³⁵ A fall in detection error probability from 40% to 10% doubles the robustly optimal reaction of conventional and unconventional policy to output and inflation; the proportionate increase in the responsiveness of QE is slightly larger than the proportionate increase in the responsiveness of the policy rate.

³⁵This generalises the result in Sargent (1999). He considers only conventional monetary policy, and finds that the interest rate becomes more sensitive to deviations in state variables as the policymaker’s preference for robustness increases.

Table 2: ‘HIGH FEAR’ OF MODEL MISSPECIFICATION REQUIRES MORE AGGRESSIVE USE OF BOTH POLICY INSTRUMENTS THAN ‘LOW FEAR’

		(a) No fear	(b) Low fear	(c) Moderate fear	(d) High fear
Risk Sensitivity	$ \sigma $	0	0.015	0.035	0.055
Entropy Penalty	θ	$+\infty$	66.7	28.6	18.2
Detection Error	$P(\sigma)$	50%	40%	25%	10%

	Conv. only (i)	Conv. & Unconv. (ii)	Conv. only (i)	Conv. & Unconv. (ii)	Conv. only (i)	Conv. & Unconv. (ii)	Conv. only (i)	Conv. & Unconv. (ii)
Conventional policy response to output	1.0081	1.0042	1.1456	1.1406	1.4442	1.4365	2.0969	2.0807
Conventional policy response to inflation	0.7084	0.7065	0.8520	0.8493	1.1724	1.1674	1.8903	1.8776
Unconventional policy response to output	-	0.0772	-	0.0879	-	0.1111	-	0.1616
Unconventional policy response to inflation	-	0.0429	-	0.0538	-	0.0783	-	0.1334

Notes: Columns labelled (i) refer to “Conventional Policy Only”; Columns labelled (ii) refer to “Conventional and Unconventional Policy.” The robustly optimal policy rule is given by $u = -\mathbb{F}x$, where $\mathbb{F}$ is a function of σ . In Column (ii) of Panel (c), for example, the unconventional policy response to output refers to $\mathbb{F}_{by}$ in the case of $|\sigma| = 0.035$.

Figure 3 illustrates this accentuation result for the full range of empirically plausible detection error probabilities (consistent with $0 \leq |\sigma| \leq 0.060$). The upper panel corresponds to the case in which only conventional policy is available; the lower panel corresponds to the baseline calibration in which non-zero values for ϕ and μ also activate unconventional policy. Greater model uncertainty translates into ‘doing more’ rather than ‘doing less’ – with both instruments.

6.2 Sensitivity

For each level of fear, Table 3 illustrates how the responsiveness of monetary policy changes following a 10% *ceteris paribus* increase in a given parameter, relative to the baseline calibration. We report elasticities for both the policymaker’s conventional and unconventional instruments, in response to deviations in both output and inflation.

Four patterns emerge. We focus mainly on the case of ‘moderate fear’ (column (b)), simply to ease exposition.

- (i) A 10% increase in the policymaker’s discount factor (β) makes him 21% more

Table 3: THE PROPORTIONATE CHANGE IN POLICY RESPONSIVENESS
FOLLOWING A PROPORTIONATE CHANGE IN EACH MODEL PARAMETER

			(a) Low fear		(b) Moderate fear		(c) High fear	
			Output	Inflation	Output	Inflation	Output	Inflation
Policymaker's Preferences	β	Conventional	17.7	24.3	20.1	27.8	24.6	33.7
		Unconventional	18.6	29.3	21.0	32.3	25.3	37.2
	Q	Conventional	3.6	4.7	5.2	7.6	10.6	14.8
		Unconventional	3.4	3.3	5.0	6.9	10.5	15.1
	Q_{yy}	Conventional	1.3	-0.2	1.9	1.5	4.3	5.5
		Unconventional	1.4	-0.4	1.9	1.7	4.4	6.1
	$Q_{\pi\pi}$	Conventional	2.3	4.9	3.3	5.9	5.7	8.5
		Unconventional	2.0	3.5	3.1	5.1	5.6	8.2
	R	Conventional	-2.1	-2.6	-0.2	-0.7	4.3	3.7
		Unconventional	-1.9	-1.1	0.0	0.8	4.6	5.2
	R_{ee}	Conventional	-2.1	-2.6	-0.2	-0.7	4.2	3.7
		Unconventional	7.7	8.7	9.9	10.7	14.9	15.6
	R_{bb}	Conventional	0.0	0.0	0.0	0.0	0.1	0.1
		Unconventional	-9.1	-9.1	-9.1	-9.1	-9.0	-9.0
Structure of Economy	δ	Conventional	-1.8	-1.6	-3.1	-2.7	-5.8	-5.3
		Unconventional	-5.8	-5.8	-7.0	-7.0	-9.6	-9.6
	λ	Conventional	10.6	-2.3	9.6	-2.2	7.6	-2.5
		Unconventional	11.4	-3.2	10.3	-2.8	8.1	-2.9
	α	Conventional	3.5	-2.8	0.9	-6.9	-5.9	-15.1
		Unconventional	3.0	-3.5	0.5	-8.2	-6.3	-16.6
Impact of Conventional Policy	γ	Conventional	-0.7	-0.4	-1.0	-1.0	-1.6	-2.0
		Unconventional	-0.8	-2.6	-1.0	-2.7	-1.5	-2.9
	τ	Conventional	-2.3	-2.0	-4.1	-3.6	-7.6	-7.0
		Unconventional	-90.8	-90.9	-91.0	-90.9	-91.3	-91.3
Impact of Unconventional Policy	ψ	Conventional	1.5	1.4	3.2	3.0	7.5	6.9
		Unconventional	95.3	95.9	98.9	99.2	107.2	107.3
	ϕ	Conventional	-1.3	-0.9	-1.6	-1.3	-2.2	-2.0
		Unconventional	97.2	97.6	96.6	96.8	95.3	95.6
	μ	Conventional	0.4	0.3	0.5	0.4	0.8	0.7
		Unconventional	-90.0	-90.0	-89.9	-89.9	-89.9	-90.0

Notes: Since we set $\beta = 0.99$ in the baseline calibration, the elasticities shown in the first row refer to a 10% increase in β from 0.90 to 0.99 rather than from 0.99 to 1.089. A 10% increase in the impact parameters associated with conventional and unconventional policy continues to respect the sign restriction imposed on the effect of unconventional policy on output (see footnote 30 in the text).

responsive with unconventional policy to an increase in the output gap and 32% more responsive with unconventional policy to an increase in the inflation gap.³⁶ The corre-

³⁶For all other parameters, the numbers in Table 3 refer to the effect of proportionate changes from the baseline calibration. However, because that baseline implies $\beta = 0.99$ – and the policymaker's discount factor can never exceed one – the first row of Table 3 instead refers to proportionate changes in policy responsiveness following a 10% increase in β from 0.90 to 0.99.

sponding elasticities for conventional policy are 20% and 28%. For a given proportionate increase in the discount factor, the rate at which policy responsiveness increases in fear (the slope of each curve in Figure 3) is also higher the higher is the original level of fear: corresponding elasticities increase monotonically between columns (a), (b) and (c). An increase in the policymaker's discount factor therefore has a similar effect to an increase in his fear of model misspecification. In our robust control framework, ambiguity over *states* is closely related to preferences over *time*. Doubt is analogous to impatience.

(ii) A 10% increase in the policymaker's aversion to asset purchase volatility (R_{bb}) prompts a 9% decrease in the responsiveness of his unconventional instrument to deviations in both output and inflation. A 10% increase in his aversion to short-term interest rate volatility (R_{ee}) prompts a 0.7% decrease in the responsiveness of his conventional instrument to inflation (and a 0.2% decrease in the case of output), combined with a compensating 11% increase in the responsiveness of his unconventional instrument to inflation (and a 10% increase in the case of output). When variance in one instrument is more costly, the policymaker substitutes towards the other.

(iii) A 10% increase in the persistence of output (λ) amplifies the policymaker's QE response to output deviations by 10%, and reduces his QE response to inflation deviations by just under 3%. A similar reaction pattern holds for his policy rate response. Because fear of model misspecification manifests itself as an over-estimation of the serial correlation of shocks (this is the sense in which, by twisting the state-transition density into a worst-case scenario, the malevolent agent is the instrument of pessimism), an increase in the inherent persistence of output in the approximating model elicits the same activism from the policymaker as an increase in his doubts.

(iv) To understand the effects of a 10% increase in each of the parameters governing the impact of unconventional policy (the fourth panel of Table 3) recall that, under the approximating model, state variables evolve according to: $x_+ = (A - B\mathbb{F})x + C\varepsilon_+$. Given our specification of A and B in Appendix B.1, given the baseline calibration of the non-policy parameters in Table 1, and given the robustly optimal solution for $\mathbb{F}$ under moderate fear (in which case $\sigma = -0.035$), this transition equation implies:

$$x_+ = \left\{ \underbrace{\begin{pmatrix} 0.8 & 0 & 0 & 0 & \frac{\tau}{\psi} \\ 0.4 & 1 & 0.2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_A - \underbrace{\begin{pmatrix} -\frac{\tau}{\psi} - 0.2 & \frac{\tau\mu}{\psi} - \phi \\ -0.2 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}}_B \underbrace{\begin{pmatrix} -1.4365 & -1.1674 & -0.2335 & 0 & -0.3636 \\ -0.1111 & -0.0783 & -0.0157 & 0 & -0.0299 \end{pmatrix}}_{\mathbb{F}(\sigma = -0.035)} \right\} x + C\varepsilon_+ \quad (19)$$

Suppose that, in equation (19), we fix the conventional policy parameters (τ, ψ) and the impact of unconventional policy on inflation (μ) at their baseline values. A 10% increase in the impact of unconventional policy on output (ϕ) , all else equal, then prompts a 97% increase in the responsiveness of asset purchases and a 1-2% decrease in the responsiveness of conventional policy to output and inflation deviations.³⁷

Suppose instead that, in equation (19), we fix the conventional policy parameters (τ, ψ) and the impact of unconventional policy on output (ϕ) at their baseline values. A 10% increase in the impact of unconventional policy on inflation (μ) , all else equal, now prompts a 90% decrease in the responsiveness of asset purchases and a 0.5% increase in the responsiveness of short-term interest rates to output and inflation deviations.³⁸

The post-crisis policymaker has a subjective view of the transmission mechanism of unconventional monetary policy. The elasticity of his response with respect to ϕ and to μ show that the relative potency of each channel has precise consequences for his robustly optimal reaction to shocks.

³⁷To interpret the signs of these proportionate changes in responsiveness, note that the relevant equations of the state transition matrix in this thought experiment imply:

$$\begin{pmatrix} y_{t+1} \\ \pi_{t+1} \end{pmatrix} = \begin{pmatrix} 0.0417 + 0.1111\phi & -0.5555 - 0.0783\phi & -0.1111 - 0.0157\phi & 0 & 0.1290 - 0.0299\phi \\ 0.1127 & 0.7665 & 0.1533 & 0 & -0.0727 \end{pmatrix} \begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix} + \begin{pmatrix} \xi_{t+1} \\ \eta_{t+1} \end{pmatrix}$$

³⁸To interpret the signs of these proportionate changes in responsiveness, note that the relevant equations of the state transition matrix in this thought experiment imply:

$$\begin{pmatrix} y_{t+1} \\ \pi_{t+1} \end{pmatrix} = \begin{pmatrix} 0.1261 - 0.0333\mu & -0.6150 + 0.0235\mu & -0.1801 + 0.0047\mu & 0 & 0.1062 + 0.0090\mu \\ 0.1127 & 0.7665 & 0.1533 & 0 & -0.0727 \end{pmatrix} \begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix} + \begin{pmatrix} \xi_{t+1} \\ \eta_{t+1} \end{pmatrix}$$

6.3 Shocks

To analyse the dynamics of the system, we must take a stand on the true data generating process. In this subsection, suppose that it is the policymaker's approximating model that is, after all, the true model of monetary transmission. This implies that his fears of misspecification are 'all in his head.'³⁹

Under the baseline calibration, the impulse response functions (IRFs) in Figure 4 show the response of the key endogenous variables of the system (output, inflation, conventional policy and unconventional policy) to a positive, one-off, unit standard deviation shock to (i) the IS curve (denoted ξ_t), (ii) the Phillips Curve (denoted η_t) and (iii) the asset market equation (denoted v_t).

State transition is governed by equation (13): $x_+ = (A - B\mathbb{F})x + C\varepsilon_+$. Figure 4 shows how the instrumental pessimism which the malevolent agent instills in the robustly optimal reaction function affects the policymaker's response to shocks. We know from Figure 3 that a stronger preference for robustness (more fear of model misspecification) provokes a more active policy response (with both the conventional and unconventional instrument). As Figure 4 shows, this precautionary activism means that deviations of output and inflation from their respective steady-states dissipate more quickly following a structural shock.

Figure 5 considers the analogous impulse responses of the worst-case conditional means w to a positive, one-off, unit standard deviation shock to ξ .⁴⁰ The worst-case conditional means of shocks are amplified as the policymaker's fears intensify. Without a preference for robustness, the policymaker is most worried about misreading a persistent shock as a temporary shock. Although the approximating model asserts that the innovations in ε are white noise, the malevolent agent's choice of w makes them positively serially correlated. The cautious policymaker therefore interprets a given innovation (ξ_t, η_t or v_t) as being larger and more persistent than he envisaged under his approximating model; his robustly optimal strategy is to adjust both of his policy instruments more aggressively in response. To insure himself, he responds to the serially uncorrelated shocks under the approximating

³⁹In Section 6.4., we will instead suppose that it is the worst-case scenario that is the true model of monetary transmission, so that the policymaker's fears of misspecification are no longer unfounded.

⁴⁰Since the inference from a unit standard deviation shock to η and v is the same as that from a shock to ξ , we present only one set of impulse response functions in Figure 5.

model as if they were positively serially correlated.

The policymaker's caution is relative to the hypothetical, worst-case scenario for the sequence of innovations. How we specify the policymaker's approximating model influences the size and serial correlation properties of these worst-case shocks, and therefore affects how the policymaker responds to actual shocks – even when his subjective model of the economy coincides with the true data generating process.

6.4 Counterfactuals

In Section 6.3, the policymaker's pessimism was unwarranted. His doubts turned out to be unfounded, even though their prevalence accentuated his policy responsiveness. Suppose now that the true model of state transition is *not* the policymaker's approximating model, but is instead the worst-case scenario inflicted by the malevolent agent.

Pessimists see good news as temporary but bad news as persistent. The context-specific pessimism instilled by the malevolent agent is instrumental in the construction of a robust decision rule for the policymaker. The larger the absolute value of σ (the smaller the value of θ), the greater the policymaker's fear of misspecification, the greater the pessimism implicit in his robust reaction function.

Given equation (14), the vector of state variables under the worst-case scenario evolves according to: $x_+ = (A - B\mathbb{F} + C\mathbb{K})x + C\varepsilon_+$. State transition now entails the most pessimistic view of the serial correlation properties characterising the shock process ε .

Figures 6 and 7 illustrate this pessimism by comparing the forecast evolution of output and inflation under the approximating model versus under the worst-case scenario.

In Figure 6, we run one representative simulation to jointly determine both the history of key endogenous variables and their subsequent forecasted paths under the approximating and distorted models. Given equations (13) and (14), and since $\mathbb{E}_t[\varepsilon_{t+1}] = 0$, the j -period-ahead projection of x as of time t under the approximating model and under the worst-case scenario is, respectively:

$$\mathbb{E}_t[x_{t+j}] = (A - B\mathbb{F})^j x_t \quad (20)$$

and

$$\mathbb{E}_t[x_{t+j}] = (A - B\mathbb{F} + C\mathbb{K})^j x_t \quad (21)$$

In the upper panel of Figure 6, the policymaker has ‘low fear’ of model misspecification; in the lower panel he has ‘high fear.’ Relative to the forecast paths under the approximating scenario, this simulation illustrates that in the worst-scenario the recession is deeper and inflation is higher. It takes between four and five quarters for this differential to close.

Figure 7 shows the evolution of worst-case counterfactuals in ‘real time.’ The policymaker in this case has ‘moderate fear’. The paths of output and inflation under the worst-case scenario are worse than under the approximating model in a precise sense: at every point in time, the worst-case paths constitute larger and longer deviations from steady-state than turn out to be the case while the approximating model is the true data generating process. The malevolent agent exponentially twists the policymaker’s approximating density according to his fear of model misspecification. This twisting produces a sequence of worst-case counterfactuals over successive periods. At each point in time, those counterfactuals impart eternal pessimism to a policymaker seeking robustness.

6.5 Losses

The worst-case forecast path is a key input into the policymaker’s robust decision process because, by responding to *it*, he ensures that the performance of his objective function does not deteriorate unacceptably under each model in a large set of alternative models – not just under his approximating model.

Suppose that in the distorted model (14), we now allow for the degree of model misspecification determining the policymaker’s optimal reaction function (σ_1) to differ from the degree of model misspecification determining the worst-case scenario that the malevolent agent inflicts on him (σ_2). State transition then follows: $x_+ = (A - B\mathbb{F}(\sigma_1) + C\mathbb{K}(\sigma_2))x + C\varepsilon_+$. For a given σ_1 , the effect of an increase in σ_2 on the value of the policymaker’s objective function captures the fragility of his robust decision rule to the set of possible models that might prevail.

Figure 8 shows that the greater the policymaker’s fear of misspecification (that is, the larger is σ_1), the less rapidly the value function associated with his decision rule

deteriorates as the true model of state transition becomes more misspecified (that is, as σ_2 increases). This is the sense in which decision rules that solve the policymaker’s max-min problem for larger values of σ_1 are more robust to model uncertainty than solutions derived under smaller values of σ_1 . That said, the more robust policy rule does *worse* than the less robust rule when the approximating model is *not* misspecified (that is, when σ_2 approaches zero).

7 Asymmetric Fear

Caballero (2010) argues that a ‘pretence of knowledge’ pervades academic models in macroeconomics. This is because “there are many instances in which our knowledge of the true structural relationship is extremely limited. In such cases, the main problem is not in how to formalize an intuition, but in the assumption that the structural relationship is known with precision. Superimposing a specific optimization paradigm is not the solution...as much of the difficulty lies precisely in not knowing which optimization problem is to be solved.” To overcome this false precision, Caballero advocates that “ultimately, for policy prescriptions, it is important to assign different weights to those [prescriptions] that follow from blocks over which we have true knowledge, and those that follow from very limited knowledge.”

The effect of unconventional monetary policy in a post-crisis economy is one instance in which our knowledge is very limited. Indeed, a central motivation for Williams (2013) is that “uncertainty regarding the effects of unconventional policies on the economy is greater than for conventional policies.” Curdia and Ferrero (2013) find that “the effects of a program like QE2 on GDP growth are smaller and more uncertain than a conventional policy move of temporarily reducing the federal funds rate.” Kashyap (2013) points out that “hundreds of thousands of person hours have been devoted to studying conventional monetary policy. Tools like LSAPs have simply not been subject to nearly as much scrutiny. It is hardly surprising that...these policies are less well understood than conventional policy.”⁴¹

Caballero (2010) does not suggest an analytical framework within which to imple-

⁴¹Kashyap goes on to say that it is “laudable...that so many policymakers are willing to admit this and it is natural to be more cautious in using such tools.” Starting from the same premise of asymmetric uncertainty, our approach challenges the optimality of this relative caution.

ment his block-weighting procedure. Our robust control approach to the problem in (IV), however, extends naturally to allow for different weights on different aspects of the approximating model, depending on the policymaker's relative confidence in different mechanisms. Allowing for asymmetric fear of model misspecification formalises precisely the same intuition that underlies Caballero's idea.

In previous sections, our treatment of uncertainty has been Knightian (over the non-denumerable *set* of distorted models) but symmetric (within a *given* distorted model). Each value of risk sensitivity σ has corresponded to a unique Lagrange multiplier θ on the entropy constraint. This has determined the penalty $\theta w'_{t+1} w_{t+1}$ associated with a specific vector of perturbations w chosen by the malevolent agent in the min-max multiplier problem (II). Since θ is scalar and the perturbation vector is defined as $w_{t+1} \equiv (w_{\xi,t+1} \quad w_{\eta,t+1} \quad w_{v,t+1})'$, the analysis so far has penalized a given distortion to each of the three state transition equations of the approximating model *uniformly*.⁴²

This restriction that fear of model misspecification be symmetric precludes cases in which the policymaker has more doubt about one channel of monetary transmission than another. In this section, we extend the analysis of Section 6 to investigate robustly optimal responses when model misspecification fear is instead treated asymmetrically.

Though he is certain of neither, the monetary policymaker now has less confidence in the effect of his unconventional instrument on the economy than he has in the effect of his conventional instrument. Rather than modelling his relative lack of confidence as a Bayesian might – that is, by attributing a higher variance to some prior on the parameter capturing the impact of asset purchases – our approach assumes unstructured uncertainty about the effect of QE. To introduce this into the model, suppose the policymaker believes that he makes a control error when conducting quantitative easing.

Recall that, in the model of Section 4, by determining the quantity of long-term debt to purchase, the central bank had perfect control over the long-term interest rate, b . This assumption allowed us to treat the long-term interest rate as the *de facto* instrument of unconventional policy. Now we relax this assumption. In this section, the monetary policymaker can only imperfectly control the interest rate on long-term government debt (b) for a given quantity of assets purchased (q). This control error ascribes Knightian

⁴²Since $\theta w'_{t+1} w_{t+1} = \theta (w_{\xi,t+1}^2 + w_{\eta,t+1}^2 + w_{v,t+1}^2)$

uncertainty to the effect of quantitative easing.⁴³

The monetary policymaker's two control variables are now e and q . His new problem is (V):

$$\begin{aligned}
& \min_{\{e_t, q_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{x'_t Q x_t + u'_t R u_t\} \\
& \text{s.t.} \\
& y_t = -\tau r_{t-1} - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t \\
& \pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \\
& e_t = \psi r_t + \mu b_t + v_t \\
& b_t = -\chi q_t + \zeta_t
\end{aligned} \tag{V}$$

The additional equation ($b_t = -\chi q_t + \zeta_t$) allows for distortions to the relationship between the quantity of long-term debt the policymaker purchases (q) and the actual price that pertains (the realized interest rate on long-term debt, b).⁴⁴ Once this additional degree of uncertainty is resolved, however, changes in short- and long-term interest rates transmit through the approximating model in (V) exactly as they did in the approximating model in (IV). For a given understanding of monetary policy transmission (the first three constraints in (V)), the policymaker's control error introduces an independent source of ambiguity.

Just as in the case of symmetric fear, we can express the policymaker's augmented approximating model as a special case of the general transition equation (1). Appendix C.1 shows this explicitly.

In $x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1}$, we now set $x_t \equiv (y_t \ \pi_t \ e_{t-1} \ b_{t-1} \ v_t \ \zeta_t)'$ since an additional equation enters the state-space form, we set $u_t \equiv (e_t \ q_t)'$ since the policy-

⁴³It is also reasonable because, in practice, many other factors apart from central bank purchases affect the yield on long-term government debt – see, for example, Bech and Lengwiler (2012).

⁴⁴There is a direct, inverse mapping between the price of debt and the rate of interest it yields. The parameter χ captures the effect of a marginal increase in the quantity of asset purchases on the long-term interest rate. Chen et. al (2012) report empirical estimates of a 3-15bp fall in long-term yields per \$100 billion of asset purchases. This effect largely operates via the term premium component of long-term interest rates.

maker's unconventional control variable is now the quantity of long-term debt purchases rather than the interest rate on long-term debt, and we let $\varepsilon_{t+1} \equiv (\xi_{t+1} \ \eta_{t+1} \ v_{t+1} \ \zeta_{t+1})'$ since the policymaker must also contend with a shock to his QE control equation. The structural coefficient matrices A , B and C are now slightly larger in dimension.⁴⁵

To allow for asymmetric fear, when modelling the policymaker's distorted state transition equation we now apply a weighting matrix Ω to the malevolent agent's vector of perturbations w_{t+1} . The fourth element of this vector is the perturbation $w_{\zeta,t+1}$ to the QE control equation, which captures the discrepancy between the approximating and distorted models of the control error that the policymaker suspects he incurs when using quantities to control prices.

The policymaker's approximating and distorted models can therefore be expressed as:

$$x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1} \quad (22)$$

$$x_{t+1} = Ax_t + Bu_t + C(\varepsilon_{t+1} + \tilde{w}_{t+1}) \quad (23)$$

where $\tilde{w}_{t+1} \equiv \Omega w_{t+1}$

in which $w_{t+1} \equiv (w_{\xi,t+1} \ w_{\eta,t+1} \ w_{v,t+1} \ w_{\zeta,t+1})'$

$$\text{and } \Omega \equiv \begin{bmatrix} \Omega_\xi & 0 & 0 & 0 \\ 0 & \Omega_\eta & 0 & 0 \\ 0 & 0 & \Omega_v & 0 \\ 0 & 0 & 0 & \Omega_\zeta \end{bmatrix} \text{ must be non-singular.}$$

If $\Omega = \mathbb{I}$ then $\tilde{w}_{t+1} = w_{t+1}$ and fear is symmetric: distortions to any of the four equations of the policymaker's approximating model are penalized equally.⁴⁶ If $\Omega \neq \mathbb{I}$

⁴⁵As Appendix C.1 shows, in state-space form, x is now a 6×1 matrix of state variables; u is a 2×1 matrix of control variables – the exchange rate (equivalently, the short-term interest rate) and the level of long-term debt purchases; ε is a 4×1 matrix of shocks to output, inflation, the exchange rate and the QE control equation; A is a 6×6 coefficient matrix capturing the autoregressive components of the system of transition equations; B is a 6×2 coefficient matrix capturing monetary policy transmission; and C is a 6×4 coefficient matrix capturing the variance and covariance of shocks. The dimensions of the new weighting matrices attached to the state variables (Q) and to the control variables (R) in the policymaker's return function are 6×6 and 2×2 respectively.

⁴⁶We forced fear of model misspecification to be symmetric in Sections 5 and 6; symmetric fear arises as a special case here in Section 7. Since the structure of the policymaker's problem (V) differs from that

then fear is asymmetric. Because the policymaker has more confidence in some aspects of his model than others, a given perturbation to an equation about which the policymaker has low conviction will be multiplied by less in the malevolent agent's minimand – and will thus be penalized by more in the policymaker's maximand – than the same perturbation to an equation about which the policymaker has high conviction.

Suppose, for example, that the policymaker fears misspecification in his model of the impact of unconventional policy *four* times more than he fears misspecification in his model of the impact of conventional policy. The state-space form of the policymaker's problem (detailed in Appendix C.1) suggests that this relative lack of confidence in QE can be captured by letting, for example, $(\Omega_\xi \ \Omega_\eta \ \Omega_v \ \Omega_\zeta) = (1 \ \frac{4}{3} \ 1 \ \frac{2}{3})$. This implies that, for a given θ and an equal perturbation to the second and fourth constraint in (V) – so that $w_{\eta,t+1} = w_{\zeta,t+1}$ – the contribution from the perturbation $w_{\zeta,t+1}$ to the increase in the malevolent agent's minimand is a *quarter* of the contribution from the perturbation $w_{\eta,t+1}$.⁴⁷

For the maximising policymaker then, the distortion to the QE control equation carries a penalty four times larger than a distortion to the Phillips curve of the same magnitude. The fact that the coefficient attached to $\theta w_{\zeta,t+1}^2$ is a fraction of that attached to $\theta w_{\eta,t+1}^2$ gives the malevolent agent more scope to perturb the approximating equation for b than the approximating equation for π . This greater penalty encodes the asymmetry in the policymaker's misspecification fears.

in (IV), setting $\Omega = \mathbb{I}$ still implies different model predictions in (V) than in (IV), but restricting the weighting matrix to identity *conceptually* reduces (V) to the case of symmetric fear imposed in (IV).

⁴⁷Given modified certainty equivalence, we set $\varepsilon_{t+1} = 0$. The perturbation in equation (23) is:

$$C\tilde{w}_{t+1} = C\Omega w_{t+1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \Omega_\xi & 0 & 0 & 0 \\ 0 & \Omega_\eta & 0 & 0 \\ 0 & 0 & \Omega_v & 0 \\ 0 & 0 & 0 & \Omega_\zeta \end{pmatrix} \begin{pmatrix} w_{\xi,t+1} \\ w_{\eta,t+1} \\ w_{v,t+1} \\ w_{\zeta,t+1} \end{pmatrix} = \begin{pmatrix} \Omega_\xi w_{\xi,t+1} \\ \Omega_\eta w_{\eta,t+1} \\ 0 \\ 0 \\ \Omega_v w_{v,t+1} \\ \Omega_\zeta w_{\zeta,t+1} \end{pmatrix}$$

and its associated penalty is:

$$\theta \tilde{w}_{t+1}' \tilde{w}_{t+1} = \theta (\Omega w_{t+1})' (\Omega w_{t+1}) = \theta w_{t+1}' \Omega \Omega' w_{t+1} = \theta (\Omega_\xi^2 w_{\xi,t+1}^2 \ \Omega_\eta^2 w_{\eta,t+1}^2 \ \Omega_v^2 w_{v,t+1}^2 \ \Omega_\zeta^2 w_{\zeta,t+1}^2)$$

Comparing this last expression in the case of $(\Omega_\xi \ \Omega_\eta \ \Omega_v \ \Omega_\zeta) = (1 \ 1 \ 1 \ 1)$ with that in the case of $(\Omega_\xi \ \Omega_\eta \ \Omega_v \ \Omega_\zeta) = (1 \ \frac{4}{3} \ 1 \ \frac{2}{3})$ leads to the argument made in the text.

More generally, suppose we have two candidates for the weighting matrix: Ω^A and Ω^B . The asymmetry in the policymaker's fear of model misspecification is greater under Ω^B than Ω^A if the diagonal elements of Ω^B are a mean-preserving spread of the diagonal elements of Ω^A . By the properties of a mean-preserving spread, this implies that the diagonal elements of Ω^B have a higher variance than the diagonal elements of Ω^A , where we can write the variance of the diagonal elements of any given weighting matrix Ω as:

$$\mathcal{V}(\Omega) = \frac{1}{\dim(\Omega)} \text{Tr} \left[\left\{ \Omega - \frac{\text{Tr}(\Omega)}{\dim(\Omega)} \mathbb{I} \right\}' \left\{ \Omega - \frac{\text{Tr}(\Omega)}{\dim(\Omega)} \mathbb{I} \right\} \right] \quad (24)$$

where $\text{Tr}(\Omega)$ denotes the trace of the square matrix Ω and $\dim(\Omega)$ denotes its dimension.

Intuitively, the greater the variance of the diagonal elements of Ω , the greater the difference in the weights attached to a given distortion of each equation in the state transition law, the greater the variation in entropy penalties across each component of the policymaker's approximating model and the greater the asymmetry in the policymaker's fear of model misspecification.⁴⁸

Suppose that the policymaker has 'moderate fear,' A reduction in the ratio of Ω_ζ to Ω_η from 0.2 to 0.1 corresponds to an increase in the asymmetry between the policymaker's doubts about unconventional policy and his doubts about conventional policy. Under the baseline calibration of Table 1 (together with $\chi = 1$), this increase in asymmetry prompts a 0.44% increase in the responsiveness of conventional policy to output, a 0.27% increase in the responsiveness of conventional policy to inflation, a 0.45% increase in the responsiveness of unconventional policy to output and a 0.38% increase in the responsiveness of unconventional policy to inflation. For a given degree of doubt over the model specification as a whole then, the robust control approach prescribes that relatively little confidence in quantitative easing should imply relatively more of it.

In Figure 9, we consider various combinations of the policymaker's overall fear of

⁴⁸The concept of 'mean-preserving spread' connects our measure of asymmetry to the concept of 'second order stochastic dominance.' In decision theory, for any pair of lotteries A and B , A second-order stochastically dominates B if and only if a decision-maker weakly prefers A to B under every weakly increasing, concave utility function u . For any pair of lotteries A and B , B is a mean-preserving spread of A if and only if $y = x + \varepsilon$ for some $x \sim A$, $y \sim B$ and ε such that $\mathbb{E}[\varepsilon|x] = 0$ for all x . If B is a mean-preserving spread of A , then A second-order stochastically dominates B . Similarly, if B is a mean-preserving spread of A , then B has a higher variance than A .

misspecification (parameterized by θ) and his relative confidence in unconventional policy (parameterized by Ω_ζ) under the baseline calibration.⁴⁹ The surface in Figure 9 shows the difference in the policymaker’s value function under the approximating model versus the worst-case scenario – eight quarters after a shock to the IS curve.⁵⁰ Asymmetric fear enhances the policymaker’s precautionary motive. For a given degree of doubt over the model specification as a whole, a lower degree of relative confidence in quantitative easing exacerbates the policymaker’s loss.

Compare this finding with Williams (2013), who concludes in the case of parameter risk that “a concern for costs of unconventional policy not captured by the model mimics (or reinforces) policy attenuation.” In the case of model uncertainty, we find that greater relative doubt about unconventional policy mimics (or reinforces) policy *accentuation*. A Bayesian approach prescribes that “the optimal strategy is to rely on the instrument associated with the least uncertainty and use alternative, more uncertain instruments only when the least uncertain instrument is employed to its fullest extent possible.” A robust control approach – which assigns different weights to blocks of the approximating model over which the policymaker has detailed knowledge versus blocks over which he has very limited knowledge – prescribes the opposite. Faced with asymmetric fear of model misspecification, the policymaker’s optimal strategy is to be more aggressive with the instrument he associates with more uncertainty.

8 The FOMC’s Pessimism

The analysis so far has been normative. Our robust control approach prescribes that, when the policymaker has symmetric fear of model misspecification, he should respond more aggressively with both his conventional and unconventional policy instrument the more he doubts his model of monetary transmission. When his fear of model misspecification

⁴⁹In this section, we revert to parameterising the policymaker’s fear of model misspecification by the Lagrange multiplier on his time-zero discounted relative entropy constraint (θ) rather than his risk sensitivity (σ). This is because the observational equivalence between the multiplier problem and the risk sensitivity approach, which we exploited in Section 6, breaks down if we restrict the class of perturbations. Introducing differential penalty terms via $\Omega \neq \mathbb{I}$ is one such class restriction.

⁵⁰Since worst-case scenarios are dynamic and time-varying (as illustrated by the real-time ‘spikes’ in Figure 7), it is misleading to identify pessimism or optimism with *average* forecast differences over time between the approximating model and the worst-case projection. Instead, we calculate forecast differences between the two at *specific* horizons – eight quarters ahead, in the case of Figure 9.

is asymmetric, he should be relatively more aggressive with the instrument associated with relatively more uncertainty. How do these normative prescriptions compare with a positive analysis of how monetary policymakers *actually* responded in the wake of the global financial crisis of 2007/08?

In this section, we examine the policy decisions made by the U.S. Federal Open Market Committee – conditional on its macroeconomic forecasts at each meeting – to infer how the FOMC’s fears of misspecification may have evolved during the three phases of QE conducted since the end of 2008.

Romer and Romer (2013) contend that concerns about the effectiveness of unconventional monetary policy dampened the Federal Reserve’s responsiveness to the economic weakness that followed the end of the recession in 2009.⁵¹ Their qualitative analysis of statements made by FOMC members is supported by the quantitative analysis of Gagnon (2009), who argues that considerably more asset purchases were justified during QE2. Despite conceding that it is still too early to pass judgement on the appropriateness of this “muted policy response,” Romer and Romer (2013) largely attribute it to the FOMC’s “unwarranted pessimism” about the stimulatory effects of quantitative easing. In response, however, Kohn (2013) argues that over the course of QE1, QE2 and QE3, the FOMC – wisely – updated its beliefs about the effectiveness of QE and amended its reaction function accordingly.

In the upper panel of Figure 10 we plot, for each quarterly publication of the FOMC’s *Summary of Economic Projections* since the beginning of 2009, the average monthly rate of long-term asset purchases conducted by the Federal Reserve over that quarter against the forecast change in the unemployment rate over the following year (the upper-left plot) and the forecast change in the rate of core inflation (the upper-right plot). In the lower panel of Figure 10, we plot the range of FOMC meeting participants’ projections for the same variables over the same horizon.⁵² Each line segment refers to forecasts made and

⁵¹The authors point out, for example, that Chairman Bernanke and Vice Chair Yellen used exactly the same language to communicate their stance on the limited effectiveness of QE: “monetary policy is no panacea.”

⁵²There were 22 sets of forecasts published between January 2009 and March 2014. These forecasts detail the individual economic projections of each FOMC meeting participant, conditional on an assumed future path of “appropriate monetary policy” that each participant deems most likely to meet the Federal Reserve’s objectives. The monthly rate of assets purchased includes Treasury bonds, agency debt and agency mortgage-backed securities. The forecast change in macroeconomic variables over the following

decisions taken in each phase of the Federal Reserve’s LSAP programme: we define QE1 as the period from Q1:2009 to Q1:2010, QE2 as the period from Q2:2010 to Q3:2012, and QE3 as the period from Q4:2012 to Q1:2014.

Recall that in Section 6 we solved numerically for our policymaker’s optimal policy function: $u = -\mathbb{F}x$. We found that the optimal responsiveness of conventional policy to output ($\mathbb{F}_{ey}$), of conventional policy to inflation ($\mathbb{F}_{e\pi}$), of unconventional policy to output ($\mathbb{F}_{by}$) and of unconventional policy to inflation ($\mathbb{F}_{b\pi}$) were each increasing in the absolute value of the policymaker’s risk sensitivity σ . Our robust control approach endogenises pessimism. It suggests that more aggressive use of quantitative easing – in response to a *given* deviation in output and in inflation – is attributable to greater fear of model misspecification.

The upper panel of Figure 10 provides tentative evidence that both $\mathbb{F}_{by}$ and $\mathbb{F}_{b\pi}$ decreased between QE1 and QE2, but increased between QE2 and QE3. The FOMC’s reaction curves for both output and inflation shifted inwards at the turn of 2010 and shifted outwards in mid-2012: a given forecast deviation in unemployment and core inflation generally prompted a less aggressive policy response during QE2 than during QE1, and a more aggressive policy response during QE3 than during QE2.

Some specific examples illustrate this pattern. In Q4:2009, the FOMC expected the unemployment rate to fall by 0.5 percentage points over the following year. It purchased an average of \$70 billion in long-term assets per month during that period. In Q4:2011, having forecast exactly the same decline in the unemployment rate over the following year, FOMC asset purchases were effectively zero. In Q1:2014, again having projected a 0.5 percentage point decline in unemployment, the FOMC purchased an average of \$51 billion in long-term assets per month.

year is the difference – between adjacent forecast intervals – of the mid-point of the central tendency of FOMC meeting participants’ projections. The central tendency excludes the three highest and the three lowest projections for each variable in each year. The range includes all participants’ projections, and simply subtracts the lowest forecast from the highest. Although the model we have described in Section 4 is in terms of output (y) and inflation (π), in practice the FOMC’s reaction function can also be understood in terms of unemployment and core inflation. The Federal Reserve’s statutory mandate is to deliver “maximum employment, stable prices, and moderate long-term interest rates.” In the March 2014 Minutes, for example, we are told: “Participants’ views of the appropriate path for monetary policy were informed by their judgements about the state of the economy, including the values of the unemployment rate and other labor market indicators that would be consistent with maximum employment, the extent to which the economy was currently falling short of maximum employment, the prospects for inflation to reach the Committee’s longer-term objective of 2 percent, and the balance of risks around the outlook.”

In Q1:2010, the FOMC expected core inflation to increase by 0.1 percentage point over the following year. It purchased an average of \$39 billion in long-term assets per month during that period. In Q2:2011, having forecast exactly the same change in core inflation over the following year, the FOMC purchased just \$3 billion in long-term assets. In Q4:2012, again having projected a 0.1 percentage point increase in core inflation, the FOMC purchased a monthly average of \$76 billion in long-term assets.

It is with this type of reasoning that Romer and Romer (2013) impute “unwarranted pessimism” about the effectiveness of QE to FOMC members after the end of the recession in 2009. According to their analysis, the fact that the FOMC reacted so aggressively in late 2012 to the same set of economic forecasts to which it reacted so cautiously in late 2011 implies that QE2 (approximately the period from mid-2010 to mid-2012) was a period of undue passivism in unconventional policy.

Our analysis, however, rationalises this passivism: a less aggressive policy response to the same deviation in output and inflation is consistent with a policymaker having less fear of model misspecification. Using robust control techniques, Ellison and Sargent (2012) find convincing evidence that the FOMC’s forecasts can be interpreted “as worst-case scenarios that inform robust policy designed to confront specification doubts.” Our approach suggests that, rather than being unwarranted, the FOMC’s pessimism during QE2 can be seen as the robustly optimal response to lower model uncertainty. The reaction curves in Figure 10 shifted inwards after QE1 because FOMC policymakers became relatively more confident in the specification of their approximating model. This interpretation is consistent with the bottom panel of Figure 10: if the dispersion of forecasts among individual participants captures the FOMC’s doubts about the effectiveness of asset purchases, fears of misspecification indeed seem to have receded steadily between QE1 and QE2. As Kohn (2013) points out, gaining “experience with the prevailing economic situation and the effects of policy” helps to “gradually clarify the costs and benefits of [unconventional] policy choices,” making the case “for slowly altering the stance of policy as you learn about its effects.”⁵³

If the range of participants’ forecasts accurately captures the degree of uncertainty implicit in the FOMC’s projections, then fears of misspecification also seem to have re-

⁵³After a forty-year career in the Federal Reserve System, Kohn contributed to his final set of FOMC participants’ projections in June 2010 – amidst QE2.

ceded (or at least stabilised) between QE2 and QE3.⁵⁴ A robust control interpretation therefore suggests that, rather than the FOMC’s relative activism during QE3 implying that its passivism during QE2 was undue, it may well have been the FOMC’s passivism during QE2 that was warranted and its return to activism during QE3 that was undue.

9 Conclusion

After the Stock Market Crash of 1929, *The Magazine of Wall Street* described uncertainty as worse than knowing the truth, no matter how bad. Amidst the Great Depression, F. Scott Fitzgerald described intelligence as the ability to hold conflicting ideas in the mind at the same time, yet retain the ability to function. These two ideas are the hallmarks of this paper. Our monetary policymaker is introspective; the economic environment is precarious.

We embed fear of model misspecification in the mind of a self-critical central banker who is required to stabilise the economy in a post-crisis climate of unknown unknowns. As an instrument for attaining a robust policy rule, the central banker contemplates the reactions of a hypothetical, malevolent agent. This malevolent agent responds to the central banker’s choice of policy function by choosing the shock process with the most pernicious serial correlation structure it can obtain. The central banker does not *actually* expect this worst-case scenario to occur. By planning against it, however, he insures himself against unknown unknowns. As Sargent puts it: “it’s the government’s ideas about things that don’t happen that influence the things that do happen.”

The conclusions of this analysis are specific to the baseline model of monetary transmission described in Section 4 and augmented in Section 7. That model of the macroeconomic effects of central bank purchases of government debt is clearly over-simplified.

⁵⁴In the Minutes to the FOMC’s March 2014 meeting, for example, it was noted that the “range of views” held by participants reflected their individual assessments of “the rate at which the headwinds that have been holding back the pace of the economic recovery would abate...and the appropriate path for monetary policy.” In recent years, FOMC meeting participants have also explicitly reported their respective “levels of uncertainty” about their economic projections. The number of individuals reporting that the uncertainty attending his/her unemployment forecast was “higher than usual” (as opposed to “broadly similar” or “lower”) fell steadily from a peak of 18 individuals in June 2012 to 2 individuals in March 2014. Self-reported uncertainty also fell steadily in the FOMC’s core inflation forecasts, with 12 individuals reporting “higher than usual” levels of uncertainty in June 2011 falling to 3 individuals in March 2014.


As such, our quantitative findings are merely illustrative. Their qualitative implications, however, are useful in thinking about how monetary policy *should* have been conducted optimally in the aftermath of the global financial crisis of 2007/08.

Contrary to a Bayesian treatment of parameter risk, robust control suggests that greater model uncertainty calls for more aggressive use of both conventional and unconventional monetary policy instruments. This precautionary activism arises because, faced with the fear that his model of monetary transmission is misspecified, the policymaker responds to the serially *uncorrelated* shocks that characterise his approximating model with the aggressiveness that would be required if those shocks were serially *correlated*.

Robustness also confronts the ‘pretence of knowledge’ syndrome in macroeconomics. Robust control allows the policymaker to express more confidence in some aspects of his model than others. With both a conventional and unconventional monetary policy tool at his disposal, we find that ‘asymmetric fear’ of model misspecification enhances the policymaker’s precautionary motive. For a given degree of uncertainty about his model of monetary transmission as a whole, the robust control approach prescribes that relatively little confidence in quantitative easing should imply relatively more of it.


More work is to be done to discern whether monetary policy *was* conducted optimally in the aftermath of the global financial crisis. The approach we have taken by no means constitutes a formal econometric analysis, but it does suggest that robust control can be used to infer a policymaker’s fear of model misspecification from empirical observation of his reaction function. The U.S. Federal Open Market Committee’s behaviour during three phases of post-crisis quantitative easing is a case in point. Rather than the FOMC’s relative passivism during QE2 constituting evidence of “unwarranted pessimism,” it can be rationalised as the robustly optimal response to less fear of model misspecification following QE1. Furthermore, given that the FOMC’s fears of misspecification seemed to recede further following QE2, it may have been the Federal Reserve’s return to activism during QE3 that was unwarranted.

Figure 1: WITH CONVENTIONAL POLICY CONSTRAINED BY THE LOWER BOUND, MAJOR CENTRAL BANKS RESORTED TO UNCONVENTIONAL EASING




Notes: The short-term interest rate captures the stance of ‘conventional’ monetary policy. The total assets held on the balance sheet of the central bank is a proxy for the stance of ‘unconventional’ monetary policy.

Figure 2: REASONABLE DETECTION ERROR PROBABILITIES
ENSURE THAT THE EXTENT OF MISSPECIFICATION FEAR IS EMPIRICALLY PLAUSIBLE




Notes: “No fear” of model misspecification corresponds to $P(\sigma = 0) = 0.5$

Figure 3: GREATER FEAR OF MODEL MISSPECIFICATION
REQUIRES MORE AGGRESSIVE USE OF BOTH POLICY INSTRUMENTS




Notes: The robustly optimal responsiveness of conventional policy to output corresponds to the element $\mathbb{F}_{ey}$ (for various values of σ). The responsiveness of conventional policy to inflation is $\mathbb{F}_{e\pi}$; the responsiveness of unconventional policy to output is $\mathbb{F}_{by}$; the responsiveness of unconventional policy to inflation is $\mathbb{F}_{b\pi}$.

Figure 4: MORE MISSPECIFICATION FEAR IMPLIES THAT THE EFFECT OF STRUCTURAL SHOCKS ON STATE VARIABLES DISSIPATES FASTER




Notes: The rows correspond to impulse response functions (IRFs) following shocks to the IS curve, the Phillips curve, and the asset market equation, respectively. The columns correspond to the responses of output, of inflation, of conventional monetary policy and of unconventional policy, respectively.

Figure 5: IN THE WORST CASE, THE POLICYMAKER INTERPRETS SHOCKS AS LARGER AND MORE PERSISTENT THAN ENVISAGED UNDER HIS APPROXIMATING MODEL




Notes: The columns correspond to the response of each worst-case conditional mean to a common shock (to the IS curve).

Figure 6: THE WORST-CASE SCENARIO ENTAILS
A DEEPER RECESSION AND HIGHER INFLATION




Notes: 'Data (under app.)' refers to the simulated history of each state variable under the approximating model;
'F'cast (under app.)' refers to its simulated forecast path under the approximating model;
'F'cast (under dist.)' refers to its simulated forecast path under the worst-case distorted model.

Figure 7: IN REAL TIME, DEPARTURES FROM STEADY-STATE THREATEN TO BE
LARGER FOR LONGER UNDER EACH WORST-CASE COUNTERFACTUAL




Notes: Each successive red 'spike' represents the worst-case counterfactual that the policymaker envisages in real time (under the distorted model) from that quarter onwards. These spikes are re-estimated quarter by quarter.

Figure 8: GREATER RISK SENSITIVITY IMPLIES
THE POLICYMAKER'S OBJECTIVE IS MORE ROBUST TO MISSPECIFICATION




Notes: The value of σ_2 regulates the degree of model misspecification determining the worst-case scenario in the malevolent agent's minimisation problem. The value of σ_1 – for which only two settings are considered above – regulates the degree of model misspecification determining the robustly optimal policy function in the policymaker's maximisation problem.

Figure 9: GREATER ASYMMETRY AMPLIFIES THE POLICYMAKER'S LOSS
JUST AS A UNIFORM INCREASE IN OVERALL FEAR OF MODEL MISSPECIFICATION



Notes: On the x -axis, a higher value implies greater asymmetry (equivalently, lower relative confidence in unconventional policy, and a lower Ω_C). On the y -axis, a higher value implies greater overall fear of model misspecification (a lower θ). On the (vertical) z -axis, a lower value implies a greater relative loss for the policymaker.

Figure 10: THE FOMC'S FORECASTS, DECISIONS AND DISAGREEMENT DURING QE1, QE2 AND QE3



Notes: See text for detailed description of variables.

A

Appendix A accompanies Section 3.

A.1 Solving the Linear-Quadratic Gaussian Robust Control Problem

By modified certainty equivalence, the policymaker's linear-quadratic Gaussian (LQG) robust control problem is deterministic. Its corresponding Bellman equation is labelled (III) in the main text:

$$\begin{aligned} -x' \mathbb{V} x &= \max_u \min_w \{-x' Q x - u' R u + \beta \theta w' w - \beta x'_+ \mathbb{V} x_+\} \\ \text{s.t.} \\ x_+ &= A x + B u + C w \end{aligned} \tag{III}$$

Postulating some unspecified (but idempotent) matrix $\mathbb{P}$, suppose the policymaker's candidate value function is $-x' \mathbb{P} x + p$. We can set $p = 0$ because – by modified certainty equivalence – there can be no stochastic element in his value function.

The solution to the inner minimisation problem in (III), for a given $\mathbb{P}$, is then a value function $-x'_+ \mathcal{D}(\mathbb{P}) x'_+$ where:

$$\begin{aligned} -x'_+ \mathcal{D}(\mathbb{P}) x'_+ &= \min_w \{\theta w' w - x'_+ \mathbb{P} x_+\} \\ \text{s.t.} \\ x_+ &= A x + B u + C w \end{aligned} \tag{A.1}$$

Substituting for x_+ from the distorted transition equation that constrains (A.1), this inner minimisation problem can be solved for $\mathcal{D}(\mathbb{P})$ by standard linear-quadratic techniques. The solution is:

$$\mathcal{D}(\mathbb{P}) = \mathbb{P} + \mathbb{P} C (\theta \mathbb{I} - C' \mathbb{P} C)^{-1} C' \mathbb{P} \tag{A.2}$$

The outer maximisation problem in (III) can then be written as:

$$\begin{aligned}
-x'_+ \mathbb{P} x'_+ &= \max_u \{-x'Qx - u'Ru - \beta x'_+ \mathcal{D}(\mathbb{P})x_+\} \\
&\text{s.t.} \\
x_+ &= Ax + Bu
\end{aligned} \tag{A.3}$$

This outer problem has the solution:

$$\mathbb{P} = Q + \beta A' \mathcal{D}(\mathbb{P}) A - \beta^2 A' \mathcal{D}(\mathbb{P}) B (R + \beta B' \mathcal{D}(\mathbb{P}) B)^{-1} B' \mathcal{D}(\mathbb{P}) A \tag{A.4}$$

The two solutions above, for $\mathcal{D}(\mathbb{P})$ and for $\mathbb{P}$, are known as algebraic Riccati equations. Iterating simultaneously over both (A.2) and (A.4) solves the decision-maker's original max-min problem (III) for $\mathcal{D}(\mathbb{P})$ and $\mathbb{P}$.⁵⁵

The robustly optimal policy function and the worst-case distortion are then given by:

$$u = -\mathbb{F}x \tag{A.5}$$

and

$$w = \mathbb{K}x \tag{A.6}$$

where

$$\mathbb{F} = \beta (R + \beta B' \mathcal{D}(\mathbb{P}) B)^{-1} B' \mathcal{D}(\mathbb{P}) A \tag{A.7}$$

and

$$\mathbb{K} = \theta^{-1} (\mathbb{I} - \theta^{-1} C' \mathbb{P} C)^{-1} C' \mathbb{P} (A - B \mathbb{F}) \tag{A.8}$$

as discussed in Section 3.3 of the main text.

⁵⁵These iterative routines are performed in MATLAB.

A.2 The observational equivalence between robust control and risk sensitivity

Suppose we begin with the most general specification. A decision-maker wishes to guard against distortions to an approximating density $\pi(\varepsilon)$, where ε is a dummy variable with the same number of elements as in the sequence of shocks $\{\epsilon_t\}$ entering the transition equation of the approximating model. The probability density of the distorted model is conditioned on an initial value of the state variable y_0 and is evaluated for the history of shocks ϵ^t . It is denoted $\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)$. As described in Section 3 of the main text, the decision-maker seeks a policy that performs well when $\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)$ is close to $\pi(\varepsilon)$. The likelihood ratio is defined by:

$$m_{t+1} = \frac{\hat{\pi}(\epsilon_{t+1} \mid \epsilon^t, y_0)}{\pi(\epsilon_{t+1})} \quad (\text{A.9})$$

Taking expectations of this likelihood ratio with respect to the approximating model yields:

$$\mathbb{E}(m_{t+1} \mid \epsilon^t, y_0) = \int \frac{\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)}{\pi(\varepsilon)} \pi(\varepsilon) d\varepsilon = 1 \quad (\text{A.10})$$

because $\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)$ is a density function.

Now suppose we define a variable as $M_0 = 1$. We then recursively construct a series $\{M_t\}$ according to $M_{t+1} = m_{t+1}M_t$. The random variable M_t is now a ratio of the joint densities of ϵ^t , conditioned on y_0 and evaluated for the history ϵ^t . Rolling the recursion forward, we can also write M_t as the factorisation of the joint density:

$$M_t = \prod_{j=0}^t m_j \quad (\text{A.11})$$

Since M_{t+1} satisfies:

$$\mathbb{E}(M_{t+1} \mid \epsilon^t, y_0) = \mathbb{E}(m_{t+1}M_t \mid \epsilon^t, y_0) = \mathbb{E}(m_{t+1} \mid \epsilon^t, y_0)M_t = M_t \quad (\text{A.12})$$

we know that M_t is a Martingale relative to the sequence of information sets generated by the shocks.

As in equation (5) of the main text, the entropy of the distortion associated with M_t is defined as the expectation of the log-likelihood ratio with respect to the distorted

distribution:

$$\int \log \left(\frac{\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)}{\pi(\varepsilon)} \dots \frac{\hat{\pi}(\varepsilon \mid \epsilon^0, y_0)}{\pi(\varepsilon)} \right) \times \hat{\pi}(\varepsilon \mid \epsilon^t, y_0) \dots \hat{\pi}(\varepsilon \mid \epsilon^0, y_0) d\varepsilon d\varepsilon = \mathbb{E}(M_t \log M_t \mid y_0) \quad (\text{A.13})$$

By definition, $\mathbb{E}(M_t \log M_t \mid y_0) \geq 0$ and we can write:

$$\mathbb{E}(M_t \log M_t \mid y_0) = \sum_{j=0}^{t-1} \mathbb{E}[M_j \mathbb{E}(m_{j+1} \log m_{j+1} \mid \epsilon^j, y_0) \mid y_0] \quad (\text{A.14})$$

where $m_{j+1} \log m_{j+1} \mid \epsilon^j, y_0$ is the conditional relative entropy of the pertrurbation to the one-step-ahead transition density. This definition of entropy, however, neglects discounting of the future. We work with a measure of *discounted* entropy of the form:

$$(1 - \beta) \sum_{j=0}^{\infty} \beta^j \mathbb{E}(M_j \log M_j \mid y_0) = \sum_{j=0}^{t-1} \beta^j \mathbb{E}[M_j \mathbb{E}(m_{j+1} \log m_{j+1} \mid \epsilon^j, y_0) \mid y_0] \quad (\text{A.15})$$

The general stochastic robust control problem can now be defined as:

$$\begin{aligned} \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{m_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \mathbb{E}[\beta^t M_t \{r(y_t, u_t) + \theta \beta \mathbb{E}(m_{t+1} \log m_{t+1} \mid \epsilon^t, y_0)\} \mid y_0] \\ \text{s.t.} \\ y_{t+1} = \varpi(y_t, u_t, \epsilon_{t+1}) \\ M_{t+1} = m_{t+1} M_t \end{aligned} \quad (\text{A.16})$$

where $y_{t+1} = \varpi(y_t, u_t, \epsilon_{t+1})$ is the state-evolution equation, the control process $\{u_t\}$ is a function of ϵ^t and y_0 , $r(y_t, u_t)$ is the decision-maker's return function, and y_0 is the initial condition for the state variable. The likelihood ratio m_{t+1} is defined as the ratio of the approximating and distorted densities as before (hence it is a function of ϵ^{t+1} and y_0) and $\mathbb{E}(m_{t+1} \mid \epsilon^t, y_0) = 1$, where the expectation is taken with respect to the distribution of the approximating model. $\theta \in [\underline{\theta}, +\infty)$ is a penalty parameter, the lower bound of which is the 'point of neurotic breakdown.'

We require a recursive formulation to solve this general stochastic robust control problem. Since the predetermined variables in the system are M_t and y_t , we postulate that the value function $\mathbb{W}(M, y)$ has a multiplicative form $M\mathbb{V}(y)$. This admits the following Bellman equation:

$$\begin{aligned} M\mathbb{V}(y) = \max_u \min_{m(\varepsilon)} \{ & r(y, u) + \beta \int (m(\varepsilon)\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta m(\varepsilon) \log m(\varepsilon)) \pi(\varepsilon) d\varepsilon \} \\ \text{s.t.} & \\ & \int m(\varepsilon) \int (\varepsilon) d\varepsilon = 1 \end{aligned} \quad (\text{A.17})$$

Equivalently, we can work with:

$$\begin{aligned} \mathbb{V}(y) = \max_u \min_{m(\varepsilon)} \{ & r(y, u) + \beta \int (m(\varepsilon)\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta m(\varepsilon) \log m(\varepsilon)) \pi(\varepsilon) d\varepsilon \} \\ \text{s.t.} & \\ & \int m(\varepsilon) \int (\varepsilon) d\varepsilon = 1 \end{aligned} \quad (\text{A.18})$$

The inner minimisation problem $\mathcal{R}$ of (A.18) is:

$$\begin{aligned} \mathcal{R}(\mathbb{V})(y, u) = \min_{m(\varepsilon)} \{ & \int (m(\varepsilon)\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta m(\varepsilon) \log m(\varepsilon)) \pi(\varepsilon) d\varepsilon \} \\ \text{s.t.} & \\ & \int m(\varepsilon) \int (\varepsilon) d\varepsilon = 1 \end{aligned} \quad (\text{A.19})$$

This problem involves a convex objective and a linear constraint; it admits the following Lagrangean:

$$\mathcal{L} = \min_{m(\varepsilon)} \int (m(\varepsilon)\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta m(\varepsilon) \log m(\varepsilon)) \pi(\varepsilon) d\varepsilon + \theta(1 + \tilde{\lambda}) \left(1 - \int m(\varepsilon) \pi(\varepsilon) d\varepsilon \right) \quad (\text{A.20})$$

where we have defined the Lagrange multiplier as $\theta(1 + \tilde{\lambda})$ for convenience.

The first-order condition with respect to $m(\varepsilon)$ is:

$$\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta(1 + \log m(\varepsilon)) - \theta(1 + \tilde{\lambda}) = 0 \quad (\text{A.21})$$

from which it follows that:

$$m(\varepsilon) = e^{-\frac{1}{\theta}\mathbb{V}(\varpi(y, u, \varepsilon))} e^{\tilde{\lambda}} \quad (\text{A.22})$$

Since $\int m(\varepsilon)\pi(\varepsilon) d\varepsilon = 1$, we can integrate (A.22) to solve for $\tilde{\lambda}$:

$$\int m(\varepsilon)\pi(\varepsilon) d\varepsilon = e^{\tilde{\lambda}} \int e^{-\frac{1}{\theta}\mathbb{V}(\varpi(y, u, \varepsilon))} \pi(\varepsilon) d\varepsilon = 1 \quad (\text{A.23})$$

We thereby define the worst-case distortion $m(\varepsilon)$ as:

$$m(\varepsilon) = \frac{e^{-\frac{1}{\theta}\mathbb{V}(\varpi(y, u, \varepsilon))}}{\int e^{-\frac{1}{\theta}\mathbb{V}(\varpi(y, u, \varepsilon))} \pi(\varepsilon) d\varepsilon} \quad (\text{A.24})$$

Letting $\theta \equiv -\sigma^{-1}$ and using (A.24), we can re-write the minimand of the inner minimisation problem $\mathcal{R}$ as:

$$\int (m(\varepsilon)\mathbb{V}(\varpi(y, u, \varepsilon)) + \theta m(\varepsilon) \log m(\varepsilon)) \pi(\varepsilon) d\varepsilon = \frac{1}{\sigma} \log \left(\int e^{\sigma\mathbb{V}(\varpi(y, u, \varepsilon))} \pi(\varepsilon) d\varepsilon \right) \quad (\text{A.25})$$

We can now return to the outer maximisation problem of (A.18), and state it as:

$$\max_u \left\{ r(y, u) + \beta \frac{1}{\sigma} \log \left(\int e^{\sigma\mathbb{V}(\varpi(y, u, \varepsilon))} \pi(\varepsilon) d\varepsilon \right) \right\} \quad (\text{A.26})$$

which admits the Bellman equation:

$$\mathbb{V}(y) = \max_u \left\{ r(y, u) + \frac{1}{\sigma} \log \left(\mathbb{E}[e^{\sigma\mathbb{V}(y_+)}] \right) \right\} \quad (\text{A.27})$$

where the subscript $+$ denotes next-period variables.

This restatement of the problem shows that solving a robust control problem by imposing a constraint on conditional relative entropy is mathematically equivalent to the policymaker having preferences (parameterised by $\sigma \equiv -\theta^{-1}$) which apply an exponential transformation to his continuation value. These preferences are known as risk sensitive preferences (following Whittle (1990)). They are a special case of the specification de-

scribed by Epstein and Zin (1989), in which preferences are no longer constrained such that the intertemporal elasticity of substitution is equal to the coefficient of risk aversion. In our case, the Lagrange multiplier on the decision-maker's entropy constraint (θ) maps directly into the parameter regulating the decision-maker's risk sensitivity (σ).

The risk sensitive formulation employs the approximating – not the distorted – model, because the policymaker by definition has no fear of model misspecification. Suppose (i) that the policymaker's return function is defined over state and control variables as: $r(y, u) = -y'Qy - u'Ru$, (ii) that the policymaker's value function is quadratic, and (iii) that his approximating model of state transition is given by: $y_{t+1} = Ay_t + Bu_t + C\varepsilon_{t+1}$. Jacobsen (1973) shows that the policymaker's Bellman equation (A.27) is then equivalent to:

$$-y'\nabla y = \max_u \{ -y'Qy - u'Ru + \beta \mathbb{E}[y'_+ \mathcal{D}(\mathbb{V})y_+] \} \quad (\text{A.28})$$

where

$$\mathcal{D}(\mathbb{V}) = \mathbb{V} - \sigma \mathbb{V}C(\mathbb{I} + \sigma C'\mathbb{V}C)^{-1}C'\mathbb{V} \quad (\text{A.29})$$

A.3 The simplification of conditional relative entropy in the LQG case

In the main text, we have used the result that conditional relative entropy in the linear-quadratic Gaussian case simplifies to $I(w_{t+1}) = \frac{1}{2}w'_{t+1}w_{t+1}$. To see why this is the case, note first that the approximating density $\pi(\varepsilon)$ is Gaussian by definition. Given modified certainty equivalence, the distorted density is only distorted with respect to its conditional mean: it too will have a Gaussian distribution, and the likelihood ratio reduces to the ratio of two Gaussian densities.

Following from Appendix A.2, assume without loss of generality that the approximating density is standard Normal: $\pi(\varepsilon) \sim N(0, \mathbb{I})$. The distorted density $\hat{\pi}(\varepsilon \mid \varepsilon^t, y_0)$ has a conditional mean of w and a variance-covariance matrix $\mathbb{I}$. From the definition of a multivariate normal density, it follows that the log-likelihood is given by:

$$\log \left(\frac{\hat{\pi}(\varepsilon \mid \varepsilon^t, y_0)}{\pi(\varepsilon)} \right) = \frac{1}{2} (-(\varepsilon - w)'(\varepsilon - w) + \varepsilon'\varepsilon) \quad (\text{A.30})$$

As discussed in Appendix A.2, the definition of conditional relative entropy requires

calculating the log-likelihood ratio expected under the distorted model:

$$\int \log \left(\frac{\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)}{\pi(\varepsilon)} \right) \hat{\pi}(\varepsilon \mid \epsilon^t, y_0) d\varepsilon = \frac{1}{2} \int (-(\varepsilon - w)'(\varepsilon - w) + \varepsilon' \varepsilon) \hat{\pi}(\varepsilon \mid \epsilon^t, y_0) d\varepsilon \quad (\text{A.31})$$

Since the distorted density is Gaussian, we know that:

$$- \int \left(\frac{1}{2}(\varepsilon - w)'(\varepsilon - w) + \varepsilon' \varepsilon \right) \hat{\pi}(\varepsilon \mid \epsilon^t, y_0) d\varepsilon = -\frac{1}{2} \text{Tr}(\mathbb{I}) \quad (\text{A.32})$$

where $\text{Tr}(\mathbb{I})$ is the trace of the identity matrix.

In addition, if we write $\varepsilon = w + (\varepsilon - w)$ then:

$$\int (\varepsilon' \varepsilon) d\varepsilon = \int \left(\frac{1}{2} w' w + \frac{1}{2} (\varepsilon - w)'(\varepsilon - w) + w'(\varepsilon - w) \right) d\varepsilon = \frac{1}{2} w' w + \frac{1}{2} \text{Tr}(\mathbb{I}) \quad (\text{A.33})$$

Combining (A.32) and (A.33), the definition of conditional relative entropy reduces to:

$$\int \log \left(\frac{\hat{\pi}(\varepsilon \mid \epsilon^t, y_0)}{\pi(\varepsilon)} \right) \hat{\pi}(\varepsilon \mid \epsilon^t, y_0) d\varepsilon = \frac{1}{2} w' w \quad (\text{A.34})$$

which is the simplification used in the main text.

B

Appendix B accompanies Section 5.

B.1 The policymaker's problem in state-space form

Begin with the policymaker's approximating model, described by equations (16)-(18) in the main text:

$$y_t = -\tau r_{t-1} - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t$$

$$e_t = \psi r_t + \mu b_t + v_t$$

Re-write the third equation in terms of r_t and substitute it into the first equation. This yields:

$$y_t = -\frac{\tau}{\psi} (e_{t-1} - \mu b_{t-1} - v_{t-1}) - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t \quad (\text{B.1})$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \quad (\text{B.2})$$

Gathering like terms and rolling forward one period yields:

$$y_{t+1} = \left(-\frac{\tau}{\psi} - \delta\right) e_t + \left(\frac{\tau\mu}{\psi} - \phi\right) b_t + \frac{\tau}{\psi} v_t + \lambda y_t + \xi_{t+1} \quad (\text{B.3})$$

$$\pi_{t+1} = \pi_t + \alpha y_t - \gamma e_t + \gamma e_{t-1} + \eta_{t+1} \quad (\text{B.4})$$

Now let the state, control and shock vectors of the system be:

$$x_t \equiv \begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix} \quad u_t \equiv \begin{pmatrix} e_t \\ b_t \end{pmatrix} \quad \varepsilon_t \equiv \begin{pmatrix} \xi_t \\ \eta_t \\ v_t \end{pmatrix}$$

In state-space form, equations (B.3) and (B.4) for y_{t+1} and π_{t+1} then constitute the following system, which is a special case of: $x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1}$. We have:

$$\underbrace{\begin{pmatrix} y_{t+1} \\ \pi_{t+1} \\ e_t \\ b_t \\ v_{t+1} \end{pmatrix}}_{x_{t+1}} = \underbrace{\begin{pmatrix} \lambda & 0 & 0 & 0 & \frac{\tau}{\psi} \\ \alpha & 1 & \gamma & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_A \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix}}_{x_t} + \underbrace{\begin{pmatrix} -\frac{\tau}{\psi} - \delta & \frac{\tau\mu}{\psi} - \phi \\ -\gamma & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}}_B \underbrace{\begin{pmatrix} e_t \\ b_t \end{pmatrix}}_{u_t} + \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}}_C \underbrace{\begin{pmatrix} \xi_{t+1} \\ \eta_{t+1} \\ v_{t+1} \end{pmatrix}}_{\varepsilon_{t+1}}$$

Given this approximating model of state transition, the vector of perturbations specialises in our model to:

$$w_{t+1} \equiv \begin{pmatrix} w_{\xi,t+1} \\ w_{\eta,t+1} \\ w_{v,t+1} \end{pmatrix}$$

Now let the co-efficient matrices in the policymaker's quadratic objective function be:

$$Q \equiv \begin{pmatrix} Q_{yy} & Q_{y\pi} & Q_{ye} & Q_{yb} & Q_{yv} \\ Q_{\pi y} & Q_{\pi\pi} & Q_{\pi e} & Q_{\pi b} & Q_{\pi v} \\ Q_{ey} & Q_{e\pi} & Q_{ee} & Q_{eb} & Q_{ev} \\ Q_{by} & Q_{b\pi} & Q_{be} & Q_{bb} & Q_{bv} \\ Q_{vy} & Q_{y\pi} & Q_{ve} & Q_{vb} & Q_{vv} \end{pmatrix} \quad R \equiv \begin{pmatrix} R_{ee} & R_{eb} \\ R_{be} & R_{bb} \end{pmatrix}$$

The policymaker's per-period return function becomes:

$$r(x_t, u_t) = - \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix}}_{x'_t} \underbrace{\begin{pmatrix} Q_{yy} & Q_{y\pi} & Q_{ye} & Q_{yb} & Q_{yv} \\ Q_{\pi y} & Q_{\pi\pi} & Q_{\pi e} & Q_{\pi b} & Q_{\pi v} \\ Q_{ey} & Q_{e\pi} & Q_{ee} & Q_{eb} & Q_{ev} \\ Q_{by} & Q_{b\pi} & Q_{be} & Q_{bb} & Q_{bv} \\ Q_{vy} & Q_{y\pi} & Q_{ve} & Q_{vb} & Q_{vv} \end{pmatrix}}_Q \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \end{pmatrix}}_{x_t} - \underbrace{\begin{pmatrix} e_t \\ b_t \end{pmatrix}}_{u'_t} \underbrace{\begin{pmatrix} R_{ee} & R_{eb} \\ R_{be} & R_{bb} \end{pmatrix}}_R \underbrace{\begin{pmatrix} e_t \\ b_t \end{pmatrix}}_{u_t}$$

which is a special case of the general return function (7) in the main text.

B.2 Detection error probabilities

Given a linear-quadratic Gaussian robust control problem with decision rules $u = -\mathbb{F}x$ for the policymaker and $w = \mathbb{K}x$ for the malevolent agent (where σ denotes the policymaker's risk sensitivity and both $\mathbb{F}$ and $\mathbb{K}$ are functions of σ), we can re-write the approximating and distorted models as:

$$x_{t+1} = \underbrace{(A - B\mathbb{F})}_{\mathbb{A}^{app}} x_t + C\varepsilon_{t+1} \tag{B.5}$$

and

$$x_{t+1} = \underbrace{(A - B\mathbb{F} + C\mathbb{K})}_{\mathbb{A}^{dist}} x_t + C\varepsilon_{t+1} \tag{B.6}$$

Conditional on the approximating model being the true data generating process, define $\mathcal{L}_{app}|app$ as the likelihood under the approximating model and define $\mathcal{L}_{dist}|app$ as the likelihood under the distorted model. If the log-likelihood ratio $\log\left(\frac{\mathcal{L}_{app}|app}{\mathcal{L}_{dist}|app}\right)$ is greater than zero, the test correctly selects the approximating model; if it is less than zero, the test erroneously selects the distorted model. Define $\mathcal{P}_{app}$ as the probability of the latter:

$$\mathcal{P}_{app} = Prob\left(\log\left(\frac{\mathcal{L}_{app}|app}{\mathcal{L}_{dist}|app}\right) < 0\right) \quad (\text{B.7})$$

Conditional on the distorted model being the true data generating process, define $\mathcal{L}_{app}|dist$ as the likelihood under the approximating model and define $\mathcal{L}_{dist}|dist$ as the likelihood under the distorted model. If the log-likelihood ratio $\log\left(\frac{\mathcal{L}_{app}|dist}{\mathcal{L}_{dist}|dist}\right)$ is less than zero, the test correctly selects the distorted model; if it is greater than zero, the test erroneously selects the approximating model. Define $\mathcal{P}_{dist}$ as the probability of the latter:

$$\mathcal{P}_{dist} = Prob\left(\log\left(\frac{\mathcal{L}_{app}|dist}{\mathcal{L}_{dist}|dist}\right) > 0\right) \quad (\text{B.8})$$

Combining the two log-likelihood tests specified above results in the detection error probability:

$$P(\sigma) = \frac{1}{2}(\mathcal{P}_{app} + \mathcal{P}_{dist}) \quad (\text{B.9})$$

Since it is not possible to calculate either $\mathcal{P}_{app}$ or $\mathcal{P}_{dist}$ for a finite data series, Anderson et al. (2003) recommend approximating $P(\sigma)$ by simulation.

We construct each likelihood according to $x_{t+1}^{app} \sim N(\mathbb{A}^{app}x_t^{app}, CC')$ and according to $x_{t+1}^{dist} \sim N(\mathbb{A}^{dist}x_t^{dist}, CC')$ by using the following definitions (noting that $p = 5$ since the vector of state variables x is five-dimensional):

$$\mathcal{L}_{app}|app = \frac{1}{(2\pi)^{\frac{Np}{2}}} \frac{1}{|CC'|^{\frac{N}{2}}} \exp\left(-\frac{1}{2} \sum_{t=1}^{T-1} (x_{t+1}^{app} - \mathbb{A}^{app}x_t^{app})' |CC'|^{-1} (x_{t+1}^{app} - \mathbb{A}^{app}x_t^{app})\right)$$

$$\mathcal{L}_{dist}|app = \frac{1}{(2\pi)^{\frac{Np}{2}}} \frac{1}{|CC'|^{\frac{N}{2}}} \exp\left(-\frac{1}{2} \sum_{t=1}^{T-1} (x_{t+1}^{dist} - \mathbb{A}^{app}x_t^{app})' |CC'|^{-1} (x_{t+1}^{dist} - \mathbb{A}^{app}x_t^{app})\right)$$

$$\mathcal{L}_{app}|dist = \frac{1}{(2\pi)^{\frac{N_P}{2}}} \frac{1}{|CC'|^{\frac{N}{2}}} \exp \left(-\frac{1}{2} \sum_{t=1}^{T-1} (x_{t+1}^{app} - \mathbb{A}^{dist} x_t^{dist})' |CC'|^{-1} (x_{t+1}^{app} - \mathbb{A}^{dist} x_t^{dist}) \right)$$

$$\mathcal{L}_{dist}|dist = \frac{1}{(2\pi)^{\frac{N_P}{2}}} \frac{1}{|CC'|^{\frac{N}{2}}} \exp \left(-\frac{1}{2} \sum_{t=1}^{T-1} (x_{t+1}^{dist} - \mathbb{A}^{dist} x_t^{dist})' |CC'|^{-1} (x_{t+1}^{dist} - \mathbb{A}^{dist} x_t^{dist}) \right)$$

To estimate $P(\sigma)$, we calculate the proportion of simulations for which the likelihood test erroneously selects one model when in fact the other has generated the data.

C

Appendix C accompanies Section 7.

C.1 The policymaker's problem in state-space form

Recall the policymaker's approximating model, augmented for the possibility of control errors when conducting quantitative easing (as in problem (V)):

$$y_t = -\tau r_{t-1} - \phi b_{t-1} - \delta e_{t-1} + \lambda y_{t-1} + \xi_t$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t$$

$$e_t = \psi r_t + \mu b_t + v_t$$

$$b_t = \chi q_t + \zeta_t$$

Using the fourth equation, we substitute for b_t in the third equation and for b_{t-1} in the first equation. Then writing the third equation in terms of r_t and substituting it into the first equation yields:

$$y_t = -\frac{\tau}{\psi} (e_{t-1} - \mu\chi q_{t-1} - \mu\zeta_{t-1} - v_{t-1}) - \phi (\chi q_{t-1} + \zeta_{t-1}) - \delta e_{t-1} + \lambda y_{t-1} + \xi_t \quad (\text{C.1})$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma(e_{t-1} - e_{t-2}) + \eta_t \quad (\text{C.2})$$

Gathering like terms and rolling forward one period yields:

$$y_{t+1} = \left(-\frac{\tau}{\psi} - \delta\right) e_t + \chi \left(\frac{\tau\mu}{\psi} - \phi\right) q_t + \frac{\tau}{\psi} v_t + \left(\frac{\tau\mu}{\psi} - \phi\right) \zeta_t + \lambda y_t + \xi_{t+1} \quad (\text{C.3})$$

$$\pi_{t+1} = \pi_t + \alpha y_t - \gamma e_t + \gamma e_{t-1} + \eta_{t+1} \quad (\text{C.4})$$

In state-space form, equations (C.3) and (C.4) for y_{t+1} and π_{t+1} then constitute the following system, which is a special case of: $x_{t+1} = Ax_t + Bu_t + C\varepsilon_{t+1}$. We have:

$$\underbrace{\begin{pmatrix} y_{t+1} \\ \pi_{t+1} \\ e_t \\ b_t \\ v_{t+1} \\ \zeta_{t+1} \end{pmatrix}}_{x_{t+1}} = \underbrace{\begin{pmatrix} \lambda & 0 & 0 & 0 & \frac{\tau}{\psi} & \frac{\tau\mu}{\psi} - \phi \\ \alpha & 1 & \gamma & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_A \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \\ \zeta_t \end{pmatrix}}_{x_t} + \underbrace{\begin{pmatrix} -\frac{\tau}{\psi} - \delta & \chi \left(\frac{\tau\mu}{\psi} - \phi\right) \\ -\gamma & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}}_B \underbrace{\begin{pmatrix} e_t \\ q_t \end{pmatrix}}_{u_t} + \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}}_C \underbrace{\begin{pmatrix} \xi_{t+1} \\ \eta_{t+1} \\ v_{t+1} \\ \zeta_{t+1} \end{pmatrix}}_{\varepsilon_{t+1}}$$

The policymaker's per-period return function becomes:

$$r(x_t, u_t) = - \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \\ \zeta_t \end{pmatrix}}_{x'_t} \underbrace{\begin{pmatrix} Q_{yy} & Q_{y\pi} & Q_{ye} & Q_{yq} & Q_{yv} & Q_{y\zeta} \\ Q_{\pi y} & Q_{\pi\pi} & Q_{\pi e} & Q_{\pi q} & Q_{\pi v} & Q_{\pi\zeta} \\ Q_{ey} & Q_{e\pi} & Q_{ee} & Q_{eq} & Q_{ev} & Q_{e\zeta} \\ Q_{qy} & Q_{q\pi} & Q_{qe} & Q_{qq} & Q_{qv} & Q_{q\zeta} \\ Q_{vy} & Q_{v\pi} & Q_{ve} & Q_{vq} & Q_{vv} & Q_{v\zeta} \\ Q_{\zeta y} & Q_{\zeta\pi} & Q_{\zeta e} & Q_{\zeta q} & Q_{\zeta v} & Q_{\zeta\zeta} \end{pmatrix}}_Q \underbrace{\begin{pmatrix} y_t \\ \pi_t \\ e_{t-1} \\ b_{t-1} \\ v_t \\ \zeta_t \end{pmatrix}}_{x_t} - \underbrace{\begin{pmatrix} e_t \\ q_t \end{pmatrix}}_{u'_t} \underbrace{\begin{pmatrix} R_{ee} & R_{eq} \\ R_{qe} & R_{qq} \end{pmatrix}}_R \underbrace{\begin{pmatrix} e_t \\ q_t \end{pmatrix}}_{u_t}$$

which is again a special case of the general return function (7) in the main text.

References

- Anderson, E., Hansen, L., and Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk and model detection. *Journal of the European Economic Association*, 1(1):68–123.
- Başar, T. S. and Bernhard, P. (1995). *H^∞ -Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Birkhäuser.
- Ball, L. M. (1999). Policy rules for open economies. *Monetary Policy Rules*, University of Chicago Press.
- Baumeister, C. and Benati, L. (2013). Unconventional monetary policy and the Great Recession: Estimating the macroeconomic effects of a spread compression at the zero lower bound. *International Journal of Central Banking*, 9(2):165–212.
- Bech, M. L. and Lengwiler, Y. (2012). The financial crisis and the changing dynamics of the yield curve. *Bank for International Settlements Papers*, 65:257–276.
- Bernanke, B. S. (2007). Remarks at the 32nd Annual Economic Policy Conference, Federal Reserve Bank of St. Louis: “Monetary policy under uncertainty”.
- Bernanke, B. S. (2012). Remarks at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming: “Monetary policy since the onset of the crisis”.
- Blinder, A. S. (1999). *Central Banking in Theory and Practice*. The MIT Press.
- Brainard, W. C. (1967). Uncertainty and the effectiveness of policy. *The American Economic Review*, 57(2):411–425.
- Bullard, J. (2013). “Four questions for current monetary policy.” Remarks at the New York Association for Business Economics luncheon.
- Caballero, R. J. (2010). Macroeconomics after the crisis: Time to deal with the pretense-of-knowledge syndrome. *Journal of Economic Perspectives*, 24(4):85–102.
- Chen, H., Curdia, V., and Ferrero, A. (2012). The macroeconomic effects of large-scale asset purchase programs. *Economic Journal*, 122(564):F289 – F315.
- Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

- Chung, H., Laforte, J.-P., Reifschneider, D., and Williams, J. C. (2012). Have we underestimated the likelihood and severity of zero lower bound events? *Journal of Money, Credit and Banking*, 44:47–82.
- Cochrane, J. H. (2011). “Is QE2 a saviour, inflator, or a dud? Business Class”: *Bloomberg View*. <http://www.bloombergview.com/articles/2011-06-02/>.
- Curdia, V. and Ferrero, A. (2013). How stimulatory are Large-Scale Asset Purchases? *Federal Reserve Bank of San Francisco Economics Letters*, 2013(22).
- Curdia, V. and Woodford, M. (2011). The central-bank balance sheet as an instrument of monetary policy. *Journal of Monetary Economics*, 58:54–79.
- D’Amico, S. and King, T. B. (2013). Flow and stock effects of Large-Scale Treasury Purchases: Evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448.
- Dennis, R. (2013). Imperfect credibility and robust monetary policy. *Scottish Institute for Research in Economics (SIRE) Discussion Papers*, 78.
- Eggertsson, G. B. and Woodford, M. (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 34(1):139–235.
- Ellison, M. and Sargent, T. J. (2012). A defense of the FOMC. *International Economic Review*, 53(4):1047–1065.
- Ellison, M. and Tischbirek, A. (2013). Unconventional government debt purchases as a supplement to conventional monetary policy. *Economics Series Working Papers, University of Oxford*, 679.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.
- Epstein, L. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behaviour of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.
- Fitzgerald, F. S. (1936). The Crack-Up. *Esquire*, February 1936.
- Friedman, M. (1961). The lag effect of monetary policy. *Journal of Political Economy*, 69(5):447–466.

- Gagnon, J., Raskin, M., Remache, J., and Sack, B. (2011). Large-scale asset purchases by the Federal Reserve: did they work? *Economic Policy Review*, pages 41–59.
- Gagnon, J. E. (2009). The world needs further monetary ease, not an early exit. *Peterson Institute for International Economics Policy Brief*, PB09(22).
- Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.
- Greenspan, A. (2004). Remarks at the Meetings of the American Economic Association, San Diego, California: “Risk and uncertainty in monetary policy”.
- Haldane, A. (2013). The banker who cried ‘simplicity.’ The Saturday Interview: *The Wall Street Journal*.
- Hamilton, J. D. and Wu, J. C. (2012). The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking*, 44:3–46.
- Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, Princeton.
- Jacobson, D. H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control*, 18(2):124–131.
- Joyce, M., Tong, M., and Woods, R. (2011). The United Kingdom’s quantitative easing policy: design, operation and impact. *Bank of England Quarterly Bulletin 2011 Q3*, 51(3):200–212.
- Kashyap, A. K. (2013). Comments on “The Ins and Outs of LSAPs”.
- Kay, J. (2012). “The other multiplier effect, or Keynes’s view of probability”: *Financial Times*.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.
- Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Hart, Schaffner and Marx, Boston.

- Kohn, D. L. (2013). Comments on “The most dangerous idea in Federal Reserve history: Monetary policy doesn’t matter”. <http://www.brookings.edu/research/papers>.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*, 42(2):215–287.
- Leitemo, K. and Söderström, U. (2004). Robust monetary policy in the New Keynesian framework. *CEPR Discussion Papers*, 4805.
- Leitemo, K. and Söderström, U. (2008). Robust monetary policy in a small open economy. *Journal of Economic Dynamics and Control*, 32(10).
- Levin, A. T. and Williams, J. C. (2003). Robust monetary policy with competing reference models. *Journal of Monetary Economics*, 50(5):945–975.
- Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 5(1):19–46.
- Romer, C. D. (1990). The Great Crash and the onset of the Great Depression. *The Quarterly Journal of Economics*, 105(3):597–624.
- Romer, C. D. and Romer, D. H. (2013). The most dangerous idea in Federal Reserve history: Monetary policy doesn’t matter. *American Economic Review*, 103(3):55–60.
- Rumsfeld, D. H. (2002). Transcript of US Department of Defense News Briefing.
- Sargent, T. and Wallace, N. (1975). Rational expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of Political Economy*, 83(2):241–54.
- Sargent, T. J. (1999). Comment. *Monetary Policy Rules*, University of Chicago Press.
- Sargent, T. J. (2001). *The Conquest of American Inflation*. Princeton University Press.
- Sargent, T. J. (2012). Remarks at Bocconi University: “The Beauty of Uncertainty”. <http://www.igier.unibocconi.it/>.
- Smets, F. and Wouters, R. (2003). An estimated dynamic stochastic general equilibrium model of the Euro Area. *Journal of the European Economic Association*, 1(5):1123–1175.

- Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.
- Svensson, L. E. O. (1997). Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review*, 41(6):1111–1146.
- von Hayek, F. A. (1974). Prize lecture: “The Pretence of Knowledge”. <http://www.nobelprize.org>.
- Wald, A. (1950). *Statistical Decision Functions*. Wiley.
- Wallace, N. (1981). A Modigliani-Miller theorem for open-market operations. *American Economic Review*, 71(3):267–274.
- Whittle, P. (1990). *Risk-Sensitive Optimal Control*. Wiley.
- Wieland, V. (2000). Monetary policy, parameter uncertainty, and optimal learning. *Journal of Monetary Economics*, 46(1):199–228.
- Williams, J. C. (2013). A defense of moderation in monetary policy. *Journal of Macroeconomics*, 38(B):137–150.

II.

THE PORTFOLIO BALANCE CHANNEL UNDER UNCERTAINTY

Adrian Paul*

Abstract

In the aftermath of the global financial crisis of 2007/08, aggregate output collapsed and macroeconomic uncertainty spiked. In response, several major central banks embarked on large-scale asset purchase programmes. The scope of these programmes extended well beyond purchases of risk-free government bonds to include the accumulation of private-sector credit instruments with increasingly uncertain returns. We construct a theoretical model in which two assets are imperfect substitutes and households' preferences over the composition of their portfolios are robust to the prevalence of Knightian uncertainty. Because relative asset quantities matter, quantitative easing has real effects. Moreover, it is households' portfolio preferences that determine the “bang for the buck” associated with central bank asset purchases. A flight to safety induced by the uncertainty accompanying a recession undermines the effectiveness of QE – just when it is needed most.

*St. Peter's College, University of Oxford. Michaelmas Term, 2017. This is the second chapter of my D.Phil. thesis: *Essays on Unconventional Monetary Policy*.

1 Introduction

“Well, the problem with QE is it works in practice, but it doesn’t work in theory.”

Ben S. Bernanke,
Chairman of the U.S. Federal Reserve,
16 January 2014.

The study of uncertainty in macroeconomics has gathered momentum since the global financial crisis of 2007/08. Theorists have found that incorporating macroeconomic uncertainty into otherwise standard models of investor behaviour resolves long-standing puzzles in asset pricing. Empiricists have found that the bankruptcy of Lehman Brothers constitutes the most acute episode of macroeconomic uncertainty for more than half a century. Policymakers have found that the pervasiveness of macroeconomic uncertainty complicates the task of business cycle stabilisation.

In this paper, we motivate a mechanism through which asset purchases by the central bank alter the relative supplies of assets available to investors and thereby affect real aggregate variables. In formalising this ‘portfolio balance’ channel, we establish a relationship between counter-cyclical macroeconomic uncertainty and the effectiveness of unconventional monetary policy.

Having cut short-term interest rates to their effective lower bound in the wake of the global financial crisis, the world’s major central banks resorted to large-scale asset purchases to stabilise their respective economies. Between 2007 and 2015, for example, the U.S. Federal Reserve, the European Central Bank, the Bank of Japan and the Bank of England expanded their balance sheets by a factor of between two and five, conducting outright purchases in the order of between 2% and 50% of GDP. Large-scale asset purchases (LSAPs) were not only concentrated in risk-free government bonds. Central banks accumulated private-sector credit instruments with diverse characteristics. The ECB, for example, bought covered bonds, asset-backed securities and corporate bonds. The U.S. Federal Reserve accumulated mortgage-backed securities. The Bank of Japan purchased commercial paper and even common stocks.

Empirical attempts to quantify the effects of these unconventional monetary policy measures are fraught with imprecision. Qualitatively, however, there is broad agreement across the macro-finance literature that large-scale asset purchases were stimulative. Av-

eraging across numerous U.S. studies, for example, suggests that each USD \$1 trillion of assets purchased by the Federal Reserve reduced long-term Treasury yields by around 40 basis points on impact, with a peak effect on the level of GDP of around 1%, and a peak effect on the rate of annual inflation of around $1/2$ percentage point.

The theoretical support for this empirical effectiveness, however, is more contentious. Former Chairman of the Federal Reserve, Ben Bernanke, famously conceded – while still in office – that “the problem with QE is it works in practice, but it doesn’t work in theory.”

The textbook mechanism most commonly cited by central bank officials is the ‘portfolio balance’ channel of monetary policy. In the theory originally espoused by Tobin (1969), if different classes of financial assets are considered imperfect substitutes in investors’ portfolios, changes in their relative supplies affect their relative prices – and therefore their relative rates of return. Central bank purchases of mortgage-backed securities, for example, raise the relative price of those securities and lower their relative yield, thereby easing the financing costs faced by indebted households. This stimulates economic activity through channels similar to those ascribed to the transmission of conventional monetary policy.

Other complementary mechanisms have also been explored in the theoretical literature, not least the signalling effects associated with the forward guidance embedded in central bank communication, and the exchange rate effects associated with post-crisis monetary policy surprises.¹ In this paper, however, we focus solely on the portfolio balance channel. It is in this aspect of the transmission mechanism of QE that modern policymakers have put the most faith in practice, but modern academics have shown the most neglect in theory.

Not only did quantitative easing debut amid a sharp decline in real economic activity, the beginning of large-scale asset purchases coincided with a sharp rise in perceived macroeconomic uncertainty.

In early 2009, IMF Chief Economist Olivier Blanchard observed:

¹See Woodford (2013), for example, on the “special role” of forward guidance in the toolkit of a central bank constrained by an effective lower bound on short-term nominal interest rates. See Glick and Leduc (2013) for a comparison of the effects of unconventional and conventional monetary policy announcements on the value of the U.S. dollar using high-frequency intraday data.

“when the economic environment is so complex as to appear nearly incomprehensible, the result is extreme prudence, if not outright paralysis, on the part of investors, consumers and firms. And this behaviour, in turn, feeds the crisis. It affects portfolio decisions. It has led to a dramatic shift away from risky assets to riskless assets, or at least assets perceived as riskless...[together with]...the realisation that many of the new complex assets were in fact much riskier than they had seemed.”

In late 2010, Federal Reserve Chairman Ben Bernanke observed:

“during the worst phase of the financial crisis, many economic actors – including investors, employers, and consumers – metaphorically threw up their hands and admitted that, given the extreme and, in some ways, unprecedented nature of the crisis, they did not know what they did not know...The profound uncertainty associated with the ‘unknown unknowns’ during the crisis resulted in panicky selling by investors, sharp cuts in payrolls by employers, and significant increases in households’ precautionary saving.”

Econometric studies comparing the forecast errors associated with key U.S. data releases confirm that subjective macroeconomic uncertainty spiked in late 2008. More generally, these studies find that over long periods of time there is a significant, negative correlation between perceived uncertainty and economic activity – a counter-cyclicity that we exploit in this paper to motivate imperfect asset substitutability.

Acknowledging uncertainty has already borne fruit in the asset pricing literature. The inability of theoretical models to replicate key features of empirical asset prices is exemplified by long-standing dilemmas such as the ‘equity premium puzzle.’ Recently, however, significant progress has been made in matching the first and second moments of observed asset prices by incorporating ambiguity aversion into investors’ preferences. Aversion to ambiguity can be interpreted as fear of model uncertainty – a fear which prompts investors to make decisions whose consequences are relatively less affected by their doubts about the true probability distribution governing future returns. In this context, both the degree of model uncertainty that investors face and their aversion to any given degree of model uncertainty fundamentally influence investors’ optimal portfolios.

In this paper, we explore how the interaction between (i) households’ information about the economy, and (ii) households’ preferences, given that information, shape the portfolio balance channel of monetary policy, and thereby affect the effectiveness of central

bank asset purchases in times of economic crisis.²

We find that when households seek to ensure that the composition of their portfolios is robust to the prevalence of Knightian uncertainty, relative asset quantities matter. Because relative asset quantities matter, quantitative easing has real effects. Moreover, it is households' portfolio preferences that determine the “bang for the buck” associated with central bank asset purchases. A flight to safety induced by the uncertainty accompanying a recession undermines the effectiveness of QE – just when it is needed most.

This paper proceeds as follows. Section 2 surveys three strands of related literature. Section 3 presents the key stylised facts which motivate the model presented in Section 4. Section 5 illustrates the effectiveness of quantitative easing. Section 6 demonstrates the importance of portfolio preferences. Section 7 concludes. Formal derivations are relegated to the Appendix.

2 Literature

This paper draws on three strands of literature: (i) the New Keynesian approach to optimal monetary policy when assets are imperfect substitutes, (ii) the portfolio choice problem faced by a mean-variance investor who is averse to ambiguity, and (iii) the empirical role of model uncertainty in the historical evolution of asset market returns.

The model in Section 4 builds on Harrison (2012), which itself is based on Andrés et al. (2004). Both papers consider optimal monetary policy in a stylised, New Keynesian model in which the assets available to investors are imperfect substitutes. This imperfect substitutability implies that central bank asset purchases – which alter the relative supplies of assets available to households – can influence the rates of return on each available asset independently, thereby affecting aggregate demand and influencing deviations of inflation from target. Harrison (2012), for example, finds that a strategy in which the central bank uses asset purchases as an unconventional complement to conventional monetary policy can improve outcomes when the economy is hit by a negative demand shock sufficiently large to drive the short-term interest rate to its effective lower bound.

²Throughout this paper, we use the terms ‘large-scale asset purchases (LSAPs)’, ‘asset purchases’, ‘quantitative easing (QE)’ and ‘unconventional monetary policy’ interchangeably. We also use the terms ‘model uncertainty’, ‘Knightian uncertainty’, ‘uncertainty’ and ‘ambiguity’ interchangeably – terms which, as we explain in Section 4, are quite distinct from ‘risk.’

The weakness of the existing literature, however, stems from the *ad hoc* nature in which imperfect asset substitutability arises. To capture the notion in Tobin (1956) and Tobin and Brainard (1963) that relative asset prices (and, therefore, relative rates of return) depend on relative asset supplies, Harrison (2012) and Andrés et al. (2004) introduce portfolio adjustment costs into households’ utility functions. But the authors do not establish why households would prefer one portfolio allocation over another, and therefore cannot explain how such portfolio adjustment costs arise. As a result, the term premium at the heart of these models – the channel through which asset purchases are transmitted to the real economy – is only loosely associated with an unmodelled concern for relative illiquidity. Moreover, the asset adjustment costs which Harrison (2012) and Andrés et al. (2004) introduce into the benchmark New Keynesian model are defined with respect to deviations from an exogenous, preferred portfolio that is not allowed to vary over time. We introduce aversion to time-varying model uncertainty to address both shortcomings of the existing literature.

In the model in Section 4, it is the incidence of negative demand shocks that prompts the policymaker to loosen monetary policy. We model negative demand shocks as shocks to the real natural interest rate, just as in Harrison (2012), but we calibrate these shocks according to Laubach and Williams (2016). The latter estimate a model in which the natural interest rate varies with the contemporaneous, trend growth rate of potential output and a time-varying, unobserved component capturing other unspecified factors. The authors show that the real natural rate of interest in the United States exhibited a moderate, secular decline during the two decades that preceded the global financial crisis, before falling sharply through 2008 and 2009.

Maccheroni et al. (2013) study the portfolio choice problem faced by a mean-variance investor concerned with model uncertainty. Based on the ‘smooth ambiguity model’ of decision-making developed by Klibanoff et al. (2005), the authors develop a preference functional which is determined by both the investor’s tastes (her aversion to risk and, separately, her aversion to ambiguity) and the quality of the investor’s information (her subjective uncertainty over prospective outcomes and models). Augmenting preferences in this way delivers a rich, analytical solution for the investor’s robustly optimal portfolio. We exploit this robust allocation to inform the portfolio adjustment costs incurred by a representative household when faced with asset purchases by the central bank, in an environment of heightened uncertainty.

Collard et al. (2011) apply the smooth ambiguity model to a single agent in a dynamic exchange economy who is uncertain about the true probability distribution governing future consumption and dividends. The authors then calibrate the agent's ambiguity aversion to match the first moment of the risk-free rate in historical U.S. data, conditioning the uncertainty that prevails in each period on the observed history of economic growth. Collard et al. (2011) find that – because ambiguity aversion endogenously accentuates conditional uncertainty in a dynamic, history-dependent fashion (ambiguity aversion increases during recessions) – modifying a standard asset-pricing model to incorporate ambiguity enables the modified model to match observed returns much more closely. In particular, time-varying uncertainty can resolve three empirical asset-pricing puzzles which have defied existing models. First, introducing ambiguity helps match the first moment of the equity premium because greater ambiguity aversion suppresses the risk-free rate while only slightly accentuating the risky rate. Second, ambiguity aversion generates conditional equity premia which are counter-cyclical. Third, ambiguity aversion implies that the volatility of conditional excess returns is appropriately pro-cyclical.

Establishing a preference for robustness creates the scope for a more formal treatment of uncertainty shocks, modelled as sudden increases in the perceived uncertainty surrounding future consumption growth. Because Collard et al. (2011) hold fixed the investor's taste parameters over time, historical movements in the authors' model-implied equity premium (that is, changes in the compensation required to tolerate unexpected equity returns) can be attributed to coincident changes in subjective macroeconomic uncertainty. It turns out that these model-based changes in subjective uncertainty closely match changes in objective measures of uncertainty, borne of a purely statistical approach.

One such statistical measure of uncertainty has been developed by Jurado et al. (2015). The authors consider 132 macroeconomic variables (including measures of real output, labour compensation and consumer prices) and define each variable's conditional uncertainty as the conditional volatility of the unforecastable component of that variable's future value. This measure of individual uncertainty is then aggregated across the entire set of variables. The result is an index of perceived economy-wide uncertainty, which serves as a proxy for the Knightian uncertainty associated with the probability distribution governing each variable's evolution. Jurado et al. (2015) emphasise that the key features of macroeconomic uncertainty – at least as identified in this fashion – are its counter-cyclicality and its persistence during recessions. These two empirical attributes corroborate the theoretical mechanism at work in Collard et al. (2011): it is ambiguity

aversion that can explain why, following a positive shock, an investor makes decisions that underplay inferred uncertainty whereas, following a negative shock, the investor behaves as if uncertainty is more severe and more persistent than implied by objective (Bayesian) inference. In this paper, we connect the counter-cyclicalities of perceived macroeconomic uncertainty to the effectiveness of unconventional monetary policy in times of crisis.

3 Stylised facts

In this section, we present the key stylised facts which motivate the model developed in Section 4.

Figure 1 suggests that the composition of the typical U.S. household’s ‘steady-state’ asset portfolio has varied considerably over the past fifty years. Although this ‘steady-state’ portfolio is invariably influenced by shifts in the *supply* of safe assets over time, to the extent that it also reflects shifts in the demand for safe assets, we use it as a proxy for the typical household’s ‘preferred’ asset portfolio.

We plot the ratio of (i) the nominal value of privately held, interest-bearing U.S. Treasury debt of residual maturity greater than one year (‘risky’ long-dated bonds), to (ii) the nominal value of privately held U.S. Treasury debt of residual maturity less than one year (‘safe’ short-dated bonds).³ This ratio, which we denote δ_t , is an indicator of the mix of asset characteristics in the representative household’s preferred savings portfolio. An increase in δ_t is indicative of an increase in desired exposure to duration risk: as δ_t rises, the representative investor becomes relatively more inclined to hold risky, long-dated debt and relatively less inclined to hold safe, short-dated securities. Between the early 1960s and the early 2010s, δ_t fluctuated considerably, between $1/2$ and $2^{1/2}$.

The average residual maturity of debt has been used in other work as a proxy for average duration risk (see, for example, Chadha et al. (2013)). Duration risk is the sensitivity of a bond’s price to changes in interest rates: the greater a bond’s duration, the greater the interest rate risk it carries. As a bond’s redemption date nears, both its duration and its residual maturity decline.

³For now, we call long-dated bonds ‘risky’ and short-dated bonds ‘safe’ in order to establish intuition. In our formal treatment of the policy problem in Section 4, however, the returns on the non-safe asset available to investors will carry uncertainty rather than risk.

Figure 2 suggests that portfolio preferences were particularly volatile during the Great Recession. Between 2007 and 2008, investors’ preferred portfolios switched sharply towards safe assets, at the expense of risky assets. This decline in δ_t demonstrates the “dramatic shift” in portfolio decisions to which the IMF’s Chief Economist alludes in the Introduction. In 2009, however, investors’ preferred portfolios reverted just as sharply towards renewed exposure to risky assets. The vertical markers in Figure 2 show that the U.S. Federal Reserve’s first program of large-scale asset purchases (‘QE1’) was conducted against the backdrop of very different private-sector portfolio preferences than the second and third LSAP programs (‘QE2’ and ‘QE3’).

Figure 3 demonstrates the counter-cyclical nature of perceived uncertainty, using the purely statistical measure constructed by Jurado et. al (2013). Deep recessions are associated with particularly sharp increases in uncertainty. This was particularly true of the Great Recession that dominated the late 2000s, but it was also a distinctive feature of the two U.S. recessions that followed the oil price shocks of 1973 and 1979.

Figure 4 confirms that model-free estimates of macroeconomic uncertainty – inferred, in this example, from the unforecastable component of U.S. data releases – are closely related to model-based estimates of macroeconomic uncertainty – inferred from the equilibrium behaviour of ambiguity-averse investors. Both measures of perceived uncertainty co-move, rising in recessions.

Simple correlations – summarised in Table 1 and Figures 5-7 – are indicative of a ‘flight to safety’ in recessions: negative shocks to the real natural interest rate have coincided with greater perceived uncertainty, which has, in turn, coincided with a stronger preference among investors for safe assets in bad times. These stylised facts concur with existing theoretical work on crises driven by uncertainty. Caballero and Krishnamurthy (2008), for example, model severe ‘flight to quality’ episodes as those involving uncertainty about the macroeconomic environment, not only risk about specific asset payoffs. This Knightian uncertainty is triggered by unusual events that prompt agents to “question their world-view” – events which culminate in market-wide capital immobility, disengagement from risk and widespread hoarding of liquidity.

Table 1: CORRELATION CO-EFFICIENTS

	Portfolio duration	Uncertainty	Natural interest rate
Portfolio duration	1		
Uncertainty	−0.57	1	
Natural interest rate	+0.39	−0.53	1

Notes: ‘Portfolio duration’ is proxied by the ratio of ‘risky’ long-dated Treasury bonds to ‘safe’ short-dated Treasury bonds in private investors’ portfolios (denoted δ_t in Figures 1, 2, 6 and 7): the lower a portfolio’s duration, the greater its safety. ‘Uncertainty’ refers to the statistical measure derived by Jurado et al. (2015). Movements in the ‘natural interest rate’ refer to cumulative, four-quarter changes in the series estimated by Laubach and Williams (2016). Correlation co-efficients are based on quarterly data between 1990 and 2015.

4 Model

In this section, we develop a model designed to illuminate the key stylised facts described in Section 3. The model is a version of Harrison (2012), modified to incorporate shocks to perceived macroeconomic uncertainty.

4.1 The household

A representative household’s utility is comprised of her consumption, her hours of work, her holdings of real money and her portfolio of financial assets. There are two financial assets available in this economy: one is risk-free, the other is not (in a sense that will be made precise below). The household’s utility is decreasing in deviations of her realised asset portfolio from her preferred asset portfolio. We first consider the determination of the latter.

4.1.1 The household’s preferred portfolio

Our formulation of the household’s preferred asset portfolio is inspired by Maccheroni et al. (2013). We begin with a general treatment of the mean-variance model under uncertainty. We then consider the special case in which a portfolio of only two assets is chosen. We use the terms ‘model uncertainty’ (or simply ‘uncertainty’) and ‘ambiguity’ interchangeably, to describe a situation in which the decision-maker does not know the probabilistic model governing possible states. This is distinct from ‘risk,’ which instead describes a situation in which the decision-maker does know the probabilistic model governing possible states, but does not know which particular state will be realised.

The general case

Suppose an investor has wealth w and considers an asset position h . She faces model uncertainty: she is unsure about the true probabilistic model $\mathcal{P}$ governing the precise return on h .

If the investor were to only face risk – that is, if she were to fully rely on a single probabilistic model $\mathcal{P}$ – her certainty equivalent $c(w + h, \mathcal{P})$ would be given by:

$$c(w + h, \mathcal{P}) = u^{-1}(\mathbb{E}_{\mathcal{P}}(u(w + h))) \quad (1)$$

where u represents her attitude towards risk, and expression (1) can be understood as the sure amount of money the investor would consider equivalent to the risky prospect $w + h$.

Since the individual also faces model uncertainty, however, she is unable to identify a single probabilistic model $\mathcal{P}$. She therefore also considers alternative models $\mathcal{Q}$, implying that her certainty equivalent $c(w + h, \mathcal{Q})$ is a variable amount of money that depends on $\mathcal{Q}$.

Suppose μ is the investor's prior probability on the space Δ of possible probabilistic models, and v represents her attitude towards model uncertainty. The rationale used to obtain the certainty equivalent (1) now leads to a modified certainty equivalent – under model uncertainty – given by:

$$C(w + h) = v^{-1}(\mathbb{E}_{\mu}(v(c(w + h, \mathcal{Q})))) = v^{-1}(\mathbb{E}_{\mu}(v(u^{-1}(\mathbb{E}(u(w + h, \mathcal{Q})))))) \quad (2)$$

where $c(w + h)$ is the random variable that associates $c(w + h, \mathcal{Q})$ with each model $\mathcal{Q}$ in the feasible set Δ .⁴

This is the ‘smooth ambiguity’ certainty equivalent of Klibanoff et al. (2005). As long as the decision-maker faces some model uncertainty, and as long as the decision-maker's reaction to model uncertainty differs from her reaction to risk, then ambiguity and risk have distinct implications, and $C(w + h)$ and $c(w + h)$ diverge. Model uncertainty itself is an information feature of the model, which implies that the support of μ is non-singleton.

⁴If there is no model uncertainty whatsoever, then the support of the prior μ is a singleton $\mathcal{P}$ and (2) reduces to (1). If the investor is ambiguity-neutral, then $v = u$ and $C(w + h) = c(w + h, \overline{\mathcal{Q}})$ so (2) again collapses to (1), with the reduced distribution $\overline{\mathcal{Q}}$ representing the entire uncertainty faced by the individual.

The reaction to a given degree of model uncertainty is a taste feature of the model, which implies that v differs from u .

Proposition 3 in Maccheroni et al. (2013) shows that the investor's certainty equivalent for the uncertain prospect $w + h$ can be approximated by:

$$C(w + h) \approx w + \mathbb{E}_{\overline{\mathcal{Q}}}(h) - \frac{1}{2}\lambda_u(w)\sigma_{\overline{\mathcal{Q}}}^2(h) - \frac{1}{2}(\lambda_v(w) - \lambda_u(w))\sigma_{\mu}^2(\mathbb{E}(h)) \quad (3)$$

where $\overline{\mathcal{Q}}$ is the reduced probability induced by the prior μ , and $\mathbb{E}(h)$ is the random variable that associates the expected value $\mathbb{E}_{\mathcal{Q}}(h)$ with each possible model $\mathcal{Q}$. The approximation in (3) is analogous to the familiar Arrow-Pratt approximation to the investor's certainty equivalent in the case where $w + h$ carries only risk.

The formal derivation of the quadratic approximation in (3) is outlined in Appendix A.

Intuitively, expression (3) implies that incorporating model uncertainty into the investor's decision problem amounts to augmenting a standard, von Neumann-Morgenstern expected utility maximiser's certainty equivalent with an ambiguity premium. This ambiguity premium is increasing in both (i) the volatility of the expected return on the investor's position in the asset, denoted $\sigma_{\mu}^2(\mathbb{E}(h))$, and (ii) the difference between the investor's attitudes towards uncertainty and her attitudes towards risk, denoted $(\lambda_v(w) - \lambda_u(w))$.

The counterpart of the quadratic approximation in (3) is an ambiguity adjustment to the classic mean-variance model. By setting $w + h = f$, $\lambda_u(w) = \lambda$, $\lambda_v(w) - \lambda_u(w) = \theta$ and $\overline{\mathcal{Q}} = \mathcal{P}$ in (3), Maccheroni et al. (2013) obtain:

$$C(f) = \mathbb{E}_{\mathcal{P}}(f) - \frac{\lambda}{2}\sigma_{\mathcal{P}}^2(f) - \frac{\theta}{2}\sigma_{\mu}^2(\mathbb{E}(f)) \quad (4)$$

The formal derivation of expression (4) is outlined in Appendix B.

The certainty equivalent specified in (4) implies that an ambiguity-averse investor's payoff from an uncertain prospect f is determined by two taste parameters, λ and θ , which are increasing in the individual's aversion to risk and to ambiguity, respectively, and one information parameter, μ , which determines the variances $\sigma_{\mathcal{P}}^2(f)$ and $\sigma_{\mu}^2(\mathbb{E}(f))$ capturing the outcome risk and model uncertainty perceived by the decision-maker.

The two-asset case

Now consider the special case of a household which decides how to allocate one unit of wealth among two assets at time zero. One asset is ‘safe’, with a constant gross return after one period of R . The other asset is ‘ambiguous’, with a gross one-period return denoted R_L . The second asset is ambiguous in the sense that the household faces model uncertainty about the true probability distribution governing R_L . Let the fraction of household wealth invested in the ambiguous asset be w . The household’s end-of-period wealth R_w , for a given choice of w , is therefore:

$$R_w = (1 - w)R + wR_L \quad (5)$$

Assuming a frictionless financial market with no transaction costs and no borrowing constraints, the modified certainty equivalent (4) implies that the household’s portfolio choice problem under uncertainty can be written as:

$$\max_w C(R_w) = \max_w \left(\mathbb{E}_{\mathcal{P}}(R_w) - \frac{\lambda}{2} \sigma_{\mathcal{P}}^2(R_w) - \frac{\theta}{2} \sigma_{\mu}^2(\mathbb{E}(R_w)) \right) \quad (6)$$

The solution to (6) delivers the following condition for the optimal fraction of wealth to be invested in the ambiguous asset:

$$w^* = \frac{\mathbb{E}_{\mathcal{P}}(R_L) - R}{\lambda \sigma_{\mathcal{P}}^2(R_L) + \theta \sigma_{\mu}^2(\mathbb{E}(R_L))} \quad (7)$$

Condition (7) implies that the household’s robustly optimal ratio of ambiguous asset holdings to safe asset holdings (let this be denoted δ) is given by:

$$\delta \equiv \frac{w^*}{1 - w^*} \equiv \frac{\mathbb{E}_{\mathcal{P}}(R_L) - R}{R - \mathbb{E}_{\mathcal{P}}(R_L) + \lambda \sigma_{\mathcal{P}}^2(R_L) + \theta \sigma_{\mu}^2(\mathbb{E}(R_L))} \quad (8)$$

The formal derivation of the optimality condition (8) is outlined in Appendix C.

Ambiguity does not affect excess returns. This creates a tractable separation between (i) the first and second moments of asset returns, and (ii) the household’s preferences and information regarding those returns.

Optimality condition (8) implies that a *ceteris paribus* increase in either the house-

hold's perception of model uncertainty, $\sigma_\mu^2(\mathbb{E}(R_L))$, or the household's aversion to a given degree of model uncertainty, θ , reduces the relative desirability of the ambiguous asset and increases the household's demand for the safe asset. In other words, a 'flight to safety' lowers the household's preferred portfolio ratio, δ .

4.1.2 The household's optimisation problem

Once her preferred ratio of ambiguous asset holdings to safe asset holdings has been established, the household maximises her utility subject to her budget constraint.

Let c denote the household's consumption of a Dixit-Stiglitz consumption bundle, let n denote hours worked, let M/P denote real money balances held and let B/B_L denote the household's actual ratio of safe asset holdings to ambiguous asset holdings. Let ϕ denote a demand shock to household preferences, and let β denote the household's discount factor.

The representative household's optimisation problem can now be expressed as:

$$\begin{aligned} & \max_{\{c_t, n_t, M_t, B_t, B_{L,t}\}_{t=0}^\infty} \mathbb{E}_0 \sum_{t=0}^\infty \beta^t \phi_t \left(\frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{n_t^{1+\psi}}{1+\psi} + \frac{\chi_m^{-1}}{1-\sigma_m^{-1}} \left(\frac{M_t}{P_t} \right)^{1-\frac{1}{\sigma_m}} - \frac{\tilde{v}}{2} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right]^2 \right) \\ & \text{s.t.} \\ & B_t + B_{L,t} + M_t = R_{t-1}B_{t-1} + R_{L,t}B_{L,t-1} + M_{t-1} + W_t n_t + T_t + D_t - P_t c_t \end{aligned} \quad (9)$$

The budget constraint stipulates that, in any given period, the sum of the household's spending on the safe asset, B , the ambiguous asset, B_L , and money, M , cannot exceed her net income. Net income is defined as the sum of the household's *ex post* returns on asset holdings in the previous period, her current wage income, W , fiscal transfers, T , and dividends, D , less her current expenditure on consumption.

The household's utility function, as described in (9), is standard except for its final term, which captures the costliness of deviations in the household's actual portfolio in any given period, $B_t/B_{L,t}$, from her preferred portfolio in that period, δ_t .⁵ If, in equilibrium, the composition of the household's actual portfolio exactly coincides with the composition of her preferred portfolio, then the final term in the household's utility function collapses to zero. By contrast, the costs associated with being overweight the ambiguous asset (un-

⁵Note that δ is defined as the household's preferred ratio of ambiguous assets to safe assets, while B/B_L denotes the household's actual ratio of safe assets to ambiguous assets.

derweight the safe asset) and the costs associated with being underweight the ambiguous asset (overweight the safe asset) are symmetric.

4.2 The government

The government comprises the fiscal authority and the monetary authority. The government's budget constraint is given by:

$$\frac{B_t - R_{t-1}B_{t-1}}{P_t} + \frac{B_{L,t}^g - R_{L,t}B_{L,t-1}^g}{P_t} + \frac{\Delta_t}{P_t} = \frac{T_t}{P_t} \quad (10)$$

which stipulates that the net issuance of government liabilities by the fiscal authority plus the change in the central bank's balance sheet finances net fiscal transfers to households.

The two financial assets in which the household can invest are both issued by the government. The safe asset, B , is modelled as a one-period bond which sells at a unit price and is redeemed at a price R in the following period. The ambiguous asset, by contrast, is subject to capital gains or losses: B_L is modelled as a consol with an *ex post* one-period return R_L .⁶ The superscript in $B_{L,t}^g$ denotes the value of the outstanding stock of ambiguous assets issued by the government. In equilibrium, this stock will be distributed between the central bank and the household (see equations (11) and (17) below). The change in the central bank's balance sheet is denoted Δ and lump-sum fiscal transfers to households are denoted T . All items in the government's budget constraint are deflated by the aggregate price index P , which is simply the price of a Dixit-Stiglitz consumption bundle as defined in equation (14) below.

We model quantitative easing as central bank purchases of some share q of the outstanding stock of ambiguous assets. The central bank's holding of the ambiguous asset is therefore given by:

$$Q_t = q_t B_{L,t}^g \quad (11)$$

Since central bank asset purchases are financed by money creation, the change in the

⁶The price of a consol at time t is denoted V_t . This implies that: (i) the *ex post* nominal return on consols, which we denote $R_{L,t}$, can be expressed as: $R_{L,t} \equiv \frac{(1+V_t)}{V_{t-1}}$, and (ii) the gross issuance of consols by the government is given by: $B_{L,t}^g \equiv V_t B_{L,t}$.

central bank's balance sheet can be expressed as:

$$\frac{\Delta_t}{P_t} = \frac{M_t - M_{t-1}}{P_t} - \left(\frac{Q_t}{P_t} - \frac{R_{L,t}Q_{t-1}}{P_t} \right) \quad (12)$$

where the first term on the right-hand side describes the expansion of the liabilities side of the central bank's balance sheet (the change in the stock of money) and the second term on the right-hand side describes the expansion of the asset side of the central bank's balance sheet (the net increase in holdings of the ambiguous asset).

Substituting (12) into (10), and denoting real quantities in lower case, the consolidated government budget constraint is given by:

$$b_t + (1 - q_t)b_{L,t}^g + m_t = \pi_t^{-1} [m_{t-1} + R_{t-1}b_{t-1} + R_{L,t}(1 - q_{t-1})b_{L,t-1}^g] + \tau_t \quad (13)$$

where the inflation rate is denoted $\pi_t \equiv \frac{P_t}{P_{t-1}}$ and real lump-sum taxes are denoted $\tau_t \equiv \frac{T_t}{P_t}$.

Since the real stock of the ambiguous asset is assumed fixed,⁷ and since we assume there is no government expenditure, the only choice variable for the fiscal authority is lump-sum transfers to households.

Net transfers are set by a simple linear rule, according to which transfers to households in any given period are decreasing in both the previous period's stock of safe assets (this ensures that debt issuance is stable) and the previous period's safe interest rate (this limits the feedback between debt-financing costs and the debt stock).⁸ The simplicity of this transfer rule reflects the narrow focus of this paper: we abstract from any discussion of fiscal policy in order to focus on the implications of model uncertainty for the effectiveness of conventional and unconventional monetary policy.

The monetary authority has two choice variables: the nominal interest rate on the safe asset (the conventional policy instrument, R) and the fraction of the ambiguous asset held on the central bank's balance sheet (the unconventional policy instrument, q).

⁷Following the previous footnote, if the real stock of consols is held fixed at $\bar{b}_L$ then the real value of the stock of consols outstanding is given by $b_{L,t}^g = \bar{b}_L V_t$.

⁸In terms of log-deviations from steady-state (see Section 4.6), we assume the following transfer rule is: $\frac{\tau}{b} \hat{\tau}_t = -\beta^{-1} \hat{R}_{t-1} - \zeta \hat{b}_{t-1}$.

4.3 Firms

The supply side of the model is standard. There is a continuum of monopolistically competitive producers indexed by $j \in (0, 1)$ creating differentiated products that form a Dixit-Stiglitz consumption bundle purchased by the representative household.

The consumption bundle is given by:

$$c_t = \left[\int_0^1 c_{j,t}^{1-\eta^{-1}} dj \right]^{\frac{1}{1-\eta^{-1}}} \quad (14)$$

where c_j is consumption of firm j 's output. Firm j 's output is produced according to a simple linear function of labour, in which A denotes productivity: $c_{j,t} = An_{j,t}$.

The real profits made by firm j are therefore given by:

$$\frac{P_{j,t}}{P_t} c_{j,t} - \omega_t n_{j,t} = \left((1+s) \frac{P_{j,t}}{P_t} - \frac{\omega_t}{A} \right) \left(\frac{P_{j,t}}{P_t} \right)^{-\eta} c_t \quad (15)$$

where s is a subsidy paid to producers in order to ensure that the steady-state level of output is efficient.⁹

Under the assumption of Calvo (1983) price-setting, with $0 \leq \alpha < 1$ denoting the probability that the producer cannot reset his price in any given period, producer j 's optimisation problem can be written as:

$$\max_{\{P_{j,t}\}_{t=0}^{\infty}} \mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k (\beta \alpha)^{k-t} \left((1+s) \frac{P_{j,t}}{P_t} - \frac{\omega_t}{A} \right) \left(\frac{P_{j,t}}{P_t} \right)^{-\eta} c_k \quad (16)$$

where Λ represents the household's stochastic discount factor.

4.4 Market clearing

To ensure that asset markets clear, the share of the ambiguous asset not purchased by the central bank must be taken up by households:

⁹We will shortly exploit the fact that this assumption implies the appropriate welfare criterion for optimal policy is a quadratic approximation of the household's utility function, as shown by Benigno and Woodford (2006).

$$b_{L,t} = (1 - q_t)b_{L,t}^g \quad (17)$$

To ensure that the goods market clears, aggregate consumption must be proportional to aggregate output:

$$c_t = \mathcal{D}_t^{-1} y_t \quad (18)$$

where $\mathcal{D}_t$ is a measure of price dispersion.

4.5 Shocks

The only shock hitting the model is the preference shock ϕ_t . In the absence of both price rigidity and imperfect asset substitutability, the preference shock has no impact on activity. The efficient level of output in our model is therefore constant. This implies that – to a first-order approximation – the ‘output gap’ is identical to the deviation of output from its steady state. Using circumflexes to denote log-deviations from steady-state, this – together with equation (18) – implies:

$$\hat{c}_t = \hat{y}_t = \hat{x}_t \quad (19)$$

where x denotes the output gap.

The absence of both price rigidity and imperfect asset substitutability also implies that the real interest rate is simply the difference between successive preference shocks.

The real *natural* interest rate in our model is therefore defined as:

$$r_t^* = -\mathbb{E}_t \left(\hat{\phi}_{t+1} - \hat{\phi}_t \right) \quad (20)$$

where r_t^* is assumed to follow a simple autoregressive process given by:

$$r_t^* = \rho r_{t-1}^* + \varepsilon_t \quad (21)$$

In the comparative statics that follow in Sections 5 and 6, an adverse shock to the economy is modelled as a negative shock to the real natural interest rate. This is characterised by

$\varepsilon_t < 0$ in equation (21).

4.6 First-order conditions

With these relationships in mind, we summarise the first-order conditions characterising optimal behaviour in each sector. This serves to illuminate the macroeconomic intuition behind the model's mechanics. The formal derivation of each first-order condition is provided in Appendix D.

The log-linearised first-order conditions of the household's constrained optimisation problem (9) can be combined to give:

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left(\frac{1}{1 + \delta_t} \hat{R}_t + \frac{\delta_t}{1 + \delta_t} \hat{R}_{L,t}^e - \mathbb{E}_t \hat{\pi}_{t+1} - r_t^* \right) \quad (22)$$

$$\hat{R}_{L,t}^e = \hat{R}_t - (1 + \delta_t) \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} (\hat{b}_t - \hat{b}_{L,t}) \quad (23)$$

$$\hat{m}_t = \frac{\sigma_m}{\sigma} \hat{x}_t - \frac{\beta \sigma_m}{1 - \beta} \hat{R}_t + \frac{\beta \sigma_m}{1 - \beta} (\delta_t) \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} (\hat{b}_t - \hat{b}_{L,t}) \quad (24)$$

The log-linearised first-order conditions of the firm's constrained optimisation problem (16) imply:

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t \quad (25)$$

in which $\kappa = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha}(\psi + \sigma^{-1})$.

The log-linearised government budget constraint (13), together with our assumed fiscal transfer rule (see footnote 8), imply:

$$\hat{b}_t + \frac{m}{b} \hat{m}_t - \delta_t q_t = - \left[\frac{m}{b} + \beta^{-1}(1 + \delta_t) \right] \hat{\pi}_t + \frac{m}{b} \hat{m}_{t-1} + (\beta^{-1} - \zeta) \hat{b}_{t-1} - \beta^{-1} \delta_t q_{t-1} \quad (26)$$

The log-linearised asset market clearing condition (17), together with the definition of *ex post* nominal returns on the ambiguous asset (see footnote 6), imply:

$$-q_t + \hat{V}_t = \hat{b}_{L,t} \quad (27)$$

$$\hat{R}_{L,t}^e = \beta \mathbb{E}_t \hat{V}_{t+1} - \hat{V}_t \quad (28)$$

where $\hat{R}_{L,t}^e \equiv \mathbb{E}_t \hat{R}_{L,t+1}$.

Intuitively, the Euler equation (22), together with the no-arbitrage asset-pricing equation (23), imply that the output gap – and, via the New Keynesian Phillips curve (25), the rate of inflation – are determined in our model by two interest rates rather than one.

In particular, the interest rate on the safe asset and the interest rate on the ambiguous asset each have an independent influence over aggregate outcomes, with the latter being determined by the relative share of the ambiguous asset absorbed from the household by the central bank. The identity of assetholders is not innocuous: whether a security is held by the central bank on the consolidated balance sheet of the government or by the household in its contemporaneous savings portfolio affects the transmission of shocks. Because relative asset quantities matter for relative asset prices, QE has real effects.

4.7 Optimal monetary policy

The monetary policymaker minimises a discounted loss function subject to the log-linearised first-order conditions (22)-(28) characterising the structure of the macroeconomy.

The monetary policymaker has two control variables: the ‘conventional’ policy rate (that is, the nominal interest rate on the safe asset in each period, R_t) and ‘unconventional’ asset purchases (that is, the fraction of the ambiguous asset held on the central bank’s balance sheet, q_t).

More precisely, the monetary policymaker’s objective can be expressed as:

$$\min_{\{R_t, q_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[\hat{x}_t^2 + \frac{\eta}{\kappa} \hat{\pi}_t^2 + \frac{\tilde{v} c^{1/\sigma-1}}{\sigma^{-1} + \psi} \left(\hat{b}_t - \hat{b}_{L,t} \right)^2 \right] \quad (29)$$

The loss function in (29) is standard, except for its last term. The monetary policymaker is not only concerned with stabilising the output gap and the deviation of inflation

from target, but also with stabilising the household’s relative holdings of each imperfectly substitutable asset available. The policymaker thereby internalises the disutility incurred by the household when the household’s actual portfolio deviates from the household’s preferred portfolio.

The formal derivation of equivalence between the the household’s utility function in (9) and the monetary policymaker’s loss function in (29) is provided in Appendix E. To solve for optimal commitment policy, we follow the approach set out in Dennis (2007).

4.8 Calibration

We calibrate our model by following recent empirical applications of standard New Keynesian models to the U.S. economy. Steady-state values are distinguished from transition values by the absence of a time subscript.

Table 2 summarises the key calibrated parameters of our model.

The productivity parameter, A , is chosen to normalise output to unity in steady state. The steady-state rate of inflation is normalised to zero, which implies $\pi = 1$. The elasticity of substitution within the household’s consumption bundle, η , follows standard calibrations of the canonical New Keynesian model. The parameter χ_m is chosen such that real money balances are a very small fraction of output in steady state. The household’s discount factor, β , implies a nominal interest rate on the ‘safe’ asset in line with the historical average yield on short-term government securities. The elasticity of money demand, σ_m , implies a unitary income elasticity. The household’s intertemporal elasticity of substitution, σ , follows Ju and Miao (2012). The slope of the Phillips curve, κ , matches Levin et al. (2010). Combined with the assumption that firms change prices on average once a year, this assumption determines the degree of price rigidity, α , and the elasticity of labour supply, ψ .

The rate of asset purchases in steady state, q , is assumed to be zero. This assumption is consistent with implementation of the efficient equilibrium. The feedback parameter in the fiscal authority’s transfer rule, ζ , implies that the stock of safe assets exhibits some persistence. The autoregressive coefficient assumed in the evolution of the real natural interest rate, ρ , also imparts a degree of persistence to r^* following the incidence of the original shock.

Table 2: CALIBRATION

Parameter	Description	Value
σ	Elasticity of intertemporal substitution	1.5
β	Household discount factor	0.9950
κ	Slope of Phillips Curve	0.024
η	Elasticity of substitution within consumption bundle	5
σ_m	Elasticity of money demand	1.5
α	Probability of firm <i>not</i> changing price	0.75
ψ	Elasticity of labour supply	0.11
ζ	Feedback parameter in fiscal transfer rule	0.5
ρ	Persistence of real natural interest rate	0.5

Notes: Selected values of key parameters. See Section 4.8 for full details.

5 The effects of quantitative easing

To demonstrate the real effects of central bank asset purchases, we consider the response of the economy to an adverse demand shock. As described in Section 4.5, an adverse demand shock is represented in the model by an exogenous shock to the real natural interest rate at time t : $\varepsilon_t < 0$.

Consider the transmission mechanism associated with each policy instrument, R and q , in turn.

In response to the negative demand shock, the policymaker optimally reduces the interest rate, R_t , on the safe asset. In the model, this reduction in the policy rate is brought about by conventional open-market operations: by intervening in the market for the safe asset, the central bank reduces its relative supply, thereby raising its price and lowering its yield. As a result of this reduction in relative supply, however, the household's actual holdings of the safe asset, B_t , must fall relative to its holdings of the ambiguous asset, $B_{L,t}$. In log-linearised terms, this implies a reduction in $(\hat{b}_t - \hat{b}_{L,t})$. Given the household's asset pricing equation (23), a reduction in $(\hat{b}_t - \hat{b}_{L,t})$ implies a *ceteris paribus* increase in the ambiguity premium, $(\hat{R}_{L,t}^e - \hat{R}_t)$.

This increase in the ambiguity premium is counterproductive.

Figure 8 plots the impulse response functions of six key endogenous variables to a negative shock to the real natural interest rate at time zero, under the assumption that the monetary policymaker only has the conventional policy instrument R_t at his disposal. As shown by the two response functions on the right-hand-side of Figure 8, the output gap

and the inflation gap endure, even five years after the incidence of the original demand shock.

Complete stabilisation of output and inflation requires elimination of the ambiguity premium. This is something that the policymaker can only achieve by using his unconventional instrument in addition to his conventional instrument. Recall that quantitative easing entails the central bank purchasing some share q_t of the ambiguous asset directly from the household. The asset market clearing condition (27) can be rearranged as:

$$\hat{b}_{L,t} = -q_t + \hat{V}_t \quad (30)$$

By increasing q_t , the policymaker can reduce $b_{L,t}$ by just enough to offset the reduction in $(\hat{b}_t - \hat{b}_{L,t})$ associated with the decline in R_t . Use of the unconventional instrument – in parallel with use of the conventional instrument – thereby mitigates any increase in the ambiguity premium $(\hat{R}_{L,t}^e - \hat{R}_t)$ and allows the monetary policymaker to achieve complete stabilisation of output and inflation.

Figure 9 plots the impulse response functions of the same six endogenous variables to the same negative shock to the real natural interest rate described in Figure 8, but now under the assumption that the monetary policymaker has both the conventional policy instrument R_t and the unconventional policy instrument q_t at his disposal. As shown by the two response functions on the right-hand-side of Figure 9, the output gap and the inflation gap are eliminated — permanently — as soon as the adverse demand shock hits.

6 The importance of portfolio preferences

According to Laubach and Williams (2016), the real natural interest rate in the United States fell from 1.75% to 0.50% between the first quarter of 2008 and the first quarter of 2010. According to Figures 1 and 2, U.S. households held around 1.1 long-term Treasury bonds for every 1 short-term Treasury bond on the eve of the Federal Reserve’s first round of large-scale asset purchases in late 2008.

Assuming this is a reasonable proxy for δ_t in our model – that is, a reasonable guide to the representative household’s preferred ratio of ambiguous asset holdings to safe asset holdings – our calibrated model suggests that the optimal policy response to a 125 basis point negative shock to the real natural interest rate when $\delta_t = 1.1$ is a sharp reduction

in the interest rate on the safe asset and the purchase of 33% of the outstanding share of the ambiguous asset.

The blue lines in Figure 10 illustrate this optimal response in more detail: intensive use of both the conventional and unconventional instrument allows the monetary policymaker to completely stabilise output and inflation.

The magnitudes of this theoretical response accord with empirical estimates of the Federal Reserve’s actual response in the aftermath of Lehman Brothers’ collapse. Gagnon et al. (2011) use the concept of ‘ten-year equivalents’ to estimate the nominal value of ten-year par Treasury debt that would have been equivalent in duration to the portfolio of securities purchased by the Federal Reserve during the implementation of LSAPs between late 2008 and March 2010. The authors conclude that the central bank’s cumulative asset purchases during that period amounted to the absorption of between a quarter and a third of the outstanding stock of ten-year equivalent Treasury debt.

Moreover, our structural model allows us to experiment with counterfactuals. Just twelve months before QE1 began, Figures 1 and 2 suggest that the ratio of long-term to short-term Treasury bonds held by the U.S. private sector was more than two times greater than it was in late 2008. Consider again the optimal monetary policy response to a 125bp negative shock to the real natural interest rate, but now under the assumption that the relative weight placed on the ambiguous asset in the representative household’s preferred portfolio is twice as high as we assumed under our original scenario.

The red lines in Figure 10 illustrate the optimal policy response under this counterfactual scenario. With $\delta_t = 2.3$ rather than $\delta_t = 1.1$, the Federal Reserve can completely offset deviations in output and inflation by purchasing just 21% of the outstanding share of the ambiguous asset rather than 33%. This suggests that, had the Federal Reserve conducted asset purchases in late-2007, the scale of the policy intervention required to stabilise the U.S. economy would have been two-thirds that of the scale required when LSAPs actually began in late-2008.

This counterfactual suggests that there are considerable costs associated with waiting, especially when it comes to the implementation of large-scale asset purchases. It also suggests that there is considerable value in using conventional and unconventional policy instruments simultaneously rather than sequentially.

Between September 2007 and April 2008, the U.S. Federal Reserve lowered the target for the federal funds rate on six separate occasions, from 5.25% to 2.00%. During this period, both regional subprime mortgage lenders and major international financial institutions came close to bankruptcy. Only after the collapse of Lehman Brothers in September 2008 did the Federal Reserve ultimately lower the federal funds rate to its effective lower bound (a target range of 0-0.25%) and proceed to purchase \$600bn in mortgage-backed securities, followed in early 2009 by the purchase of \$300bn in longer-dated Treasury debt.

In hindsight, this chronology is unsurprising. Conventional wisdom dictates that interest rate policy and balance sheet policy are to be used lexicographically – that is, central banks should resort to “unconventional” tools only after “conventional” tools have been exhausted. Yet Figures 8-10 suggest that the policymaker’s optimal strategy in response to a negative shock to demand involves early and active absorption of ambiguous assets from household portfolios *in parallel* with sharp reductions in the policy rate. To the extent that this strategy can pre-empt – or even mitigate – a ‘flight to safety’ among investors, the policymaker is better able to stabilise output and inflation.¹⁰

As early as August 2007, the Federal Reserve had begun to express concerns about the incipient subprime lending crisis: “financial market conditions have deteriorated, and tighter credit conditions and increased uncertainty have the potential to restrain economic growth going forward.” By December 2008, the Federal Reserve confirmed that it would “purchase large quantities of agency debt and mortgage-backed securities to provide support to the mortgage and housing markets,” and that it “stands ready to expand its purchases of agency debt and mortgage-backed securities as conditions warrant.” The FOMC was already in the process of “evaluating the potential benefits of purchasing longer-term Treasury securities.”

Fifteen months passed between the Federal Reserve’s first response with conventional policy and the Federal Reserve’s first response with unconventional policy. By the end of that period, our model suggests that a “dramatic shift” in portfolio preferences had already undermined the ability of central bank purchases of risky assets to stimulate real activity in the private sector.

¹⁰This is a conclusion reached in a different context by Ellison and Tischbirek (2013).

6.1 Portfolio preferences as “bang for the buck”

The asset pricing equation (23) implies the following ambiguity premium:

$$\hat{R}_{L,t}^e - \hat{R}_t = -(1 + \delta_t) \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} (\hat{b}_t - \hat{b}_{L,t}) \quad (31)$$

It is this spread between the yield on the ambiguous asset and the yield on the safe asset that offers the central bank traction to influence the output gap – via the Euler equation (22), and to control the deviation of inflation from target – via the Phillips Curve (25).

Combining equations (22) and (23) with the asset market clearing condition (27), and considering them alongside equation (25), we arrive at the following compact representation of the key determinants of output and inflation in our model:

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left\{ \left(\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - r_t^* \right) - \omega \delta_t \left(q_t + \hat{b}_t - \hat{V}_t \right) \right\} \quad (32)$$

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t \quad (33)$$

In equations (32) and (33), κ and ω denote the following convolutions of parameters:

$$\kappa = \frac{(1 - \alpha)(1 - \beta\alpha)}{\alpha} (\psi + \sigma^{-1}) \quad (34)$$

$$\omega = \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} \quad (35)$$

Differentiating output with respect to asset purchases in (32) yields:

$$\frac{\partial \hat{x}_t}{\partial q_t} = \sigma \omega \delta_t > 0 \quad (36)$$

The derivative (36) implies that the lower the value of δ_t , the lower the semi-elasticity of output with respect to marginal changes in the share of the ambiguous asset absorbed by the central bank. Equivalently, the lower the value of δ_t , the greater the share of the ambiguous asset that must be purchased by the central bank to eliminate the ambiguity premium, and the less effective the marginal increase in q_t in eliminating output and inflation deviations. In intuitive terms, the greater the representative household’s

relative preference for the safe asset, the smaller the “bang for the buck” associated with unconventional monetary policy.

Moreover, our choice-theoretic specification of the robustly optimal portfolio implies that the representative household’s relative preference for the safe asset is increasing in her subjective uncertainty about the state of the world. The optimality condition (8) dictates that the greater the household’s perception of model uncertainty, $\sigma_\mu^2(\mathbb{E}(R_L))$, and/or the greater the household’s aversion to a given degree of model uncertainty, θ , the lower the value of δ_t . Put differently, the greater the ambiguity that prevails and/or the greater the household’s fear of a given degree of ambiguity, the greater the preference for safety. Not only do households’ portfolio preferences determine the “bang for the buck” of central bank asset purchases, but a ‘flight to safety’ induced by the uncertainty accompanying a recession undermines the effectiveness of unconventional monetary policy – just when it is needed most.

6.2 Safety traps

This result is an alternative interpretation of the ‘safety trap’ described by Landau (2014) and inspired by Caballero and Farhi (2014). In that formulation, both as a sign and a consequence of extreme risk aversion, economic agents develop a strong inclination in certain states to hold money and government bonds at the expense of other real or financial assets. In order to accumulate more safe assets, agents must reduce their consumption or cutback on investment: a form of precautionary saving which depresses private-sector activity, and from which there is little escape other than via a transfer of risk from private to public balance sheets.

In Caballero and Farhi (2014), safe asset shortages operate via the interest rate. An excess demand for safe assets drives the risk-free natural interest rate lower. If inflation is already subdued, and if nominal interest rates are already stuck at their effective lower bound, it becomes impossible for the monetary policymaker to attain a real interest rate low enough to stimulate aggregate demand. Inflation therefore falls further, which increases the real interest rate and leads to further disinflation.

In Landau (2014), a safety trap accelerates this self-reinforcing dynamic. A dramatic shift in investors’ desire for safety can accelerate the decline in the risk-free interest rate, exacerbating the gap between effective and equilibrium rates of return. An increase in demand for safe assets therefore acts as an endogenous tightening of monetary conditions.

In our model, δ_t determines the sensitivity of output to asset purchases. This formalises the effect of a safety trap on the effectiveness of unconventional monetary policy under uncertainty. As Landau (2014) puts it (albeit informally), a strong preference for safe assets can inhibit the apparent arbitrage between riskless and risky assets and has the potential to significantly impair – or indeed paralyse – the process of portfolio rebalancing.

7 Conclusion

Two key macroeconomic ideas, advanced almost four decades before the global financial crisis began, became the theoretical basis for a practical policy response that remains the subject of debate almost a decade after the global financial crisis ended.

First, Tobin (1969) argued that when assets are imperfect substitutes, an increase in the relative supply of any one asset not only forces the rate of return on that particular asset to change, it also prompts the rates of return on all other assets to adjust in order to induce the public to hold the incremental addition to supply. As a result, “there is no reason to think that the impact [of monetary policy] will be captured in any single exogenous or intermediate variable, whether it is a monetary stock or a market interest rate.” Rather, the impact of monetary policy must be inferred from changes in the entire term structure of the complete set of interest rates at which borrowing takes place in the real economy.

Second, Tobin (1969) suggested that “changes can occur, and undoubtedly do, in the portfolio preferences – asset demand functions – of the public, the banks, and other sectors. These preferences are based on expectations, estimates of risk, attitudes towards risk, and a host of other factors.” As a result, both the private-sector’s tastes and the private-sector’s information have indirect effects on relative rates of return beyond the direct effects that stem from relative asset supplies.

Our model combines these two ideas by embedding a choice-theoretic optimal asset portfolio within the consumption-savings problem faced by a representative household. The representative household’s decision problem, in turn, is embedded within a general equilibrium framework which incorporates the behaviour of firms and the government. Optimal behaviour in each sector jointly determines aggregate output and inflation.


Our model looks beyond the riskiness of asset returns to incorporate the presence of a

more pervasive form of unstructured uncertainty.¹¹ In particular, a household’s optimal allocation among available assets is a function of both her attitude towards a given degree of uncertainty and the degree of uncertainty that she perceives in the macroeconomic environment.

During the global financial crisis of 2007/08, aggregate output collapsed and macroeconomic uncertainty spiked. In response, several major central banks embarked on large-scale asset purchase programmes. The scope of these programmes extended well beyond purchases of risk-free government bonds to include the accumulation of private-sector credit instruments with increasingly uncertain returns. Our model suggests that in times of crisis, “when the economic environment is so complex as to appear nearly incomprehensible,” the “extreme prudence, if not outright paralysis” that characterises the portfolio preferences of investors and consumers tends to undermine the effectiveness of unconventional monetary policy — just when it is needed most.

¹¹Elsewhere, Tobin (1969) quips: “It is now more than a decade ago that I participated in the modest endeavour of doubling the number of parameters of investors’ probability estimates involved in economists’ analyses of asset choice. This extension from one moment to two was never advertised as the complete job or the final word.”


Figure 1: THERE HAS BEEN CONSIDERABLE VARIATION IN HOUSEHOLDS' 'PREFERRED PORTFOLIOS' OVER TIME



Notes: δ_t denotes the ratio of (i) the private sector's holdings of Treasury securities of maturity *greater* than one year, to (ii) the private sector's holdings of Treasury securities of maturity *less* than one year. Shaded bars indicate US recessions.

Source: Post-1975: Table FD-5 of the Treasury Bulletin; Pre-1975: Federal Reserve's Annual Statistical Digest, and Banking and Monetary Statistics.


Figure 2: PORTFOLIO PREFERENCES WERE PARTICULARLY VOLATILE DURING THE GLOBAL FINANCIAL CRISIS



Notes: δ_t denotes the ratio of (i) the private sector's holdings of Treasury securities of maturity *greater* than one year, to (ii) the private sector's holdings of Treasury securities of maturity *less* than one year. Red lines indicate commencement of three phases of Federal Reserve LSAPs.


Source: Table FD-5 of the Treasury Bulletin.

Figure 3: STATISTICAL ESTIMATES OF PERCEIVED MACROECONOMIC UNCERTAINTY
ARE COUNTERCYCLICAL




Notes: The statistical index of uncertainty is taken from Jurado, Ludvigson and Ng (2015). Shaded bars indicate US recessions.

Figure 4: MODEL-FREE AND MODEL-IMPLIED MEASURES OF PERCEIVED
MACROECONOMIC UNCERTAINTY CO-MOVE, RISING IN RECESSIONS




Notes: The statistical index of uncertainty is taken from Jurado, Ludvigson and Ng (2015). (We plot the trailing 12-month average of the authors' monthly series). Model-implied uncertainty is taken from Collard, Mukerji, Sheppard and Tallon (2011). (We plot the posterior conditional variance of the distribution of the latent state driving consumption and dividends – distorted according to ambiguity aversion. The authors denote this variance $\widehat{Var}(x_t)$). Both series are shown as (signed) proportionate deviations from their respective means.

Figure 5: NEGATIVE DEMAND SHOCKS COINCIDE WITH GREATER MACROECONOMIC UNCERTAINTY




Notes: The statistical index of uncertainty is taken from Jurado, Ludvigson and Ng (2013). The shock to the real natural interest rate is identified by the cumulative, four-quarter change in the real natural interest rate estimated by Laubach and Williams (2016). Shaded bars indicate US recessions.

Figure 6: GREATER MACROECONOMIC UNCERTAINTY IS ASSOCIATED WITH DE-RISKING OF HOUSEHOLD PORTFOLIOS




Notes: δ_t denotes the ratio of (i) the private sector's holdings of Treasury securities of maturity *greater* than one year, to (ii) the private sector's holdings of Treasury securities of maturity *less* than one year. The statistical index of uncertainty is taken from Jurado, Ludvigson and Ng (2015). Shaded bars indicate US recessions.

Figure 7: FLIGHT TO SAFETY: HOUSEHOLDS PREFER SAFE ASSETS IN BAD TIMES




Notes: δ_t denotes the ratio of (i) the private sector's holdings of Treasury securities of maturity *greater* than one year, to (ii) the private sector's holdings of Treasury securities of maturity *less* than one year. The shock to the real natural interest rate is identified by the cumulative, four-quarter change in the real natural interest rate estimated by Laubach and Williams (2016). Shaded bars indicate US recessions.

Figure 8: WITH ONLY THE CONVENTIONAL INSTRUMENT AVAILABLE, THE POLICYMAKER CANNOT STABILISE OUTPUT AND INFLATION




Notes: Each panel depicts the response of a key endogenous variable to a one-off negative shock to the real natural interest rate (r^*) at time zero, assuming that the unconventional policy instrument is de-activated. R is the interest rate on the safe asset (the conventional policy instrument). R_L is the interest rate on the ambiguous asset. All three interest rates are expressed in percent. q is the fraction of the ambiguous asset held on the central bank's balance sheet (the unconventional policy instrument). x is the output gap (in percent). π is the deviation of inflation from target (in percent). The horizontal axis is in quarters.

Figure 9: WITH BOTH THE CONVENTIONAL AND UNCONVENTIONAL INSTRUMENT, THE POLICYMAKER CAN PERFECTLY STABILISE BOTH OUTPUT AND INFLATION



Notes: Each panel depicts the response of a key endogenous variable to a one-off negative shock to the real natural interest rate (r^*) at time zero, assuming that both the conventional policy instrument and the unconventional policy instrument are available. R is the interest rate on the safe asset (the conventional policy instrument). R_L is the interest rate on the ambiguous asset. All three interest rates are expressed in percent. q is the fraction of the ambiguous asset held on the central bank's balance sheet (the unconventional policy instrument). x is the output gap (in percent). π is the deviation of inflation from target (in percent). The horizontal axis is in quarters.

Figure 10: A ‘FLIGHT TO SAFETY’ NEGATES THE EFFECTIVENESS OF QE



Notes: Each panel depicts the response of a key endogenous variable to a one-off 125bp negative shock to the real natural interest rate (r^*) at time zero, assuming that both the conventional policy instrument and the unconventional policy instrument are available. R is the interest rate on the safe asset (the conventional policy instrument). R_L is the interest rate on the ambiguous asset. All three interest rates are expressed in percent. q is the fraction of the ambiguous asset held on the central bank’s balance sheet (the unconventional policy instrument). When δ is “high” rather than “low”, the representative household’s preferred portfolio is skewed towards the ambiguous asset. The horizontal axis is in quarters.

A

This appendix accompanies Section 4.1.1. It follows Maccheroni et al. (2013).

A.1 Mathematical preliminaries

Given a probability space $(\Omega, \mathcal{F}, P)$, let $L^2 = L^2(\Omega, \mathcal{F}, P)$ be the Hilbert space of square integrable random variables on Ω . Let $L^\infty = L^\infty(\Omega, \mathcal{F}, P)$ be the subset of L^2 consisting of its almost surely bounded elements. Given an interval $I \subseteq \mathbb{R}$, set:

$$L^\infty(I) = \{f \in L^\infty : \text{essinf} f, \text{esssup} f \in I\} \quad (37)$$

Let $E_P(X)$ and $\sigma_P^2(X)$ denote the mean and variance of a random variable $X \in L^2$, and let $\sigma_P(X, Y)$ denote the covariance between two random variables $X, Y \in L^2$.

The set of probability measures Q on $\mathcal{F}$ that have square integrable density $q = dQ/dP$ with respect to P can be identified with the closed and convex subset of L^2 given by:

$$\Delta = \{q \in L_+^2 : \int_{\Omega} q(\omega) dP(\omega) = 1\} \quad (38)$$

This implies the following lemma, in which the density $\bar{q}$ is known as the ‘barycenter’ of μ and is denoted $\int_{\Delta} q d\mu(q)$.

Lemma A.1: Given a Borel probability measure μ on Δ with bounded support, there exists, for all $X \in L^2$, a unique $\bar{q} \in \Delta$ such that:

$$\int_{\Omega} X(\omega) \bar{q}(\omega) dP(\omega) = \int_{\Delta} \left(\int_{\Omega} X(\omega) q(\omega) dP(\omega) \right) d\mu(q) \quad (39)$$

When restricted to indicator functions 1_A of elements of $\mathcal{F}$, (39) can be rewritten, for all $A \in \mathcal{F}$, as:

$$\bar{Q}(A) = \int_{\Delta} Q(A) d\mu(Q) \quad (40)$$

As discussed in Section 2.1 of Maccheroni, Marinacci and Ruffino (2013), μ can therefore be interpreted as a lottery whose outcomes are all possible models. These possible models, in turn, can be interpreted as lotteries determining states.

A.2 Decision theory

Given any nonsingleton interval $I \subseteq \mathbb{R}$ of monetary outcomes, consider decision-makers who behave according to the smooth model of decision-making under ambiguity, as developed by Klibanoff, Marinacci and Mukerji (2005). Such decision-makers rank prospects via the functional $V : L^\infty(I) \rightarrow \mathbb{R}$ which is defined, for all $f \in L^\infty(I)$, as:

$$V(f) = \int_{\Delta} \phi \left(\int_{\Omega} u(f(\omega)) q(\omega) dP(\omega) \right) d\mu(q) \quad (41)$$

where μ is a Borel probability measure on Δ with bounded support, and $u : I \rightarrow \mathbb{R}$ and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ are continuous and strictly increasing functions. The functional $V : L^\infty(I) \rightarrow \mathbb{R}$ is well defined, with $V(L^\infty(I)) = \phi(u(I))$.

The certainty equivalent function $C : L^\infty \rightarrow I$ induced by V is defined by $V(C(f)) = V(f)$ for all prospects $f \in L^\infty(I)$. Specifically:

$$C(f) = u^{-1} \left(\phi^{-1} \left(\int_{\Delta} \phi \left(\int_{\Omega} u(f(\omega)) q(\omega) dP(\omega) \right) d\mu(q) \right) \right) \quad (42)$$

Monetary certainty equivalents in the face of risky financial prospects can be captured by setting $v = \phi \circ u : I \rightarrow \mathbb{R}$, which implies that (42) can be written as:

$$C(f) = v^{-1} \left(\int_{\Delta} v \left(u^{-1} \left(\int_{\Omega} u(f(\omega)) q(\omega) dP(\omega) \right) \right) d\mu(q) \right) \quad (43)$$

In (43), while u represents the decision-maker's attitude towards risk, the function v represents her attitude towards model uncertainty. Model uncertainty augments state uncertainty. Together, model uncertainty and state uncertainty determine the ambiguity faced by the decision-maker in ranking prospects $f : \Omega \rightarrow \mathbb{R}$. The concavity of the function ϕ characterises ambiguity aversion.

Since:

$$\lambda_\phi(u(w)) = \frac{\phi''}{\phi'} = \frac{1}{u'(w)} (\lambda_v(w) - \lambda_u(w)) \quad (44)$$

Maccheroni, Marinacci and Ruffino (2013) conclude that ambiguity aversion amounts to the difference between λ_v and λ_u .

A.3 Quadratic approximation

Let $w \in \text{int } I$ be a scalar measure of current wealth, which includes the degenerate random variable $w1_\Omega$. Given any prospect $h \in L^\infty$ such that $w + h \in L^\infty(I)$, we focus on the certainty equivalent $C(w + h)$ of $w + h$, that is:

$$C(w + h) = v^{-1} \left(\int_{\Delta} v \left(u^{-1} \left(\int_{\Omega} u(w + h)q \, dP(\omega) \right) \right) d\mu(q) \right) \quad (45)$$

For all $h \in L^\infty$ consider:

$$\begin{aligned} E(h) : q &\mapsto \int_{\Omega} hq \, dP \\ \sigma^2(h) : q &\mapsto \int_{\Omega} h^2q \, dP - \left(\int_{\Omega} hq \, dP \right)^2 \\ \sigma_\mu^2(E(h)) &:= \int_{\Delta} \left(\int_{\Omega} h(\omega)q(\omega) \, dP(\omega) \right)^2 d\mu(q) - \left(\int_{\Delta} \left(\int_{\Omega} h(\omega)q(\omega) \, dP(\omega) \right) d\mu(q) \right)^2 \end{aligned} \quad (46)$$

The variance $\sigma_\mu^2(E(h))$ reflects uncertainty about the expectation $E(h)$ which, in turn, derives from model uncertainty.

The second-order approximation of the certainty equivalent (45) is given by the following proposition (see Proposition 3 of Maccheroni, Marinacci and Ruffino (2013)):

Proposition A.3: Let μ be a Borel probability measure with bounded support on Δ , and let $u, v : I \rightarrow \mathbb{R}$ be twice differentiable with $u', v' > 0$. Then, for all $h \in L^\infty$ such that $w + h \in L^\infty(I)$, we have:

$$C(w + h) = w + E_{\overline{Q}}(h) - \frac{1}{2}\lambda_u(w)\sigma_{\overline{Q}}^2(h) - \frac{1}{2}(\lambda_v(w) - \lambda_u(w))\sigma_\mu^2(E(h)) + R_2(h) \quad (47)$$

where:

$$\lim_{t \rightarrow 0} \frac{R_2(th)}{t_2} = 0 \quad (48)$$

and, if $\mathcal{F}$ is finite, then $R_2(h) = o(\|h\|^2)$ as $h \rightarrow 0$ in L^2 . $\square$

The variance $\sigma_Q^2(h)$, in the third term on the right-hand side of (47), can be decomposed along two dimensions in the following way:

$$\sigma_Q^2(h) = E_\mu(\sigma^2(h)) + \sigma_\mu^2(E(h)) \quad (49)$$

State uncertainty, which exists within each model, determines the average variance $E_\mu(\sigma^2(h))$. Model uncertainty determines the variance of averages $\sigma_\mu^2(E(h))$.

Given (49), the approximation in (47) can be rearranged – given Arrow-Pratt coefficients for u and v – as:

$$C(w+h) = w + E_{\bar{Q}}(h) - \frac{\lambda_u(w)}{2} E_\mu(\sigma^2(h)) - \frac{\lambda_v(w)}{2} \sigma_\mu^2(E(h)) + R_2(h) \quad (50)$$

According to the certainty equivalent given by (50), risk aversion determines the decision-maker's reaction to the average variance of returns, while ambiguity aversion determines her reaction to the variance of average returns. If the decision-maker is risk averse at w , then $\lambda_u(w) \geq 0$. If the decision-maker is ambiguity averse at w , then $\lambda_v(w) - \lambda_u(w) \geq 0$.

The quadratic approximation developed in this section holds exactly when the decision-maker's preferences exhibit constant absolute aversion to both risk and model uncertainty, and when the prospect h follows a normal distribution with unknown mean and known variance.

B Robust mean-variance preferences

This appendix accompanies Section 4.1.1. It follows Maccheroni et al. (2013).

B.1 Incorporating uncertainty

Given the quadratic approximation in Appendix A.3, standard mean-variance preferences can be generalised to account for model uncertainty.

Consider a decision-maker who ranks prospects f in L^2 via the robust mean-variance functional $C : L^2 \rightarrow \mathbb{R} \cup \{-\infty\}$ given by:

$$C(f) = E_{\bar{Q}}(f) - \frac{\lambda}{2} \sigma_{\bar{Q}}^2(f) - \frac{\theta}{2} \sigma_\mu^2(E(f)) \quad (51)$$

where λ and θ are strictly positive coefficients and μ is a Borel probability measure on Δ with bounded support and barycenter $\bar{Q}$.

Under the assumption that $\bar{Q} = P$, $C(f)$ is always finite, (51) can be rewritten as:

$$C(f) = E_P(f) - \frac{\lambda}{2}\sigma_P^2(f) - \frac{\theta}{2}\sigma_\mu^2(E(f)) \quad (52)$$

C The generic portfolio problem

This appendix accompanies Section 4.1.1. It follows Maccheroni et al. (2013).

C.1 The robustly optimal allocation

Consider the one-period optimisation problem of a decision-maker who must decide how to allocate one unit of wealth among $n + 1$ assets at time zero. The gross return on asset i after one period is denoted $r_i \in L^2$ for $i = 1, \dots, n$. As such, the $(n \times 1)$ vector of returns on the first n assets is denoted $\mathbf{r}$. The return on the $(n + 1)^{th}$ asset is risk-free and is set equal to a constant $\mathbf{r}_f$. The $(n \times 1)$ vector of portfolio weights indicating the fraction of wealth invested in each asset is denoted $\mathbf{w}$.

The end-of-period wealth $\mathbf{r}_w$ induced by a choice of $\mathbf{w}$ is therefore given by:

$$\mathbf{r}_w = \mathbf{r}_f + \mathbf{w} \cdot (\mathbf{r} - \mathbf{r}_f \mathbf{1}) \quad (53)$$

where $\mathbf{1}$ is the n -dimensional unit vector.

Assuming frictionless financial markets with no transactions costs and no restrictions on either borrowing or short-selling, the decision-maker's portfolio problem can be written as:

$$\max_{\mathbf{w} \in \mathbb{R}^n} C(\mathbf{r}_w) = \max_{\mathbf{w} \in \mathbb{R}^n} \left(\mathbb{E}_P(\mathbf{r}_w) - \frac{\lambda}{2}\sigma_P^2(\mathbf{r}_w) - \frac{\theta}{2}\sigma_\mu^2(\mathbb{E}(\mathbf{r}_w)) \right) \quad (54)$$

Let:

$$\mathbf{m} = [E_P(\mathbf{r}_1 - \mathbf{r}_f), \dots, E_P(\mathbf{r}_n - \mathbf{r}_f)]^T \quad (55)$$

$$\Sigma_P = [\sigma_p(\mathbf{r}_i, \mathbf{r}_j)]_{i,j=1}^n \quad (56)$$

$$\Sigma_\mu = [\sigma_\mu(E(\mathbf{r}_i), E(\mathbf{r}_j))]_{i,j=1}^n \quad (57)$$

$$\Xi = \lambda \Sigma_P + \theta \Sigma_\mu \quad (58)$$

These simplifications imply that (54) becomes:

$$\max_{\mathbf{w} \in \mathbb{R}^n} \left(\mathbf{r}_f + \mathbf{w} \cdot \mathbf{m} - \frac{\lambda}{2} \mathbf{w}^T \Sigma_P \mathbf{w} - \frac{\theta}{2} \mathbf{w}^T \Sigma_\mu \mathbf{w} \right) \quad (59)$$

which is equivalent to:

$$\max_{\mathbf{w} \in \mathbb{R}^n} \left(\mathbf{w} \cdot \mathbf{m} - \frac{1}{2} \mathbf{w}^T \Xi \mathbf{w} \right) \quad (60)$$

From (60) it becomes clear that the optimal solution $\hat{\mathbf{w}}$ satisfies $\Xi \hat{\mathbf{w}} = \mathbf{m}$.

The optimality condition associated with (60) can therefore be interpreted as:

$$(\lambda \cdot \text{Var}_P[\mathbf{r}] + \theta \cdot \text{Var}_\mu[E(\mathbf{r})]) \cdot \hat{\mathbf{w}} = E_P(\mathbf{r} - \mathbf{r}_f \cdot \mathbf{1}) \quad (61)$$

where:

$\text{Var}_P[\mathbf{r}] \equiv [\sigma_p(\mathbf{r}_i, \mathbf{r}_j)]_{i,j=1}^n$ is the variance-covariance matrix of returns under P ,

$\text{Var}_\mu[E(\mathbf{r})] \equiv [\sigma_\mu(E(\mathbf{r}_i), E(\mathbf{r}_j))]_{i,j=1}^n$ is the variance-covariance matrix of expected returns under μ ,

and

$E_P(\mathbf{r} - \mathbf{r}_f \cdot \mathbf{1}) \equiv [E_P(\mathbf{r}_i - \mathbf{r}_f)]_{i=1}^n$, is the vector of expected excess returns under P ,

If $n = 1$, there is only one ambiguous asset. Equation (61) then reduces to:

$$\hat{w} = \frac{E_P(r) - r_f}{\lambda \sigma_P^2(r) + \theta \sigma_\mu^2(E(r))} \quad (62)$$

as in Section 4.1.1 of the main text.

D

This appendix accompanies Section 4. It follows Harrison (2012).

D.1 The household

The household optimisation problem considered in Section 4.1 is given by:

$$\begin{aligned} & \max_{\{c_t, n_t, M_t, B_t, B_{L,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \phi_t \left(\frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{n_t^{1+\psi}}{1+\psi} + \frac{\chi_m^{-1}}{1-\sigma_m^{-1}} \left(\frac{M_t}{P_t} \right)^{1-1/\sigma_m} - \frac{\tilde{v}}{2} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right]^2 \right) \\ & \text{s.t.} \\ & B_t + B_{L,t} + M_t = R_{t-1}B_{t-1} + R_{L,t}B_{L,t-1} + M_{t-1} + W_t n_t + T_t + D_t - P_t c_t \end{aligned} \quad (63)$$

Let μ denote the Lagrange multiplier on the budget constraint in (63). The first-order conditions associated with the problem are then given by:

$$\frac{\phi_t}{c_t^{1/\sigma}} = \mu_t P_t \quad (64)$$

$$\phi_t n_t^\psi = W_t \mu_t \quad (65)$$

$$\phi_t \chi_m^{-1} \left(\frac{M_t}{P_t} \right)^{-1/\sigma_m} \frac{1}{P_t} - \mu_t + \beta \mathbb{E}_t \mu_{t+1} = 0 \quad (66)$$

$$-\mu_t + \phi_t \tilde{v} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right] \frac{\delta_t B_t}{B_{L,t}^2} + \beta \mathbb{E}_t \mu_{t+1} R_{L,t+1} = 0 \quad (67)$$

$$-\phi_t \tilde{v} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right] \frac{\delta_t}{B_{L,t}} - \mu_t + \beta R_t \mathbb{E}_t \mu_{t+1} = 0 \quad (68)$$

Let the real Lagrange multiplier be defined as:

$$\Lambda_t \equiv P_t \mu_t \quad (69)$$

Let real money balances and bond holdings be defined as:

$$m_t \equiv \frac{M_t}{P_t} \quad (70)$$

$$b_t \equiv \frac{B_t}{P_t} \quad (71)$$

$$b_{L,t} \equiv \frac{B_{L,t}}{P_t} \quad (72)$$

Let inflation be defined as:

$$\frac{P_t}{P_{t-1}} \equiv \pi_t \quad (73)$$

Combining (64) and (68) implies that the Euler equation for consumption is given by:

$$-\tilde{v} \left[\delta_t \frac{b_t}{b_{L,t}} - 1 \right] \frac{\delta_t}{b_{L,t}} - \frac{1}{c_t^{1/\sigma}} + \beta R_t \mathbb{E}_t \pi_{t+1}^{-1} \frac{\phi_{t+1}/\phi_t}{c_{t+1}^{1/\sigma}} = 0 \quad (74)$$

A log-linear approximation of equation (74) implies:

$$\hat{c}_t = \mathbb{E}_t \hat{c}_{t+1} - \sigma \left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} + \mathbb{E}_t (\phi_{t+1} - \phi_t) \right] + \frac{\tilde{v} \delta_t c^{1/\sigma} \sigma}{b_L} \left[\hat{b}_t - \hat{b}_{L,t} \right] \quad (75)$$

Similarly, the labour supply condition (65) can be log-linearised to give:

$$\psi \hat{n}_t = \hat{w}_t - \sigma^{-1} \hat{c}_t \quad (76)$$

The first-order condition for the ambiguous asset (67) can be written as:

$$-\Lambda_t + \phi_t \tilde{v} \left[\delta_t \frac{b_t}{b_{L,t}} - 1 \right] \frac{\delta_t b_t}{b_{L,t}^2} + \beta \mathbb{E}_t \left[\frac{\Lambda_{t+1}}{\pi_{t+1}} R_{L,t+1} \right] = 0 \quad (77)$$

Log-linearising (77), noting that $\Lambda = 1/c^{1/\sigma}$, implies:

$$\hat{\Lambda}_t = \frac{\tilde{v}c^{1/\sigma}}{b_L} \left[\hat{b}_t - \hat{b}_{L,t} \right] + \mathbb{E}_t \left[\hat{\Lambda}_{t+1} - \hat{\pi}_{t+1} + \hat{R}_{L,t+1} \right] \quad (78)$$

The first-order condition for the safe asset can be written as:

$$-\phi_t \tilde{v} \left[\delta_t \frac{b_t}{b_{L,t}} - 1 \right] \frac{\delta_t}{b_{L,t}} - \Lambda_t + \beta R_t \mathbb{E}_t \left[\pi_{t+1}^{-1} \Lambda_{t+1} \right] = 0 \quad (79)$$

Log-linearising (79) implies:

$$\hat{\Lambda}_t = \mathbb{E}_t \left[\hat{\Lambda}_{t+1} - \mathbb{E}_t \hat{\pi}_{t+1} \right] + \hat{R}_t - \tilde{v} \frac{c^{1/\sigma} \delta_t}{b_L} \left[\hat{b}_t - \hat{b}_{L,t} \right] \quad (80)$$

Equating expressions (78) and (80) for $\hat{\Lambda}_t$ implies:

$$\mathbb{E}_t \hat{R}_{L,t+1} = \hat{R}_t - (1 + \delta_t) \tilde{v} \frac{c^{1/\sigma}}{b_L} \left[\hat{b}_t - \hat{b}_{L,t} \right] \quad (81)$$

A money demand relationship follows from equation (66), whereby:

$$\phi_t \chi_m^{-1} m_t^{-1/\sigma_m} - \Lambda_t + \beta \mathbb{E}_t \pi_{t+1}^{-1} \Lambda_{t+1} = 0 \quad (82)$$

Equation (82) implies that:

$$\hat{m}_t = \frac{\sigma_m}{\sigma} \hat{c}_t - \frac{\beta \sigma_m}{1 - \beta} \hat{R}_t + \frac{\beta \sigma_m}{1 - \beta} \tilde{v} \frac{c^{1/\sigma} \delta_t}{b_L} \left[\hat{b}_t - \hat{b}_{L,t} \right] \quad (83)$$

Note that:

$$R_{L,t} = \frac{1 + V_t}{V_{t-1}} \quad (84)$$

Given equation (84), log-linearising implies:

$$\beta^{-1} \left[\hat{R}_{L,t} + \hat{V}_{t-1} \right] = \hat{V}_t \quad (85)$$

and

$$\hat{R}_{L,t} = \beta \hat{V}_t - \hat{V}_{t-1} \quad (86)$$

D.2 The government

The consolidated government's budget constraint in real terms, as described in Section 4.2, is given by:

$$\frac{B_t - R_{t-1}B_{t-1}}{P_t} + \frac{B_{L,t}^g - R_{L,t}B_{L,t-1}^g}{P_t} + \frac{\Delta_t}{P_t} = \frac{T_t}{P_t} \quad (87)$$

where:

$$\frac{\Delta_t}{P_t} = \frac{M_t - M_{t-1}}{P_t} - \left(\frac{Q_t}{P_t} - \frac{R_{L,t}Q_{t-1}}{P_t} \right) \quad (88)$$

Asset purchases are characterised by:

$$Q_t = q_t B_{L,t}^g \quad (89)$$

Together, (87) and (89) imply that the government's consolidated budget constraint can be written as:

$$b_t + (1 - q_t)b_{L,t}^g + m_t = \pi_t^{-1} [m_{t-1} + R_{t-1}b_{t-1} + R_{L,t}(1 - q_{t-1})b_{L,t-1}^g] + \tau_t \quad (90)$$

where:

$$\tau_t \equiv \frac{T_t}{P_t} \quad (91)$$

The rule governing transfers to households is assumed to be:

$$\frac{\tau}{b} \hat{\tau}_t = -\zeta \hat{b}_{t-1} - \beta^{-1} \hat{R}_{t-1} \quad (92)$$

and the quantity of long-term bonds is held fixed in real terms, implying that:

$$b_{L,t}^g = \bar{b}_C V_t \quad (93)$$

Log-linearising (and linearising with respect to q) implies that:

$$\hat{b}_t + \frac{m}{b} \hat{m}_t + \left[\frac{m}{b} + R + \frac{R_L \bar{b}_L}{b} \right] \hat{\pi}_t - \frac{R_L \bar{b}_L}{b} \hat{R}_{L,t} - \frac{\bar{b}_L}{b} q_t = (R - \zeta) \hat{b}_{t-1} + \frac{m}{b} \hat{m}_{t-1} - R_L \frac{\bar{b}_L}{b} q_{t-1} \quad (94)$$

Since, in steady state, $R = R_L = \beta^{-1}$ and $\delta = \frac{\bar{b}_L}{b}$, equation (94) can be rewritten as:

$$\hat{b}_t + \frac{m}{b} \hat{m}_t - \delta q_t = - \left[\frac{m}{b} + \beta^{-1} (1 + \delta) \right] \hat{\pi}_t + \frac{m}{b} \hat{m}_{t-1} + (\beta^{-1} - \zeta) \hat{b}_{t-1} - \beta^{-1} \delta q_{t-1} \quad (95)$$

D.3 Firms

The first-order condition for a producer, as described in Section 4.3, resetting her price at time t is:

$$\mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k (\beta \alpha)^{k-t} \left((1 - \eta) \frac{1+s}{P_k} + \eta \frac{w_k}{P_{j,t} A} \right) \left(\frac{P_{j,t}}{P_t} \right)^{-\eta} c_k = 0 \quad (96)$$

Define the price set by firm j – relative to the aggregate price level – as:

$$p_{j,t} \equiv \frac{P_{j,t}}{P_t} \quad (97)$$

Define the relative inflation factor as:

$$\Pi_{t,k} \equiv \begin{cases} \frac{P_k}{P_t} = \Pi_k \times \Pi_{k+1} \times \dots \times \Pi_{t+1}, & \text{for } k \geq t+1 \\ 1, & \text{for } k=t \end{cases} \quad (98)$$

Then we can re-write equation (96) as:

$$\mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k (\beta \alpha)^{k-t} \left((1 - \eta) \frac{(1+s)p_{j,t}}{\Pi_{t,k}} + \eta \frac{w_k}{A} \right) \left(\frac{p_{j,t}}{\Pi_{t,k}} \right)^{-\eta} c_k = 0 \quad (99)$$

Since all firms are identical, all firms that can change prices at time t will choose the same price, denoted p_t^* . Therefore, (99) becomes:

$$\mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k (\beta \alpha)^{k-t} \left((1-\eta) \frac{(1+s)p_t^*}{\Pi_{t,k}} + \eta \frac{w_k}{A} \right) \left(\frac{p_t^*}{\Pi_{t,k}} \right)^{-\eta} c_k = 0 \quad (100)$$

The aggregate price is:

$$P_t = \left[\int_0^1 P_{j,t}^{1-\eta} dj \right]^{\frac{1}{1-\eta}} = \left[\sum_{k=0}^{\infty} (1-\alpha) \alpha^k (P_{t-k}^*)^{1-\eta} \right]^{\frac{1}{1-\eta}} \quad (101)$$

where the second equality is a result of grouping firms into cohorts (according to the time at which they last reset their price) and noting that the mass of firms that have not reset their price since date $t-k$ is given by $(1-\alpha)\alpha^k$.

From (101), the aggregate price level can be written as:

$$P_t = [\alpha P_{t-1}^{1-\eta} + (1-\alpha)(P_t^*)^{1-\eta}]^{\frac{1}{1-\eta}} \quad (102)$$

which implies:

$$1 = \alpha \left(\frac{1}{\pi_t} \right)^{1-\eta} + (1-\alpha)(p_t^*)^{1-\eta} \quad (103)$$

Log-linearising the pricing equation gives:

$$\mathbb{E}_t \sum_{k=0}^{\infty} (\beta \alpha)^{s-t} \left(\hat{p}_t^* - \hat{\Pi}_{t,k} - \hat{w}_k \right) = 0 \quad (104)$$

which, by using the law of iterated (conditional) expectations, can be rearranged to give:

$$\hat{p}_t^* = (1-\beta\alpha)\hat{w}_t + \beta\alpha\mathbb{E}_t\hat{\pi}_{t+1} + \beta\alpha\mathbb{E}_t\hat{p}_{t+1}^* \quad (105)$$

Linearising the expression for the aggregate price level yields:

$$\hat{p}_t^* = \frac{\alpha}{1-\alpha} \hat{\pi}_t \quad (106)$$

Using this information in the log-linearised pricing equation implies:

$$\hat{\pi}_t = \frac{(1 - \beta\alpha)(1 - \alpha)}{\alpha} \hat{w}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} \quad (107)$$

Combined with the aggregate labour supply equation and the condition for market clearing (as discussed in Appendix D.4), the Phillips Curve (107) can be rewritten as:

$$\hat{\pi}_t = \frac{(1 - \beta\alpha)(1 - \alpha)}{\alpha} \left(\psi + \frac{1}{\sigma} \right) + \beta \mathbb{E}_t \hat{\pi}_{t+1} \quad (108)$$

D.4 Market clearing

Goods market clearing requires that:

$$c_t = \mathcal{D}_t^{-1} y_t \quad (109)$$

where $\mathcal{D}_t$ is a measure of price dispersion, as defined in Appendix E.1.

As discussed in Section 4.5, the only shock hitting the model is the preference shock ϕ_t . In the absence of both price rigidity and imperfect asset substitutability, the preference shock has no impact on activity. The efficient level of output is therefore constant. To a first-order approximation, the ‘output gap’ is therefore identical to the deviation of output from its steady state, which implies:

$$\hat{c}_t = \hat{y}_t = \hat{x}_t \quad (110)$$

D.5 Solution

Collecting the first-order conditions derived above, we now summarise the model’s solution. Equations (111)-(117) below are simply a replica of equations (22) to (28) in Section 4.6 of the main text.

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left(\frac{1}{1 + \delta_t} \hat{R}_t + \frac{\delta_t}{1 + \delta_t} \hat{R}_{L,t}^e - \mathbb{E}_t \hat{\pi}_{t+1} - r_t^* \right) \quad (111)$$

$$\hat{R}_{L,t}^e = \hat{R}_t - (1 + \delta_t) \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} (\hat{b}_t - \hat{b}_{L,t}) \quad (112)$$

$$\hat{m}_t = \frac{\sigma_m}{\sigma} \hat{x}_t - \frac{\beta \sigma_m}{1 - \beta} \hat{R}_t + \frac{\beta \sigma_m}{1 - \beta} (\delta_t) \tilde{v} c^{1/\sigma} (\bar{b}_L)^{-1} (\hat{b}_t - \hat{b}_{L,t}) \quad (113)$$

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t \quad (114)$$

$$\hat{b}_t + \frac{m}{b} \hat{m}_t - \delta_t q_t = - \left[\frac{m}{b} + \beta^{-1} (1 + \delta_t) \right] \hat{\pi}_t + \frac{m}{b} \hat{m}_{t-1} + (\beta^{-1} - \zeta) \hat{b}_{t-1} - \beta^{-1} \delta_t q_{t-1} \quad (115)$$

$$-q_t + \hat{V}_t = \hat{b}_{L,t} \quad (116)$$

$$\hat{R}_{L,t}^e = \beta \mathbb{E}_t \hat{V}_{t+1} - \hat{V}_t \quad (117)$$

E

This appendix accompanies Section 4. It follows Harrison (2012).

E.1 Per-period utility

The household's per-period utility function is:

$$U_t = \phi_t \left(\frac{c_t^{1-\frac{1}{\sigma}}}{1 - \frac{1}{\sigma}} - \frac{n_t^{1+\psi}}{1 + \psi} + \frac{\chi_m^{-1}}{1 - \sigma_m^{-1}} \left(\frac{M_t}{P_t} \right)^{1-\frac{1}{\sigma_m}} - \frac{\tilde{v}}{2} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right]^2 \right) \quad (118)$$

Since the preference shock is exogenous to policy, and the model is calibrated to ensure that the quantity of money in circulation is negligible, we can set M_t to zero in (118) when constructing the policymaker's loss function. As such:

$$U_t \approx \frac{c_t^{1-\frac{1}{\sigma}}}{1 - \frac{1}{\sigma}} - \frac{n_t^{1+\psi}}{1 + \psi} - \frac{\tilde{v}}{2} \left[\delta_t \frac{B_t}{B_{L,t}} - 1 \right]^2 \quad (119)$$

A second-order approximation to (119) yields:

$$U_t \approx c^{1-\frac{1}{\sigma}} \left[\hat{c}_t - \frac{1}{2\sigma} \hat{c}_t^2 \right] - n^{1+\psi} \left[\hat{n}_t + \frac{\psi}{2} \hat{n}_t^2 \right] - \frac{\tilde{v}}{2} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \quad (120)$$

As explained in Appendix D.4, the market-clearing condition for goods is:

$$c_t = \mathcal{D}_t^{-1} y_t \quad (121)$$

$\mathcal{D}_t \equiv \int_0^1 (P_t(i)/P_t)^{-\eta} di$ is the price dispersion term associated with staggered price setting. It implies:

$$\hat{c}_t = \hat{y}_t - \hat{\mathcal{D}}_t \quad (122)$$

The household's utility function (120) can therefore be rewritten as:

$$U_t \approx c^{1-\frac{1}{\sigma}} \left[\hat{y}_t - \hat{\mathcal{D}}_t - \frac{1}{2\sigma} (\hat{y}_t - \hat{\mathcal{D}}_t)^2 \right] - n^{1+\psi} \left[\hat{n}_t + \frac{\psi}{2} \hat{n}_t^2 \right] - \frac{\tilde{v}}{2} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \quad (123)$$

Since the price dispersion term $\mathcal{D}_t$ is only a second-order term, and the production function implies $\hat{y}_t = \hat{n}_t$, (123) can be written as

$$U_t = \left(c^{1-\frac{1}{\sigma}} - n^{1+\psi} \right) \hat{y}_t - \frac{1}{2} \left(\frac{c^{1-\frac{1}{\sigma}}}{\sigma} + \psi n^{1+\psi} \right) \hat{y}_t^2 - c^{1-\frac{1}{\sigma}} \hat{\mathcal{D}}_t - \frac{\tilde{v}}{2} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \quad (124)$$

The steady-state relationship between labour supply and consumption, under the assumption that subsidies paid to firms are set to eliminate the distortion associated with monopolistic competition, can be written as:

$$n^\psi = wc^{-1/\sigma} = Ac^{-1/\sigma} \quad (125)$$

Steady-state market clearing, given that steady-state price dispersion $\mathcal{D} = 1$ is given by:

$$c = y = An \quad (126)$$

Together, (125) and (126) imply:

$$n^{1+\psi} = c^{1-1/\sigma} \quad (127)$$

which allows us to write the household's per-period utility function as:

$$U_t = -\frac{1}{2}c^{1-\frac{1}{\sigma}} \left(\frac{1}{\sigma} + \psi \right) - c^{1-\frac{1}{\sigma}} \hat{\mathcal{D}}_t - \frac{\tilde{v}}{2} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \quad (128)$$

Price dispersion in equilibrium is given by:

$$\mathcal{D}_t \equiv \int_0^1 \left(\frac{P_t(i)}{P_t} \right)^{-\eta} di = \alpha \mathcal{D}_{t-1} \pi_t^n + (1-\alpha) p_t^{*- \eta} \quad (129)$$

Substituting for the optimal price p_t^* , we can re-write (129) as:

$$\mathcal{D}_t = \alpha \mathcal{D}_{t-1} \pi_t^n + (1-\alpha) \left[\frac{1 - \alpha \pi_t^{\eta-1}}{1-\alpha} \right]^{\frac{\eta}{\eta-1}} \quad (130)$$

Taking a second-order Taylor expansion of (130) yields:

$$\begin{aligned} \mathcal{D}_t &\approx \alpha (\mathcal{D}_{t-1} + \eta \hat{\pi}_t) + (1-\alpha) \left[\frac{-\alpha \eta \hat{\pi}_t}{1-\alpha} \right] + \frac{\alpha \eta (\eta-1)}{2} \hat{\pi}_t^2 + \frac{1}{2} \left[\frac{\alpha^2}{1-\alpha} - \alpha \eta (\eta-2) \right] \hat{\pi}_t^2 \\ &\approx \alpha \hat{\mathcal{D}}_{t-1} + \frac{\alpha \eta}{2(1-\alpha)} \hat{\pi}_t^2 \end{aligned} \quad (131)$$

E.2 The present value of losses

The loss function to be minimised can be defined as the present discounted value of all future per-period losses:

$$\begin{aligned} \mathcal{L} &= -2 \sum_{t=0}^{\infty} \beta^t U_t \\ &= \sum_{t=0}^{\infty} \beta^t \left(c^{1-\frac{1}{\sigma}} \left(\frac{1}{\sigma} + \psi \right) \hat{y}_t^2 + 2c^{1-\frac{1}{\sigma}} \hat{\mathcal{D}}_t + \frac{\tilde{v}}{2} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \right) \end{aligned} \quad (132)$$

Noting that:

$$\begin{aligned}
\sum_{t=0}^{\infty} \beta^t \hat{\mathcal{D}}_t &= \alpha \sum_{t=0}^{\infty} \beta^t \hat{\mathcal{D}}_{t-1} + \sum_{t=0}^{\infty} \beta^t \frac{\alpha \eta}{2(1-\alpha)} \hat{\pi}_t^2 \\
&= \alpha \hat{\mathcal{D}}_{t-1} + \alpha \beta \sum_{t=0}^{\infty} \beta^{t-1} \hat{\mathcal{D}}_{t-1} + \sum_{t=0}^{\infty} \beta^t \frac{\alpha \eta}{2(1-\alpha)} \hat{\pi}_t^2 \\
&= \alpha \hat{\mathcal{D}}_{t-1} + \alpha \beta \sum_{t=0}^{\infty} \beta^t \hat{\mathcal{D}}_t + \sum_{t=0}^{\infty} \beta^t \frac{\alpha \eta}{2(1-\alpha)} \hat{\pi}_t^2
\end{aligned} \tag{133}$$

implies that:

$$\sum_{t=0}^{\infty} \beta^t \hat{\mathcal{D}}_t = \frac{\alpha}{1-\alpha\beta} \hat{\mathcal{D}}_{t-1} + \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \frac{\alpha \eta}{(1-\alpha\beta)(1-\alpha)} \hat{\pi}_t^2 \tag{134}$$

Substituting into (132) implies:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left(c^{1-\frac{1}{\sigma}} \left(\frac{1}{\sigma} + \psi \right) \hat{y}_t^2 + c^{1-1/\sigma} \frac{\alpha \eta}{(1-\alpha\beta)(1-\alpha)} \hat{\pi}_t^2 + \tilde{v} \left[\hat{b}_t - \hat{b}_{L,t} \right]^2 \right) + \frac{\alpha}{(1-\alpha\beta)(1-\beta)} \hat{\mathcal{D}}_{t-1} \tag{135}$$

where the final term in (135) is independent of policy and can therefore be ignored. Normalising such that the weight on output gap deviations is unity implies:

$$\mathcal{L} \propto \sum_{t=0}^{\infty} \beta^t \left[\hat{x}_t^2 + \frac{\eta}{\kappa} \hat{\pi}_t^2 + \frac{\tilde{v} c^{1/\sigma-1}}{\sigma^{-1} + \psi} \left(\hat{b}_t - \hat{b}_{L,t} \right)^2 \right] \tag{136}$$

which is the loss function characterising the monetary policymaker's objective in Section 4.7 of the main text.

References

- Andrés, J., López-Salido, J., and Nelson, E. (2004). Tobin’s imperfect asset substitution in optimizing general equilibrium. *Journal of Money, Credit and Banking*, 36(4):665–691.
- Benigno, P. and Woodford, M. (2006). Linear-quadratic approximation of optimal policy problems. *NBER Working Paper*, 12672.
- Bernanke, B. S. (2010). “Implications of the financial crisis for economics,” Remarks at The Bendheim Centre for Finance, Princeton University.
- Bernanke, B. S. (2014). Transcript of “Central banking after the Great Recession: Lessons learned and challenges ahead,” Remarks at The Brookings Institution.
- Blanchard, O. (2009). (Nearly) nothing to fear but fear itself. *The Economist*, 29 January.
- Caballero, R. and Farhi, E. (2014). The safety trap. *National Bureau of Economic Research Working Paper*, 19927.
- Caballero, R. and Krishnamurthy, A. (2008). Collective risk management in a flight to quality episode. *Journal of Finance*, 63(5):2195–2230.
- Calvo, G. (1983). Staggered prices in a utility-maximising framework. *Journal of Monetary Economics*, 12:383–398.
- Chadha, J., Turner, P., and Zampolli, F. (2013). The interest rate effects of government debt maturity. *BIS Working Paper*, 415.
- Collard, S., Sheppard, S. M. K., and Tallon, J. (2011). Ambiguity and the historical equity premium. *Oxford University Economics Dept. Working Paper*, 550.
- Dennis, R. (2007). Optimal policy in rational expectations models: new solution algorithms. *Macroeconomic Dynamics*, 11(1):31–55.
- Dennis, R. (2013). Imperfect credibility and robust monetary policy. *Scottish Institute for Research in Economics (SIRE) Discussion Papers*, 78.
- Ellison, M. and Tischbirek, A. (2013). Unconventional government debt purchases as a supplement to conventional monetary policy. *Economics Series Working Papers, University of Oxford*, 679.

- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.
- Gagnon, J., Raskin, M., Remache, J., and Sack, B. (2011). Large-scale asset purchases by the Federal Reserve: Did they work? *Economic Policy Review*, pages 41–59.
- Glick, R. and Leduc, S. (2013). The effects of unconventional and conventional U.S. monetary policy on the dollar. *Federal Reserve Bank of San Francisco Working Paper*, 2013-11.
- Harrison, R. (2012). Asset purchase policy at the effective lower bound for interest rates. *Bank of England Working Paper*, No. 444.
- Ju, N. and Miao, J. (2012). Ambiguity, learning and asset returns. *Econometrica*, 80(2):559–591.
- Jurado, K., Ludwigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73:1849–1892.
- Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Hart, Schaffner and Marx, Boston.
- Landau, J. P. (2014). Why is euro inflation so low? *VoxEU.org*, 2 December.
- Laubach, T. and Williams, J. C. (2016). Measuring the natural rate of interest. *Business Economics*, 51:257–267.
- Levin, A., Lopez-Salido, D., Nelson, E., and Yun, T. (2010). Limitations on the effectiveness of forward guidance at the zero lower bound. *International Journal of Central Banking*, 6(1):143–189.
- Maccheroni, F., Marinacci, M., and Ruffino, D. (2013). Alpha as ambiguity: Robust mean-variance portfolio analysis. *Econometrica*, 81(3):1075–1113.
- Tobin, J. (1956). Liquidity preference as behaviour towards risk. *Review of Economic Studies*, 25(2):65–86.
- Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of Money, Credit and Banking*, 1(1):15–29.

- Tobin, J. and Brainard, W. C. (1963). Financial intermediaries and the effectiveness of monetary controls. *American Economic Review*, 53(2):383–400.
- Vayanos, D. and Vila, J.-L. (2009). A preferred-habitat model of the term structure of interest rates. *London School of Economics, mimeo*.
- Woodford, M. (2013). “Forward guidance by inflation-targeting central banks,” Paper presented at Sveriges Riksbank.

III.

DEBT DURATION

AS A POLICY INSTRUMENT

Adrian Paul*

Abstract

It is the net composition of the government's *consolidated* balance sheet – which incorporates asset purchases by the monetary authority as well as debt issuance by the fiscal authority – that determines the duration of public debt held by the private sector, and therefore influences the real economy. Seen from this perspective, the transmission channel from ‘unconventional’ monetary policy to output and inflation is not especially unconventional. We estimate a structural vector autoregression with time-varying parameters based on a sample of U.S. data stretching back to 1975. We find that an exogenous increase in the average maturity of government debt held in private hands reduces inflation and increases unemployment. These effects are three or four times larger in the depths of the Great Recession than in the midst of the Great Moderation – a finding which underscores the importance of imperfect asset substitutability, but also suggests that the joint behaviour of the U.S. Treasury and the Federal Reserve in the aftermath of the global financial crisis of 2007/08 amounted to quantitative *tightening* rather than quantitative easing.

*St. Peter's College, University of Oxford. Michaelmas Term, 2017. This is the third chapter of my D.Phil. thesis: *Essays on Unconventional Monetary Policy*.

1 Introduction

“The president announced: ‘I’m scheduled to go into hospital tomorrow for a gall bladder operation. You wouldn’t raise the discount rate while I’m in the hospital, would you?’ Martin waited for a second, and said: ‘No, Mr. President, we’ll wait until you get out of the hospital’...

...When he heard about the rate increase, Lyndon Johnson was as angry as he could possibly be.”

An exchange between U.S. President Lyndon B. Johnson
and U.S. Federal Reserve Chairman William McChesney Martin,
October 1965.

Having cut short-term nominal interest rates to their effective lower bound in the wake of the global financial crisis, major central banks such as the U.S. Federal Reserve and the Bank of England turned to “unconventional” monetary policy measures to stabilise their economies after 2008. In large part, these unconventional measures involved large-scale central bank purchases of long-term government debt: an innovation which blurred the distinct lines of separation between fiscal policy and monetary policy that had underpinned central bank independence in the decades preceding the crisis.¹

Economic theory suggests that it is the net composition of the government’s *consolidated* balance sheet – which incorporates asset purchases by the monetary authority as well as debt issuance by the fiscal authority – that matters.² The ‘portfolio balance’ chan-

¹In a speech as Chairman of the Federal Reserve in September 2010, Ben Bernanke argued that “macroeconomic modelling must accommodate the possibility of unconventional monetary policies, a number of which have been used during the crisis. Earlier work on this topic relied primarily on the example of Japan; now, a number of data points can be used. For example, the experience of the United States and the United Kingdom with large-scale asset purchases could be explored to improve our understanding of the effects of such transactions on longer-term yields and how such effects can be incorporated into modern models of the term structure of interest rates.”

²Greenwood et al. (2015), for example, study optimal government debt maturity in a model in which investors derive monetary services from holding riskless short-term securities. In *The New Palgrave Dictionary of Economics*, Leeper and Nason (2008) define a related concept – the consolidated government budget constraint – in the following way: “The government budget constraint is an accounting identity linking the monetary authority’s choices of money growth or nominal interest rate and the fiscal authority’s choices of spending, taxation, and borrowing at a point in time and across time. The intertemporal links create a rich set of possible outcomes from standard macro policy experiments. Taking the government budget constraint seriously can overturn some widely held beliefs about policy effects.”

nel of monetary policy, for example, which originates in the work of Tobin (1956) and was adopted as the dominant rationale for quantitative easing (QE) throughout the crisis, implies that it is ultimately the composition of assets in the hands of private investors that determines the transmission of policy stimulus to real economic activity.³

All else equal, central bank purchases of long-term debt in the secondary market serve to absorb duration from the portfolios of private-sector bondholders. In theory, at least, this is stimulative because it pushes the private sector into other, riskier assets in search of yield. As demand is squeezed out of relatively safe government bonds into relatively risky assets such as corporate bonds and equities, the prices of those risky assets rise and their yields fall. This reduces the cost of finance in the real economy, and thereby raises output and inflation.

But the same change in the composition of private-sector portfolios can be brought about by changes in the pattern of asset supply (via government bond issuance) just as it can by changes in the pattern of asset demand (via central bank purchases). Seen from this perspective, quantitative easing is observationally equivalent to government debt management. Despite this equivalence, academic studies of the macroeconomic effects of post-crisis monetary policies designed to remove duration risk have typically ignored the potentially confounding effects of long-standing fiscal policies designed to inject duration risk. If the relevant policy instrument is, in fact, the net duration risk borne by the private sector, the actions taken by one branch of government (the central bank) cannot be accurately assessed without acknowledging the decisions taken by the other branch of government (the treasury).

More fundamentally, the primacy of the government’s *consolidated* balance sheet suggests there is little unconventional about the “unconventional” monetary policy pursued after 2008. As Turner (2014) points out, government debt management was a key component of monetary policy in the decades that followed the Second World War.⁴

³The portfolio balance channel was not only the rationale for QE in the United States. When the European Central Bank embarked on public-sector asset purchases in March 2015, Executive Board Member Benoît Cœuré asserted that “the proper working of the portfolio rebalancing channel...is at the heart of our asset purchase programmes.”

⁴Indeed, Wallace (1981) begins the exposition of his seminal “irrelevance theorem” for central bank market operations by stating: “Monetary policy determines the composition of the government portfolio.” In 1930, Keynes argued strongly that the British government was mistaken in thinking that its “sound money” policy of lengthening the maturity of gilts would not offset the decline in short-term interest

In this paper, we assess the U.S. experience over a forty-year period which spans both pre-crisis and post-crisis policy regimes. We estimate a structural vector autoregression with parameters that are allowed to vary over time. We aim to identify the ‘portfolio balance’ channel associated with duration absorption, independently of the ‘signalling’ channel associated with forward guidance. The former captures the macroeconomic effects of shifting the risks borne by the private sector; the latter captures the effects of shifting the private sector’s beliefs about the future path of policy. Both channels have been credited with averting a second Great Depression in the aftermath of Lehman Brothers’ collapse.

This paper proceeds as follows. Section 2 surveys three strands of related literature. Section 3 presents post-crisis quantitative easing as a case study. Section 4 explains the formal econometric methodology that we adopt in Section 5. Section 6 presents the key results. Section 7 concludes.

2 Literature

This paper is related to three strands of literature.

The first strand of literature documents the extent to which asset purchases by major central banks in the wake of the global financial crisis of 2007/08 were unwound by the debt issuance undertaken by their counterparts in the treasury.

Greenwood et al. (2014) argue that, throughout the aftermath of the global financial crisis, monetary and fiscal policies in the United States were “pushing in opposite directions.” The U.S. Federal Reserve’s attempts – via successive rounds of quantitative easing – to reduce the supply of long-term government debt held by private investors were offset by the U.S. Treasury’s attempts to increase the maturity of interest-bearing debt issued by the government.

Greenwood et al. (2014) find that QE by the Federal Reserve between 2008 and 2014 reduced the supply of ten-year duration-equivalent debt by 15½ percentage points of GDP. Over the same period, active maturity extension by the Treasury increased the supply of ten-year equivalent debt by 5½ percentage points of GDP. In effect, a fiscal

rates that motivated Britain’s abandonment of the Gold Standard.

policy of “reverse quantitative easing” served to undo 35% of the duration absorption implied by post-crisis monetary policy. Combined with the price impact of asset purchase announcements, as estimated by Williams (2014), the authors conclude that the Federal Reserve’s actions lowered the term premium on ten-year U.S. government debt by 137bps while the Treasury’s actions simultaneously raised the term premium by 48bps.

If the Treasury’s policy of active maturity extension contributed to a significant increase in the supply of duration for a given stock of debt, Greenwood et al. (2014) argue that the surge in the overall stock of U.S. debt between 2008 and 2014 boosted the supply of duration yet further. The positive debt demand shock implied by quantitative easing was dwarfed by the positive debt supply shock implied by fiscal expansion.

Yet this sharp increase in the net supply of government debt is difficult to reconcile with the significant declines in interest rates documented around each QE announcement. Greenwood et al. (2014) conjecture that the prospect of Federal Reserve asset purchases had both a direct, mechanical effect on risk premia as well as an indirect, expectational effect on risk-neutral yields. Both the announcement of purchases themselves, as well as the language used to describe the outlook which necessitated them, conveyed enough additional information about the likely evolution of future monetary policy (both conventional and unconventional) to warrant a reduction in the required rate of return on long-dated government debt even in the face of a surge in its net supply.

In an analysis of the first nine months of the European Central Bank’s asset purchase programme, Andrade et al. (2016) find only “limited evidence” that greater issuance of long-term bonds by national Euro area governments offset the ECB’s attempts to withdraw duration from the portfolios of private investors. The authors find that, between March and December 2015, government net issuance increased the supply of ten-year equivalent debt by 1.9% of Euro area GDP, while ECB asset purchases reduced the supply of ten-year equivalent debt by 4.5% of Euro area GDP. At least in its earliest phases, Andrade et al. (2016) argue that the ECB’s asset purchase programme was “able to remove duration risk from the economy, despite some increase in the supply of such risk induced by national debt management policies.”

Haldane et al. (2016) argue that, in the United Kingdom, the maturity of assets purchased by the Bank of England closely tracked the maturity of outstanding government debt. This observation supports the view that the monetary and fiscal authorities in the United Kingdom did not work across purposes to the same extent as they did in the

United States. Indeed, on the eve of QE in the U.K. in early 2009, the Governor of the Bank of England wrote to the Chancellor of the Exchequer asserting that “if the [Asset Purchase] facility were to be used to purchase gilts, it would be important that the Government’s debt management policy remain consistent with the aims of monetary policy.” The Governor went on to emphasise that the government “should not alter its [gilt] issuance strategy as a result of the transactions undertaken through the Asset Purchase Facility for monetary policy purposes.” The Chancellor concurred, restating the U.K. Debt Management Office’s objective “to minimise, over the long-term, the costs of meeting the Government’s financing needs, taking into account risk, whilst ensuring that debt management is consistent with the aims of monetary policy.” No such co-ordination seems to have taken place in the United States.

The second strand of literature investigates the macroeconomic consequences of duration absorption. One set of empirical papers focuses solely on large-scale asset purchases by major central banks after the global financial crisis began. Another set of empirical papers exploits the evolution of official-sector balance sheets over several decades before 2008.

Adopting the first approach, Andrade et al. (2016) find evidence that the removal of duration risk was a key channel through which ECB QE stimulated the real economy. The authors find that the decline in yields on announcement of the ECB’s asset purchase programme was larger the longer the maturity of the bonds targeted. The authors’ calibrated macroeconomic model implies that an increase in the average maturity of ECB public-sector purchases from eight to eleven years would have lead to a 50bp increase in peak Euro area inflation.

Weale and Wieladek (2016) compare the impact of asset purchase announcements on output and inflation in the U.S and the U.K. by estimating a vector autoregression on monthly data from 2009 to 2014. The authors find that a central bank announcement to purchase government bonds in the order of 1% of GDP prompted a 0.6% increase in real U.S GDP (0.25% in the U.K.) and a 0.6% increase in U.S. consumer prices (0.3% in the U.K.). Pascual and Wieladek (2016) conduct a similar exercise for the Euro area, concluding based on monthly data from 2012 to 2016 that, in the absence of the first round of ECB QE, real Euro area GDP and core consumer prices would have been 1.3% and 0.9% lower, respectively. Lloyd (2016) attempts to separately identify the ‘signalling’ channel and the ‘portfolio balance’ channel of U.S. quantitative easing by estimating a

vector autoregression on monthly data from 2008 to 2013. Lloyd (2016) argues that signalling effects were particularly powerful at one and two-year horizons. Moreover, these signalling effects generated three or four times as much stimulus as the effects attributable to portfolio rebalancing.

Adopting the second approach, Chadha et al. (2013) estimate a single-equation OLS regression based on a pre-crisis sample starting in 1976. The authors find that reducing the average maturity of U.S. Treasury debt held outside the Federal Reserve by one year typically reduced the five-year forward yield on ten-year U.S. government debt by between 130 and 150 basis points. Chadha et al. (2013) conclude that “these large estimates are consistent with the existence of significant portfolio balance effects” – even in normal times.

Baumeister and Benati (2013) estimate a structural vector autoregression with time-varying parameters to identify the macroeconomic effects of a ‘pure spread’ shock. A pure spread shock, which compresses long-term bond yields but leaves short-term interest rates unchanged, is designed to replicate central bank asset purchases in an environment in which the policy rate is stuck at its effective lower bound. Conditional on existing estimates of the impact of QE on long-term bond yields in the U.S. and the U.K. (estimates which are derived from single-equation regressions and/or financial market event studies, and which the authors simply take as given), Baumeister and Benati (2013) map their estimates of the growth and inflation impact of yield curve compression since the mid-1970s into estimates of the growth and inflation impact of large-scale asset purchases after 2008. They conclude that the unconventional monetary policy pursued by the Federal Reserve and the Bank of England mitigated the risk of a pronounced deflation and averted a collapse in output comparable to that which took place during the Great Depression.

Blattner and Joyce (2016) consider the average maturity of government debt outstanding in the Euro area’s four largest economies between 2000 and 2013 – adjusted for official foreign holdings. Purchases by foreign central banks pre-crisis serve as a proxy for purchases by the ECB after 2015. First, the authors estimate the latent factors that explain the shape of a synthetic Euro area yield curve. Second, the authors estimate a time-invariant vector autoregression which relates those latent factors to a set of macroeconomic and financial variables, including a measure of net debt supply. Blattner and Joyce (2016) argue that the ECB’s asset purchase programme constituted a shock to the net supply of government bonds which reduced Euro area ten-year bond yields by around

30bps in 2015. By implication, the authors argue that this removal of duration risk raised output above potential and increased inflation by 0.2 percentage points and 0.3 percentage points, respectively, in 2016.

Appealing to a longer history, Kuttner (2006) estimates a single-equation OLS regression using data from 1964 to 2002. He finds that, in aggregate, the composition of privately held debt had only small effects on the slope of the yield curve over that period, but changes in the long-term securities held specifically by the Federal Reserve did tend to have statistically significant and economically meaningful effects on term premia. The author contends that this result may owe to changes in the central bank’s portfolio being “more exogenous” with respect to term premia, especially given that changes in the overall volume of privately held government debt include the (potentially endogenous) issuance of new Treasury securities.

Despite relying on a handful of discrete events, Swanson (2011) also adopts a historical approach, drawing a parallel between the second round of asset purchases conducted by the Federal Reserve in late 2010 (“QE2”) and a similar programme undertaken by the Kennedy administration and the Federal Reserve in 1961 (“Operation Twist”). The author identifies six specific announcements over the course of Operation Twist that are likely to have had a major effect on financial markets. He argues that the cumulative effect of these six announcements was highly statistically significant but of little economic significance, amounting to a reduction in long-term Treasury yields of only 15 basis points. The effects of Operation Twist seem to have diminished yet further once the U.S. authorities moved away from Treasury securities towards private-sector credit instruments.

The third strand of literature develops techniques for estimating time-varying structural vector autoregressions. In a seminal methodological contribution, Primiceri (2005) introduces two sources of variation into a small-scale multiple equation model of the economy. Within such a framework, the coefficients on all endogenous variables and the variance-covariance matrix of additive shocks to those endogenous variables are both allowed to vary over time. Primiceri (2005) and Del Negro and Primiceri (2015) develop a multivariate, stochastic volatility strategy to model the law of motion of the variance-covariance matrix of innovations, proposing an efficient Markov chain Monte Carlo algorithm to numerically evaluate the posterior probability distributions of the model’s key parameters.

As opposed to Cogley (2003) and Cogley and Sargent (2005) (in which the variance

of endogenous variables changes over time, but the simultaneous relationships among them are assumed to be time-invariant), and as opposed to Boivin (1999) (in which the simultaneous relationships among endogenous variables are time-varying, but the variance of their innovations is constant), Primiceri (2005) distinguishes between changes in the typical size of the model’s exogenous shocks and changes in the typical response of the model’s endogenous variables.

In this paper, we combine the three strands of literature described above. In particular, the few empirical studies that do exploit pre-crisis variation in the government’s consolidated balance sheet typically estimate its effect on long-dated government bond yields, or the term premia embedded therein. Only via the imputed, off-model effects of such yield curve compression do these studies have anything to say about the indirect impact of duration absorption on the wider economy. We resist this reliance on rules of thumb. By contrast, in an econometric framework which allows for simultaneous relationships between the maturity composition of private sector portfolios and the forward guidance implicit in central bank communication – and which allows the effects of shocks to the government’s consolidated balance sheet to vary across policy regimes – we directly estimate the effect of duration absorption on the joint evolution of the macroeconomic variables uppermost in the minds of monetary and fiscal policymakers alike: output and inflation.

3 Post-crisis QE: A case study

Documenting the tension between monetary and fiscal policy during the decade that followed the 2007/08 financial crisis is an instructive starting point for our analysis. In subsequent sections, we consider a longer history of macroeconomic performance, within the context of a more formal methodological framework.

In the United States, the Treasury publishes quarterly data on the average maturity of total outstanding marketable debt. The Federal Reserve publishes weekly data on the factors affecting reserve balances, including the nominal value of government securities – by maturity category – held by the central bank outright.⁵ Finally, the Treasury publishes

⁵The maturity distribution of the Federal Reserve’s SOMA portfolio is reported in Table 1.19 of the Statistical Supplement to the *Federal Reserve Bulletin*.

monthly data on the average residual maturity of U.S. Federal debt held in private hands.⁶

In the Euro area and the United Kingdom, however, neither the monetary authority nor the fiscal authority publishes analogous data on the composition of debt held outside the official sector. We therefore construct our own estimates of the average maturity of government debt held in private hands by consolidating the balance sheets of the government and of the central bank.⁷

The ECB publishes monthly data (for each of the nineteen Euro area member states, and for the Euro area in aggregate) on the nominal value of outstanding government debt, and its average residual maturity. In January 2015, the ECB also began publishing monthly data on the nominal value of ECB holdings of government debt under the Public Sector Purchase Programme, as well as their average residual maturity.

It is straightforward to find comparable monthly data on the nominal value of outstanding government debt in the U.K., and its average residual maturity. In March 2009, the Bank of England began publishing monthly data on the nominal value of government debt held under the Asset Purchase Facility. But the Bank of England does not publish a precise measure of the average residual maturity of those holdings, except for having restated in market notices published alongside each round of QE that “the Bank intends to purchase evenly across the three gilt maturity sectors.” These three sectors are defined as: gilts with a residual maturity of 3 to 7 years, gilts with a residual maturity of 7 to 15 years, and gilts with a residual maturity of more than 15 years. We construct our own database of the purchase date, the redemption date, the nominal value and the precise residual maturity of each U.K. government security bought by the Bank of England (by ISIN identifier) since asset purchases began in March 2009.

3.1 Public duration in private hands

With these data in hand, we calculate the average maturity of government debt held in private hands as the residual that remains after netting out the weighted average residual

⁶The Treasury data on the maturity distribution of privately held public debt are reported in Table FD-5 of the *Treasury Bulletin*, and available post-1975 through Haver Analytics. Data pre-1975 are reported in the Federal Reserve’s *Annual Statistical Digest*, and *Banking and Monetary Statistics*.

⁷We refer to a single, hypothetical Euro area fiscal authority as “the government.” In reality, of course, the Euro area ‘government’ is made up of 19 individual Euro area governments, each with its own balance sheet, its own financing needs, and its own funding costs.

maturity of central bank purchases of government debt from the weighted average residual maturity of the total stock of debt issued by the government.

Let l_{outst} be the average maturity of government debt outstanding. Let D_{outst} be the nominal value of government debt outstanding. Assuming that all outstanding debt is held either by the central bank ('CB') or by the private sector ('private'), we have:

$$l_{outst} \cdot D_{outst} = l_{CB} \cdot D_{CB} + l_{private} \cdot D_{private} \quad (1)$$

and

$$D_{outst} \equiv D_{CB} + D_{private} \quad (2)$$

Together, equations (1) and (2) imply that the average maturity of public debt in the hands of the private sector is given by:

$$l_{private} = \frac{l_{outst} \cdot D_{outst} - l_{CB} \cdot D_{CB}}{D_{private}} = \frac{l_{outst} \cdot D_{outst} - l_{CB} \cdot D_{CB}}{D_{outst} - D_{CB}} \quad (3)$$

For a given nominal value and a given maturity of debt outstanding (that is, for given D_{outst} and l_{outst}), withdrawal of duration by the central bank (that is, an increase in l_{CB}) implies a reduction in the duration held by the private sector (that is, a decrease in $l_{private}$). To the extent that quantitative easing affects the real economy via the portfolio balance channel, this absorption of duration by the central bank is stimulative because it pushes the private sector into other, riskier assets in search of yield. As demand is squeezed out of relatively safe government bonds into relatively risky assets such as corporate bonds and equities, the prices of those risky assets rise and their yields fall. This reduces the cost of finance facing firms and households, and thereby raises output and inflation.

3.2 Quantitative tightening

Figures 1-3 plot the evolution of l_{outst} , l_{CB} and $l_{private}$ in the United States, the United Kingdom and the Euro area since quantitative easing began in each jurisdiction. Our estimates suggest that, to varying degrees, each fiscal authority's strategy of debt issuance and each monetary authority's strategy of duration absorption ultimately pushed in opposite directions. By the end of the post-crisis phase of "unconventional" macroeconomic policy, the average residual maturity of public debt in private hands was *higher* – not

lower – in each major economy than it had been at the outset.

This is not to say that large-scale asset purchases by the central bank were always ineffective. In the United States between late 2011 and late 2012, for example, even while the U.S. Treasury actively extended the average maturity of government debt from 62 months to 65 months, asset purchases by the U.S. Federal Reserve (which raised the maturity of the central bank’s bond portfolio from 61 to 85 months) served to reduce the average maturity of public debt in private hands from 60 to 54 months. Similarly, in the United Kingdom between mid-2009 and mid-2010, the average maturity of outstanding U.K. government debt was broadly unchanged, but the Bank of England’s asset purchases (which raised the maturity of the central bank’s gilt portfolio from 133 months to 167 months) reduced the average maturity of public debt held in private hands from 178 months to 166 months.

Yet these phases of genuine, net maturity withdrawal were invariably short-lived. Once one accounts for both the securities purchased by the central bank and the securities issued by the treasury, the decade after the crisis was one of quantitative *tightening* rather than quantitative easing. The average residual maturity of private-sector portfolios in the U.S. was 18 months higher in early 2017 than it had been in late 2008; in the U.K., it was 24 months higher in early 2017 than it had been in early 2009; in the Euro area, it was 3 months higher in early 2017 than it had been in early 2015.

In the remainder of this paper, we investigate the macroeconomic consequences of a given degree of duration absorption in the context of a more formal econometric model. We describe our empirical methodology in detail in Section 4. We apply it to forty years of U.S. data in Section 5.

4 Empirical methodology

4.1 Estimation

We identify the real effects of duration absorption by estimating a time-varying structural vector autoregression, following the methodology developed by Primiceri (2005). Such an approach admits time variation in both (i) the coefficients capturing the joint dynamics of the model’s endogenous variables, and (ii) the variance-covariance matrix capturing heteroskedasticity in the model’s exogenous innovations.

Consider the following generic model:

$$\begin{aligned} y_t &= c_t + B_{1,t}y_{t-1} + \dots + B_{k,t}y_{t-k} + u_t & t = 1, \dots, T \\ \text{Var}(u_t) &= \Omega_t. \end{aligned} \tag{4}$$

At each point in time, y_t is a $n \times 1$ vector of observed endogenous variables. c_t is an $n \times 1$ vector of time-varying intercepts. Each $B_{i,t}$, for $i = 1, \dots, k$, is an $n \times n$ matrix of time-varying coefficients on the i^{th} lag of the vector of endogenous variables y . Finally, u_t captures the model's additive innovations. These innovations represent heteroskedastic shocks which are unobservable. The variance-covariance matrix of the vector of shocks u_t is denoted Ω_t .

Suppose A_t is a lower triangular matrix capturing the contemporaneous relationships between endogenous variables:

$$A_t = \begin{bmatrix} 1 & 0 & \dots & 0 \\ \alpha_{2,1,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \alpha_{n,1,t} & \dots & \alpha_{n,n-1,t} & 1 \end{bmatrix}$$

Suppose Σ_t is a diagonal matrix capturing the stochastic volatility of structural shocks:

$$\Sigma_t = \begin{bmatrix} \sigma_{1,t} & 0 & \dots & 0 \\ 0 & \sigma_{2,t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \sigma_{n,t} \end{bmatrix}$$

Consider the triangular reduction of Ω_t given by:

$$A_t \Omega_t A_t' = \Sigma_t \Sigma_t' \tag{5}$$

Equation (5) implies that (4) can be re-written as:

$$\begin{aligned} y_t &= c_t + B_{1,t}y_{t-1} + \dots + B_{k,t}y_{t-k} + A_t^{-1} \Sigma_t \varepsilon_t, \\ \text{Var}(\varepsilon_t) &= \mathbb{I}_n. \end{aligned} \tag{6}$$

Define B_t as a matrix stacking all coefficient vectors on the right-hand side of (6). Now (6) can be re-written as:

$$\begin{aligned} y_t &= X_t' B_t + A_t^{-1} \Sigma_t \varepsilon_t, \\ X_t' &= \mathbb{I}_n \otimes [1, y_{t-1}', \dots, y_{t-k}'] \end{aligned} \quad (7)$$

The fact that A_t varies over time in (7) allows each additive innovation to each endogenous variable in the model to have a time-varying impact on every other endogenous variable.

We assume the model's time-varying parameters follow random walks. Define α_t as a vector stacking all elements of A_t that are neither zero nor unity. Define σ_t as a vector of all diagonal elements of Σ_t . The dynamics of the model's parameters are then given by:

$$\begin{aligned} B_t &= B_{t-1} + \nu_t \\ \alpha_t &= \alpha_{t-1} + \zeta_t \\ \log \sigma_t &= \log \sigma_{t-1} + \eta_t \end{aligned} \quad (8)$$

We assume all innovations in (7) and (8) follow a joint normal distribution. Their variance-covariance matrix is given by:

$$\mathbb{V} = Var \left(\begin{bmatrix} \varepsilon_t \\ \nu_t \\ \zeta_t \\ \eta_t \end{bmatrix} \right) = \begin{bmatrix} \mathbb{I}_n & 0 & 0 & 0 \\ 0 & \mathbb{Q} & 0 & 0 \\ 0 & 0 & \mathbb{S} & 0 \\ 0 & 0 & 0 & \mathbb{W} \end{bmatrix} \quad (9)$$

The specification of the variance-covariance matrix of shocks in (9) embodies some important assumptions. In particular, we assume that $Var(\varepsilon_t)$ is an $n \times n$ identity matrix. We assume that $\mathbb{Q}$, $\mathbb{S}$ and $\mathbb{W}$ are all positive definite. Further, we assume that $\mathbb{S}$ is block diagonal: the parameters capturing contemporaneous relationships within one block of endogenous variables evolve independently of the parameters within another block.

Following Primiceri (2005), we use a Bayesian approach to evaluate the posterior probability distributions of interest. Bayesian methods allow us to combine *a priori* information about key aspects of the model (the prior distribution) with *ex post* realisations of the data (the likelihood function) to obtain more informed estimates of the model's defining parameters (the posterior distribution). In particular, we focus on (i) the historical

evolution of the unobservable states B^T , A^T and Σ^T (where T denotes the sequence of all realisations from $t = 1$ to T), and (ii) the hyperparameters characterising the variance-covariance matrix $\mathbb{V}$. Rather than drawing from the high-dimensional *joint* posterior distribution of the entire set of model parameters, we draw from the low-dimensional *conditional* posterior distribution of each parameter by means of a Gibbs sampler – one of several standard methods for Markov chain Monte Carlo (MCMC) simulation.

For convenience, we assume that the initial state vectors characterising the time-varying lag coefficients, the contemporaneous relationships, the log volatilities and the hyperparameters are mutually independent. The priors for the hyperparameters ($\mathbb{Q}$, $\mathbb{W}$ and the blocks of $\mathbb{S}$) are each assumed to follow an inverse-Wishart distribution. The priors for the initial states of the time-varying lag coefficients, the contemporaneous relationships and the log standard errors (B_0 , α_0 and $\log \sigma_0$) are each assumed to follow a Gaussian distribution. Together with the state equations in (8), these assumptions imply Gaussian priors on the entire sequences of B_t , α_t and $\log \sigma_t$ from $t = 1$ to T — conditional on $\mathbb{Q}$, $\mathbb{W}$ and $\mathbb{S}$. We provide more detail on the priors adopted in our specific empirical application in Section 5.2.

A Gibbs sampler allows us to simulate the distribution of the model’s parameters conditional on the observed data. We use the algorithm developed by Shephard (1994) and Del Negro and Primiceri (2015). Conditional on A^T and Σ^T , the state-space representation of B^T in (8) is linear and Gaussian. The conditional posterior distribution of B^T is therefore a product of Gaussian densities, and B^T can thus be drawn from the standard simulation smoother developed by Carter and Kohn (1994). Similarly, conditional on B^T and Σ^T , the state-space representation of A^T in (8) is linear and Gaussian. The conditional posterior distribution of A^T is therefore also a product of Gaussian densities, and can be drawn in the same way as in the case of B^T . By contrast, the state-space representation of Σ^T in (8) necessitates a more involved simulation smoother – one which first transforms the non-linear, non-Gaussian state-space form of Σ^T into a linear counterpart which is approximately Gaussian. Finally, since the conditional posterior distribution of $\mathbb{V}$ is the product of inverse-Wishart distributions, as evident from (9), simulating the conditional posterior distribution of the model’s hyperparameters is straightforward. We provide more detail on the MCMC algorithm used in our specific empirical application in Appendix A and Appendix B.

4.2 Identification

The methodology described in Section 4.1 allows us to estimate the time-varying posterior probability distributions of the parameters B_t and Ω_t . These parameters characterise the ‘reduced form’ vector autoregression specified in (4). For ease of exposition, we re-write (4) below as:

$$\begin{aligned} y_t &= X_t' B_t + u_t \\ \text{Var}(u_t) &= \Omega_t. \end{aligned} \tag{10}$$

Now consider the following ‘structural’ vector autoregression:

$$y_t = X_t' B_t + \Xi_t \varepsilon_t \tag{11}$$

Given the posterior distribution of Ω_t , we can recover the posterior distribution of Ξ_t at each point in time by solving – for each draw of Ω_t – the system of equations given by:

$$\Xi_t \Xi_t' = \Omega_t \quad t = 1, \dots, T \tag{12}$$

Solving (12) is useful because it allows us to recover structural shocks from estimated reduced-form errors. But solving (12) requires assumptions to achieve identification.

Throughout this paper, we assume that Ξ_t is lower triangular, implying (just as in (5)) that the solution to (12) is given by:

$$\Xi_t = A_t^{-1} \Sigma_t \tag{13}$$

This identifying assumption implies that each endogenous variable in the model can respond contemporaneously to those variables ordered *above* it in y , but can only respond with a lag to those variables ordered *below* it. We provide more detail on the causal ordering assumed in our specific empirical application in Section 5.3.

As discussed in Primiceri (2005), the fact that the elements of Ξ_t vary over time is the key distinction between a time-varying structural VAR and a time-invariant structural VAR. Consider a time-invariant VAR with three variables, for example, in which the vector y consists of y_1 , y_2 and y_3 (in that order). There is only one source of shocks to the third equation: the third element ε_3 of the vector of additive innovations ε . In a time-varying VAR with three variables, however, (9) implies there are at least three

independent sources of shocks to the third equation: (i) the usual, additive innovation ε_3 (just as in the time-invariant case), (ii) an innovation to the contemporaneous response of y_3 to y_1 and y_2 , denoted ζ , and (iii) an innovation to the lagged response of y_3 to y_1 and y_2 , denoted ν .

Shocks of the first type are assumed to be mutually independent. Shocks of the second type are assumed to be correlated within equation, but uncorrelated between equations (since $\mathbb{S}$ is block diagonal). Shocks of the third type are potentially correlated both within and between equations. The virtue of a time-varying structural VAR is therefore its ability to accommodate innovations which cannot be neatly orthogonalised within a multivariate system – an attribute which is particularly useful for modelling the impact of policy shocks on macroeconomic variables.

5 Duration absorption: An application

We apply the techniques developed by Primiceri (2005) to estimate a structural vector autoregression with time-varying parameters. We focus on the macroeconomic performance of the United States, before and after the global financial crisis of 2007/08. We aim to identify the real effects of the ‘portfolio balance’ channel associated with duration absorption, as distinct from the real effects of the ‘signalling’ channel associated with forward guidance.

5.1 Data

Our quarterly U.S. dataset covers the period from Q1:1975 to Q2:2017. Our model contains five endogenous variables. Based on *ex ante* lag length tests such as the Hannan-Quinn information criterion and the Bayesian information criterion, our model incorporates two lags. In terms of the notation used in Section 4, we set $T = 169$, $n = 5$ and $k = 2$. Although some analyses of the effects of macroeconomic policy consider vector autoregressions which include many more variables,⁸ the time-varying nature of our specification imposes computational constraints on the viable size of y_t . In particular, the number of parameters to be estimated increases exponentially as each additional endogenous variable is included in the model.

⁸See, for example, Christiano et al. (1999), Bernanke et al. (2005), and Lenza et al. (2010).

5.1.1 The economy

The non-policy block of our model consists of two variables: activity and inflation.

To capture economic activity, we take the quarterly average of monthly, seasonally adjusted data on the unemployment rate among all U.S. civilians aged 16 and over. To capture inflation, we calculate the annual growth rate of quarterly, seasonally adjusted data on the implicit U.S. GDP deflator (a chain-type index of the aggregate price level). Maximising employment and stabilising prices are the two key pillars of the ‘dual mandate’ bestowed on the Federal Reserve by the U.S. Congress.

5.1.2 Policy

The policy block of our model consists of three variables: one measure of “conventional” policy and two measures of “unconventional” policy.

To capture conventional policy, we take the quarterly average of the effective federal funds rate. The federal funds rate is a benchmark measure of the interest rate charged on overnight loans between banks in the United States. It is the primary instrument under the control of the Federal Open Market Committee. To capture unconventional policy, we distinguish between the signalling effects associated with forward guidance and the portfolio balance effects associated with duration absorption. A prominent example of duration absorption is the purchase of long-dated securities by the Federal Reserve in the secondary market, financed either by the issuance of reserves or the sale of short-dated holdings. A prominent example of forward guidance is the language used by Federal Reserve officials to communicate with participants in financial markets, in either time-contingent or state-contingent terms. Neither central bank asset purchases (“open market operations”) nor central bank forward guidance (“open mouth operations”) are unique to the post-crisis period, but both aspects of policy have become the subject of ever more scrutiny since the collapse of Lehman Brothers.

The zero lower bound

Having cut the effective federal funds rate to zero in December 2008, the Federal Reserve not only announced changes to the conduct *current* monetary policy, it also offered more explicit guidance on the likely evolution of *future* monetary policy. In March 2009, for example, the Federal Open Market Committee not only announced that it would purchase \$300bn of long-term Treasury securities (“QE1”), it also indicated that it

expected the federal funds rate to remain between 0% and 0.25% for “an extended period” of time. In November 2010, the FOMC announced that it would purchase an additional \$600bn of long-term Treasuries (“QE2”). In August 2011, the Committee indicated that it expected to keep the federal funds rate unchanged “at least through mid-2013.” In September 2011, the FOMC announced that it would sell \$400bn of short-term Treasuries to buy an additional \$400bn of long-term Treasuries (“Operation Twist”). In January 2012, the Committee indicated that it expected to keep the federal funds rate unchanged “at least through late 2014.” In December 2012, the FOMC not only announced that it would purchase long-term Treasuries indefinitely – at a rate of \$45bn per month (“QE3”) – it also indicated that it now expected to keep the federal funds rate unchanged for at least as long as the unemployment rate remained above 6.5% and inflation expectations remained subdued.

Throughout this post-crisis period, empirical estimates suggest that monetary policy was incrementally loosened despite the federal funds rate being stuck at its effective lower bound. According to the ‘shadow’ policy rate estimated by Wu and Xia (2016), for example, model-based short-term yields fell by the equivalent of two percentage points between January 2009 and December 2012, from 0.61% to -1.43%.

At the end of each subsequent year, however, the Federal Reserve indicated that it would incrementally tighten monetary policy. In December 2013, the FOMC announced that it would taper its purchases of longer-term Treasuries to a rate of \$40bn per month. In December 2014, the FOMC promised that it would be “patient” in its approach to the withdrawal of monetary accommodation. In December 2015, for the first time in seven years, the FOMC raised the federal funds rate, from a target range of between 0% and 0.25% to a range of between 0.25% and 0.50%.

In June 2017, having increased the federal funds rate for a second time, the FOMC confirmed that it would soon begin the process of “policy normalisation,” publishing detailed plans of how it would phase out reinvestments of the proceeds of maturing Treasury securities as they rolled off the Federal Reserve’s \$4¹/₂ trillion balance sheet.

Although it is clear that U.S. monetary policy was incrementally loosened – and subsequently tightened – even absent any change in the actual federal funds rate, it is far less clear how actual changes in current policy compare with expected changes in future policy in their respective contribution to the post-crisis recovery. Identifying portfolio balance effects independently of signalling effects is the objective of this section.

Signalling effects

To capture signalling effects, we appeal to the decomposition of Treasury yields proposed by Adrian et al. (2013). In an affine term structure model, which estimates yields as linear functions of five pricing factors, the authors impose no-arbitrage restrictions to ensure that the time-series and the cross-section of bond returns are mutually consistent. Such an approach enables us to decompose U.S. Treasury yields into two distinct components: (i) the expected future path of short-term risk-neutral yields, and (ii) the contemporaneous risk premia required to compensate holders of long-dated government debt.

We measure the signal embedded in central bank behaviour as the difference between the current one-year risk-neutral yield, which we denote ι_1 , and the implied one-year risk-neutral yield one year ahead, which we denote ι_1^1 . The decomposition proposed by Adrian et al. (2013) gives us direct estimates of ι_1 and ι_2 . To infer ι_1^1 , we enforce the following no-arbitrage relationship:

$$(1 + \iota_2)^2 = (1 + \iota_1)(1 + \iota_1^1) \quad (14)$$

The no-arbitrage condition (14) is standard. It implies that, assuming all assets are risk-free and no opportunities for riskless profit can survive in equilibrium, it must be the case that investing \$1 for two years at a spot (net) annual interest rate of ι_2 is equivalent to investing \$1 for one year at a spot annual interest rate of ι_1 and then reinvesting one year later (for a further year) at a forward annual interest rate of ι_1^1 .

Re-writing (14) implies:

$$\iota_1^1 = \frac{(1 + \iota_2)^2}{1 + \iota_1} - 1 \quad (15)$$

At time t , we can then define f_t as:

$$f_t = \iota_{1,t}^1 - \iota_{1,t} \quad (16)$$

In any given quarter, f_t captures the expected change in one-year risk-neutral yields over the course of the subsequent year. If, as in some of the episodes described above, the Federal Reserve were to surprise financial markets not by cutting the federal funds

rate itself (because, for instance, it is already stuck at zero) but rather by announcing that it will not increase the federal funds rate for at least two years, this type of forward guidance would imply a decline in f_t , as investors reduce the probability they attach to a prospective rate increase over the subsequent 12 to 24 months.

Portfolio balance effects

To capture portfolio balance effects, we define d_t as the average residual maturity of U.S. Treasury debt held outside the official sector in any given quarter.

As Kuttner (2006) notes, there has been considerable variation in the maturity composition of privately held debt since the end of the Second World War. Relative to the 1950s and 1960s, debt management policy during the 1970s was stable and predictable, with no pronounced shifts in the composition of debt held either by the Treasury or by the Federal Reserve. The average residual maturity of the Federal Reserve’s portfolio tended to mirror the maturity distribution of outstanding Treasury debt, increasing as the Treasury issued more long-term debt in the late 1970s and the late 1990s. The 1980s, however, were a notable exception: the Federal Reserve reduced the average maturity of its portfolio, even as the U.S. Treasury increased the maturity of its outstanding debt. In the early 2000s, Treasury Secretary Lawrence Summers announced that the Treasury would buy back as much as \$30 billion in long-term bonds, in an attempt to reduce the government’s overall debt servicing costs. Soon enough, the Treasury announced it would stop issuing 30-year bonds altogether.⁹

Figure 4 plots the net effect of the policies pursued by the fiscal and monetary authorities in the U.S. over the past forty years.¹⁰ Between 1975 and 1990, the average residual maturity of public debt in private hands rose steadily, from around 30 months to close to 75 months. Having dipped in the mid-1990s, the average maturity of privately held debt fell back to around 58 months in the years leading up to the global financial

⁹Volcker (2002) recalls a similar adventure: “In 1953, the Treasury got aggressive and issued some 30 year bonds – the ‘3.25s of 78-83.’ The Treasury market was somewhat disturbed, and the economy went into a recession. Whether the Treasury’s aggressive debt issuance was a contributing factor was much discussed, and a long argument ensued about whether the Federal Reserve should intervene directly by conducting open market operations in the long-term market or whether such intervention should be left to debt management, with the Fed operating with ‘bills only.’ ”

¹⁰The series capturing the average residual maturity of U.S. public debt in private hands in Figure 4 is identical to that in Figure 1. Figure 4 merely illustrates the full history of the series since the mid-1970s.

crisis. We discussed the evolution of d_t during the post-crisis period in detail in Section 3.2: despite two bouts of genuine duration absorption – in 2008/09 and in 2011/12 – the average residual maturity of government debt held in private hands *increased* by 18 months between Q4:2008 and Q1:2017.

Three policy instruments

Including three endogenous variables in the policy block of our model is, for the most part, an attempt at parsimony. However, a very different approach – informed by event studies in financial markets – reinforces our modelling strategy. Swanson (2017) extends the high-frequency methodology developed by Gurkaynak et al. (2005). He finds that the systematic impact of FOMC announcements on asset prices between 1991 and 2015 can be explained by just three factors: changes in the federal funds rate, changes in forward guidance and changes in large-scale asset purchases. Because we are more interested in measuring the ultimate impact of asset purchases on the real economy rather than the intermediate impact of asset purchases on asset prices, our empirical approach is unrelated to that of Swanson (2017). Moreover, we differ fundamentally on the metric that best characterises a shock to macroeconomic policy.¹¹ But the fact that a careful diagnosis of asset price reactions points to similar policy variables as those which we include to understand the joint response of output and inflation suggests that the econometric constraints imposed by time-varying tractability are unlikely to incur large economic costs in terms of lost explanatory power.

5.2 Priors

Our results are based on 5,000 iterations of a Gibbs sampler, with the first 1,000 burn-in draws discarded to ensure convergence. The first forty observations (covering the first ten years of data, from Q1:1975 to Q1:1985) are used to calibrate the prior distributions of the parameter space, as in Primiceri (2005).

The mean and variance of B_0 , A_0 and $\log \sigma_0$ are derived from the OLS estimates $\hat{B}_{OLS}$, $\hat{A}_{OLS}$ and $\hat{\sigma}_{OLS}$ that minimise the sum of squared residuals in a time-invariant

¹¹One criticism of our identification strategy is that – in the absence of data on expectations of future net bond supply – we restrict our analysis to realised data on government bond issuance and central bank asset purchases. One way to overcome this constraint would be to adopt a narrative-based approach similar in spirit to the identification of fiscal shocks in Ramey (2011).

vector autoregression estimated on our initial, ten-year subsample. To complete our specification of the inverse-Wishart prior distributions of the hyperparameters $\mathbb{Q}$, $\mathbb{W}$ and $\mathbb{S}$, we follow the degrees of freedom and scale matrices adopted in the existing literature.¹²

Our priors are generally diffuse and largely uninformative. In summary:

$$\begin{aligned}
B_0 &\sim \mathcal{N}(\hat{B}_{OLS}, 4\sigma_{\hat{B}_{OLS}}^2) \\
A_0 &\sim \mathcal{N}(\hat{A}_{OLS}, 4\sigma_{\hat{A}_{OLS}}^2) \\
\log \sigma_0 &\sim \mathcal{N}(\log \hat{\sigma}_{OLS}, \mathbb{I}_n) \\
\mathbb{Q} &\sim \mathcal{IW}(40k_{\mathbb{Q}}^2\sigma_{\hat{B}_{OLS}}^2, 40) \\
\mathbb{W} &\sim \mathcal{IW}(4k_{\mathbb{W}}^2\mathbb{I}_n, 4) \\
\mathbb{S}_1 &\sim \mathcal{IW}(4k_{\mathbb{S}}^2\sigma_{\hat{A}_{1,OLS}}^2, 4) \\
\mathbb{S}_2 &\sim \mathcal{IW}(3k_{\mathbb{S}}^2\sigma_{\hat{A}_{2,OLS}}^2, 3)
\end{aligned} \tag{17}$$

where $\mathbb{S}_1$ and $\mathbb{S}_2$ denote the policy and non-policy blocks of $\mathbb{S}$, while $\hat{A}_{1,OLS}$ and $\hat{A}_{2,OLS}$ denote the corresponding blocks of $\hat{A}_{OLS}$. We set $k_Q = 0.01$, $k_S = 0.1$ and $k_W = 1$.

5.3 Shocks

Our interest lies primarily in identifying the macroeconomic impact of ‘duration shocks,’ defined in our model as exogenous changes in the average residual maturity of government debt held by the private sector. Duration shocks are exogenous in the sense that they are independent of the joint behaviour of the other endogenous variables in our model: inflation, unemployment, actual changes in the policy rate and perceived changes in the future path of monetary policy.

More specifically, our identifying assumption dictates that the maturity composition of private-sector portfolios affects inflation and unemployment with a lag of at least one quarter, and that duration shocks affect neither actual short-term interest rates nor expected future short-term interest rates within a given quarter.

We impose our identifying assumption by modelling the contemporaneous relationships between each endogenous variable in y_t in lower triangular form. As we showed in Section

¹²See, for example, Cogley (2003) and Cogley and Sargent (2005).

4, pre- and post-multiplying the variance-covariance matrix of the model's innovations Ω_t by a lower triangular matrix of contemporaneous coefficients A_t allows us to recover structural shocks from reduced-form errors. This Cholesky decomposition has long been used to identify the macroeconomic effects of conventional monetary policy shocks.¹³

The five variables in y_t are ordered as follows:

$$y_t = [\pi_t \quad u_t \quad i_t \quad f_t \quad d_t]'$$
 (18)

As described in Section 5.1, π_t is the inflation rate, u_t is the unemployment rate, i_t is the current federal funds rate, f_t is our estimate of the expected future change in short-term interest rates, and d_t is our proxy for the duration risk borne by the private sector. In y_t , the non-policy block is ordered above the policy block; within the policy block, d_t is ordered last.

These restrictions are consistent with the view that: (i) current monetary policy, as well as expected future monetary policy, responds contemporaneously to macroeconomic shocks, but only affects the macroeconomy with a transmission lag, and (ii) issuance decisions by debt managers are not scrutinised closely enough in financial markets for changes in portfolio duration to have an instantaneous impact on the expectations of investors in short-term interest rate futures. Indeed, Greenwood et al. (2014) provide event-study evidence to suggest that, during the crisis, the financial market impact of a U.S. Treasury refunding announcement (which reveals news on the average maturity of prospective debt supply) was typically only half the size of the financial market impact of an equivalent Federal Reserve asset purchase announcement (which reveals news on the average maturity of prospective debt demand). Modelling the economy such that the response of inflation and unemployment to a duration shock is restricted to be zero on impact (but left unrestricted afterwards) also reduces the risk that we conflate a shock to the average residual maturity of debt in private hands with a shock to the stance of debt-financed fiscal policy.

¹³See, for example, Sims (1980), Christiano et al. (1999), Peersman and Smets (2001) and Primiceri (2005).

6 Results

6.1 The volatility of shocks

First, we consider the posterior mean of the time-varying standard deviation of each shock identified in our model.

Panels (a) and (b) of Figure 5 plot the time-varying standard deviation of identified shocks to the model’s non-policy variables: inflation and unemployment.

The volatility of shocks to inflation was particularly elevated in the mid-1980s, following two pronounced oil price spikes in 1973 and 1979. The volatility of shocks to inflation was low and stable between the late 1980s and the mid-2000s, but rose sharply with the onset of the global financial crisis in 2007/08, before normalising somewhat in the early 2010s. That said, the volatility of shocks to inflation was still higher in 2017 than it was in 1987. The evolution of the standard deviations in Panel (a) is consistent with the conventional narrative surrounding the inflation performance of the U.S. economy after the Second World War: the recovery from the Great Inflation of the 1970s, the conquest of inflation under Federal Reserve Chairman Volcker in the 1980s, the anchoring of inflation expectations pre-2007, and the deflationary shock stemming from the global financial crisis post-2008.

Panel (b) also accords with the conventional narrative. The volatility of shocks to unemployment began trending lower in the mid-1980s, halving between 1990 and 2000 before reaching a trough in 2004. This twenty-year period of relative stability has since become known as the Great Moderation. The global financial crisis had a similar impact on the volatility of shocks to unemployment as it had on the volatility of shocks to inflation, inducing a sharp increase in 2007/08 followed by a gradual normalisation. That said, a decade on from the Great Recession, the volatility of shocks to unemployment was still higher than it had been in the final stages of the Greenspan era.

Based on their historical averages, exogenous shocks to inflation were around twice as volatile as exogenous shocks to unemployment over the period from 1985 to 2017.

Panels (c), (d) and (e) of Figure 5 plot the time-varying standard deviation of identified shocks to the model’s policy variables: the contemporaneous short-term interest rate, expected future changes in the short-term interest rate, and the average residual maturity of public debt in private hands.

The volatility of shocks to the contemporaneous policy rate was also particularly elevated in the mid-1980s, coinciding with Paul Volcker’s tenure as Chairman of the Federal Reserve. The volatility of shocks to the contemporaneous policy rate fell throughout the subsequent Great Moderation, to a level in 2005 that was a quarter of its level in 1985. Consistent with the immediate Federal Reserve response to the collapse of Lehman Brothers, and the subsequent incidence of the zero lower bound, the volatility of shocks to the contemporaneous policy rate rose abruptly in 2008, before falling back shortly afterwards. Between 2010 and 2017, the volatility of shocks to the federal funds rate remained low and stable relative to recent U.S. history.

The volatility of shocks to expected future changes in the policy rate rose in each of the three recessions in our sample – 1991, 2001 and 2008 – but it was after the 2001 recession that the increase in volatility was most pronounced. This is consistent with the FOMC’s decision, in mid-1999, to begin including language about the future stance of monetary policy in its post-meeting statements. As Rudebusch and Williams (2008) note, for example, in early 2000 the FOMC’s directional guidance – “the Committee...is tilted toward the possibility of a firming in the stance of monetary policy” – was replaced with more circumspect language describing the “balance of risks” around the Committee’s dual mandate. In August 2003, the FOMC indicated that “policy accommodation can be maintained for a considerable period.” In May 2004, however, the FOMC argued that “policy accommodation can be removed at a pace that is likely to be measured.” The volatility of shocks to expected future changes in the policy rate also rose in 2013, albeit temporarily. This is consistent with the sharp reaction in financial markets to Chairman Bernanke’s guidance, in the spring and summer of that year, that the Federal Reserve would likely soon start moderating the pace of its large-scale asset purchases, conditional on continued strength in the macroeconomic data. This episode became known as the “taper tantrum,” owing to investors’ hypersensitivity to news about the prospective withdrawal of post-crisis policy stimulus.

The volatility of shocks to the average residual maturity of public debt in private hands was low and stable throughout the Great Moderation, but rose sharply with the onset of the financial crisis, to a level four times higher in Q4:2008 than it had been, on average, between 1985 and 2000. This is consistent with the Federal Reserve’s first round of large-scale asset purchases. Only in 2012 did the volatility of shocks to private-sector duration revert to its pre-crisis mean. In the context of our model, this increase in the volatility of identified duration shocks can be attributed to neither ‘conventional’ changes

in the federal funds rate nor ‘unconventional’ use of forward guidance. Nor can it be attributed to changes in the private sector’s behaviour, as captured by the autoregressive equations for inflation and unemployment estimated in the non-policy block of our model. Our identification scheme suggests that, although U.S. debt management policy cultivated a reputation of being “regular and predictable” over several decades since the 1970s, the maturity composition of the net supply of Treasury debt systematically responded to extraneous factors between 2007 and 2011.¹⁴

Based on their historical averages, exogenous shocks to the residual maturity of public debt in private hands were three to four times more volatile than exogenous shocks to our other policy variables between 1985 and 2017 – even abstracting from the 2007-2011 spike in the standard deviation of the former.

6.2 The effects of duration shocks

We now consider the impulse responses of inflation, unemployment, the contemporaneous short-term interest rate and the expected future short-term interest rate to a positive, one standard deviation shock to the average residual maturity of public debt in private hands at various points in time. This is our standardised measure of a ‘duration shock.’

In Figures 6 and 7, we plot the impact of a duration shock at each of the three peaks in the U.S. business cycle since the beginning of the Greenspan era: (i) Q3:1990, (ii) Q1:2001, and (iii) Q4:2007. Figure 6 compares median impulse response functions; Figure 7 plots the 16% and 84% confidence bands associated with each median response.

An exogenous increase in the average residual maturity of public debt in private hands reduces inflation and increases unemployment. In response to the fall in inflation and the rise in unemployment, contemporaneous short-term interest rates fall to stabilise the economy (albeit with a lag) and expected future short-term interest rates rise.

The response of the economy to shocks to portfolio duration exhibits considerable

¹⁴In a monograph on the modern history of U.S. Treasury debt management, Garbade (2015) notes that “Treasury officials are less likely to be criticized for breaching the principle of regular and predictable issuance when changing the auction calendar and/or offering amounts if market participants anticipated the actions on the basis of the Treasury’s objectives and the contemporaneous economic environment. In a perfect world, the Treasury’s objectives would be so well-known that its actions would be fully anticipated from the economic data.”

nonlinearity over time. A one standard deviation shock to the duration of private-sector portfolios in Q3:1990 reduces annual inflation by 2bps after three years; the same shock in Q1:2001 reduces inflation by 10bps. A one standard deviation shock to the duration of private-sector portfolios in Q3:1990 increases the unemployment rate by 3bps after four years; the same shock in Q4:2007 increases the unemployment rate by 6bps.

In Figures 8 and 9, we conduct an analogous exercise, but with the incidence of the duration shock coinciding with the beginning of each round of large-scale asset purchases by the Federal Reserve: (i) Q1:2009, which corresponds to the start of QE1; (ii) Q4:2010, which coincides with the start of QE2; and (iii) Q3:2012, which marks the start of QE3. Figure 8 compares median impulse response functions; Figure 9 plots the 16% and 84% confidence bands associated with each median response.

We know from Figure 4 that, between January 2009 and December 2010, the average residual maturity of public debt in private hands *rose*, from 47 to 57 months, despite Federal Reserve purchases of \$300bn of longer-term Treasury securities. Between July 2012 and March 2017, the average residual maturity of public debt in private hands rose further, from 56 to 64 months, despite Federal Reserve purchases of a further \$1.7trn of longer-term Treasury securities. Figures 8 and 9 imply that these episodes of duration injection exacerbated unemployment and undermined inflation. A one standard deviation shock to the duration of private-sector portfolios in Q1:2009 reduces the annual rate of inflation by 15bps and increases the unemployment rate by 6bps after four years. The same shock in Q3:2012 has a smaller impact but is nevertheless contractionary, reducing the annual rate of inflation by 4bps and increasing the unemployment rate by 2bps after four years.

In Figures 10 and 11, we compare the impact of a duration shock in a representative quarter in each of three very different macroeconomic regimes in recent U.S. history: (i) the midst of the Great Moderation, Q1:1997, (ii) the nadir of the global financial crisis, Q2:2009, and (iii) the beginning of policy normalisation, Q2:2017. The first date marks the halfway point in the 120-month expansion of the 1990s, which the National Bureau of Economic Research classifies as the longest recession-free period in the United States since the Civil War era.¹⁵ The second date marks the trough in U.S. GDP after an 18-month contraction that began in December 2007 and intensified after September 2008.

¹⁵See Bernanke (2004).

The third date marks the second post-crisis increase in the federal funds rate, and the formal publication of Federal Reserve plans for balance sheet “normalisation.”¹⁶ Figure 10 compares median impulse response functions; Figure 11 plots the 16% and 84% confidence bands associated with each median response.

A one standard deviation shock to the average maturity of public debt in private hands in the midst of the Great Moderation reduces annual inflation by around 6bps, increases the unemployment rate by 1bp, and increases expected future short-term interest rates by 1bp after three years. By contrast, an equivalent injection of duration into private-sector portfolios at the nadir of the financial crisis implies a reduction in three-year-ahead inflation that is almost three times as large, an increase in unemployment that is five times as large, and an increase in expected future interest rates that is twice as large.

Overall, our empirical analysis suggests that there is considerable time variation in the macroeconomic effects of exogenous shocks to the maturity of public debt in private-sector portfolios.

6.3 State dependence

In theory, imperfect substitutability among financial assets is a pre-requisite for the viability of the portfolio balance channel. Only if investors have different ‘preferred habitats’ along the spectrum of available assets will changes in the relative supplies of those assets provoke changes in their relative prices. And only by influencing relative prices will the absorption or injection of duration risk by policymakers affect the real economy.

As we discussed in the Introduction, this theoretical proposition became the practical rationale for quantitative easing in the aftermath of the global financial crisis. The argument proceeded as follows. By purchasing assets from the non-bank private sector in return for central bank reserves, QE would increase sellers’ holdings of broad money. Insofar as broad money was seen as an imperfect substitute for the assets purchased by the central bank, those sellers would seek to rebalance their portfolios by buying riskier assets, such as corporate bonds, in place of government securities. The sellers of those riskier assets would, in turn, seek to rebalance their own portfolios, and so on, trigger-

¹⁶In June 2017, the FOMC raised the target range for the federal funds rate to between 1% and 1.25% and formalised its plans to “begin implementing a balance sheet normalisation program this year, provided that the economy evolves broadly as anticipated.”

ing a sequence of asset reallocation. During this process, asset prices would rise – yields would fall – until private-sector investors were indifferent between the relative supply of broad money and the relative supply of financial assets in the economy. The increase in risky asset prices would increase the net wealth of asset-holders, while the decline in risky yields would reduce the borrowing costs faced by households and firms. Together, these portfolio balance effects would stimulate real activity and boost inflation.

A corollary of this argument is that portfolio balance effects are strongest when asset heterogeneity is greatest. Asset heterogeneity is greatest when risk premia are most pervasive. An investor who sells ten-year U.S. Treasury debt to the Federal Reserve at a time when the market price of risk is high would be far less likely to settle for the return on the central bank reserves she receives in return – and therefore far more likely to rebalance her portfolio by buying U.S. corporate debt – than she would be had she sold ten-year U.S. Treasury debt to the Federal Reserve at a time when the market price of risk was low. In other words, the more imperfect the perceived substitutability among assets, the greater the search for yield, and the more stimulative the resultant compression of risk premia.¹⁷

Figures 6-11 corroborate the theoretical case for state-dependent portfolio balance effects. Over a thirty-year sample spanning both pre- and post-crisis policy regimes, the macroeconomic effects of a shock to duration were most acute at the nadir of the financial crisis. Only during the period between 2008 and 2010 did a one standard deviation duration shock have a three-year-ahead inflation impact in excess of 10bps and a three-year-ahead impact on the unemployment rate in excess of 5bps. Between Q3:2008 and Q4:2010, the term premium on U.S. Treasury debt remained in line with its long-run historical average, but the market price of risk spiked to unprecedented levels. The VIX – an index of the expected near-term volatility embedded in equity options – rose to a level three times its historical average in Q4:2008. The TED spread – the difference between the ‘risk-free’ yield on Treasury bills and the counterparty risk embedded in Eurodollar futures – rose to a level six times its historical average in Q4:2008. The Corporate BBB spread – the difference between the yield on BBB-rated corporate debt and spot Treasury

¹⁷Based on U.S. data to 2012, Carpenter et al. (2013) find that the private-sector investors – primarily domestic households – from whom the Federal Reserve bought the bulk of its assets indeed rebalanced their portfolios toward more risky assets during this period. Based on U.K. data to 2012, Joyce et al. (2014) find that asset purchases by the Bank of England indeed prompted life insurance companies and pension funds to shift their portfolios away from gilts towards corporate bonds.

rates – rose to a level four times its historical average in Q1:2009. Portfolio balance effects indeed appear to have been strongest when risk premia were most pervasive.

The irony, however, is that – at the nadir of the financial crisis, just when they were at their most potent – these portfolio balance effects were *contractionary* rather than expansionary. The *joint* behaviour of the U.S. Treasury and the Federal Reserve during the Great Recession culminated in the injection of duration risk into private-sector portfolios rather than its absorption. The Treasury’s pursuit of active maturity extension dominated the Federal Reserve’s pursuit of duration withdrawal. Changes in the composition of the U.S. government’s consolidated balance sheet amounted to quantitative *tightening* rather than quantitative easing. Relative to a no-shock counterfactual, the net effect of fiscal and monetary policy was therefore to exacerbate the collapse in employment and to amplify the risk of deflation.

6.4 Balance sheet normalisation

Notwithstanding its potency in times of crisis, Figures 6-11 suggest that an exogenous shock to the duration of private-sector portfolios has meaningful macroeconomic effects even in normal times.

By 2013, the U.S. economy had begun to recover from the trauma of 2008. In late 2015, the Federal Reserve raised the federal funds rate for the first time in seven years. In early 2017, Chair Yellen warned that waiting too long before withdrawing more monetary stimulus “could risk a nasty surprise [of] too much inflation” if the economy were allowed “to run markedly and persistently hot.”¹⁸ Along with a second rate increase, the FOMC formalised its intention to “gradually reduce the Federal Reserve’s holdings of Treasury and agency securities once normalisation of the level of the federal funds rate is well under way.”

Our results rationalise this strategy of balance sheet normalisation. By reducing the extent to which it reinvests principal payments from maturing government securities held on its balance sheet, the Federal Reserve would – all else equal – add to the average residual maturity of public debt in private-sector portfolios. Figures 10 and 11 suggest that this duration injection would curb inflation and curtail employment. A one standard

¹⁸See Yellen (2017a) and Yellen (2017b).

deviation duration shock in Q2:2017 reduces the annual rate of inflation by 5bps after one year and by 4bs after three years, while increasing the civilian unemployment rate by 1bp after one year and by 3bps after three years.¹⁹

7 Conclusion

More than half a century ago, Tobin (1963) advanced the notion that “the government – comprising both the Federal Reserve and the Treasury – should continuously adjust the maturity structure of the debt, seeking to minimize its net cost, while achieving the required restraint on aggregate demand.” Although this argument was ultimately eclipsed by the allure of central bank independence in the 1990s, it did in fact underpin macroeconomic policymaking in the United States through much of the post-war experience. According to Volcker (2002), for example, the 1950s and 1960s were decades during which “debt management was considered to be an active ‘third leg’ of policy.”


Seen from this perspective, the “unconventional” monetary policies adopted by major central banks in response to the global financial crisis of 2007/08 are not especially unconventional. Large-scale purchases of long-term government bonds by a monetary authority in the secondary market constitute little more than an attempt by the government to adjust the maturity structure of its debt, with the objective of achieving the required stimulus to aggregate demand. Similarly, the monetary authority’s subsequent sale of those long-term government bonds – in pursuit of “balance sheet normalisation” – is another attempt by the government to adjust the maturity structure of its debt, now with the opposite objective of achieving the required restraint on aggregate demand.

In this paper, we exploit the equivalence between central bank purchases and treasury issuance to estimate the macroeconomic impact of attempts to adjust the maturity structure of public debt in private hands at different points in time. In particular, controlling for the signalling effects associated with central bank communication, we find that an

¹⁹All else is unlikely to be equal. In the minutes of its May 2017 meeting, the Treasury Borrowing Advisory Committee (TBAC) commented on “the demand for ultra-long debt.” Despite press reports suggesting that issuance of a 50-year or 100-year U.S. Treasury bond was imminent, “the Committee recommended that further work be done to study these demand dynamics to get a better sense of where an ultra-long bond might price.” The TBAC recommended that the Treasury “consider issuing a zero coupon 50-year bond, and coupon maturities between 10- and 30-years, preferably the re-introduction of the 20-year,” noting the “preferable attributes of stripped 30-year bonds to meet a similar duration as a 100-year coupon bond.”

exogenous increase in the average residual maturity of Treasury debt in private-sector portfolios reduces inflation and increases unemployment. These effects are three or four times larger in the depths of the Great Recession than in the midst of the Great Moderation – a finding which underscores the importance of imperfect asset substitutability, but also suggests that the *joint* behaviour of the U.S. Treasury and the Federal Reserve in the aftermath of the global financial crisis of 2007/08 amounted to quantitative *tightening* rather than quantitative easing.

Figure 1: IN THE U.S., WHILE THE CENTRAL BANK ABSORBED DURATION AFTER 2008, THE TREASURY INJECTED DURATION




Notes: Each line indicates the average residual maturity of interest-bearing debt attributable to each agent in the economy. “CB” refers to the assets held on the balance sheet of the U.S. Federal Reserve. “TREASURY” refers to the total stock of debt issued by the U.S. Treasury. “PRIVATE” refers to the portfolio of the typical U.S. household in the private sector.

Figure 2: IN THE U.K., THERE WAS RELATIVELY MORE CO-ORDINATION BETWEEN THE FISCAL AND MONETARY AUTHORITIES




Notes: Each line indicates the average residual maturity of interest-bearing debt attributable to each agent in the economy. “CB” refers to the assets held on the balance sheet of the Bank of England. “TREASURY” refers to the total stock of debt issued by HM Treasury. “PRIVATE” refers to the portfolio of the typical U.K. household in the private sector.

Figure 3: IN THE EURO AREA, THE CENTRAL BANK’S ATTEMPTS TO ABSORB DURATION AFTER 2015 WERE ALSO OFFSET BY THE INJECTION OF DURATION BY NATIONAL FISCAL AUTHORITIES




Notes: Each line indicates the average residual maturity of interest-bearing debt attributable to each agent in the economy. “CB” refers to the assets held on the balance sheet of the European Central Bank. “TREASURY” refers to the total stock of debt issued by Euro area governments. “PRIVATE” refers to the portfolio of the typical Euro area household in the private sector.

Figure 4: IN THE U.S., THERE HAS BEEN CONSIDERABLE HISTORICAL VARIATION IN THE MATURITY COMPOSITION OF PRIVATELY HELD GOVERNMENT DEBT




Notes: “PRIVATE” refers to the portfolio of the typical U.S. household in the private sector. Shading denotes U.S. recessions. Figure 4 is the historical counterpart to the private-sector data shown in Figure 1.

Figure 5: THE TIME-VARYING STANDARD DEVIATION OF IDENTIFIED SHOCKS TO THE MODEL'S ENDOGENOUS VARIABLES




Notes: Each panel plots the posterior mean of the time-varying standard deviation of identified shocks to the model's non-policy variables (panels (a) and (b)) and to the model's policy variables (panels (c), (d) and (e)). The abbreviation "inf" refers to the inflation rate, π_t . The abbreviation "unemp" refers to the unemployment rate, u_t . The abbreviation "int" refers to the contemporaneous federal funds rate, i_t . The abbreviation "expdint" refers to the expected future change in the short-term interest rate, f_t . The abbreviation "prdur" refers to the duration risk borne by the private sector, d_t .

Figure 6: THE IMPACT OF A DURATION SHOCK AT EACH PEAK IN THREE MAJOR U.S. BUSINESS CYCLES




Notes: Each panel plots the median impulse response (IR) of inflation (inf), unemployment (unemp) and the policy rate (int) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q3:1990 (1990.50), Q1:2001 (2001.00) and Q4:2007 (2007.75). The horizontal axis denotes quarters.

Figure 7: THE IMPACT OF A DURATION SHOCK AT EACH PEAK IN THREE MAJOR U.S. BUSINESS CYCLES




Notes: Each panel plots the impulse response (IR) of inflation (inf), unemployment (unemp), the policy rate (int) and forward rates (expdint) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q3:1990 (1990.50), Q1:2001 (2001.00) and Q4:2007 (2007.75). The solid line illustrates the median impulse response; the dotted lines illustrate 16% and 84% confidence bands. The horizontal axis denotes quarters.

Figure 8: THE IMPACT OF A DURATION SHOCK AT THE BEGINNING OF EACH ROUND OF LARGE-SCALE ASSET PURCHASES BY THE U.S. FEDERAL RESERVE




Notes: Each panel plots the median impulse response (IR) of inflation (inf), unemployment (unemp) and the policy rate (int) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q1:2009 (2009.00), Q4:2010 (2010.75) and Q3:2012 (2012.50). The horizontal axis denotes quarters.

Figure 9: THE IMPACT OF A DURATION SHOCK AT THE BEGINNING OF EACH ROUND OF LARGE-SCALE ASSET PURCHASES BY THE U.S. FEDERAL RESERVE




Notes: Each panel plots the impulse response (IR) of inflation (inf), unemployment (unemp), the policy rate (int) and forward rates (expdint) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q1:2009 (2009.00), Q4:2010 (2010.75) and Q3:2012 (2012.50). The solid line illustrates the median impulse response; the dotted lines illustrate 16% and 84% confidence bands. The horizontal axis denotes quarters.

Figure 10: THE IMPACT OF A DURATION SHOCK DURING THREE DIFFERENT MACROECONOMIC REGIMES IN RECENT U.S. HISTORY



Notes: Each panel plots the median impulse response (IR) of inflation (inf), unemployment (unemp) and the policy rate (int) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q1:1997 (1997.00), Q2:2009 (2009.25) and Q2:2017 (2017.00). The horizontal axis denotes quarters.

Figure 11: THE IMPACT OF A DURATION SHOCK DURING THREE DIFFERENT MACROECONOMIC REGIMES IN RECENT U.S. HISTORY



Notes: Each panel plots the impulse response (IR) of inflation (inf), unemployment (unemp), the policy rate (int) and forward rates (expdint) to a positive, one standard deviation shock to the average residual maturity of public debt in private hands (pridur) at three points in time: Q1:1997 (1997.00), Q2:2009 (2009.25) and Q2:2017 (2017.00). The solid line illustrates the median impulse response; the dotted lines illustrate 16% and 84% confidence bands. The horizontal axis denotes quarters.

A

In this Appendix, we outline the basic Markov chain Monte Carlo algorithm first developed by Primiceri (2005), and subsequently modified by Del Negro and Primiceri (2015).

A.1 Step 1: Drawing coefficient states

Let $M^\tau = [m'_1, m'_2, \dots, m'_\tau]'$ denote the history of a generic matrix of variables M_t from $t = 1$ to τ , where m_t is a column vector constructed with the time-varying elements of M_t . Let $p(\cdot)$ denote a generic density function, with $\mathcal{N}$ denoting the Gaussian distribution.

Conditional on A^T , Σ^T and $\mathbb{V}$, the observation equation (7) is linear, with Gaussian innovations of known variance. As shown in Fruhwirth-Schnatter (1994) and Carter and Kohn (1994), the density $p(B^T \mid y^T, A^T, \Sigma^T, \mathbb{V})$ can be factored as:

$$p(B^T \mid y^T, A^T, \Sigma^T, \mathbb{V}) = p(B_T \mid y^T, A^T, \Sigma^T, \mathbb{V}) \prod_{t=1}^{T-1} p(B_t \mid B_{t+1}, y^t, A^T, \Sigma^T, \mathbb{V}) \quad (19)$$

in which

$$\begin{aligned} B_t \mid B_{t+1}, y^t, A^T, \Sigma^T, \mathbb{V} &\sim \mathcal{N}(B_{t|t+1}, P_{t|t+1}), \\ B_{t|t+1} &= \mathbb{E}(B_t \mid B_{t+1}, y^t, A^T, \Sigma^T, \mathbb{V}), \\ \Psi_{t|t+1} &= \text{Var}(B_t \mid B_{t+1}, y^t, A^T, \Sigma^T, \mathbb{V}). \end{aligned}$$

It is straightforward to draw each B_t . This is because $B_{t|t+1}$ and $\Psi_{t|t+1}$ can be computed using the forward and backward recursion reported in Appendix B, as applied to the state-space model given by (7) and (8). In particular, the last recursion of the Kalman filter provides the mean and variance of the posterior distribution of B_T , that is, $B_{T|T}$ and $\Psi_{T|T}$. A value drawn from this distribution is then used in the backward recursion to obtain $B_{T-1|T}$ and $\Psi_{T-1|T}$. A value drawn from this distribution is then used in the backward recursion to obtain $B_{T-2|T}$ and $\Psi_{T-2|T}$, and so on.

A.2 Step 2: Drawing covariance states

In (7) we wrote the system of equations as:

$$\begin{aligned}
y_t &= X_t' B_t + A_t^{-1} \Sigma_t \varepsilon_t, \\
X_t' &= \mathbb{I}_n \otimes [1, y_{t-1}', \dots, y_{t-k}']
\end{aligned} \tag{20}$$

This can be re-written as:

$$A_t(y_t - X_t' B_t) = A_t \hat{y}_t = \Sigma_t \varepsilon_t \tag{21}$$

In (21), $\hat{y}_t$ is observable if B^T is taken as given. Since A_t is a lower triangular matrix with ones on the main diagonal, (21) can be rewritten as:

$$\hat{y}_t = Z_t \alpha_t + \Sigma_t \varepsilon_t \tag{22}$$

In (22), α_t is as defined in Section 4.1, and Z_t denotes the following matrix:

$$Z_t = \begin{bmatrix} 0 & \dots & \dots & 0 \\ -\hat{y}_{1,t} & 0 & \dots & 0 \\ 0 & -\hat{y}_{[1,2],t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -\hat{y}_{[1,2,\dots,n-1],t} \end{bmatrix}$$

in which $\hat{y}_{[1,2,\dots,i],t}$ denotes the row vector $[\hat{y}_{1,t}, \hat{y}_{2,t}, \dots, \hat{y}_{i,t}]$.

The model described in (22) has a Gaussian state-space representation, but it is non-linear because the dependent variable ($\hat{y}_t$) on the left-hand side of the observation equation (22) also appears as an independent variable on the right-hand side (via Z_t). As such, the vector $[\hat{y}_t, \hat{\alpha}_t]$ is not jointly normal, implying that the conditional distributions cannot be computed using the standard Kalman filter recursion described in Appendix A.1. However, since we assume that $\mathbb{S}$ is block diagonal, the Kalman filter and the backward recursion can be applied on an equation-by-equation basis. Indeed, in the i^{th} equation, $\hat{y}_{i,t}$ is not among the variables explaining $\hat{y}_{i,t}$: due to the triangular structure of Z_t , we can treat $\hat{y}_{[1,2,\dots,i-1],t}$ as pre-determined.

An equation-by-equation approach, under the assumption that $\mathbb{S}$ is block diagonal, allows us to follow a similar technique as in Appendix A.1 to recursively recover:

$$\begin{aligned}\alpha_{i,t|t+1} &= \mathbb{E}(\alpha_{i,t} \mid \alpha_{i,t+1}, y^t, B^T, \Sigma^T, \mathbb{V}), \\ \Lambda_{i,t|t+1} &= \text{Var}(\alpha_{i,t} \mid \alpha_{i,t+1}, y^t, B^T, \Sigma^T, \mathbb{V})\end{aligned}$$

in which $\alpha_{i,t}$ is the i^{th} block of α_t , corresponding to the coefficients of the i^{th} equation in (22). As in Appendix A.1, $\alpha_{i,t}$ can be drawn recursively from $p(\alpha_{i,t} \mid \alpha_{i,t+1}, y^t, B^T, \Sigma^T, \mathbb{V})$, which is $\mathcal{N}(\alpha_{i,t|t+1}, \Lambda_{i,t|t+1})$.

A.3 Step 3: Drawing volatility states

Re-write (20) instead as:

$$A_t(y_t - X_t' B_t) = y_t^* = \Sigma_t \varepsilon_t \quad (23)$$

y_t^* is observable if A^T and B^T are taken as given. The model described in (23) is also non-linear, but can be linearised by squaring and taking logarithms of every measurement equation it contains. Letting $y_{i,t}^{**} = \log((y_{i,t}^*)^2 + \bar{c})$, letting $e_{i,t} = \log(\varepsilon_{i,t}^2)$ and letting $s_{i,t} = \log(\sigma_{i,t})$, this leads to the following state-space form:²⁰

$$\begin{aligned}y_t^{**} &= 2s_t + e_t, \\ s_t &= s_{t-1} + \eta_t\end{aligned} \quad (24)$$

in which e and η are uncorrelated. Though linear, the system in this form is non-Gaussian because the innovations in the measurement equations are distributed as a $\log\chi^2(1)$. In order to further transform (24) into a Gaussian system, we use the ‘mixture of normals’ approximation of the $\log\chi^2(1)$ distribution, as described by Kim, Shephard and Chib (1998). Since $\text{Var}(\varepsilon_t) = \mathbb{I}_n$ from (6), this implies that the variance-covariance matrix of e_t is diagonal, allowing us to use the same (independent) ‘mixture of normals’ approximation for any element of e .

Let $\theta^T = [\theta_1, \theta_2, \dots, \theta_T]'$ denote a matrix of indicator variables that selects which member of the ‘mixture of normals’ approximation to use for each element of e at each point in time. Conditional on $A^T, B^T, \mathbb{V}$ and θ^T , the system in (24) has an approximately linear,

²⁰Following Primiceri (2005), the offset constant $\bar{c} = 0.001$ is used to make the estimation procedure more robust to the fact that $y_{i,t}^2$ may be very small.

Gaussian state-space form. This allows us to follow a similar technique as in Appendix A.1 and Appendix A.2 to recursively recover:

$$\begin{aligned} s_{t|t+1} &= \mathbb{E}(s_t \mid s_{t+1}, y^t, A^T, B^T, \mathbb{V}, \theta^T), \\ \Phi_{t|t+1} &= \text{Var}(s_t \mid s_{t+1}, y^t, A^T, B^T, \mathbb{V}, \theta^T) \end{aligned}$$

As in Appendix A.1 and Appendix A.2, s_t can then be drawn recursively from $p(s_t \mid s_{t+1}, y^t, A^T, B^T, \mathbb{V}, \theta^T)$, which is $\mathcal{N}(s_{t|t+1}, \Phi_{t|t+1})$. Conditional on y^{**T} and the resultant s^T , we can sample the new θ^T to be used in the next iteration.

A.4 Step 4: Drawing hyperparameters

The hyperparameters of the model are the diagonal blocks of $\mathbb{V}$. As described in (9), we denote these: $\mathbb{Q}$, $\mathbb{W}$ and $\mathbb{S}$. Since $\mathbb{S}$ is block diagonal, we further decompose $\mathbb{S}$ into $\mathbb{S}_1$ and $\mathbb{S}_2$, with each corresponding to the parameters of separable sets of equations. Conditional on B^T , Σ^T , A^T and y^T , each diagonal block of $\mathbb{V}$ follows an inverse-Wishart distribution, independent of the other blocks.

This implies that:

$$\begin{aligned} p(\mathbb{Q}, \mathbb{W}, \mathbb{S} \mid y^T, A^T, B^T, \Sigma^T) &= p(\mathbb{Q} \mid y^T, A^T, B^T, \Sigma^T) \cdot p(\mathbb{W} \mid y^T, A^T, B^T, \Sigma^T) \\ &\cdot p(\mathbb{S}_1 \mid y^T, A^T, B^T, \Sigma^T) \cdot p(\mathbb{S}_2 \mid y^T, A^T, B^T, \Sigma^T) \quad (25) \end{aligned}$$

Conditional on B^T , Σ^T , A^T and y^T , it is straightforward to draw from these inverse-Wishart posterior distributions using the sampling procedure developed by Gelman et al. (1995), because the innovations are observable.

A.5 Summary of steps

In summary, our algorithm proceeds in seven steps:²¹

²¹Following Del Negro and Primiceri (2015), we switch the order of steps (d) and (e) in Primiceri (2005). This is because the original algorithm was found to yield draws from the incorrect posterior distribution of model parameters: see Appendix A.5 of Primiceri (2005) and Section 3 of Del Negro and Primiceri (2015) for more details.

- a. Initialise A^T , Σ^T , Θ^T and $\mathbb{V}$.
- b. Sample B^T from $p(B^T \mid y^T, A^T, \Sigma^T, \mathbb{V})$.
- c. Sample A^T from $p(A^T \mid y^T, B^T, \Sigma^T, \mathbb{V})$.
- d. Sample θ^T from $p(\theta^T \mid y^T, A^T, \Sigma^T, \mathbb{V})$.
- e. Sample Σ^T from $p(\Sigma^T \mid y^T, A^T, B^T, \mathbb{V}, \theta^T)$.
- f. Sample $\mathbb{V}$ by sampling $\mathbb{Q}$, $\mathbb{W}$ and $\mathbb{S}$ from $p(\mathbb{Q}, \mathbb{W}, \mathbb{S} \mid y^T, A^T, B^T, \Sigma^T)$.
- g. Return to (b.).

B

In this Appendix, we outline the technique used for Gibbs sampling in state-space models.

B.1 The Gibbs sampler

Consider the generic measurement equation given by:

$$y_t = H_t \beta_t + \varepsilon_t \quad (26)$$

Consider the generic transition equation given by:

$$\beta_t = F \beta_{t-1} + u_t \quad (27)$$

in which:

$$\begin{pmatrix} \varepsilon_t \\ u_t \end{pmatrix} \sim i.i.d. \mathcal{N} \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} R_t & 0 \\ d & Q \end{pmatrix} \right]$$

Now, let:

$$\begin{aligned} \beta_{t|s} &= \mathbb{E}(\beta_t \mid Y^s, H^s, R^s, Q), \\ V_{t|s} &= Var(\beta_t \mid Y^s, H^s, R^s, Q), \end{aligned} \quad (28)$$

Given $\beta_{0|0}$ and $V_{0|0}$, a standard Kalman filter delivers:

$$\begin{aligned}
\beta_{t|t-1} &= F\beta_{t-1|t-1} \\
V_{t|t-1} &= FV_{t-1|t-1}F' + Q \\
K_t &= V_{t|t-1}H_t'(H_tV_{t|t-1}H_t' + R_t)^{-1} \\
\beta_{t|t} &= \beta_{t|t-1} + K_t(y_t - H_t\beta_{t|t-1}) \\
V_{t|t} &= V_{t|t-1} - K_tH_tV_{t|t-1}
\end{aligned} \tag{29}$$

From (29), it is clear that the final elements of the recursion are $\beta_{T|T}$ and $V_{T|T}$. These are, respectively, the mean and the variance of the normal distribution from which we draw β_T . This draw of β_T and the other remaining outputs of the filter are now used for the first step of the backward recursion, which follows the updating formulae given by:

$$\begin{aligned}
\beta_{t|t+1} &= \beta_{t|t} + V_{t|t}F'V_{t+1|t}^{-1}(\beta_{t+1} - F\beta_{t|t}) \\
V_{t|t+1} &= V_{t|t} - V_{t|t}F'V_{t+1|t}^{-1}FV_{t|t} \mid t)
\end{aligned} \tag{30}$$

In the first instance, (30) provides $\beta_{T-1|T}$ and $V_{T-1|T}$, which are used to make a draw of β_{T-1} . This backward recursion continues until $t = 0$.

References

- Adrian, T., Crump, R., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138.
- Andrade, P., Breckenfelder, J., de Fiore, F., Karadi, P., and Tristani, O. (2016). The ECB’s Asset Purchase Programme: An early assessment. *ECB Working Paper*, 1956.
- Baumeister, C. and Benati, L. (2013). Unconventional monetary policy and the Great Recession: Estimating the macroeconomic effects of a spread compression at the zero lower bound. *International Journal of Central Banking*, 9(2):165–212.
- Bernanke, B. S. (2004). “The Great Moderation,” Remarks at the meetings of the Eastern Economic Association, Washington, DC.
- Bernanke, B. S. (2010). “Implications of the financial crisis for economics,” Remarks at The Bendheim Centre for Finance, Princeton University.
- Bernanke, B. S., Boivin, J., and Elias, P. (2005). Measuring the effects of monetary policy: A factor augmented vector autoregressive FAVAR approach. *The Quarterly Journal of Economics*, 120(1):387–422.
- Blattner, T. and Joyce, M. A. S. (2016). Net debt supply shocks in the euro area and the implications for QE. *ECB Working Paper*, 1957.
- Boivin, J. (1999). The Fed’s conduct of monetary policy: Has it changed and does it matter? *Columbia University, mimeo*.
- Bremner, R. P. (2004). *Chairman of the Fed*. Yale University Press, New Haven, CT.
- Carpenter, S., Demiralp, S., Ihrig, J., and Klee, E. (2013). Analyzing Federal Reserve asset purchases: From whom does the Fed buy? *Finance and Economics Discussion Series*, 2.
- Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81:541–553.
- Chadha, J., Turner, P., and Zampolli, F. (2013). The interest rate effects of government debt maturity. *BIS Working Paper*, 415.
- Christiano, L. J., Eichenbaum, M., and Evans, C. (1999). Monetary shocks: What have we learned and to what end? *Handbook of Macroeconomics*, 1A:65–148.

- Coeure, B. (2015). “Embarking on public sector asset purchases,” Remarks at the Second International Conference on Sovereign Bond Markets, Frankfurt.
- Cogley, T. (2003). How fast can the new economy grow?: A Bayesian analysis of the evolution of trend growth. *University of California, Davis, mimeo*.
- Cogley, T. and Sargent, T. J. (2005). Drift and volatilities: Monetary policy outcomes in the post-WWII U.S. *Review of Economic Dynamics*, 8(2):262–302.
- Del Negro, M. and Primiceri, G. E. (2015). Time varying structural vector autoregressions and monetary policy: A corrigendum. *The Review of Economic Studies*, 82(4):1342–1345.
- Garbade, K. D. (2015). *Treasury debt management under the rubric of regular and predictable issuance: 1983-2012*. The Federal Reserve Bank of New York, New York, N.Y.
- Greenwood, R., Hanson, S., Rudolph, J., and Summers, L. (2014). Government debt management at the zero lower bound. *Hutchins Center Working Paper*, 5.
- Greenwood, R., Hanson, S., and Stein, J. (2015). A comparative-advantage approach to government debt maturity. *Journal of Finance*, 70(4):1683–1722.
- Gurkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.
- Haldane, A. G., Roberts-Sklar, M., Wieladek, T., and Young, C. (2016). QE: The story so far. *Bank of England Working Paper*, (624).
- Joyce, M. A. S., Liu, Z., and Tonks, I. (2014). Institutional investor portfolio allocation, quantitative easing and the global financial crisis. *Bank of England Working Paper*, 210.
- Kuttner, K. (2006). Can central banks target bond prices? *NBER Working Paper*, (12454).
- Leeper, E. M. and Nason, J. M. (2008). *The New Palgrave Dictionary of Economics*. Edited by Steven N. Durlauf and Lawrence E. Blume, 2nd edition.
- Lenza, M., Pill, H., and Reichlin, L. (2010). Monetary policy in exceptional times. *Economic Policy*, 25:295–339.

- Lloyd, S. P. (2016). Unconventional monetary policy and the interest rate channel: Signalling and portfolio rebalancing. *University of Cambridge, mimeo*.
- Pascual, G. and Wieladek, T. (2016). The European Central Bank’s QE: A new hope. *CEPR Discussion Papers*, (11309).
- Peersman, G. and Smets, F. (2001). The monetary transmission mechanism in the euro area: more evidence from VAR analysis. *ECB Working Paper*, 91.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72:821–852.
- Ramey, V. (2011). Identifying government spending shocks: It’s all in the timing. *Quarterly Journal of Economics*, 126.
- Rudebusch, G. and Williams, J. C. (2008). Revealing the secrets of the temple: The value of publishing central bank interest rate projections. *Asset Prices and Monetary Policy*.
- Shephard, N. (1994). Partial non-Gaussian state space. *Biometrika*, 81:115–131.
- Sims, C. (1980). Macroeconomics and reality. *Econometrica*, 48(1):1–48.
- Swanson, E. (2011). Let’s twist again: A high-frequency event-study analysis of Operation Twist and its implications for QE2. *Brookings Papers on Economic Activity*, 42(1):151–207.
- Swanson, E. T. (2017). Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. *NBER Working Papers*, 23311.
- Tobin, J. (1956). Liquidity preference as behaviour towards risk. *Review of Economic Studies*, 25(2):65–86.
- Tobin, J. (1963). An essay on the principles of debt management. *Fiscal and debt management policies*.
- Treasury, H. M. (2009). *Debt and reserves management report 2009-10*. United Kingdom Debt Management Office, April edition.
- Turner, P. (2014). The exit from non-conventional monetary policy: What challenges? *BIS Working Papers*, 448.

- Volcker, P. A. (2002). “Monetary policy transmission: Past and future challenges,” Remarks at The Conference on Financial Innovation and Monetary Transmission, Federal Reserve Bank of New York.
- Wallace, N. (1981). A Modigliani-Miller theorem for open-market operations. *American Economic Review*, 71(3):267–274.
- Weale, M. and Wieladek, T. (2016). What are the macroeconomic effects of asset purchases? *Journal of Monetary Economics*, 79:81–93.
- Williams, J. C. (2014). Monetary policy at the zero lower bound: Putting theory into practice. *Hutchins Center Working Paper*, 3.
- Wu, J. C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.
- Yellen, J. L. (2017a). “The economic outlook and the conduct of monetary policy,” Remarks at the Stanford Institute for Economic Policy Research, Stanford University, Stanford, California.
- Yellen, J. L. (2017b). “The goals of monetary policy and how we pursue them,” Remarks at the Commonwealth Club, San Francisco, California.