BANK OF ENGLAND

Staff Working Paper No. 678
Optimal quantitative easing

Richard Harrison

September 2017

BANK OF ENGLAND

Staff Working Paper No. 678

Optimal quantitative easing

Richard Harrison[(1)]

Abstract

I study optimal monetary policy in a simple New Keynesian model with portfolio adjustment costs. Purchases of long-term debt by the central bank (quantitative easing; 'QE') alter the average portfolio return and hence influence aggregate demand and inflation. The central bank chooses the short-term policy rate and QE to minimise a welfare-based loss function under discretion. Adoption of QE is rapid, with large-scale asset purchases triggered when the policy rate hits the zero bound, consistent with observed policy responses to the Global Financial Crisis. Optimal exit is gradual. Despite the presence of portfolio adjustment costs, a policy of 'permanent QE' in which the central bank holds a constant stock of long-term bonds does not improve welfare.

Key words: Quantitative easing, optimal monetary policy, zero lower bound.

JEL classification: E52, E58.

---

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

The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees. I am grateful to seminar participants at the Bank of England, Birkbeck and NIESR and to Yunus Aksoy, Peter Sinclair, Ron Smith, Stephen Wright and an anonymous referee for helpful comments on earlier drafts. The first draft of this paper was written during a visit to University College London and their hospitality is gratefully acknowledged.

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

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

ISSN 1749-9135 (on-line)

1 Introduction

Central bank purchases of long-term government debt – often called quantitative easing (QE) – have been deployed as a monetary policy tool since the depth of the Great Recession, when short-term policy rates became constrained at their effective lower bounds. The widespread use of an unconventional monetary policy instrument has spawned much research. Perhaps surprisingly, however, there has been relatively little investigation of the optimal conduct of monetary policy when QE is a policy instrument, though recent contributions include Darracq Pariès and Kühl (2016), Harrison (2012), Reis (2015b) and Woodford (2016).

In this paper I study the optimal use of QE alongside the short-term policy rate. I use a version of a textbook New Keynesian model,[1] extended to include a bond market friction, following Andrés et al. (2004) and Harrison (2012). The representative household faces portfolio adjustment costs when allocating its assets between short-term and long-term bonds. These adjustment costs create a wedge between returns on short-term and long-term bonds that can be influenced by changes in the relative supplies of assets, thus providing a role for QE as a policy instrument. The adjustment cost specification captures 'flow effects' of QE purchases: the effects of *changes* in the stock of long-term bonds held by the central bank. The portfolio adjustment costs are calibrated to match estimates of those effects by D'Amico and King (2013).

The policymaker in the model acts under discretion to minimise a loss function derived from a quadratic approximation to the welfare of the representative household. In addition to the standard New Keynesian terms in inflation and the output gap, the loss function includes terms in the quantitative easing instrument. These arise because QE has traction via welfare-reducing portfolio adjustment costs borne by households.

The model is solved using projection methods, accounting for the non-linearities generated by the zero bound on the short-term interest rate and the possibility that bounds may also apply to the QE instrument (for example, the central bank's holdings of long-term debt must be non-negative).

I study entry into and exit from a 'QE regime', defined as a period during which the central bank holds a positive stock of long-term bonds on its balance sheet. I find that entry into QE regimes can be rapid, with large scale asset purchases commencing as soon as the short-term policy rate hits the zero bound. Exit from QE is slower in order to mitigate the costs of changes in portfolios.

Relative to the case in which the only policy instrument is the short-term policy rate, use of QE reduces the welfare costs of fluctuations by around 50%. In this 'active QE' case, the central bank holds a positive stock of long-term bonds on average and the average long-term interest rate is below the average short-term rate.

These observations suggest that a policy of 'permanent QE' in which the central bank is instructed to hold a constant stock of long-term bonds on its balance sheet may mitigate the effects of the zero bound on the short-term interest rate, by increasing the average value of the short-term interest rate. I show that this is *not* the case.

While permanent QE does succeed in 'twisting' the term structure on average (the

[1] See, for example, Galí (2008) and Woodford (2003).

long-term rate falls and the short-term policy rate rises), this has little effect on average inflation expectations. Raising average inflation expectations requires agents to expect that the central bank will cushion the effects of future deflationary shocks by purchasing assets if those shocks are sufficiently large to force the short-term policy rate to the zero bound. A permanent QE policy does not have this property. The welfare gains of 'active QE' are therefore generated by an expectation effect.

I also study the effects of delegation schemes by allowing the central bank to use both instruments, but instructing the central bank to minimise a loss function that differs from the one derived from household welfare. As in similar analysis using textbook New Keynesian models, I find that a very small increase in the inflation target does improve welfare, but increasing average inflation beyond a small amount generates welfare costs that outweigh the benefits associated with hitting the zero bound less frequently.

Allowing active use of QE but instructing the central bank to target a positive average quantity of long-term bonds on its balance sheet does not improve welfare relative to the case in which the central bank may freely choose the scale of QE. As in the case of permanent QE, this result stems from the fact that the most powerful effects of QE arise from the expectations that it will be deployed when necessary, rather than the direct effects of central bank asset holdings on long-term bond returns.

The rest of this paper is organised as follows. Section 2 discusses the 'portfolio balance effect' through which QE operates in my model and relates it to the broader literature on QE. Section 3 presents the model. Section 4 analyses the optimal policy problem of a central bank tasked with using the short-term interest rate and QE to minimise a welfare-based loss function in a time-consistent manner. The results from the baseline parameterisation of the model are presented in Section 5. Section 6 examines the effects of delegating alternative loss functions to the central bank. Section 7 assesses the robustness of the results to alternative assumptions about key parameter values and Section 8 concludes.

2 The portfolio balance transmission mechanism

In an oft-quoted remark, former FOMC Chairman Ben Bernanke argued that "the trouble with QE is that it works in practice, but not in theory".[2] In this section, I argue that the so-called 'portfolio balance' mechanism has become the predominant channel through which most monetary policymakers believe that quantitative easing affects asset prices and the wider economy.

When quantitative easing was introduced as a response to the global financial crisis, there was uncertainty among policymakers over the channels through which the policy might operate and skepticism among academics over whether it would have any effect at all.[3] For example, Benford et al. (2009) document several possible channels through

[2] The comment was made during a discussion session at the Brookings Institution: Bernanke (2014).

[3] I focus on quantitative easing measures of the type introduced by several central banks in response to the Global Financial Crisis. The Bank of Japan introduced a range of (somewhat different) balance sheet measures much earlier, given that it encountered the zero bound in the late 1990s.

which quantitative easing might stimulate spending and inflation.[4] Academic skepticism over the likely effects of the policy was typified by the analysis of Eggertsson and Woodford (2003), who demonstrated that a change in the composition of households' portfolios would have no effect on equilibrium asset prices or allocations in a widely studied benchmark model.

A wide range of studies provided evidence that the quantitative easing policies enacted in response to the financial crisis increased asset prices and reduced longer-term interest rates.[5] Other studies attempted to estimate the macroeconomic effects of these changes in asset prices and yields, with a general consensus that QE interventions generated increases in output and inflation.[6]

Alongside the accumulating evidence, economists explored possible theoretical mechanisms that could give rise to such effects. From an asset pricing perspective, King (2015) notes that the neutrality results of Eggertsson and Woodford (2003) rely on the (common) assumption of an additively time separable utility function. This implies that the stochastic discount factor used to price assets depends only on consumption allocations across time. A broader class of utility functions imply that the stochastic discount factor also depends on the return on wealth (or the average portfolio return). In such cases, shifts in the composition of agents' portfolios can affect the average portfolio return and hence individual rates of return via the stochastic discount factor. King (2015) demonstrates that Epstein-Zinn-Weil preferences[7] and the 'preferred habitat' investor framework set out by Vayanos and Vila (2009) fit into the wider class of models in which portfolio composition affects asset prices.

The Vayanos and Vila (2009) model is an important contribution, as it provides a link with the strand of the macroeconomics literature, described below, to which the present paper contributes. The model features two types of agents, one of which has preferences for assets of a particular maturity which give rise to a downward sloping demand curve for the asset. The second type of agent is an arbitrageur, trading in all assets. The interaction of the two agents gives rise to an equilibrium in which changes in the supply of an asset of a particular maturity affects the price of that asset (through the downward sloping demand of preferred habitat investors) and the prices of other assets with similar maturities (through the effect of arbitrage).

In macroeconomics, there is a long tradition of studying the effects of portfolio allocations on asset prices (and vice versa), dating back at least to the work of James Tobin and coauthors.[8] The key assumption underpinning the theory was that the relative demand

[4] As well as the from the portfolio balance effect discussed in this section, the authors argue that the expansion of bank reserves generated by asset purchases may create conditions that encourage greater bank lending and that asset purchases may help to anchor inflation expectations close to target by signalling the central bank's commitment to returning inflation to that target.

[5] Notable examples include D'Amico and King (2013), Greenwood and Vayanos (2010, 2014), Joyce et al. (2011) and Krishnamurthy and Vissing-Jorgensen (2012).

[6] See, among many others, Baumeister and Benati (2013), Lenza et al. (2010), Kapetanios et al. (2012), Pesaran and Smith (2016) and Weale and Wieladek (2016).

[7] This specification of preferences has become a benchmark model for the case in which the elasticity of intertemporal substitution is distinct from the coefficient of relative risk aversion. See Epstein and Zin (1989) and Weil (1990).

[8] See, for example, Tobin (1956, 1969) and Tobin and Brainard (1963).

for alternative asset classes would depend on their relative prices or returns, because of imperfect substitutability:

> [A]ssets are assumed to be imperfect substitutes for each other in wealth-owners' portfolios. That is, an increase in the rate of return on any one asset will lead to an increase in the fraction of wealth held in that asset, and to a decrease or at most no change in the fraction held in every other asset. (Tobin and Brainard, 1963)

These models assumed a (primitive) relationship between relative yields and relative asset demands. Frankel (1985) showed that this type of asset demand could be derived as the solution to a Markowitz portfolio problem.[9] As King (2015) notes, this approach does not incorporate rational expectations, because the portfolio problem does not account for the fact that future asset prices (which determine the rates of return on some assets) will be determined in the same way as current asset prices.

The seminal work of Andrés et al. (2004) embedded portfolio adjustment costs into a New Keynesian rational expectations model to provide a more microfounded treatment of imperfect substitutability. The model echoes the finance approach of Vayanos and Vila (2009) and also features two types of agents.[10] Unconstrained households have access to both short-term and long-term bonds, paying a portfolio adjustment cost when investing in the latter. Arbitrage by these households equates the returns (accounting for adjustment costs) of the two bonds. Constrained households only have access to long-term bonds. The consumption of constrained households is influenced by changes in the price of long-term bonds, which can be driven by changes in their relative supply via the portfolio adjustment costs paid by unconstrained households.

The Andrés et al. (2004) model has been modified and extended in several directions. Harrison (2012) builds a representative agent model in which all households face portfolio adjustment costs. In such a setting, aggregate demand depends on the average returns of short-term and long-term bonds as in Andrés et al. (2004). However, there is no heterogeneity, so that the effect of long-term returns on aggregate demand depends on the (average) shares of long-term and short-term debt held by households rather than on the fraction of constrained households. Harrison (2012) argues that the representative agent framework is more tractable, in particular facilitating welfare analysis. Ellison and Tischbirek (2014) uses an indirect utility argument to directly impose portfolio balance terms in the asset pricing equations of the banks who manage portfolios on the behalf of households.[11]

Chen et al. (2012) develop a medium-scale model based on the Andrés et al. (2004) setup, estimate it on US data and use it to study the effects of the FOMC's Large Scale

[9]The investor's objective function is the expected return on the portfolio, less a term in the covariance across returns that captures risk aversion.

[10]Note, however, that the portfolio adjustment cost role for QE is somewhat different from the role generated by the effects on portfolio risk studied in the finance context by, for example, King (2015).

[11]In some ways, this approach has more similarities with the early models of Tobin and others, though the indirect utility approach does deliver cross equation restrictions on the asset pricing relationships.

Asset Purchase programmes.[12] Carlstrom et al. (2017) adopt a market segmentation approach, but assume leveraged financial intermediaries provide the channel through which households invest in long-term government debt.

Once QE programmes had been implemented and their effects observed, a consensus among monetary policymakers on the portfolio balance transmission channel seemed to emerge. For example, Bernanke (2010) argues that:

> The channels through which the Fed's purchases affect longer-term interest rates and financial conditions more generally have been subject to debate. I see the evidence as most favorable to the view that such purchases work primarily through the so-called portfolio balance channel, which holds that once short-term interest rates have reached zero, the Federal Reserves purchases of longer-term securities affect financial conditions by changing the quantity and mix of financial assets held by the public.

Of course, while there may be near consensus among monetary policymakers on the transmission channel of QE, the portfolio balance effect is not without critique.[13] Thornton (2014) challenges the empirical evidence on the effects of QE, finding little evidence of that QE operations had economically important effects on long-term bond yields.[14]

One alternative theory for the efficacy of QE is that it contains signals about the likely path for the short-term policy rate (Bauer and Rudebusch, 2014); another is that changes in the composition of the central bank balance sheet can be used to reduce the risk exposure of private agents (Farmer and Zabczyk, 2016). Other authors have focused on the liabilities side of the central bank balance sheet, arguing that QE operates through the expansion of central bank reserves associated with asset purchases (Aksoy and Basso, 2014; Reis, 2015b). My model can be seen as complementary to these models in many respects as it is possible that QE operates via several channels. However, my focus on a portfolio balance mechanism is prompted by the views of monetary policymakers cited above.

3 The model

This section provides an overview of the model which is based on Harrison (2012). More details of the derivation and the modifications relative to Harrison (2012) are presented in Appendix A.

[12] Canzoneri and Diba (2005) and Canzoneri et al. (2008, 2011) have explored models in which government bonds provide liquidity services so that the mix of assets held by private agents affects relative returns via liquidity premia. This complementary strand of the literature does not focus on the effects of quantitative easing *per se*.

[13] The discussion here focuses on quantitative easing operations in which the central bank purchases long-term government debt (the focus of this paper). In response to the financial crisis, some central banks also engaged in the purchase of private debt instruments. Such policies may be expected to operate through different channels and represent a complementary line of research. Important contributions to this research include Cúrdia and Woodford (2009), Del Negro et al. (2017) and Gertler and Karadi (2011).

[14] These findings chime with the argument that Federal Reserve purchases of US government debt constituted such a small fraction of total debt holdings that any portfolio balance effects would likely be very small (Bauer and Rudebusch, 2014; Cochrane, 2011).

Following Harrison (2012), I focus on a simple extension to the canonical New Keynesian model in order to highlight the marginal implications of introducing a new friction (portfolio adjustment costs) and hence the possibility of using an additional monetary policy instrument relative to a widely studied benchmark model. Using a relatively small scale model also facilitates the use of projection methods to solve the model, thus accounting for the *risk* that bounds on policy instruments may bind.

Several papers have studied QE using larger models featuring similar portfolio frictions: for example, Chen et al. (2012), Darracq Pariès and Kühl (2016), De Graeve and Theodoridis (2016), Hohberger et al. (2017) and Priftis and Vogel (2016). However, all of these papers assume that agents' expectations satisfy a certainty equivalence assumption.[15] With the exception of Darracq Pariès and Kühl (2016) and Quint and Rabanal (2017), these papers do not consider the optimal design of QE policies. Neither Darracq Pariès and Kühl (2016) nor Quint and Rabanal (2017) consider potential bounds on the QE instrument or use a welfare-based loss function.

3.1 Short-term and long-term bonds

There are two assets in the economy: short-term and long-term nominal government bonds. Following Woodford (2001), I model long-term government bonds as infinite maturity instruments, paying a geometrically declining coupon. Specifically, a bond issued at date t pays nominal coupons χ^s in dates $t+1+s$, $s \geq 0$. This modelling assumption is convenient because it implies that a one dollar holding of a bond issued j periods ago is equivalent to a χ^j dollar holding of a bond issued today. The fact that the values of long-term bonds issued at different dates can be linked in this way means that it is possible to write budget constraints in terms of a single bond price and a single stock of long-term bonds.[16]

As an illustration, consider first the budget constraint of a representative household:

$$V_t \widetilde{B}^h_{L,t} + B^h_t = (1+\chi V_t)\, \widetilde{B}^h_{L,t-1} + R_{t-1} B^h_{t-1} + W_t n_t + T_t + D_t - P_t c_t - \Psi_t \tag{1}$$

The right hand side of the budget constraint captures income from working n_t hours at wage W_t, net transfers/taxes T_t from the government and dividends D_t from firms, less spending on consumption goods c_t at price P_t and portfolio adjustment costs Ψ (discussed in Section 3.2). The household decision problem will be analyzed in detail below: here I focus on the role of short and long-term bonds.

The household holds one-period bonds B^h, which pay a gross rate of return R. The budget constraint with respect to short-term bonds is standard: bonds purchased at date $t-1$ mature in date t with a nominal payoff of R_{t-1} per bond.

The household also holds long-term bonds, where $\widetilde{B}^h_L$ denotes the nominal bonds held, measured in terms of the equivalent quantity of newly issued bonds. V is the nominal value (price) of each bond. The right hand side of the budget constraint contains the current value of existing holdings of the long-term bond. The quantity of long-term bonds purchased at all previous dates can be summarised in terms of a quantity

[15] While the rational expectations assumption specifies that shocks are zero in expectation, the certainty equivalence assumption specifies that these shocks are assumed (by agents) to be zero with certainty.

[16] See Woodford (2001) and Chen et al. (2012) for further discussion.

of bonds (newly) issued in the previous period by virtue of the pricing relationship discussed above. The bond holdings from the previous period $\widetilde{B}^h_{L,t-1}$ pay a coupon of 1 per bond and have a value equal to χV_t, reflecting the fact that the quantity $B_{L,t-1}$ of date $t-1$ issued bonds has the same value as a quantity $\chi B_{L,t-1}$ of date t issued bonds.

The budget constraint can be conveniently re-written in terms of the one-period return on long-term bonds:

$$B^h_{L,t} + B^h_t = R^1_{L,t} B^h_{L,t-1} + R_{t-1} B^h_{t-1} + W_t n_t + T_t + D_t - P_t c_t - \Psi_t \tag{2}$$

where:

$$B^h_{L,t} \equiv V_t \widetilde{B}^h_{L,t}$$
$$R^1_{L,t} \equiv \frac{1 + \chi V_t}{V_{t-1}}$$

This formulation treats the choice variables of the household as the *value* of long-term bond holdings. Because households take bond prices as given this is isomorphic to the original formulation, but simplifies the subsequent algebra. Similarly, the one period return is simply a definition expressed in terms of other asset prices which simplifies the derivation. While the one-period return on long-term bonds is a sufficient statistic to characterize household behavior in the model, it is possible to map the implications for the one-period return back to bond yields that are more readily compared with the data, as shown below.

3.2 Households

The optimisation problem of the representative household is

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\frac{1}{\sigma}} - 1}{1 - \frac{1}{\sigma}} - \frac{\phi_t n_t^{1+\psi}}{1+\psi} \right\}$$

where c is consumption and n is hours worked. A preference shock ϕ_t is included and will serve as the 'demand shock' that generates a persistent decline in the natural real interest rate considered in the simulation experiments examined below.

Maximisation is subject to the budget constraint (2), including an explicit formulation of portfolio adjustment costs, Ψ:[17]

$$\begin{aligned} B^h_{L,t} + B^h_t = {} & R^1_{L,t} B^h_{L,t-1} + R_{t-1} B^h_{t-1} + W_t n_t + T_t + D_t - P_t c_t \\ & - \frac{\tilde{\nu} P_t \left(b^h + b^h_L\right)}{2} \left[\delta \frac{B^h_t}{B^h_{L,t}} - 1\right]^2 - \frac{\tilde{\xi} P_t \left(b^h + b^h_L\right)}{2} \left[\frac{B^h_t / B^h_{L,t}}{B^h_{t-1} / B^h_{L,t-1}} - 1\right]^2 \end{aligned} \tag{3}$$

The portfolio adjustment costs have two components. The first component is a function of the deviation of the households 'portfolio mix', $\frac{B^h_t}{B^h_{L,t}}$ from their desired level, δ^{-1}.

[17] Here b^h and b^h_L denote the steady state *real* levels of short-term and long-term bonds.

These adjustment costs are intended to capture 'stock effects': shifts in the supply of these assets can have a direct effect on their price. Following Andrés et al. (2004), δ is set equal to the steady-state ratio of long-term bonds to short-term bonds so that these portfolio costs are zero at the non-stochastic steady state.

The second component of the portfolio adjustment costs is a function of the *change* in the household's portfolio mix. This adjustment cost is motivated by the empirical evidence that changes in asset supplies associated with the auctions that implement asset purchases have an effect on the prices of assets purchased and their close substitutes (see D'Amico and King, 2013). In the context of my model, such 'flow' effects may be interpreted in terms of frictions in adjusting portfolios including transactions costs.

The tractability of this type of adjustment costs has led to their adoption in a range of monetary models.[18] In reality, transactions costs are likely to be low, so the portfolio adjustment costs in the model are a stand in for a broader range of frictions. Andrés et al. (2004) argue that they represent a perception by households that longer-term bonds are riskier than short-term bonds, such that households' require a greater quantity of liquid assets (in their model, money) as compensation. Cast in this way, these costs may be better suited to inclusion in the utility function. However, Harrison (2012) demonstrates that taking this approach gives rise to isomorphic expressions for the model equations and welfare functions. Moreover, frictions in financial intermediation can give rise to very similar behavioural equations (Carlstrom et al., 2017; Harrison, 2011).

As shown in Appendix A, the household's first order conditions with respect to consumption, short-term bonds and long-term bonds can be log-linearised around the steady state and combined to give:

$$\hat{c}_t = \mathbb{E}_t \hat{c}_{t+1} - \sigma \left[\frac{1}{1+\delta} \hat{R}_t + \frac{\delta}{1+\delta} \mathbb{E}_t \hat{R}^1_{L,t+1} - \mathbb{E}_t \hat{\pi}_{t+1} \right] \tag{4}$$

$$\begin{aligned} \mathbb{E}_t \hat{R}^1_{L,t+1} =& \hat{R}_t - \nu (1+\delta) \left[\hat{b}^h_t - \hat{b}^h_{L,t} \right] \\ & - \xi \delta^{-1} (1+\delta) \left[\Delta \left(\hat{b}^h_t - \hat{b}^h_{L,t} \right) - \beta \mathbb{E}_t \Delta \left(\hat{b}^h_{t+1} - \hat{b}^h_{L,t+1} \right) \right] \end{aligned} \tag{5}$$

where $\nu \equiv \tilde{\nu}(1+\delta)$ and $\xi \equiv \tilde{\xi}(1+\delta)$ and $\hat{z}_t \equiv \ln(z_t/z)$ denotes the log-deviation of variable z_t from its non-stochastic steady state, z.[19]

The Euler equation (4) demonstrates that aggregate demand is driven by a weighted average of the interest rates on short-term and long-term bonds. The pricing equation for long-term bonds (5) indicates that aggregate demand therefore also depends on the household's relative holdings of short-term and long-term bonds. An increase in the household's relative holdings of short-term bonds acts like a reduction in the short-term real interest rate and boosts demand. According to the interpretation of Andrés et al. (2004), an increase in relative holdings of short-term bonds represents an increase in households' (marginal) liquidity. So a shift towards short-term bond holdings reduces

[18] See, for example, Andrés et al. (2004), Chen et al. (2012), Gertler and Karadi (2013) and Darracq Pariès and Kühl (2016).

[19] Real valued debt stocks are denoted using lower case letters, so that $b^h_t \equiv B^h_t / P_t$ and $b^h_{L,t} \equiv B^h_{L,t} / P_t$.

the wedge between the rates of return on long-term and short-term bonds, as shown in equation (5).

3.3 Firms

There is a set of monopolistically competitive producers indexed by $j \in (0,1)$ that produce differentiated products that form a Dixit-Stiglitz bundle that is purchased by households. Preferences over differentiated products are given by

$$y_t = \left[\int_0^1 y_{j,t}^{1-\eta_t^{-1}} dj \right]^{\frac{1}{1-\eta_t^{-1}}}$$

where y_j is firm j's output. The elasticity of demand among consumption varieties η_t is assumed to be time-varying, which generates a 'cost push' shock in the Phillips curve that characterises log-linear pricing decisions.

Firms produce using a constant returns production function in the single input (labour):

$$y_{j,t} = A n_{j,t}$$

where A is a productivity parameter.

The real profit of producer j is:

$$\frac{(1+s)\, P_{j,t}}{P_t} y_{j,t} - w_t n_{j,t} = \left((1+s) \frac{P_{j,t}}{P_t} - \frac{w_t}{A} \right) \left(\frac{P_{j,t}}{P_t} \right)^{-\eta_t} y_t$$

where s is a subsidy paid to producers in order to ensure that the steady-state level of output is efficient. This assumption permits the use of a quadratic approximation of the household utility function as the appropriate welfare criterion (see Benigno and Woodford (2006)).

Under a Calvo (1983) pricing scheme the objective function for a producer that is able to reset prices is:

$$\max \mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k \left(\beta \alpha \right)^{k-t} \left((1+s) \frac{P_{j,t}}{P_k} - \frac{w_k}{A} \right) \left(\frac{P_{j,t}}{P_k} \right)^{-\eta_t} y_k$$

where Λ represents the household's stochastic discount factor and $0 \leq \alpha < 1$ is the probability that the producer is *not* allowed to reset its price each period.

Well-known manipulations (presented in Appendix A.2) lead to a New Keynesian Phillips Curve:

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{y}_t + u_t$$

where

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

and u_t is the cost push shock (a linear function of log-deviations of the demand elasticity from steady state). The cost push shock is assumed to evolve according to a first order autoregressive process:

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u \tag{6}$$

3.4 Government and monetary policies

To focus on the role of monetary policy, fiscal policy is highly simplified. There is no government spending and net transfers to households are lump sum.

Given the specification of the long-term bond discussed in Section 3.1, the nominal government budget constraint is:

$$B_t + V_t \widetilde{B}_{L,t} = R_{t-1} B_{t-1} + (1 + \chi V_t)\, \widetilde{B}_{L,t-1} + Z_t - P_t \tau_t$$

where B and $\widetilde{B}_L$ represent stocks of short-term and long-term debt, Z denotes net asset purchases by the central bank and τ represents net tax/transfer payments from/to households. The inclusion of Z reflects the assumption that QE is financed by the central government, as discussed below.

Applying the same change of variables as above allows the constraint to be expressed in terms of the value of long-term bonds and their one-period return:

$$B_t + B_{L,t} = R_{t-1} B_{t-1} + R^1_{L,t} B_{L,t-1} + Z_t - P_t \tau_t \tag{7}$$

The government implements the following debt issuance policies:

$$\frac{B_t}{P_t} \equiv b_t = b > 0, \ \forall t \tag{8}$$

$$\frac{B_{L,t}}{P_t} \equiv b_{L,t} = \delta b, \ \forall t \tag{9}$$

These issuance policies will ensure that – absent QE operations by the central bank – household achieve their desired (or target) portfolio positions. Conditional on these issuance policies and QE by the central bank, net transfers to households T are pinned down by the government budget constraint (7). In particular, changes in the value of the total government debt stock are transferred to/from households (lump sum) in order to keep the overall value of debt constant over time (a form of balanced budget financing).

Net purchases of long-term government bonds by the central bank are:[20]

$$Z_t = V_t \widetilde{Q}_t - (1 + \chi V_t)\, \widetilde{Q}_{t-1}$$

where the quantity of long-term bonds purchased by the central bank is denoted by $\widetilde{Q}$ and it is assumed that coupon payments are paid to the central bank. The now familiar change of variables (so that $Q_t \equiv V_t \widetilde{Q}_t$) can be used to express Z in terms of the value of long-term bonds purchased by the central bank:

$$Z_t = Q_t - R^1_{L,t} Q_{t-1} \tag{10}$$

I define the policy instrument as the fraction of the market value of long-term bonds purchased by the central bank, denoted q:

$$Q_t = q_t B_{L,t}$$

[20] In a model with money, the net expansion in the monetary base would also be included in this expression. Here, QE is financed by a loan from the central government, which must ultimately be financed by lump sum taxes on households.

Monetary policy is therefore conducted using the short-term nominal interest rate (R) and the fraction of long-term bonds held on the central bank's balance sheet (q). Section 4 considers the case in which monetary policy is set according to optimal discretion.

Quantitative easing, by its nature, is a prime candidate for study from the perspective of monetary and fiscal policy interactions. For example, Del Negro and Sims (2015) and Benigno and Nistico (2015) take such an approach to examine the potential importance of the government and central bank intertemporal budget constraints. My assumptions abstract from these considerations entirely.

Two aspects of my assumptions are intended to make quantitative easing an exclusively monetary policy operation. By assuming that the fiscal authority adopts a 'neutral' debt management policy (so that relative bond supplies are kept always in line with the desired holdings of households), the monetary policymaker has maximal control over the debt stocks actually held by households. By assuming that the total value of debt is fixed, I ensure that fiscal policy is active.[21] In particular, losses and gains on the central bank's asset portfolios are immediately financed/rebated to private agents via lump sum taxes/transfers. As Benigno and Nistico (2015) point out, such a setup implies that only the consolidated government/central bank budget constraint matters for allocations. As a result, the only non-neutrality from QE operates through the portfolio balance channel.

To the extent that fiscal policy does not deliver the optimal mix of assets for households on average, my model would imply a role for QE in normal times (away from the zero bound). However, my assumptions are an attempt to capture the key elements of institutional arrangements in practice. For example, government treasury departments (or their agents) are tasked with actively manage the maturity structure of government debt. Their mandate is typically expressed in terms of achieving favourable financing conditions for the government and ensuring adequate liquidity in government debt markets. In the context of my model, debt issuance in line with household portfolio preferences would (other things equal) minimise portfolio adjustment costs and hence the (social) costs of financing a given debt stock.

My assumptions also require that debt management policy remains unchanged when monetary policy uses QE when constrained by the zero bound. There is an active debate on the extent to which US government debt issuance may have offset some effects of FOMC asset purchases (see, for example, Greenwood et al., 2015). However, my assumptions are consistent with the institutional arrangements for QE in the United Kingdom, where the Debt Management Office was instructed to ensure that debt management operations "be consistent with the aims of monetary policy" including the asset purchases implemented by the Bank of England's Monetary Policy Committee.[22]

[21] Tax revenues are adjusted to hold the debt stock constant which ensures that the government's intertemporal budget constraint is satisfied.

[22] The quotation is from the letter from the Chancellor of the Exchequer to the Governor of the Bank of England, 3 March 2009: http://www.bankofengland.co.uk/monetarypolicy/Documents/pdf/chancellorletter050309.pdf.

3.5 Market clearing

Market clearing for short term bonds implies that:

$$\frac{B_t}{P_t} \equiv b_t^h = b$$

Market clearing for long-term bonds requires that:

$$\frac{Q_t}{P_t} + b_{L,t}^h = \frac{B_{L,t}}{P_t}$$

where $b_{L,t}^h \equiv B_{L,t}^h / P_t$.

Combining the government debt issuance policy with the specification of the QE instrument q gives:

$$b_{L,t}^h = (1 - q_t)\,\delta^{-1} b$$

In log-linear terms we have:

$$\hat{b}_t^h - \hat{b}_{L,t}^h = -\hat{b}_{L,t}^h = q_t \tag{11}$$

where the equation is linearised (rather than log-linearised) with respect to q.

Equation (11) shows that asset purchases (q) influence the quantity of long bonds available to households and hence long-term bond yields via (5).

3.6 Model equations

As shown in Appendix A, the log-linearised model can be reduced to an Euler equation for the output gap ($\hat{x}$) and a Phillips curve:

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \gamma q_t + \xi q_{t-1} + \beta \xi \mathbb{E}_t q_{t+1} - r_t^* \right] \tag{12}$$

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t \tag{13}$$

where $\gamma \equiv \nu\delta + \xi(1+\beta)$ and the 'natural rate of interest' is $r_t^* \equiv -\mathbb{E}_t\left(\hat{\phi}_{t+1} - \hat{\phi}_t\right)$.

As shown in Appendix C, the yield to maturity of the long-term bond is given by:

$$\hat{\mathcal{R}}_t = \chi\beta \mathbb{E}_t \hat{\mathcal{R}}_{t+1} + (1-\chi\beta) \begin{pmatrix} \hat{R}_t - \delta^{-1}(1+\delta)\gamma q_t \\ +\xi\delta^{-1}(1+\delta) q_{t-1} + \beta\xi\delta^{-1}(1+\delta)\mathbb{E}_t q_{t+1} \end{pmatrix} \tag{14}$$

The shock processes are:

$$r_t^* = \rho_r r_{t-1}^* + \varepsilon_t^r \tag{15}$$

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u \tag{16}$$

where $\varepsilon_t^r \sim N(0, \sigma_r^2)$ and $\varepsilon_t^u \sim N(0, \sigma_u^2)$.

Table 1: Parameter values

	Description	Value
σ	Elasticity of intertemporal substitution	1
κ	Slope of Phillips curve	0.0516
β	Discount factor	0.9918
ρ_r	Autocorrelation of natural rate	0.85
σ_r	Standard deviation of natural rate	0.25
ρ_u	Autocorrelation of cost push shock	0
σ_u	Standard deviation of cost push shock	0.154
η	Consumption bundle elasticity of substitution	7.66
α	Calvo probability of *not* changing price	0.855
ψ	Inverse labour supply elasticity	1
χ	Long-term bond coupon decay rate	0.975
δ	Steady-state ratio of long-term bonds to short-term bonds	0.3
$b + b_L$	Total debt stock (relative to GDP)	2
ν	Elasticity of long-term bond rate with respect to portfolio mix	0.35
ξ	Elasticity of long-term bond rate with respect to change in portfolio mix	3.2
$\underline{q}$	Lower bound on QE	0
$\bar{q}$	Upper bound on QE	0.5

3.7 Parameter values

Table 1 shows the baseline parameter values.[23]

The key parameters of the aggregate demand and pricing equations are σ and κ. Setting $\sigma = 1$ is a standard assumption in the literature. Many studies that examine optimal policy at the zero bound use a much higher value (Adam and Billi, 2006; Bodenstein et al., 2012; Levin et al., 2010, among others, use a value of 6 or more). As Levin et al. (2010) point out, such calibrations are often required to generate significant effects on output at the zero bound under optimal *commitment* policy in a canonical New Keynesian model. As I focus on the case of optimal discretionary policy, the zero bound has substantial effects even with a value for σ that is more in line with empirical evidence (such as that presented by Guvenen, 2006). The slope of the Phillips curve ($\kappa = 0.0516$), though larger than the values used in similar studies (typically around 0.02–0.024), is consistent with my choice of a lower value for σ, given the values for the other parameters.[24]

The value of β is chosen to be consistent with a real interest rate of 3.35% in the non-

[23] The productivity parameter A is chosen to normalise output to unity in the steady state.

[24] Given the assumed elasticity of disutility of labour supply ($\psi = 1$), achieving this value of κ requires setting $\alpha = 0.855$. This high degree of price stickiness is consistent with estimates from macroeconomic models such as Smets and Wouters (2005). More plausible estimates of average contract length can be obtained by adopting more flexible formulations of the demand for alternative product varieties as demonstrated by Smets and Wouters (2007). The value of $\eta = 7.66$ is commonly used in the canonical New Keynesian model (see, for example, Adam and Billi, 2006; Bodenstein et al., 2012).

stochastic steady state. As shown by Adam and Billi (2007), as β increases, the steady-state real interest rate falls and so the chances of encountering the zero bound (and the costs associated with hitting it) increase.[25]

Other parameters that are important in determining the incidence of the zero bound are those governing the shock processes. The process for the natural real interest rate is assumed to be persistent, with $\rho_r = 0.85$, following Levin et al. (2010). The standard deviation of the shock is roughly in line with the value used by Adam and Billi (2006) in their 'RBC calibration' and the values of the parameters governing the cost push shock are also taken from that calibration.

The parameters related to long-term and short-term bonds deserve particular attention. The value of χ is chosen to imply that the long-term bond has a duration of between 7 and 8 years in the non-stochastic steady state (see Appendix C). This corresponds to the average duration of 10-year US Treasuries at the time of the first large scale asset purchase programme (D'Amico and King, 2013). I therefore interpret the long-term bond as a 10-year bond for the purposes comparing the model predictions with the data. The steady-state ratio of long-term to short-term bonds (δ) is set to 0.3 on the basis of the data presented in D'Amico and King (2013)[26] and is also consistent with the longer-term data presented by Kuttner (2006).[27]

The values for the parameters governing the portfolio adjustment costs, ν and ξ are designed to capture the empirical effects of quantitative easing, defined as 'stock effects' and 'flow effects' by D'Amico and King (2013). To arrive at these parameter values the model was solved on a grid of $\{\nu, \xi\}$ pairs and the values that generated stock and flow effects closest to those estimated by D'Amico and King (2013) selected. This procedure and the results are discussed in Section 5.1. Andrés et al. (2004) estimate a parameter similar to ν (relating the long-term bond premium to household's relative holdings of money and long bonds) using US data. Their estimate implies a value of ν of around 0.035, though the long-term rate in that study is a three-year bond, a somewhat shorter maturity than the focus of my model. The evidence presented in Bernanke et al. (2004) would, using a simple back of the envelope calculation, suggest a much larger value for $\nu \approx 2$.[28] Of course, such calculations ignore the fact that asset purchases will have effects

[25] My calibration implies a lower non-stochastic steady-state real interest rate than previous studies, including Adam and Billi (2007). Nevertheless, this value may be considered rather high, even by pre-crisis standards. The calibration is best thought of as an assumption about the non-stochastic steady-state *nominal* interest rate, because the efficient inflation rate in the model is zero.

[26] The ratio can be inferred from the data on the dollar amounts and percentages of stock purchased displayed in D'Amico and King (2013, Fig 1) where short-term bonds are interpreted as those with an outstanding maturity of six years or less.

[27] Kuttner (2006, Figure 3) shows that the average fraction of short-term (less than five year maturity) bonds held by the private sector was around 25% over the period from 1965 to 2006, suggesting a value of δ around one third. Debt management strategies differ across countries, of course: the data underlying Figure 1 suggests $\delta > 1$ for the United Kingdom, for example.

[28] This is calculated using a steady-state version of equation (14), assuming that an asset purchase operation of size $q = 0.1$ is permanent and there are no effects on short-term rates long-term bond returns. To see this note that, since $\gamma \equiv \nu\delta + \xi(1+\beta)$, equation (14) can be written as
$\hat{\mathcal{R}}_t = \chi\beta\mathbb{E}_t\hat{\mathcal{R}}_{t+1} + (1-\chi\beta)\left(\hat{R}_t - (1+\delta)\nu q_t - \xi\delta^{-1}(1+\delta)\Delta q_t + \beta\xi\delta^{-1}(1+\delta)\mathbb{E}_t\Delta q_{t+1}\right)$. A 'steady state' version of the equation sets $\hat{z}_t = \hat{z}, \forall t$ so that $\hat{\mathcal{R}} = \hat{R} - (1+\delta)\nu q$ and hence $\frac{\partial\hat{\mathcal{R}}}{\partial q} = -(1+\delta)\nu$. Since $\hat{\mathcal{R}}$ is

on other asset prices (in particular, expected short-term rates). The simulation approach discussed in Section 5.1 attempts to overcome these issues.

Finally, the parameters $\underline{q}$ and $\bar{q}$ represent the lower and upper bounds on the scale of QE operations that the central bank may undertake. Recall that q represents the fraction of the total quantity of outstanding long-term bonds held by the central bank. Under the assumption that the central bank cannot issue long-term bonds that are perfect substitutes for long-term government bonds, $q_t \geq 0$, and I set $\underline{q} = 0$.

It must also be the case that $\bar{q} \leq 1$, since the central bank cannot purchase more than 100% of the existing bonds. There may be practical reasons why the upper bound on asset purchases is less than 1, for example if there are some financial institutions that must hold long-term safe assets for regulatory purposes. In addition, if the central bank balance sheet is considered independently from the government, then the size of the balance sheet may be limited by a solvency constraint.[29]

Figure 1: Approximate measure of q for the United Kingdom

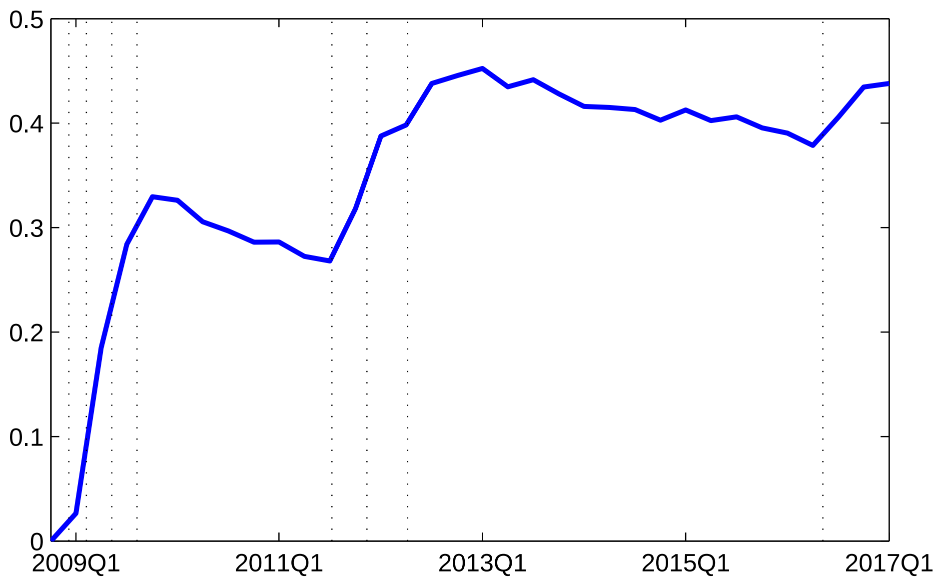

Notes: The figure shows the ratio of the value of the Bank of England's asset purchase facility (APF) to the value of outstanding medium-term and long-term UK government debt. Vertical dotted lines indicate dates at which the Monetary Policy Committee voted to increase the size of the APF.
Sources: Bank of England; UK Debt Management Office.

These types of friction are not explicitly incorporated in the model, but I set $\bar{q} = 0.5$ as a way of capturing them. This value is based on QE in the United Kingdom, which purchased around half of the long-term debt stock (Figure 1).[30] Quantitative easing programmes in the United States, while substantial, represented a much smaller share of the long-term government debt market. Section 7.2 examines the robustness of the results

measured in quarterly units, we require $0.1 \times (1 + \delta)\, \nu = 0.25$ which implies $\nu \approx 2$.

[29]The issue depends on the financing agreement between the central bank and government, as explored by Benigno and Nistico (2015).

[30]Figure 1 is consistent with the results in Daines et al. (2012), who estimate that the first phase of QE in the United Kingdom purchased around 30% of the long-term government debt stock.

to the assumed value of $\bar{q}$ and the implications for profits and losses associated with the central bank's asset portfolio.

4 The monetary policy problem

In this section, I consider the optimal use of QE alongside the short-term policy rate. I assume that the monetary policymaker sets both instruments to minimise a loss function based on an approximation to the utility of the representative household. As in Harrison (2012), this loss function includes terms in the QE instrument (q), reflecting the fact that the portfolio frictions that give QE traction impose costs on households.

Indeed, Alla et al. (2016) argue that welfare-based loss functions for models that feature a wide range of unconventional policy instruments (for example, including foreign exchange intervention) should include terms in the variability of those instruments for this reason. Using an ad hoc loss function to study the optimal use of QE (as in, for example, Darracq Pariès and Kühl, 2016) may fail to capture the full welfare costs of policy actions, which may in turn determine some of the policy prescriptions.

Appendix B demonstrates that a loss function based on a quadratic approximation to household utility is given by:

$$\mathcal{L} = \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left(\omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2 + \omega_q q_t^2 + \omega_{\Delta q} \left(q_t - q_{t-1} \right)^2 \right) \tag{17}$$

where the weights are related to the model parameters according to:

$$\begin{aligned} \omega_x &\equiv \left(\psi + \sigma^{-1} \right) \\ \omega_\pi &\equiv \frac{\alpha \eta}{(1-\alpha\beta)(1-\alpha)} \\ \omega_q &\equiv \tilde{\nu} \left(b^h + b_L^h \right) \\ \omega_{\Delta q} &\equiv \tilde{\xi} \left(b^h + b_L^h \right) \end{aligned}$$

The loss function specifies that the policymaker seeks to stabilise the output gap, inflation and the extent of its quantitative easing policy. The first two terms in parentheses appear in the welfare-based loss function of the canonical New Keynesian model.[31] The third and fourth terms appear because of the introduction of imperfect substitutability between assets. This additional friction can be mitigated by stabilising the relative supplies of assets and the rate at which portfolio shares change. Given the assumption that the maturity structure of government debt issuance is matched to the preferred portfolio mix of households, deviations in the relative supplies of assets are due entirely to quantitative easing decisions (q_t) by the central bank.

The policymaker minimises the loss function (17) subject to (12), (13) and the relevant

[31] See Woodford (2003).

constraints on the policy instruments:

$$\hat{R}_t \geq 1 - \beta^{-1} \tag{18}$$
$$q_t \geq \underline{q} \tag{19}$$
$$q_t \leq \bar{q} \tag{20}$$

I assume that there is no commitment technology that allows the policymaker to make credible promises about future policy actions. Examining the case in which policy operates under 'discretion' is motivated by two considerations. The first is that, for the class of models I consider, the zero lower bound on nominal interest rates does not pose a substantial problem if the policymaker is able to make commitments about how the short-term policy rate will be set in the future (Eggertsson and Woodford, 2003; Adam and Billi, 2006).[32] The second reason is that many central bankers have expressed doubts over their ability to credibly commit to future policy actions (Nakata, 2015).[33]

In this discretionary setting, the policymaker at date t is treated as a Stackelberg leader with respect to both private agents at date t and policymakers (and private agents) in dates $t+i, i \geq 1$. I seek a Markov perfect policy in which optimal decisions are a function only of the relevant state variables in the model ($\{u_t, r_t^*, q_{t-1}\}$). Under this interpretation, the policymaker understands that future policymakers will choose allocations according to time-invariant Markovian policy functions and therefore that its current policy decisions affect future outcomes through their impact on the endogenous state variable (q).

Appendix E shows that the first order conditions of the policymaker's problem are given by:

$$0 = \omega_\pi \hat{\pi}_t - \lambda_t^\pi \tag{21}$$
$$0 = \omega_x \hat{x}_t + \kappa \lambda_t^\pi - \lambda_t^x \tag{22}$$
$$\begin{aligned} 0 = & \ \omega_q q_t + \omega_{\Delta q}(q_t - q_{t-1}) + \beta \frac{\partial \mathbb{E}_t \mathcal{L}_{t+1}}{\partial q_t} + \beta \frac{\partial \mathbb{E}_t \pi_{t+1}}{\partial q_t} \lambda_t^\pi \\ & + \left[\frac{\partial \mathbb{E}_t x_{t+1}}{\partial q_t} + \sigma \frac{\partial \mathbb{E}_t \pi_{t+1}}{\partial q_t} + \sigma\gamma - \beta\sigma\xi \frac{\partial \mathbb{E}_t q_{t+1}}{\partial q_t} \right] \lambda_t^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}} \end{aligned} \tag{23}$$
$$0 = -\sigma \lambda_t^x - \lambda_t^R \tag{24}$$

where $\lambda_t^x, \lambda_t^\pi, \lambda_t^R, \lambda_t^{\underline{q}}, \lambda_t^{\bar{q}}$ are the Lagrange multipliers on the constraints (12), (13), (18), (19) and (20) respectively.[34]

The first order condition for quantitative easing (23) shows that the policymaker accounts for the fact in which the choice of QE today may have effects on welfare and future

[32] Levin et al. (2010) point out that if aggregate demand is very sensitive to real interest rates, then the zero bound can be costly, even under commitment. Harrison (2012) studies optimal quantitative easing under commitment in a model with such a calibration.

[33] This evidence is consistent with the observation that, in general, QE was used as a policy tool somewhat earlier than explicit forward guidance. Moreover, even when forward guidance was deployed, there was much debate over the extent to which it represented a commitment by policymakers (see, for example, Plosser, 2012).

[34] Appendix E also reports the required Kuhn-Tucker conditions for the multipliers on the inequality constraints.

outcomes because future policymakers will inherit the stock of QE chosen today. In the case that the optimal level of QE is an interior solution (that is, $q_t \in (\underline{q}, \bar{q})$), (23) can be written as:

$$q_t = \frac{\omega_{\Delta q}}{\omega_q + \omega_{\Delta q}} q_{t-1} - \frac{\beta}{\omega_q + \omega_{\Delta q}} \frac{\partial \mathbb{E}_t \mathcal{L}_{t+1}}{\partial q_t} - \frac{\beta}{\omega_q + \omega_{\Delta q}} \frac{\partial \mathbb{E}_t \pi_{t+1}}{\partial q_t} \omega_\pi \pi_t$$
$$- \frac{1}{\omega_q + \omega_{\Delta q}} \left[\frac{\partial \mathbb{E}_t x_{t+1}}{\partial q_t} + \sigma \frac{\partial \mathbb{E}_t \pi_{t+1}}{\partial q_t} + \sigma\gamma - \beta\sigma\xi \frac{\partial \mathbb{E}_t q_{t+1}}{\partial q_t} \right] (\omega_x x_t + \kappa \omega_\pi \pi_t)$$

which shows that current QE will be larger if the policymaker inherits a larger initial stock of QE and if additional QE reduces losses in the next period (the first two terms on the right hand side). The policymaker's current choice of QE will also affect current losses via the effects on private agents' expectations and hence current decisions. The third term, for example, captures the effect of QE on current losses via the effect of QE on inflation expectations and hence current inflation choices through the Phillips curve (13).

This expression for optimal QE also highlights the importance of 'flow effects' of portfolio changes. If these flow effects are absent, then $\tilde{\xi} = \xi = \omega_{\Delta q} = 0$. Moreover, q ceases to be a state variable in the model, so that current choices of QE have no effect on expectations or future losses. In this case, for an interior solution for q_t, the first order condition becomes:

$$q_t = -\frac{\sigma\gamma}{\omega_q} (\omega_x x_t + \kappa \omega_\pi \pi_t)$$

so that the choice of depends only on the current output gap and inflation. This condition balances the marginal cost of QE ($\omega_q q_t$) with the marginal benefits of improved output gap and inflation stabilisation via the effect of QE through the IS equation, (12). Crucially, for this equation to hold, the IS equation must be an active constraint on policy choices. For this to be the case, we require that λ_t^x be non-zero which (from (24) and the Kuhn-Tucker conditions) requires that the zero bound on the short-term policy rate must be binding. A corollary of this observation is that when the zero bound is not binding, the policymaker does not use the QE instrument: in this case we have $\omega_x x_t + \kappa \omega_\pi \pi_t = 0$ and hence $q_t = 0$. The logic of this result is simple. The policymaker has access to two instruments that affect the output gap in the same way, but one is costly to operate (since q appears in the loss function). When the zero bound on the short-term interest rate is not binding, the policymaker is unconstrained in her ability to choose the output gap by her choice of the short term interest rate and hence will choose not to use the costly instrument.

4.1 Solution approach

To capture the distortions created by the zero bound on the short-term interest rate, I solve the model using projection methods. The algorithm is an extension of a time iteration algorithm to solve for equilibrium policy functions as in Coleman (1990). I specify a grid for the state vector $\{u_t, r_t^*, q_{t-1}\}$ formed as a tensor product of three linearly spaced vectors. The vector for q_{t-1} is defined on the range $[\underline{q}, \bar{q}]$ and the grids for u_t and r_t are specified across ± 4 standard deviations. Expectations are computed using Gauss-Hermite quadrature using five nodes for each shock (ε^r and ε^u) and linear interpolation of the policy functions. Appendix E.3 describes the solution algorithm in more detail.

5 Results

In this section I discuss the results from the baseline model. I first examine the model's ability to replicate the effects of QE actions on asset prices that have been reported in the literature. I then examine the behaviour of the model and in particular asset purchases/sales as the economy enters or leaves a recession. Finally, I compare my results with the statements by monetary policymakers about their plans for asset purchases and sales.

5.1 Stock and flow effects

To compare the model's implications for the effects of asset purchases on long-term bond prices with the empirical evidence presented in D'Amico and King (2013), I focus on the estimated 'stock effects' and 'flow effects' of asset purchases reported by those authors.

Figure 2: Model implied estimates of 'stock effects' and 'flow effects'

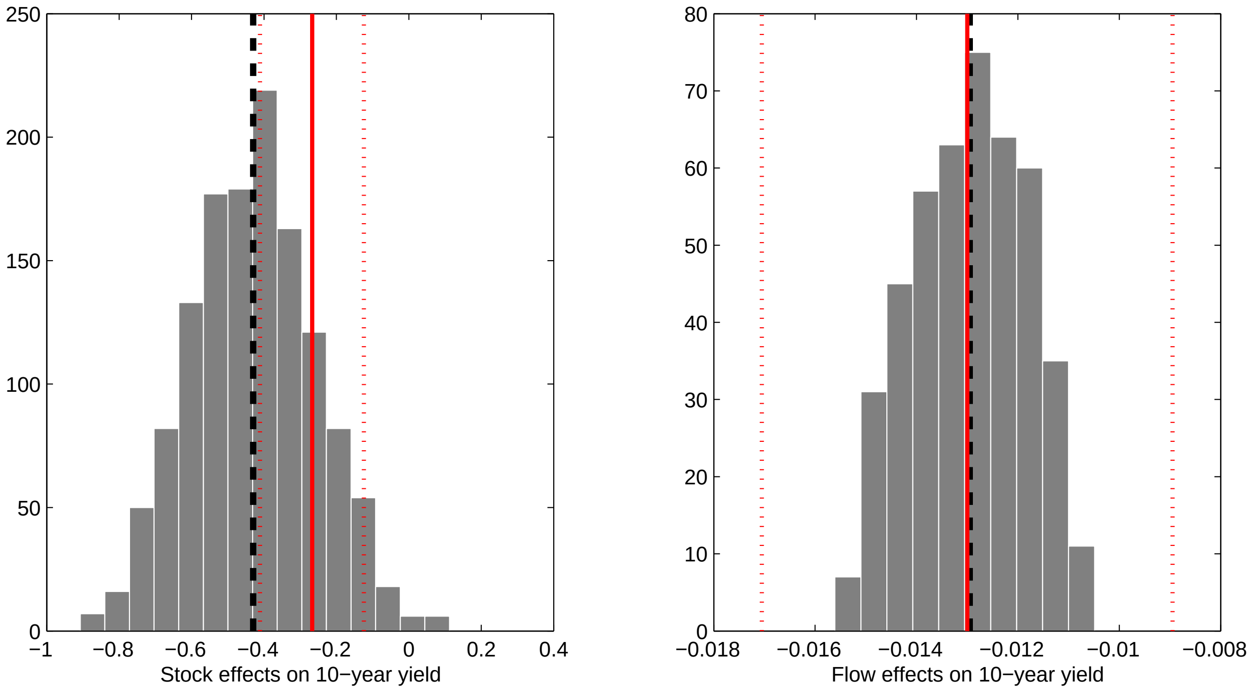

Notes: The left panel shows the distribution of the surprise movements in long-term bond yields for QE surprises amounting to approximately 10% of the debt stock. This is calibrated to match the size of from the FOMC's QE1 programme. The dashed black line is the mean of the distribution and the solid red line is the estimated effect on long-term bond yields attributed to QE1 D'Amico and King (2013). The right panel shows the distribution of surprise movements in the yield differential $G_t = \hat{\mathcal{R}}_t - \hat{\mathcal{R}}^a_t$, where $\hat{\mathcal{R}}^a_t$ is defined in equation (25), for surprise movements in QE of a similar size to the individual QE1 operations. The dashed black line shows the mean of the distribution and the solid red line (against the left hand axis) shows the estimate of flow effects from D'Amico and King (2013). In both panels, the dotted red lines show 90% confidence intervals around the mean estimate.

To uncover the model's implications for stock effects and flow effects, I simulate the model for 100,000 periods and relate surprise movements in QE to surprises in long-term bond yields.[35] By focusing on surprises I mimic the empirical approach of D'Amico and King (2013), which attempts to control for all predictable asset price movements. However, in my model all surprises are generated by optimal policy responses to unforeseen macroeconomic shocks, whereas it is possible that D'Amico and King (2013) estimate the effects of a surprise innovation to a non-optimal QE policy rule (that is, a 'QE policy shock').

Stock effects are computed by isolating the set of QE surprises of a similar magnitude to the FOMC's QE1 programme ($0.10 \leq q_t - \mathbb{E}_{t-1}q_t \leq 0.12$) and recording the corresponding surprise movements in long-term bond rates $\hat{\mathcal{R}}_t - \mathbb{E}_{t-1}\hat{\mathcal{R}}_t$. The mean of the distribution of these surprises (in annualised units), shown as the thick black dashed line in the left panel of Figure 2, is -0.43 implying a substantially larger effect than the estimate of -0.27 presented in D'Amico and King (2013) (the solid red line) and just outside the 90% confidence interval (dotted red line).

To estimate flow effects, I examine the surprise on the difference in yields between the long-term bond and the price of that bond when agents do not face costs of adjusting their portfolio mix ($\tilde{\xi} = \xi = 0$). The yield to maturity of that hypothetical bond is given by:

$$\hat{\mathcal{R}}_t^a = \chi\beta\mathbb{E}_t\hat{\mathcal{R}}_{t+1}^a + (1-\chi\beta)\left(\hat{R}_t - (1+\delta)\,\nu q_t\right) \tag{25}$$

and the yield differential is defined as $G_t = \hat{\mathcal{R}}_t - \hat{\mathcal{R}}_t^a$. This definition is designed to be the closest match to the effects estimated by D'Amico and King (2013). If estimated accurately, the flow effects in D'Amico and King (2013) reflect reactions to surprises in the maturity composition of asset purchases when the New York Fed enacted the purchases. My model does not incorporate the full maturity structure of government debt, but focusing on marginal effects on yields of *changes* in households' portfolio mix is the closest analogue.

Flow effects are computed by isolating the set of QE surprises of a similar magnitude to individual QE auctions ($0.006 \leq q_t - \mathbb{E}_{t-1}q_t \leq 0.008$) and recording the corresponding surprise movements in the long-term yield differential $G_t - \mathbb{E}_{t-1}G_t$. This distribution is compared to D'Amico & King's estimate of the flow effects of -0.013% shown as the red line in the right hand panel of Figure 2, with 90% confidence intervals indicated by the dotted red lines.[36] The model matches the estimated flow effects well.

These estimation results reflect the use of a loss function that weights the mean squared deviation from the mean estimated (stock and flow) effects according to the inverse of the confidence intervals. Because the confidence interval for the flow effects is quite narrow (relative to the mean estimate) the flow effect receives a relatively high weight in the estimation process.

[35] A simulation of 110,000 periods is used with the first 10,000 periods discarded.

[36] Appendix C explains how these confidence intervals are constructed from the results in D'Amico and King (2013).

5.2 QE entry and exit

In this section, I study the properties of the model and in particular its predictions for asset purchases when the economy enters and leaves a recession. Figure 3 plots 'slices' of the policy functions conditioned on particular values for the cost push state, u_t, and the lagged value of the QE instrument, q_{t-1}. In all cases I condition on $u_t = 0$. I then consider policy functions conditional on $q_{t-1} = \underline{q} = 0$ and $q_{t-1} = \bar{q} = 0.5$. By conditioning on the minimum and maximum levels of QE, I can study conditions of 'entry' into and 'exit' from periods in which the central bank holds assets on its balance sheet. Both of these cases are compared to a variant in which the central bank is not allowed to implement QE (that is $q_t = 0, \forall t$).[37] Given the conditioning assumptions, all policy function 'slices' show how optimal outcomes are affected by the natural real interest rate, $r^{\star}$, holding u_t and q_{t-1} constant.

Figure 3: Policy function comparison

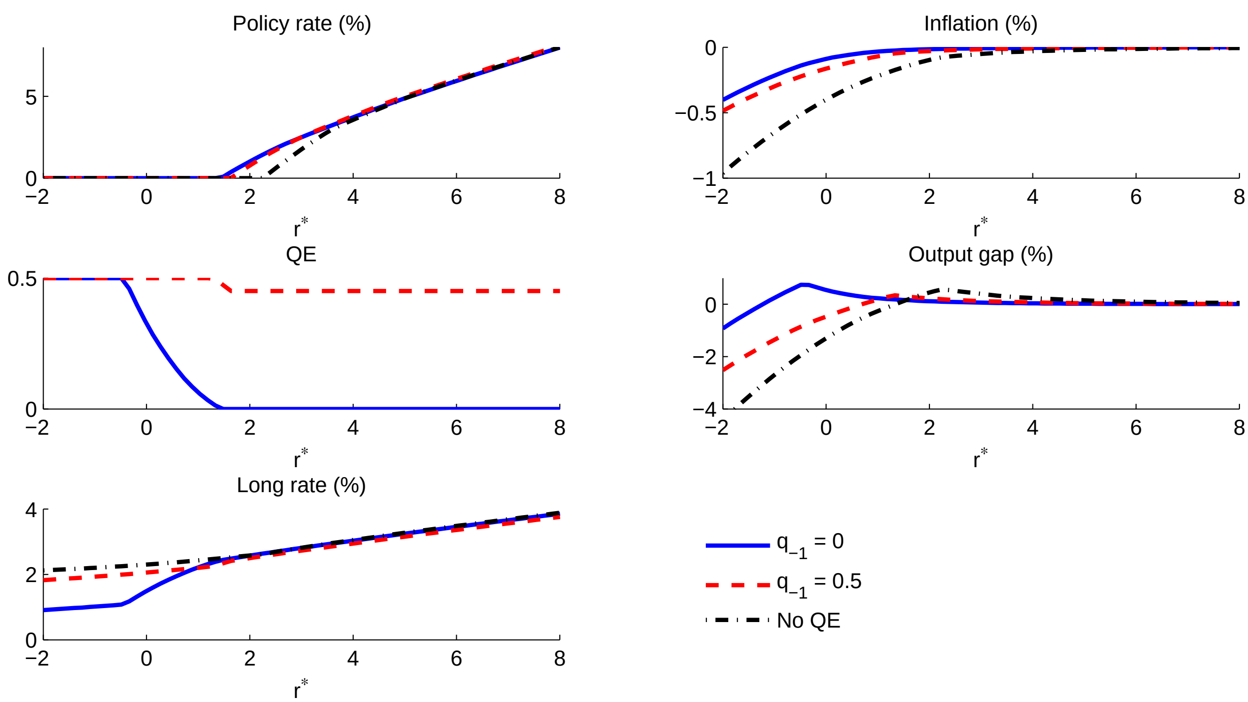

Notes: 'Slices' of policy functions for alternative model variants. The solid blue lines are slices of the policy functions conditional $\{u_t, q_{t-1}\} = \{0, 0\}$. The dashed red lines are slices of the policy functions conditional on $\{u_t, q_{t-1}\} = \{0, 0.5\}$. The dot-dash black lines are policy functions conditional on $u_t = 0$ for a version of the model in which the policymaker does not use QE (so $q_t = 0, \forall t$).

Figure 3 demonstrates that when QE is not used (black dot-dash lines), the policy functions have the same qualitative features as those presented in Adam and Billi (2007).[38]

[37] Formally, this case corresponds to a situation in which the policymaker acts to minimise the welfare-based loss function using only the short-term nominal interest rate as the policy instrument. This setup therefore corresponds to the case of discretionary policy subject to the zero lower bound in the canonical New Keynesian model, as studied by Adam and Billi (2007).

[38] The policy functions are quantitatively different because different parameter values are used.

Low realisations of r^* are associated with the policy rate at the zero bound and negative outcomes for the output gap and inflation. The fact that agents understand that policy will be constrained in this way for low realisations of r^* reduces inflation expectations for realisations of r^* that are low enough to imply a substantial risk of hitting the zero bound. This effect implies that the policy rate hits the zero bound when the natural rate is positive (around 2%). Moreover, the downward skew in the distribution of future inflation outcomes induces the policymaker to generate a positive output gap for values of r^* are slightly above the value that induces the policy rate to hit the zero bound. As described by Adam and Billi (2007), this is the optimal response to the effect of low inflation expectations on inflation.[39]

When QE is used, the recessionary consequences of low realisations of r^* are mitigated relative to the case in which QE is not used. This follows from the fact that QE can be used to ease monetary conditions when the short-term policy rate is constrained by the zero bound. The policy functions conditioned on $q_{t-1} = 0$ (solid blue lines) show that the policymaker does not make substantial use of QE until the short-term policy rate is constrained by the zero bound. This follows from the first order condition for QE in the case that the policymaker is unconstrained in her instrument settings:[40]

$$q_t = \frac{\omega_{\Delta q}}{\omega_q + \omega_{\Delta q}} q_{t-1} - \frac{\beta}{\omega_q + \omega_{\Delta q}} \frac{\partial \mathbb{E}_t \mathcal{L}_{t+1}}{\partial q_t} - \frac{\beta}{\omega_q + \omega_{\Delta q}} \frac{\partial \mathbb{E}_t \pi_{t+1}}{\partial q_t} \omega_\pi \pi_t \tag{26}$$

When the short-term policy rate is unconstrained, active use of QE does not affect current allocations (since the short-term policy rate can be adjusted to deliver the unconstrained optimal allocations). Moreover, setting $0 \leq q_t \leq \bar{q}$ reduces the scope for subsequent stimulus in the event of bad shocks arriving (such that the short-term policy rate is constrained by the zero bound). So the effects of choosing $0 \leq q_t \leq \bar{q}$ on future losses are positive and the effects on expected future inflation are negative. Taken together, these observations imply that it is not optimal to engage in QE until the short-term policy rate has hit the zero bound.

When the policy rate is constrained by the zero bound, the optimal level of QE rises for lower realisations of r^*. Higher QE reduces the long-term interest rate and provides additional monetary stimulus, hence reducing the recessionary effects of these realisations of r^*. Importantly, the anticipation of additional monetary easing via QE when the short-term policy rate is constrained by the zero bound supports inflation expectations for 'low' values of r^*. As a result, the policy rate becomes constrained for values of r^* below around 1.5% compared with a value of around 2% when QE is not used (black dash-dot lines).

Nevertheless, the policy functions in which QE is used as a policy instrument exhibit a similar tradeoff between inflation and the output gap for very low realisations of r^*. This reflects the fact that as r^* reaches very low levels, it becomes optimal to purchase the maximum possible quantity of assets, $q_t = \bar{q} = 0.5$. In such states, further easing in the event of future recessionary shocks is not possible and the same downward skew in

[39]That is, the policymaker pursues the targeting rule $\hat{\pi}^t = -\frac{\omega_x}{\omega_\pi}\hat{x}_t$ when away from the zero bound.

[40]That is, $\lambda_t^x = \lambda_t^{\bar{q}} = \lambda_t^{q} = 0$.

future inflation outturns described above for the standard New Keynesian model once again emerges.

Finally, comparisons of the policy functions for the cases in which $q_{t-1} = 0$ and $q_{t-1} = 0.5$ (dashed red lines) reveal the importance of 'flow effects' in influencing long-term interest rates, monetary conditions and hence output and inflation. When $q_{t-1} = 0$, setting $q_t > 0$ generates both 'stock effects' and 'flow effects' on the long-term rate. For extremely low realisations of r^*, setting $q_t = 0.5$ reduces the long-term by around 100bp compared to setting $q_t = 0.5$ when $q_{t-1} = 0.5$.[41]

To shed further light on how initial conditions effect outcomes as the economy enters or exits a recession, Figures 4 and 5 present 'modal' simulations for alternative initial conditions. In each case, the simulation traces out the outcomes in the event that the sequence of cost push and natural rate shocks are equal to their most likely value of zero (that is, $\varepsilon_t^u = \varepsilon_t^r = 0, t = 2, \ldots$). The alternative paths represent outcomes for different initial conditions for the exogenous states and QE holdings.

Figure 4: Modal simulation of a severe recessionary scenario

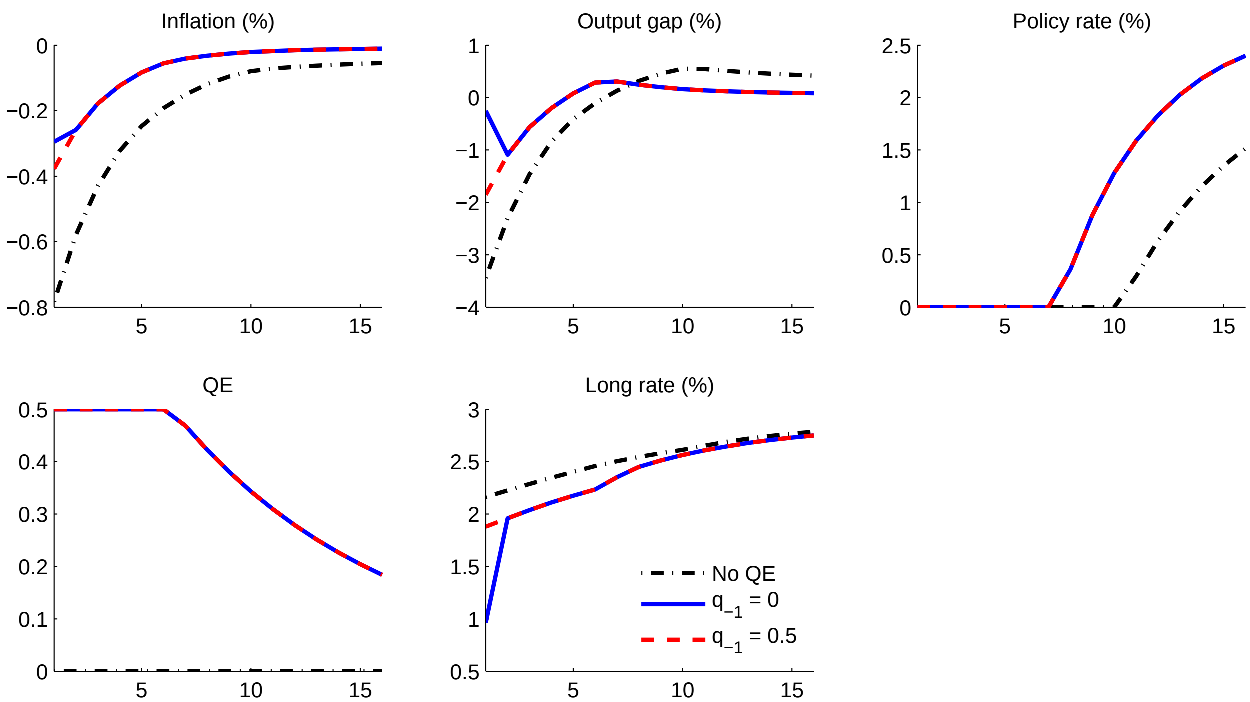

Notes: Each simulated path is computed under the assumptions that the sequence of shocks is equal to the most likely value ($\varepsilon_t^u = \varepsilon_t^r = 0, t = 2, \ldots$). The values of the exogenous state variables in period 1 are $u_1 = 0$ and $r_1^* = -2.25\%$ (in annualised units). The solid blue lines correspond to the case in which the initial stock of QE is $q_0 = 0$. The red dashed lines correspond to the case in which the initial stock of QE is $q_0 = 0.5$. The dash-dotted black lines show the case in which the policymaker does not use QE (so $q_t = 0, \forall t$).

Figure 4 shows the case in which the initial condition for the natural rate of interest is extremely low ($r_1^* = -2.25\%$, measured as an annualised rate)[42] and the short-term

[41] Compare the solid blue and red dashed lines in the bottom left panel of Figure 3.

[42] In these units, the steady-state value of the natural rate is approximately 3.35%.

policy rate is constrained by the zero bound. When the initial condition for QE is zero ($q_0 = 0$, solid blue lines), the policy maker sets QE to its maximium level immediately. The flow effects from this action are sufficient to generate a substantial initial fall in long-term interest rates.

In contrast, a policymaker who experiences the same recessionary state ($r_1^* = -2.25\%$), but inherits a maximal stock of QE ($q_0 = \bar{q} = 0.5$, red dashed lines) does not have recourse to further policy loosening. So in this case, inflation and the output gap are more negative in period 1. However, from period 2 onwards, the outcomes from these two simulations are identical, because the endogenous state variable is identical in period 2 (that is, $q_1 = 0.5$ in both cases). Compared to the case in which QE is not used (black dash-dotted lines), the additional stimulus from asset purchases allows the short-term policy rate to liftoff from the zero bound several quarters earlier.

Figure 5: Modal simulation of a mild recessionary scenario

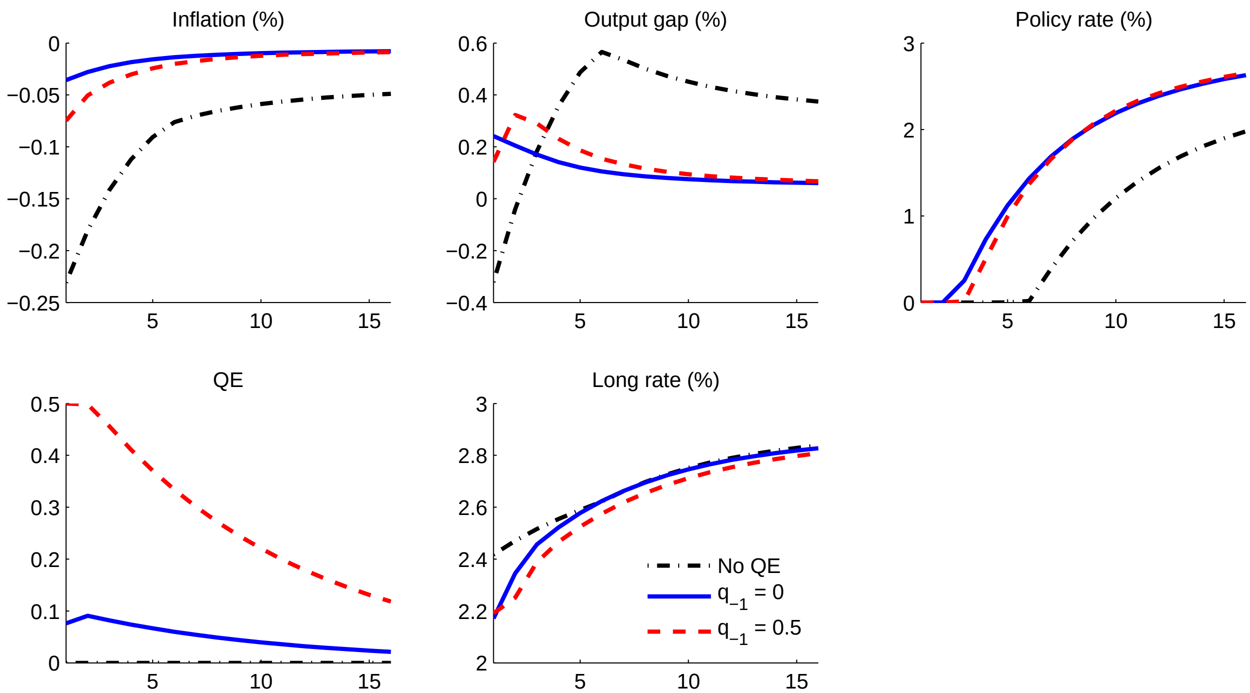

Notes: Each simulated path is computed under the assumptions that the sequence of shocks is equal to the most likely value ($\varepsilon_t^u = \varepsilon_t^r = 0, t = 2, \ldots$). The values of the exogenous state variables in period 1 are $u_1 = 0$ and $r_1^* = 0.55\%$ (in annualised units). The solid blue lines correspond to the case in which the initial stock of QE is $q_0 = 0$. The red dashed lines correspond to the case in which the initial stock of QE is $q_0 = 0.5$. The dashed-dotted black lines show the case in which the policymaker does not use QE (so $q_t = 0, \forall t$).

Figure 5 shows the case of a much milder recessionary state, so that $r_1^* = 0.55\%$ (in annualised units). In the case in which the policymaker does not inherit any assets on its balance sheet ($q_0 = 0$, solid blue lines), it is optimal to engage in a small-scale QE operation that is unwound slowly. When the policymaker inherits a large stock of assets ($q_0 = \bar{q} = 0.5$, red dashed lines), it is optimal to start unwinding the stock after one period. The unwinding of the large stock of assets tightens monetary conditions, both through the flow effects of asset sales and the stock effect of reducing the size of the balance sheet. This explains why the trajectory of long term interest rates is similar despite the different paths for QE.[43]

In the case of a large inherited balance sheet ($q_0 = 0.5$), exit from the zero bound is delayed by one quarter relative to the case in which the policymaker does not inherit any assets on its balance sheet. Moreover, when the policymaker is unwinding a large initial stock of assets, it has, on average, less capacity to respond to a future shock that constrains QE at its upper bound. As a result, the tradeoff between weaker inflation and stronger output is more acute in this case. Compared with the case in which QE is not used (black dashed-dotted lines), however, the tradeoff is managed much more effectively and, once again, liftoff from the zero bound occurs earlier.

Figure 6: Distribution of changes in QE (Δq_t)

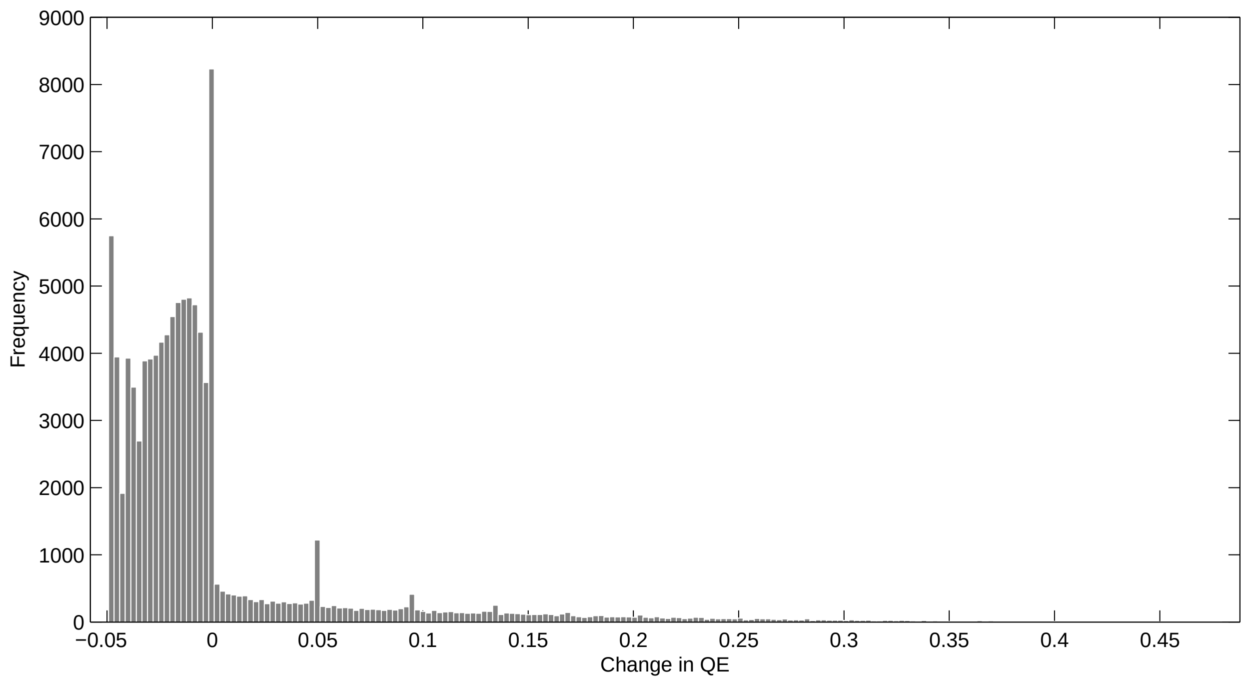

Notes: The histogram records the distribution of outcomes for the change in QE (Δq_t) from a stochastic simulation of 100,000 periods.

The analysis of Figures 4 and 5 suggests that there is a skew in the distribution of QE policy actions: it is more common to observe large asset purchases than large sales. This is because large scale purchases can be triggered by a large recessionary shock when the policy rate is constrained by the zero bound but exit from QE typically occurs slowly and at least partially during periods in which the short-term policy rate is unconstrained by the zero bound. Figure 6 confirms this intuition by plotting the distribution of changes in QE (Δq_t) from a stochastic simulation of the model. The distribution exhibits an upward skew.[44]

[43]Mechanically, the long rate is calibrated to a duration of around 8 years and the fact that QE is substantially unwound after four years in both simulations implies that the effects on the long rate of the different paths will be relatively small.

[44]One observation from Figure 6 is that there are 'spikes' in the distribution of QE changes. This reflects

The discussion of the policy functions in Figure 3 noted the tradeoff between the output gap and inflation that occurs as the limits of policy accommodation are reached (that is, when the policy rate is at the lower bound and $q \approx \bar{q}$). Figure 5 demonstrated that this tradeoff may still be present even when policy is relatively unconstrained. Using simple New Keynesian models (without QE), Evans et al. (2016) examine the typical size of the output gap and deviation of inflation from target at the point of liftoff from the zero bound. Figure 7 explores this issue in the context of my model.

Figure 7: Distributions of variables at liftoff

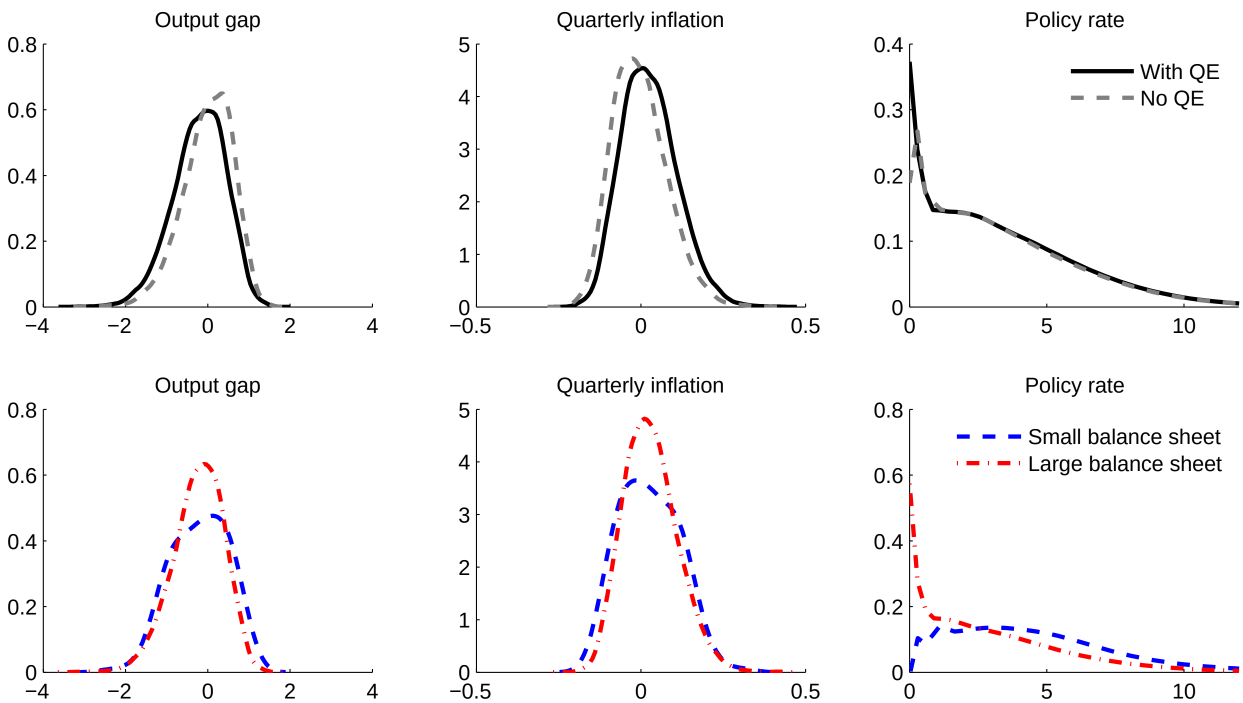

Notes: The top row shows kernel based estimates of the distributions of the output gap, inflation and the policy rate for the baseline version of the model in which QE is used (solid black lines) and a case in which QE is not used (dashed grey lines). The distributions are computed using a simulation of 100,000 periods and selecting those periods in which the current policy rate is positive and in which the policy rate in the preceding was at the zero bound. The bottom row shows distributions conditional on the policymaker having a 'large' balance sheet immediately prior to liftoff ($0.4 \leq q_{t-1} \leq 0.5$, dash-dotted red lines) and conditional on a 'small' balance sheet ($0 \leq q_{t-1} \leq 0.1$, dashed blue lines).

The top row of Figure 7 shows the distributions of the output gap, inflation and the short-term policy rate in liftoff quarters (defined as those in which the policy rate is positive, but was equal to the lower bound in the previous period). The solid black lines show

the fact that unwinding existing QE stocks is in many cases a near deterministic process. To see this, recall the first order condition for QE when the policy instruments are unconstrained, (26). When the policy rate is unconstrained, the effects of current QE decisions on expectations are likely to be small (because there are many future states of the world in which the short-term policy rate will be unconstrained and current QE decisions will have no impact in those states). In the limiting case where QE has no effect on future outcomes, the first order condition implies $q_t = \frac{\omega_{\Delta q}}{\omega_q + \omega_{\Delta q}} q_{t-1}$. For states in which the effects of current QE on future decisions are small, we will observe $q_t \approx \frac{\omega_{\Delta q}}{\omega_q + \omega_{\Delta q}} q_{t-1}$ and the implied changes in QE will 'bunch' around the values implied by a deterministic unwind of QE.

the distribution in the baseline model, with active use of QE and the dashed grey lines show the distributions when QE is not used ($q_t = 0, \forall t$). These distributions show that, when QE is used as an active policy tool, liftoff of the short-term policy rate will tend to occur with a more negative output gap and smaller inflation overshoot when compared to the case in which QE is not used.[45] The use of QE allows the policymaker to lift off before the output gap has closed (on average).[46]

The bottom row of Figure 7 shows the distributions in liftoff quarters for cases in which the central bank's pre-liftoff balance sheet is 'small' ($0 \leq q_{t-1} \leq 0.1$, dashed blue lines) and 'large' ($0.4 \leq q_{t-1} \leq 0.5$, dash-dotted red lines). The distributions for the output gap and inflation in these cases are similar, with little difference in the means.[47] However, it is notable that the variance of the distributions is larger when the policymaker lifts off with a small initial balance sheet. This reflects the fact that such liftoff episodes are likely to occur in relative benign situations (in which previous shocks have not required substantial use of QE). So these liftoff episodes are more likely to correspond to cases in which policy is unlikely to be constrained in the near future. The distribution of the short-term policy rate supports this reasoning: lifting off with a large balance sheet is more likely to be associated with a smaller initial rate rise.

5.3 Discussion

The results from the baseline model suggest that QE is not actively used until the short-term policy rate hits the zero lower bound. It also suggests that asset purchases of a very large scale occur fairly frequently. These predictions seem consistent with the implementation of QE in early 2009 in both the United States and the United Kingdom. In those cases, the initial purchases of long-term government debt were large and occurred at or soon after the short-term policy rate hit the effective lower bound.

However, it is less obvious that the model's predictions for exit from QE are consistent with the presumptive exit strategies announced by real world policymakers. In the United Kingdom and United States, policymakers expect to begin unwinding QE only after the short-term policy rate has been increased from the zero bound.[48] The Bank of England's MPC has stated that unwinding of QE will not begin until the short-term policy rate has reached levels that make it possible to respond to negative shocks by reducing the policy rate rather than increasing the level of QE.[49] In contrast, the results from my model suggest that it is optimal to start reducing the stock of QE *before* the short-term policy rate has been increased from the zero lower bound.[50]

[45] When QE is used, the average output gap is -0.21 and the average inflation rate is 0.02. Without QE the means are both zero to three decimal places.

[46] Inspection of Figure 4 indicates that for deep recessionary shocks this is likely to be because liftoff will be somewhat later after the impact of the shock if QE is not used as a policy instrument.

[47] For the 'large' ('small') balance sheet cases the mean output gap is -0.23 (-0.20) and the mean inflation rate is 0.03 (0.03).

[48] See Federal Open Market Committee (2011, p3) and Monetary Policy Committee (2015, p34).

[49] In Monetary Policy Committee (2015, p34) the MPC estimates the appropriate level of Bank Rate to be around 2%.

[50] Harrison (2012) and Darracq Pariès and Kühl (2016) reach similar conclusions: optimal policy behaviour implies that QE is halted and begins to unwind at or before the date of liftoff. However, both

What might explain these differences?

One key difference is that the QE policy variable (q) in the model represents the fraction of the stock of long-term government bonds held by the central bank, rather than the absolute size of the asset stock held by the central bank. Recall that Figure 1 plots a crude approximation of q for the United Kingdom. The figure shows that q rises following MPC decisions to increase the stock of assets purchased (dashed vertical lines). But for periods during which the asset stock was held constant and total government debt rose, q was falling.[51] So actual policy behaviour is indeed broadly consistent with the model's predictions: a fixed central bank asset stock when government debt is rising constitutes a reduction in q. In the model, because government debt is assumed to be fixed, q can only be reduced by active sales of assets.

Another consideration is that the policymaker in the model minimises a loss function based on the utility function of agents in the model. This incorporates the fact that holding assets on the central bank balance sheet imposes costs on households by shifting their portfolio mix away from its desired level. The mandates of real-world central banks in reality do not incorporate such costs explicitly. Rather, they resemble a so-called 'flexible inflation targeting' loss function which accounts for the costs of output gap and inflation variability.

Moreover, monetary policymakers have stressed the relative uncertainty over the effect of QE on aggregate demand and inflation relative to the effects of movements in the short-term policy rate.[52] This gives rise to a preference to use the short-term policy rate as the 'marginal instrument' to set the overall stance of monetary policy. This sentiment may be strengthened further by the possibility that asset sales may generate rather different effects from asset purchases.

In the model, the costs of portfolio misallocation are assumed to be quadratic, suggesting symmetry in the marginal effects of QE tightening and loosening. However, these effects will not be symmetric under optimal policy. That is because increasing the level of QE reduces the remaining scope for policy loosening via flow effects whereas many instances in which QE is reduced will have negligible effects on outcomes because the short-term policy rate is unconstrained.[53] Nevertheless, the model abstracts from any uncertainty over the impact of policy actions on outcomes: the aforementioned asymmetries are perfectly understood by agents in the model.

6 Welfare and alternative delegation schemes

The results of Section 5 suggest that active use of QE improves welfare by allowing the policymaker to use an additional instrument to offset the effects of shocks on output and

of these studies assume that the policymaker is has access to a commitment technology and adopt a perfect foresight methodology. In addition, only Harrison (2012) accounts for bounds on the the QE instrument.

[51]Greenwood et al. (2015) argue present evidence of a similar effect for the United States.

[52]One source of uncertainty over the model's predictions is that the factors that gave rise to large effects from initial asset purchases may have been related to the particular state of financial stress during the period in which they were implemented. In contrast, the model assumes that the portfolio adjustment costs that give QE traction are structural.

[53]See the discussion on page 31.

inflation.[54] Table 2 confirms this by reporting the means of key variables for a simulation of 100,000 periods.[55] The mean of the period loss (that is, $\omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2 + \omega_q q_t^2 + \omega_{\Delta q} (q_t - q_{t-1})^2$) is also reported.

Table 2: Statistics from model simulations

Mean (%)	Baseline	No QE	No ZLB
Qtly inflation	-0.03	-0.10	0.00
Output gap	-0.00	-0.01	-0.00
Policy rate	3.30	2.92	3.31
10-year rate	2.79	2.92	3.31
QE	0.28	0	0
Loss	3.49	7.25	2.52

Table 2 shows the results from the baseline version of the model (with active use of QE), a 'no QE' version in which the policymaker sets $q_t = 0, \forall t$ and a 'no ZLB' version in which the zero bound on the short-term policy rate is ignored. Results from this variant represent the best achievable outcomes for the policymaker, conditional on her inability to commit to future policy actions.

Comparing the case in which the policymaker does not use QE with the variant in which the zero bound is ignored reveals that the presence of the zero bound reduces mean outcomes for inflation, the output gap and the short-term policy rate. These effects are sufficient to almost triple the average loss.[56]

When QE is actively used, the downward skew in the distributions of the output gap and inflation are reduced. Relative to the case in which QE is not used, losses are more than halved. This improved performance is associated with an average level of QE of 0.28, a higher level of the short-term policy rate and a lower average level of the long-term rate.

These observations suggest that it may be possible to mimic the outcomes from active use of QE by mandating that the central bank holds a fixed fraction of the stock of long-term bonds on its balance sheet at all times. This type of 'permanent QE' implies that the central bank sets $q_t = q^*, \forall t$, where $\underline{q} \leq q^* \leq \bar{q}$. Such a policy might be expected to reduce average long-term nominal interest rates so that a higher short-term nominal interest rate is required, on average, to deliver inflation at target. A higher average short-term nominal interest rate should in turn reduce the frequency with which the short-term policy rate is constrained by the zero lower bound and therefore improve the policymaker's ability to stabilise the economy.

The above logic is similar to the argument for increasing the inflation target: a higher average inflation rate increases the level of the short-term policy rate consistent inflation at target and reduces the frequency with which the short-term policy rate is constrained

[54] Quint and Rabanal (2017) also consider the welfare implications of using unconventional policies in normal times. However, they assume that conventional monetary policy is set using a Taylor-type feedback rule.

[55] A simulation of 110,000 periods is produced and the first 10,000 periods are discarded.

[56] By using (17) to compare losses, I am ignoring the 'terms independent of policy' that generate costly fluctuations in potential output. The scale of the changes in losses therefore represent an upper bound on the changes in welfare.

by the zero bound. So I also compare the results from 'permanent QE' policies with the case in which the policymaker sets $q_t = 0, \forall t$, but sets the short-term policy rate to minimise the period loss function $\omega_x \hat{x}_t^2 + \omega_\pi (\hat{\pi}_t - \pi^*)^2$ where $\pi^* \geq 0$ is the inflation target delegated to the central bank.

Figure 8: Mean outcomes under 'permanent QE' and alternative inflation targets

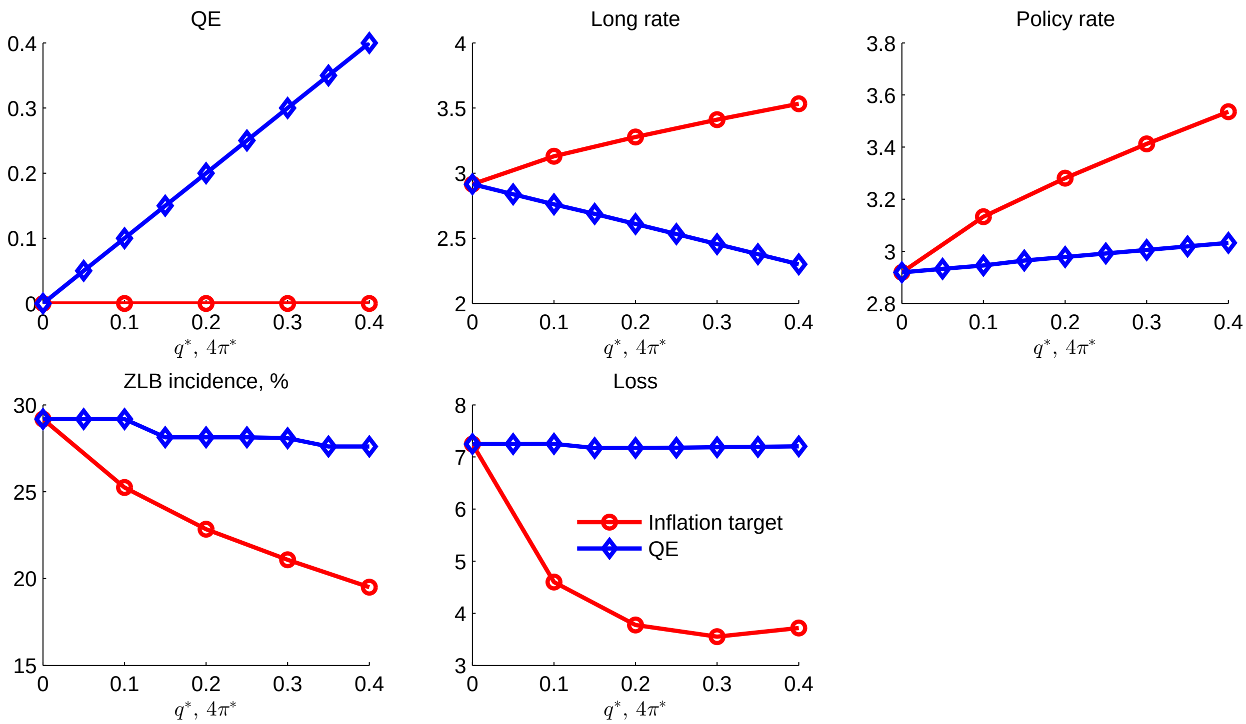

Notes: Each panel reports mean outcomes from a simulation of 100,000 periods for alternative policy specifications. The bottom left panel shows the frequency with which the short-term policy rate is constrained by the zero bound for each of the policy specifications. The blue lines with diamond markers show the outcomes from the case in which the policymaker sets $q_t = q^*, \forall t$ for alternative values of q^* shown on the x-axis. The red lines with circle markers show the case in which the policymaker sets $q_t = 0, \forall t$, but sets the short-term policy rate to minimise the period loss function $\omega_x \hat{x}_t^2 + \omega_\pi (\hat{\pi}_t - \pi^*)^2$. The values of the annualised inflation target $4\pi^*$ are shown on the x-axis.

Figure 8 shows the results of these experiments. Up to a point, increasing the inflation target reduces losses and is associated with a lower frequency of the short-term policy rate being constrained by the zero bound. Beyond this point, while the incidence of the zero bound continues to fall, losses start to increase because the higher level of average inflation is sufficiently costly.

In contrast, 'permanent QE' policies do not improve welfare. As predicted, these policies succeed in 'twisting' the term structure so that the long-term rate falls and the short-term policy rate rises as q^* is increased. However, the strength of these effects is limited and the frequency of zero bound incidents falls only marginally as q^* is increased.

What accounts for these results? By prohibiting active use of QE in response to shocks, a 'permanent QE' policy influences the term structure only through 'stock effects'. These effects are determined by ν which is small relative to the parameter determining flow effects, ξ. A policy of permanent, but fixed, QE therefore has less traction over the term structure than a policy of adjusting the level of QE in response to shocks. Moreover, increasing the inflation target has a effect on inflation expectations which is absent for permanent QE policies. Increasing the inflation target raises inflation expectations and nominal yields directly, which mitigates the downward skew in the distribution of inflation outcomes. This improves the policymaker's ability to stabilise outcomes (as measured by the delegated loss function) both at and away from the zero bound. A permanent QE policy does not have this effect on inflation expectations and as a result the only effect on welfare comes through the frequency with which the short-term policy rate is constrained by the zero bound.

To examine potential gains from more general delegation schemes – when active use of QE is permitted – I assume that the central bank is delegated the following loss function:

$$\widetilde{\mathcal{L}} = \sum_{t=0}^{\infty} \beta^t \left(\omega_x \hat{x}_t^2 + \omega_\pi \left(\hat{\pi}_t - \pi^* \right)^2 + \omega_q \left(q_t - q^* \right)^2 + \omega_{\Delta q} \left(q_t - q_{t-1} \right)^2 \right) \tag{27}$$

where π^* and $q_* \in (0, \bar{q})$ (with $\bar{q} \leq 1$) are inflation and QE targets delegated to the policymaker. The loss function coincides with the utility-based benchmark when $\pi^* = q^* = 0$.

Figure 9 shows the results of experiments using (27). Once again increasing the inflation target, to a point, reduces losses. However, the optimal increase in the inflation target is relatively small, reflecting the fact that active QE is quite effective at offsetting the negative skew in inflation outcomes. As a result, there is less benefit from increasing average inflation expectations and the costs of higher inflation in states of the world when the policymaker is unconstrained offset these benefits quickly as π^* is increased.

Figure 9 also illustrates that mandating the central bank to target a higher level of assets on its balance sheet does not improve outcomes. Losses increase with q^* because the central bank is forced to hold assets (which imposes costs on households) even when the short-term policy rate is unconstrained.[57]

Despite a mild increase in the short-term policy rate (and a fall in the long-term rate) generated by higher average stock effects, the frequency with which the short-term policy rate is constrained by the zero bound actually *rises* as q^* is increased. This reflects the importance of flow effects in determining the effects of QE on output and inflation. To see this, consider the case in which the policymaker inherits a stock of assets equal to the target level ($q_{t-1} = q^* > 0$). Suppose a recessionary shock arrives that constrains the short-term policy rate at the zero bound and necessitates active use of QE. In this case, the maximum 'firepower' that policymaker can deploy through QE is $\bar{q} - q^*$. A higher value of q^* therefore limits the ability of the policymaker to reduce long-term rates via flow effects by increasing QE and so the ability to stabilise the economy at the zero bound is reduced.

[57]Recall that the social welfare function is given by (17), which penalises any $q_t > 0$.

Figure 9: Mean outcomes with active QE under alternative delegated loss functions

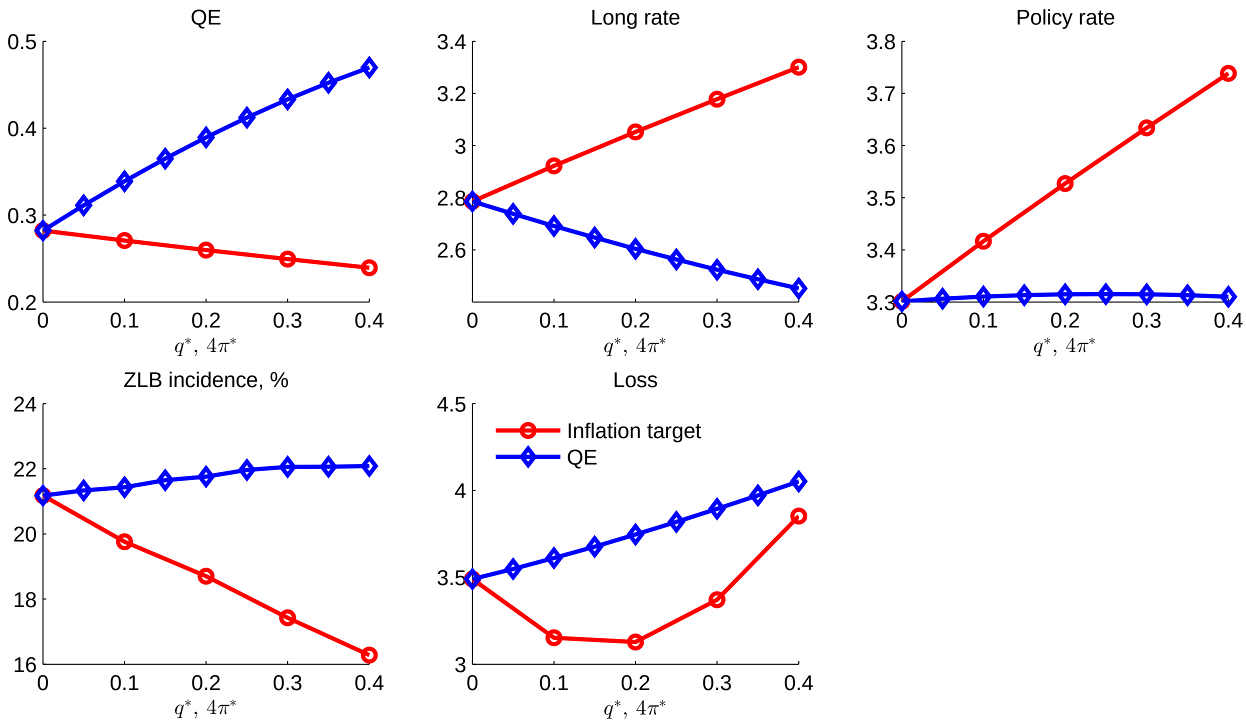

Notes: Each panel reports mean outcomes from a simulation of 100,000 periods for alternative policy specifications. The bottom left panel shows the frequency with which the short-term policy rate is constrained by the zero bound for each of the policy specifications. The blue lines with diamond markers show the outcomes from the case in which the policymaker minimises the loss function (27) for alternative values of q^* shown on the x-axis (with $\pi^* = 0$). The red lines with circle markers show the case in which the policymaker minimises the loss function (27) for values of the annualised inflation target $4\pi^*$ shown on the x-axis (with $q^* = 0$).

7 Robustness analysis

In this section, I consider the robustness of the results presented in Section 5 to alternative assumptions for key parameter values in the model and also to the assumption about the maximal level of QE ($\bar{q}$).

7.1 Alternative parameter values

To assess robustness to the choice of parameter values, I focus on those parameters that are most important for the transmission of monetary policy actions.

I consider the case in which the interest elasticity of demand is smaller than in the baseline case by setting $\sigma = 0.5$ following Eggertsson and Woodford (2003). The interest elasticity of demand is a key parameter because it affects the extent to which changes in both short-term and long-term interest rates affect the output gap. A smaller interest elasticity reduces the power of monetary policy, but also reduces the extent to which monetary conditions are tightened when the zero bound binds. In this case, the slope of the Phillips curve (κ) is held fixed to the baseline value of 0.0516 by setting $\alpha = 0.8805$.

I also consider a case in which the Phillips curve is flatter. Setting $\alpha = 0.9$ reduces the Phillips curve slope to $\kappa = 0.024$. This is the value used by Eggertsson and Woodford (2003) and Levin et al. (2010) in their studies of optimal commitment policy at the zero bound. Other things equal, the flatter Phillips curve specification mitigates the downward drag on inflation expectations near the zero bound.

Flattening the IS and Phillips curves is likely to improve the policymaker's ability to stabilise the economy. In a similar vein, I also consider a case in which the standard deviation of the shock to the natural rate of interest is smaller ($\sigma_r = 0.225$ rather than 0.25 as in the baseline specification). This alternative calibration implies that the zero bound on the short-term policy rate will be less often (and less severely) binding.

Finally, I consider a case in which there are no 'flow effects' ($\xi = 0$). This case is of interest given the uncertainty over the size of flow effects. For example, D'Amico and King (2013) note that flow effects seem to be somewhat short-lived (though are persistent in my model) and Kandrac and Schlusche (2013) find evidence that flow effects from the Fed's LSAP2 and MEP operations were smaller than those for LSAP1.[58]

Figure 10: Policy functions for alternative parameter values

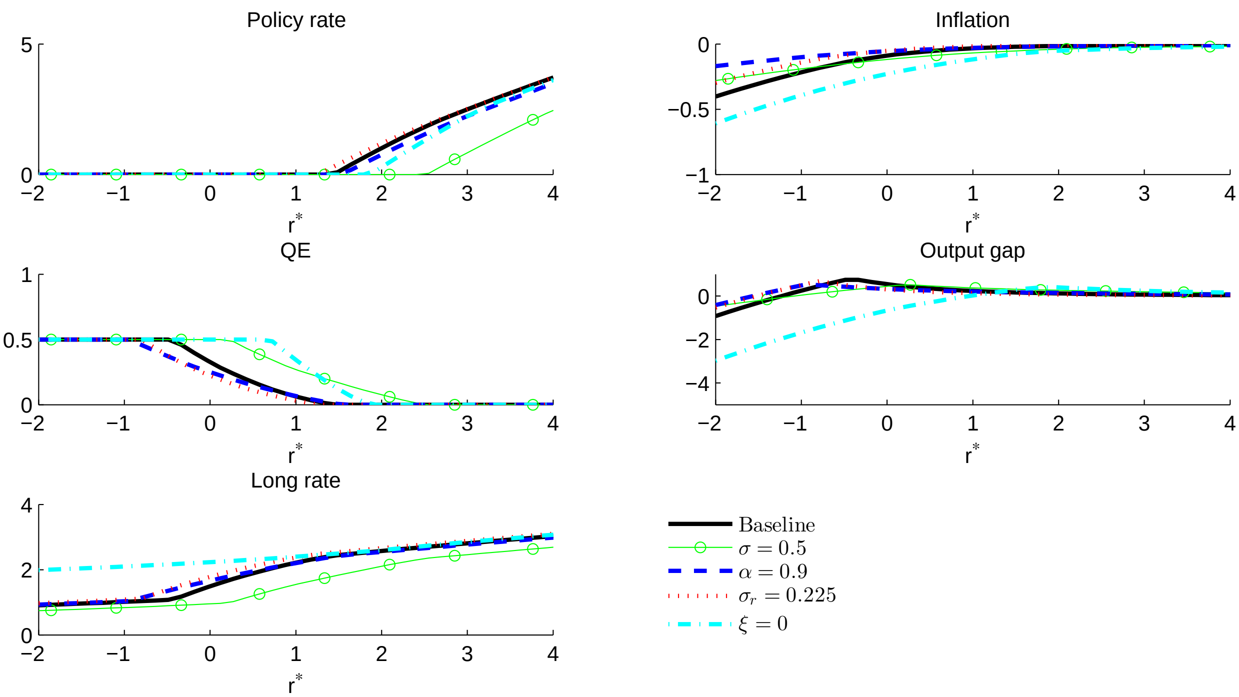

Notes: 'Slices' of policy functions for alternative model variants. All slices of the policy functions are conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. See the text for a full description of the alternative model parameterisations.

Figure 10 plots 'slices' of the policy functions for key variables as functions of the natural real interest rate, r^* conditional on a zero cost push shock state and zero inherited

[58] In this case, the policy problem becomes static, as in the simple New Keynesian model with no portfolio balance effects ($\nu = \xi = 0$), because the existing stock of QE does not affect the choices of future policymakers.

QE. Table 3 summarises statistics for model variables and losses for each of the model variants.

Table 3: Model statistics for alternative parameterisations

Variant	Baseline	$\sigma = 0.5$	$\alpha = 0.9$	$\sigma_r = 0.225$	$\xi = 0$
Mean (%)					
Qtrly inflation	-0.03	-0.06	-0.03	-0.02	-0.06
Output gap	-0.01	-0.01	-0.01	-0.00	-0.01
Policy rate	3.29	3.28	3.33	3.34	3.11
10-year rate	2.78	2.49	2.75	2.86	2.92
QE	0.28	0.43	0.32	0.27	0.10
QE gain (%)	51.89	16.75	12.62	20.74	40.19

Notes: QE gain is percentage difference in loss when QE is used relative to the case in which $q_t = 0, \forall t$.

As expected, the gain from active use of QE is smaller for all of the variants considered (Table 3).[59] When the IS curve and Phillips curves are flatter ($\sigma = 0.5$ and $\alpha = 0.9$), the welfare costs of hitting the zero bound in the absence of QE are smaller so the gains from using QE are smaller. When the variance of natural rate shocks is smaller ($\sigma_r = 0.225$) the likelihood of hitting the zero bound is lower, so the benefits of QE are once again reduced. In the absence of flow effects from QE ($\xi = 0$), agents recognise that future QE actions will have a smaller effect on long-term bond yields. This weakens the expectations channel through which QE helps to support inflation expectations and hence reduces the ability of QE to stimulate spending in the event of future negative demand shocks.

These results are consistent with the policy functions in Figure 10. For the cases in which the IS and Phillips curves are flatter ($\sigma = 0.5$ and $\alpha = 0.9$) the policy functions for the output gap and inflation are generally closer to zero relative to the baseline parameterisation. This reflects the fact that monetary policy is better able to stabilise the economy in light of the smaller downward skews in expected inflation and output gap realisations associated with the zero bound on the short-term policy rate.[60] When $\sigma = 0.5$ the policy functions for the instruments show that more aggressive policy is required to deliver the better stabilisation outcomes. That follows from the fact that, when the IS curve is flatter, larger changes in both the short-term and long-term interest rate are required to achieve a given change in the output gap.

Finally, Figure 10 shows that the absence of flow effects from QE ($\xi = 0$) generates substantially worse outcomes than the baseline parameterisation. The policy function for the long-term interest rate is much flatter because the strong effects of flow effects on depressing long-term yields are absent. As a result, outcomes for both inflation and the output gap are much worse when the economy is at or in the vicinity of the zero bound on the short-term rate.

[59] The 'QE gain' is given by $100 \times \left(1 - \frac{\bar{\mathcal{L}}}{\bar{\mathcal{L}}_{q=0}}\right)$ where $\bar{\mathcal{L}}$ is the mean welfare-based loss and $\bar{\mathcal{L}}_{q=0}$ is the mean welfare-based loss computed under a policy in which $q_t = 0, \forall t$.

[60] Those skews are smaller precisely because the IS or Phillips curves are flatter.

7.2 The upper bound on QE

As noted in Section 3.4, the model makes no distinction between the central bank and government balance sheets (or budget constraints). When the central bank holds long-term bonds on its balance sheet, it faces the risk that the value of those bonds may fall if the long-term interest rate rises. In the model, the government implicitly stands ready to cover any losses incurred on the central bank's portfolio by means of a transfer (or capital injection) funded by levying (lump sum) taxes on households. In practice, doubts over whether the government would guarantee such unconditional support to the central bank have been regarded as a limit on the scale of asset purchases by the central bank.

For example, Dennis Lockhart, President of the Federal Reserve Bank of Atlanta argues that:

> A second perspective on limits [on monetary policy] might reference statutory or self-imposed limits that central banks observe. These might encompass limits on how far the central bank can or should go in addressing what are fiscal concerns. Monetary policymakers have tried to avoid interventions that put taxpayers at risk of loss. (Lockhart, 2012)

Such considerations motivate the imposition of the upper bound on the scale of quantitative easing that the central bank may undertake, $\bar{q} \in (0, 1)$. In this section, I explore the implications of alternative assumptions about $\bar{q}$ for the results presented in Section 5 and for the implicit revaluations of the central bank's portfolio.

Figure 11 shows representations of the policy functions for key variables, under alternative assumptions about $\bar{q}$. Each policy function is plotted holding both the cost push shock and the inherited stock of QE equal to zero.[61] This representation is convenient for assessing the conditions under which a QE regime is entered (the range of values for the natural real interest rate r^* for which asset purchases are initiated) and how the scale of asset purchases is influenced by the natural real interest rate.

The results show that, unsurprisingly, restricting the maximum scale of asset purchases inhibits the policymaker's ability to stabilise output and inflation for low realisations of the natural real interest rate. When $\bar{q} = 0.25$, the policy rate hits the zero lower bound at a (slightly) higher value of r^*. Moreover, the scale of asset purchases is *larger* over the range $r^* \in (0.5, 1.5)$ when the maximal scale of QE is *smaller*. This means that a higher $\bar{q}$ implies a larger 'bang for buck' of a given scale of asset purchases. The reason for this result is that, as explained in Section 6, agents recognise that a policymaker with a larger $\bar{q}$ has more 'firepower' remaining. Agents therefore expect that outcomes will be better stabilised in future bad states. This mitigates the drag on inflation and output gap expectations generated by the presence of the zero bound on the short-term interest rate and the upper bound on QE.

The fact that output and inflation are better stabilised implies that welfare is higher at all points on the slice of the policy function plotted in Figure 11. Table 4 shows that this is also true on average. Increasing the upper bound on QE reduces welfare losses and keeps inflation and the output gap closer to zero on average. Moreover, comparing the

[61] That is, conditioned on $\{u_t, q_{t-1}\} = \{0, 0\}$.

Figure 11: Policy functions for alternative $\bar{q}$

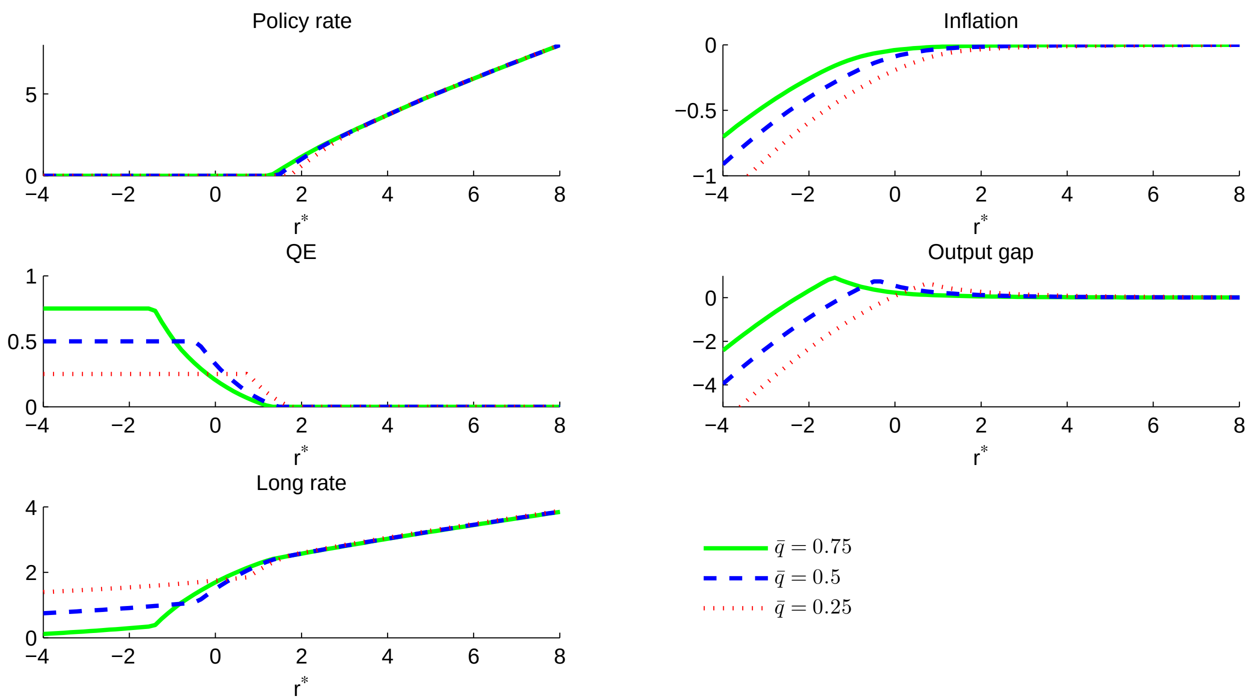

Notes: 'Slices' of policy functions for alternative assumptions about the upper bound on asset purchases, $\bar{q}$. Each slice is conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$.

Table 4: Model statistics for alternative $\bar{q}$

Mean (%)	$\bar{q} = 0.75$	$\bar{q} = 0.5$	$\bar{q} = 0.25$
Qtrly inflation	-0.02	-0.03	-0.05
Output gap	-0.00	-0.01	-0.01
Policy rate	3.37	3.29	3.16
10-year rate	2.74	2.78	2.86
QE	0.35	0.28	0.17
Loss	3.08	3.49	4.34

results for $\bar{q} = 0.75$ with those of the baseline specification ($\bar{q} = 0.5$) reveals that improved stabilisation is achieved with only slightly higher average QE holdings and very slightly lower long-term bond rates.

As noted, the reason why $\bar{q} = 0.75$ delivers better stabilisation performance even with relatively low average QE holdings by the central bank is the recognition that the scale of the central bank's portfolio *can* be expanded to a substantial level in particularly bad states. The fact that the central bank will (optimally) hold a larger quantity of long-term government debt in some states exposes its portfolio to larger interest rate risk compared to a central bank operating under a lower $\bar{q}$ constraint.

To explore the extent of the interest rate risk, I calculate the size of the revaluation of the central bank's portfolio as a fraction of steady-state GDP. Appendix D derives the following expression for the revaluation effect:

$$\mathcal{K}_t \approx \frac{\delta\left(b+b_L\right)}{1+\delta}\left[\hat{R}^1_{L,t}-\hat{R}_{t-1}\right]q_{t-1} \tag{28}$$

based on the assumption that purchases of long-term debt are financed by issuing interest-bearing reserves.[62]

Calculating the distribution of $\mathcal{K}_t$ from simulations of the model provides a way to assess the interest rate risk associated with alternative assumptions about $\bar{q}$. Without a fuller treatment of the central bank budget constraint it is not possible to infer from this distribution the likelihood of the central bank paying a negative dividend to the government. However, very large revaluation effects make it more likely that dividends will be negative in some states.[63]

To assess the likelihood of such an event, I calculate the probability of $\mathcal{K} < -0.005$: the probability that capital losses exceed 0.5% of GDP. The choice of critical value is intended to proxy for the fact that, in general, central banks generate seigniorage revenue from the issuance of non-interest bearing currency. A generous estimate for seigniorage revenue as a proportion of GDP is 0.5%, so I interpret the estimated probability as approximating that of a negative dividend payment: the portfolio revaluation exceeds the likely flow of seigniorage revenue.[64]

Figure 12 plots the histograms of $\mathcal{K}_t$ for 100,000 period simulations of the models with $\bar{q} \in \{0.75, 0.5, 0.25\}$. The bottom right panel shows kernel estimates of the underlying distributions of $\mathcal{K}_t$. While there is a notable skew in the distributions, the case in which $\bar{q} = 0.75$ has a particularly long left tail.

Indeed, when $\bar{q} = 0.75$, negative dividends occur with a frequency of almost 10%, and even the baseline assumption of $\bar{q} = 0.5$ implies a frequency of around 7%. Reducing $\bar{q}$ to 0.25 reduces the risk to less than 1%. In the absence of a full articulation of the central bank budget constraint and balance sheet, the revaluation effects considered here can only provide an indication of the effects on central bank profitability. However, Benigno (2017, equation (35)) demonstrates that central bank profits will be equal to the revaluation effect computed above in the special case that the nominal interest rate in the previous period is at the zero bound (that is $R_{t-1} = 1$).

Figure 13 plots estimates of the distributions of central bank profits, conditional on being constrained by the zero bound in the preceding period. Once again, allowing the

[62] Reserves are assumed to earn the same return as the short-term government bond. A formal model of this type of approach is presented in Bassetto and Messer (2013), which is used to study the feasibility of alternative central dividend policies (see also, Hall and Reis, 2014). Since my model abstracts from the central bank balance sheet and budget constraint, I am restricted to much simpler, indicative, exercises.

[63] As made clear by Reis (2015a), negative dividends would only create solvency problems for the central bank under an extremely strict dividend policy (generating what he calls 'period insolvency' in the event that the government insists that all positive net income is transferred to the government as a dividend). Moreover, my calculations represent an upper bound on the extent to which balance sheet risks present a problem for central bank solvency because they assume that the central bank's asset portfolio is marked to market.

[64] Reis (2015a) argues that the steady-state seigniorage ratio is 0.23% in the model presented by Del Negro and Sims (2015). The data in Aisen and Veiga (2008) give average ratios of 0.3% and 0.4% for the United States and United Kingdom respectively.

Figure 12: Distributions of portfolio revaluation $\mathcal{K}$ for alternative $\bar{q}$

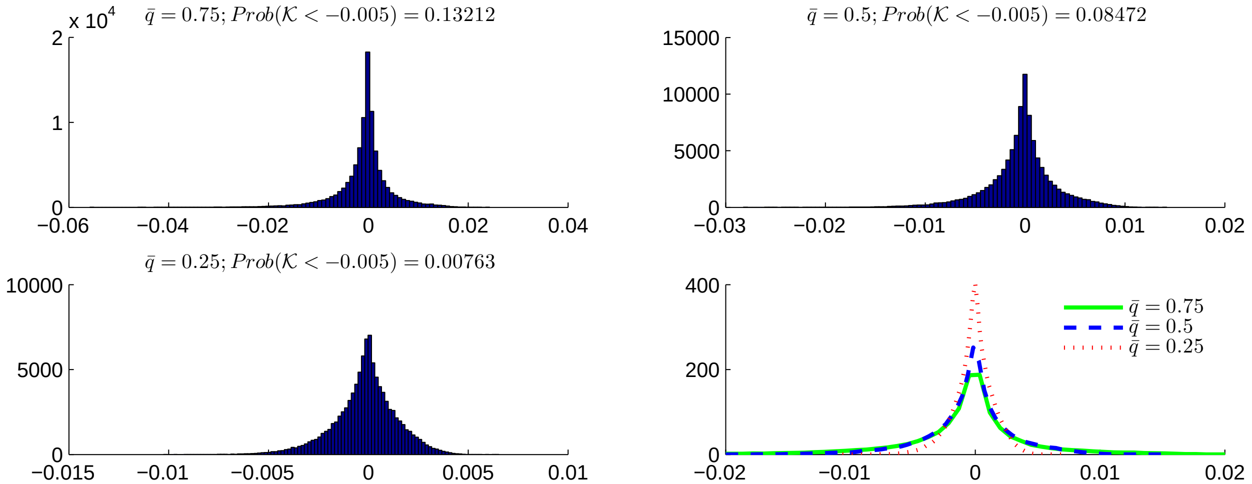

Notes: Histograms of $\mathcal{K}$ computed using equation (28) from a 100,000 period simulation of the model with $\bar{q} \in \{0.25, 0.5, 0.75\}$. The bottom right panel compares kernel estimates of the underlying densities of $\mathcal{K}$.

Figure 13: Central bank profits, conditional on being at the zero bound

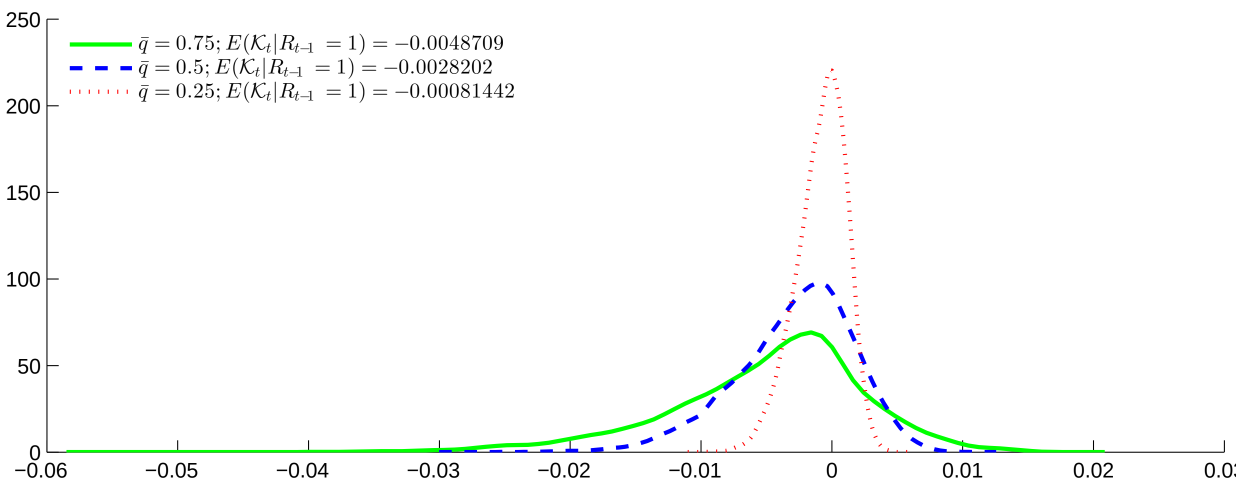

Notes: Kernel estimates of the the distributions of central bank profits, conditional on being at the ZLB in the previous period ($R_{t-1} = 1$) for variants of the model with $\bar{q} \in \{0.25, 0.5, 0.75\}$. Results are based on 100,000 period simulations.

central bank to undertake larger asset purchases results in a larger left tail of losses. Conditional on being constrained by the zero bound, losses are 0.5% of steady state GDP when $\bar{q} = 0.75$ compared with less than 0.1% when $\bar{q} = 0.25$.[65] Of course, a higher $\bar{q}$ also allows the central bank to post larger profits on its asset portfolio during a period in which the ZLB is binding, as Figure 13 also shows. However, the debate on policy constraints imposed by the central bank's balance sheet has focused on the extent to which

[65] The frequency of losses is roughly the same for the alternative values of $\bar{q}$: although the distribution of profits is narrower for lower values of $\bar{q}$, the frequency with which the ZLB constraint binds is also higher. These effects roughly cancel out.

the government will stand ready to make transfers to the central bank (or forgo dividend payments).

8 Conclusion

I study the optimal use of quantitative easing alongside the short-term policy rate using a textbook New Keynesian model extended to include portfolio adjustment costs. The existence of these costs implies both that QE can influence long-term rates and that its use has welfare costs. Reducing long-term rates can increase aggregate demand when the economy is in a recessionary state in which the short-term policy rate is constrained by the zero bound. In such cases, the welfare costs of portfolio distortion may be outweighed by the benefits of increased aggregate demand. Indeed, relative to the case in which the only policy instrument is the short-term policy rate, use of QE reduces the welfare costs of fluctuations by around 50%.

The model predicts that entry into QE regimes (a period during which the central bank holds a positive stock of long-term bonds) can be rapid, with asset purchases commencing as soon as the short-term policy rate hits the zero bound. Exit from QE is slower. Both of these findings are consistent with some aspects of real-world QE policies.

The model also predicts that exit from a QE regime (sales of previously accumulated long-term bonds by the central bank) occurs before the short-term policy rate lifts off from the zero bound. This contrasts with real-world policy statements and indeed the actions of the FOMC following liftoff of the federal funds rate. One reason for this apparent difference is that the measure of quantitative easing in the model represents the share of long-term government debt held by the central bank. In the model, this ratio can only be reduced via active sales of assets held by the central bank. A real-world analogue of this ratio can fall if the central bank holds its asset portfolio constant and government debt rises, as has been observed in several economies.

References

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

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

Aisen, A. and F. J. Veiga (2008): "The political economy of seigniorage," *Journal of Development Economics*, 87, 29–50.

Aksoy, Y. and H. S. Basso (2014): "Liquidity, term spreads and monetary policy," *The Economic Journal*, 124, 1234–1278.

Alla, Z., R. A. Espinoza, and A. R. Ghosh (2016): "Unconventional Policy Instruments in the New Keynesian Model," *IMF Working Paper*.

Andrés, J., J. López-Salido, and E. Nelson (2004): "Tobin's imperfect asset substitution in optimizing general equilibrium," *Journal of Money, Credit and Banking*, 36, 665–91.

Bassetto, M. and T. Messer (2013): "Fiscal consequences of paying interest on reserves," *Fiscal Studies*, 34, 413–436.

Bauer, M. D. and G. D. Rudebusch (2014): "The Signaling Channel for Federal Reserve Bond Purchases," *International Journal of Central Banking*.

Baumeister, C. and L. Benati (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*.

Benford, J., S. Berry, K. Nikolov, C. Young, and M. Robson (2009): "Quantitative Easing," *Bank of England Quarterly Bulletin*, 49, 90–100.

Benigno, P. (2017): "A Central Bank Theory of Price Level Determination," *CEPR Discussion Papers*.

Benigno, P. and S. Nistico (2015): "Non-neutrality of Open-market Operations," *CEPR Discussion Paper No. DP10594*.

Benigno, P. and M. Woodford (2006): "Linear-quadratic approximation of optimal policy problems," Tech. rep., *NBER Working Paper No. 12672*.

Bernanke, B., V. Reinhart, and B. Sack (2004): "Monetary policy alternatives at the zero bound: an empirical assessment," *Brookings Papers on Economic Activity*, 2, 1–78.

Bernanke, B. S. (2010): "Opening remarks: the economic outlook and monetary policy," in *Proceedings: Economic Policy Symposium, Jackson Hole*, Federal Reserve Bank of Kansas City, 1–16.

——— (2014): "A Conversation: The Fed Yesterday, Today and Tomorrow," *interview by Liquat Ahamed at the Brookings Institution*, 14.

BODENSTEIN, M., J. HEBDEN, AND R. NUNES (2012): "Imperfect credibility and the zero lower bound," *Journal of Monetary Economics*, 59, 135–149.

CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–98.

CANZONERI, M., R. CUMBY, B. DIBA, AND D. LÓPEZ-SALIDO (2008): "Monetary Aggregates and Liquidity in a Neo-Wicksellian Framework," *Journal of Money, Credit and Banking*, 40, 1667–1698.

CANZONERI, M., R. CUMBY, B. DIBA, AND D. LOPEZ-SALIDO (2011): "The role of liquid government bonds in the great transformation of American monetary policy," *Journal of Economic Dynamics and Control*, 35, 282–294.

CANZONERI, M. B. AND B. T. DIBA (2005): "Interest rate rules and price determinacy: The role of transactions services of bonds," *Journal of Monetary Economics*, 52, 329–343.

CARLSTROM, C. T., T. S. FUERST, AND M. PAUSTIAN (2017): "Targeting Long Rates in a Model with Segmented Markets," *American Economic Journal: Macroeconomics*, 9, 205–42.

CHEN, H., V. CÚRDIA, AND A. FERRERO (2012): "The Macroeconomic Effects of Large-scale Asset Purchase Programmes," *Economic Journal*, 122, F289–F315.

COCHRANE, J. H. (2011): "Inside the black box: Hamilton, Wu, and QE2," *mimeo*.

COLEMAN, W. J. (1990): "Solving the stochastic growth model by policy-function iteration," *Journal of Business & Economic Statistics*, 8, 27–29.

CÚRDIA, V. AND M. WOODFORD (2009): "Conventional and unconventional monetary policy," Tech. rep., *Federal Reserve Bank of New York Staff Report No. 404*.

DAINES, M., M. JOYCE, AND M. TONG (2012): "QE and the gilt market: a disaggregated analysis," .

D'AMICO, S. AND T. B. KING (2013): "Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply," *Journal of Financial Economics*, 108, 425–448.

DARRACQ PARIÈS, M. AND M. KÜHL (2016): "The optimal conduct of central bank asset purchases," *ECB Working Paper*.

DE GRAEVE, F. AND K. THEODORIDIS (2016): "Forward guidance, quantitative easing, or both?" *National Bank of Belgium Working Paper*.

DEL NEGRO, M., G. EGGERTSSON, A. FERRERO, AND N. KIYOTAKI (2017): "The great escape? A quantitative evaluation of the Fed's liquidity facilities," *The American Economic Review*, 107, 824–857.

Del Negro, M. and C. A. Sims (2015): "When does a central bank's balance sheet require fiscal support?" *Journal of Monetary Economics*, 73, 1–19.

Eggertsson, G. and M. Woodford (2003): "The zero interest rate bound and optimal monetary policy," *Brookings Papers on Economic Activity*, 1, 139–211.

Ellison, M. and A. Tischbirek (2014): "Unconventional government debt purchases as a supplement to conventional monetary policy," *Journal of Economic Dynamics and Control*, 43, 199 – 217, the Role of Financial Intermediaries in Monetary Policy Transmission.

Epstein, L. G. and S. E. Zin (1989): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica: Journal of the Econometric Society*, 937–969.

Evans, C., J. Fisher, F. Gourio, and S. Krane (2016): "Risk management for monetary policy near the zero lower bound," *Brookings Papers on Economic Activity*, 2015, 141–219.

Farmer, R. E. and P. Zabczyk (2016): "The theory of unconventional monetary policy," *Bank of England Staff Working Paper No. 613*.

Federal Open Market Committee (2011): "Minutes of the Federal Open Market Committee, June 21–22," *Federal Reserve Board of Governors*.

Frankel, J. A. (1985): "Portfolio crowding-out, empirically estimated," *The Quarterly Journal of Economics*, 100, 1041–1065.

Galí, J. (2008): *Monetary policy, inflation, and the business cycle*, Princeton University Press.

Gertler, M. and P. Karadi (2011): "A model of unconventional monetary policy," *Journal of Monetary Economics*, 58, 17–34.

——— (2013): "QE 1 vs. 2 vs. 3...: A framework for analyzing large-scale asset purchases as a monetary policy tool," *International Journal of Central Banking*, 9, 5–53.

Greenwood, R., S. G. Hanson, J. S. Rudolph, and L. Summers (2015): "Debt Management Conflicts between the U.S. Treasury and the Federal Reserve," in *The $13 Trillion Question: How America Manages Its Debt*, ed. by D. Wessel, Brookings Institution Press, chap. 2, 43–89.

Greenwood, R. and D. Vayanos (2010): "Price pressure in the government bond market," *The American Economic Review*, 100, 585–590.

——— (2014): "Bond supply and excess bond returns," *Review of Financial Studies*, 27, 663–713.

Guvenen, F. (2006): "Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective," *Journal of Monetary Economics*, 53, 1451 – 1472.

HALL, R. E. AND R. REIS (2014): "Maintaining central-bank solvency under new-style central banking," .

HARRISON, R. (2011): "Asset purchase policies and portfolio balance effects: a DSGE analysis," in *Interest rates, prices and liquidity*, ed. by J. Chadha and S. Holly, Cambridge University Press, chap. 5.

——— (2012): "Asset purchase policy at the effective lower bound for interest rates," *Bank of England Working Paper No. 444*.

HOHBERGER, S., R. PRIFTIS, AND L. VOGEL (2017): "The macroeconomic effects of quantitative easing in the Euro area : evidence from an estimated DSGE model," *European University Institute Working Paper*.

JOYCE, M., A. LASAOSA, I. STEVENS, M. TONG, ET AL. (2011): "The financial market impact of quantitative easing in the United Kingdom," *International Journal of Central Banking*, 7, 113–161.

KANDRAC, J. AND B. SCHLUSCHE (2013): "Flow effects of large-scale asset purchases," *Economics Letters*, 121, 330–335.

KAPETANIOS, G., H. MUMTAZ, I. STEVENS, AND K. THEODORIDIS (2012): "Assessing the Economy-wide Effects of Quantitative Easing," *Economic Journal*, 122, F316–F347.

KING, T. B. (2015): "A portfolio-balance approach to the nominal term structure," *mimeo*.

KRISHNAMURTHY, A. AND A. VISSING-JORGENSEN (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120, 233–267.

KUTTNER, K. (2006): "Can central banks target bond prices?" Tech. rep., *NBER Working Paper No. 12454*.

LENZA, M., H. PILL, AND L. REICHLIN (2010): "Monetary policy in exceptional times," *Economic Policy*, 25, 295–339.

LEVIN, A., D. LÓPEZ-SALIDO, E. NELSON, AND T. YUN (2010): "Limitations on the effectiveness of forward guidance at the zero lower bound," *International Journal of Central Banking*, 6, 143–89.

LOCKHART, D. P. (2012): "Monetary Policy Limits: Federal Reserve Actions and Tools," *Speech at Institute of Regulation and Risk, Tokyo*, 5.

MONETARY POLICY COMMITTEE (2015): "Inflation Report, November," *Bank of England*.

NAKATA, T. (2015): "Credibility of Optimal Forward Guidance at the Interest Rate Lower Bound," *FEDS notes*.

PESARAN, M. H. AND R. P. SMITH (2016): "Counterfactual analysis in macroeconometrics: An empirical investigation into the effects of quantitative easing," *Research in Economics*, 70, 262–280.

PLOSSER, C. I. (2012): "Perspectives on Monetary Policy," *Official Monetary and Financial Institutions Forum (OMFIF) Golden Series Lecture London, England*.

PRIFTIS, R. AND L. VOGEL (2016): "The Portfolio Balance Mechanism and QE in the Euro Area," *The Manchester School*, 84, 84–105.

QUINT, D. AND P. RABANAL (2017): "Should Unconventional Monetary Policies Become Conventional?" *International Monetary Fund*.

REIS, R. (2015a): "Different Types of Central Bank Insolvency and the Central Role of Seignorage," NBER Working Papers 21226, National Bureau of Economic Research, Inc.

——— (2015b): "QE in the future: the central banks balance sheet in a fiscal crisis," *mimeo*.

SMETS, F. AND R. WOUTERS (2005): "Comparing Shocks and Frictions in US and Euro Area Business Cycles: A Bayesian DSGE Approach," *Journal of Applied Econometrics*, 20, 161–183.

——— (2007): "Shocks and frictions in US business cycles," *American Economic Review*, 97, 586–606.

THORNTON, D. L. (2014): "QE: is there a portfolio balance effect?" *Federal Reserve Bank of St. Louis Review*, 96, 55–72.

TOBIN, J. (1956): "Liquidity preference as behavior towards risk," *Review of Economic Studies*, 25, 65–86.

——— (1969): "A general equilibrium approach to monetary theory," *Journal of Money, Credit and Banking*, 1, 15–29.

TOBIN, J. AND W. C. BRAINARD (1963): "Financial intermediaries and the effectiveness of monetary controls," *American Economic Review*, 53, 383–400.

VAYANOS, D. AND J.-L. VILA (2009): "A preferred-habitat model of the term structure of interest rates," *NBER working paper*.

WEALE, M. AND T. WIELADEK (2016): "What are the macroeconomic effects of asset purchases?" *Journal of Monetary Economics*, 79, 81 – 93.

WEIL, P. (1990): "Nonexpected utility in macroeconomics," *The Quarterly Journal of Economics*, 105, 29–42.

WOODFORD, M. (2001): "Fiscal Requirements for Price Stability," *Journal of Money, Credit and Banking*, 669–728.

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

——— (2016): "Quantitative easing and financial stability," *CEPR Discussion Paper No. DP11287*.

A Model derivation

My model modifies Harrison (2012) in several respects. I model the long-term government bond as an infinitely-lived security paying a geometrically declining coupon rather than a pure consol to better approximate the behavior of long-term interest rates. I simplify the behavior of fiscal policy to focus exclusively on the role of monetary policy. I assume that the portfolio friction is in the form of adjustment costs rather than within the utility function and I also allow portfolio adjustment costs to depend on changes in households' portfolio mix (between short-term and long-term bonds) as a way to capture 'flow effects' of asset purchases on bond yields. Finally, I abstract from base money, which reduces the scale of the model without affecting the main conclusions.

A.1 Households

The optimisation problem considered in Section 3.2 is

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\frac{1}{\sigma}} - 1}{1 - \frac{1}{\sigma}} - \frac{\phi_t n_t^{1+\psi}}{1+\psi} \right\}$$

subject to

$$\begin{aligned} B_{L,t}^h + B_t^h =& R_{L,t}^1 B_{L,t-1}^h + R_{t-1} B_{t-1}^h + W_t n_t + T_t + D_t - P_t c_t \\ & - \frac{\tilde{\nu} P_t \left(b^h + b_L^h\right)}{2} \left[\delta \frac{B_t^h}{B_{L,t}^h} - 1\right]^2 \\ & - \frac{\tilde{\xi} P_t \left(b^h + b_L^h\right)}{2} \left[\frac{B_t^h}{B_{t-1}^h} \frac{B_{L,t-1}^h}{B_{L,t}^h} - 1\right]^2 \end{aligned} \tag{A.1}$$

The first-order conditions for the optimisation problem are:

$$c_t^{-\frac{1}{\sigma}} = \mu_t P_t \tag{A.2}$$

$$\phi_t n_t^{\psi} = W_t \mu_t \tag{A.3}$$

$$\begin{aligned} 0 =& - \mu_t - \mu_t \frac{\tilde{\nu} \delta P_t \left(b^h + b_L^h\right)}{B_{L,t}^h} \left[\delta \frac{B_t^h}{B_{L,t}^h} - 1\right] \\ & - \mu_t \frac{\tilde{\xi} P_t \left(b^h + b_L^h\right)}{B_{t-1}^h} \frac{B_{L,t-1}^h}{B_{L,t}^h} \left[\frac{B_t^h}{B_{t-1}^h} \frac{B_{L,t-1}^h}{B_{L,t}^h} - 1\right] + \beta R_t \mathbb{E}_t \mu_{t+1} \\ & + \beta \mathbb{E}_t \mu_{t+1} \frac{\tilde{\xi} P_{t+1} \left(b^h + b_L^h\right) B_{t+1}^h}{\left(B_t^h\right)^2} \frac{B_{L,t}^h}{B_{L,t+1}^h} \left[\frac{B_{t+1}^h}{B_t^h} \frac{B_{L,t}^h}{B_{L,t+1}^h} - 1\right] \end{aligned} \tag{A.4}$$

$$
\begin{aligned}
0 = & - \mu_t + \mu_t \frac{\tilde{\nu}\delta P_t \left(b^h + b^h_L\right) B^h_t}{\left(B^h_{L,t}\right)^2} \left[\delta \frac{B^h_t}{B^h_{L,t}} - 1\right] \\
& + \mu_t \frac{\tilde{\xi} P_t \left(b^h + b^h_L\right) B^h_t}{B^h_{t-1}} \frac{B^h_{L,t-1}}{\left(B^h_{L,t}\right)^2} \left[\frac{B^h_t}{B^h_{t-1}} \frac{B^h_{L,t-1}}{B^h_{L,t}} - 1\right] \\
& + \beta \mathbb{E}_t R^1_{L,t+1} \mu_{t+1} - \beta \mathbb{E}_t \mu_{t+1} \frac{\tilde{\xi} P_{t+1} \left(b^h + b^h_L\right) B^h_{t+1}}{B^h_t} \frac{1}{B^h_{L,t+1}} \left[\frac{B^h_{t+1}}{B^h_t} \frac{B^h_{L,t}}{B^h_{L,t+1}} - 1\right] \qquad \text{(A.5)}
\end{aligned}
$$

where μ is the Lagrange multiplier on the nominal budget constraint (A.1).

Let the real Lagrange multiplier be defined as:

$$\Lambda_t \equiv P_t \mu_t$$

and real bond holdings and inflation as

$$
\begin{aligned}
b^h_t & \equiv \frac{B^h_t}{P_t} \\
b^h_{L,t} & \equiv \frac{B^h_{L,t}}{P_t}
\end{aligned}
$$

The first order conditions for short-term and long-term bond holdings, (A.4) and (A.5) can be written in terms of real-valued variables as:

$$
\begin{aligned}
0 = & - \Lambda_t - \Lambda_t \frac{\tilde{\nu}\delta \left(b^h + b^h_L\right)}{b^h_{L,t}} \left[\delta \frac{b^h_t}{b^h_{L,t}} - 1\right] \\
& - \Lambda_t \frac{\tilde{\xi} \left(b^h + b^h_L\right)}{b^h_{t-1}} \frac{b^h_{L,t-1}}{b^h_{L,t}} \left[\frac{b^h_t}{b^h_{t-1}} \frac{b^h_{L,t-1}}{b^h_{L,t}} - 1\right] + \beta R_t \mathbb{E}_t \Lambda_{t+1} \pi^{-1}_{t+1} \\
& + \beta \mathbb{E}_t \Lambda_{t+1} \frac{\tilde{\xi} \left(b^h + b^h_L\right) b^h_{t+1}}{\left(b^h_t\right)^2} \frac{b^h_{L,t}}{b^h_{L,t+1}} \left[\frac{b^h_{t+1}}{b^h_t} \frac{b^h_{L,t}}{b^h_{L,t+1}} - 1\right] \qquad \text{(A.6)}
\end{aligned}
$$

$$
\begin{aligned}
0 = & - \Lambda_t + \Lambda_t \frac{\tilde{\nu}\delta \left(b^h + b^h_L\right) b^h_t}{\left(b^h_{L,t}\right)^2} \left[\delta \frac{b^h_t}{b^h_{L,t}} - 1\right] \\
& + \Lambda_t \frac{\tilde{\xi} \left(b^h + b^h_L\right) b^h_t}{b^h_{t-1}} \frac{b^h_{L,t-1}}{\left(b^h_{L,t}\right)^2} \left[\frac{b^h_t}{b^h_{t-1}} \frac{b^h_{L,t-1}}{b^h_{L,t}} - 1\right] \\
& + \beta \mathbb{E}_t R^1_{L,t+1} \Lambda_{t+1} \pi^{-1}_{t+1} - \beta \mathbb{E}_t \Lambda_{t+1} \frac{\tilde{\xi} \left(b^h + b^h_L\right) b^h_{t+1}}{b^h_t} \frac{1}{b^h_{L,t+1}} \left[\frac{b^h_{t+1}}{b^h_t} \frac{b^h_{L,t}}{b^h_{L,t+1}} - 1\right] \qquad \text{(A.7)}
\end{aligned}
$$

The following steady state relationships are useful for the subsequent log-linearisation.

$$\frac{b_L^h}{b^h} = \delta$$
$$\frac{b^h}{b^h + b_L^h} = (1+\delta)^{-1}$$
$$\frac{b_L^h}{b^h + b_L^h} = 1 - (1+\delta)^{-1} = \frac{\delta}{1+\delta}$$

Combining (A.2) and (A.6) creates an Euler equation for consumption:

$$\begin{aligned} c_t^{-\frac{1}{\sigma}} =& \beta R_t \mathbb{E}_t c_{t+1}^{-\frac{1}{\sigma}} \pi_{t+1}^{-1} - c_t^{-\frac{1}{\sigma}} \frac{\tilde{\nu}\delta\left(b^h + b_L^h\right)}{b_{L,t}^h} \left[\delta \frac{b_t^h}{b_{L,t}^h} - 1\right] \\ &- c_t^{-\frac{1}{\sigma}} \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{b_{t-1}^h} \frac{b_{L,t-1}^h}{b_{L,t}^h} \left[\frac{b_t^h}{b_{t-1}^h} \frac{b_{L,t-1}^h}{b_{L,t}^h} - 1\right] \\ &+ \beta \mathbb{E}_t c_{t+1}^{-\frac{1}{\sigma}} \pi_{t+1}^{-1} \frac{\tilde{\xi}\left(b^h + b_L^h\right) b_{t+1}^h}{\left(b_t^h\right)^2} \frac{b_{L,t}^h}{b_{L,t+1}^h} \left[\frac{b_{t+1}^h}{b_t^h} \frac{b_{L,t}^h}{b_{L,t+1}^h} - 1\right] \end{aligned}$$

which can be log-linearised to give:

$$\begin{aligned} \hat{c}_t =& \mathbb{E}_t \hat{c}_{t+1} - \sigma \left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1}\right] + \sigma \tilde{\nu} \delta (1+\delta) \left[\hat{b}_t^h - \hat{b}_{L,t}^h\right] \\ &+ \sigma \tilde{\xi} (1+\delta) \left[\Delta\left(\hat{b}_t^h - \hat{b}_{L,t}^h\right) - \beta \mathbb{E}_t \Delta\left(\hat{b}_{t+1}^h - \hat{b}_{L,t+1}^h\right)\right] \end{aligned} \quad \text{(A.8)}$$

The first order conditions for labour supply (A.3) and consumption (A.2) can be combined and log-linearised to give

$$\psi \hat{n}_t + \hat{\phi}_t = \hat{w}_t - \sigma^{-1} \hat{c}_t \quad \text{(A.9)}$$

Log-linearising the first order condition for long-term bonds (A.7) gives:

$$\begin{aligned} \hat{\Lambda}_t =& \tilde{\nu} (1+\delta) \left[\hat{b}_t^h - \hat{b}_{L,t}^h\right] + \mathbb{E}_t \left[R_{L,t+1}^1 + \hat{\Lambda}_{t+1} - \hat{\pi}_{t+1}\right] \\ &+ \tilde{\xi} \delta^{-1} (1+\delta) \left[\Delta\left(\hat{b}_t^h - \hat{b}_{L,t}^h\right) - \beta \mathbb{E}_t \Delta\left(\hat{b}_{t+1}^h - \hat{b}_{L,t+1}^h\right)\right] \end{aligned}$$

Log-linearising the first order condition for short-term bonds (A.6) gives:

$$\begin{aligned} -\hat{\Lambda}_t =& -\mathbb{E}_t \left[\hat{R}_t + \hat{\Lambda}_{t+1} - \hat{\pi}_{t+1}\right] + \tilde{\nu} \delta (1+\delta) \left[\hat{b}_t^h - \hat{b}_{L,t}^h\right] \\ &+ \tilde{\xi} (1+\delta) \left[\Delta\left(\hat{b}_t^h - \hat{b}_{L,t}^h\right) - \beta \mathbb{E}_t \Delta\left(\hat{b}_{t+1}^h - \hat{b}_{L,t+1}^h\right)\right] \end{aligned}$$

Adding the previous two equations gives:

$$\begin{aligned} 0 =& \mathbb{E}_t \left[R_{L,t+1}^1 - \hat{R}_t\right] + \tilde{\nu} (1+\delta)^2 \left[\hat{b}_t^h - \hat{b}_{L,t}^h\right] \\ &+ \tilde{\xi} \delta^{-1} (1+\delta)^2 \left[\Delta\left(\hat{b}_t^h - \hat{b}_{L,t}^h\right) - \beta \mathbb{E}_t \Delta\left(\hat{b}_{t+1}^h - \hat{b}_{L,t+1}^h\right)\right] \end{aligned}$$

or

$$\begin{aligned}\mathbb{E}_t R^1_{L,t+1} =& \hat{R}_t - \tilde{\nu}\left(1+\delta\right)^2\left[\hat{b}^h_t - \hat{b}^h_{L,t}\right] \\ & -\tilde{\xi}\delta^{-1}\left(1+\delta\right)^2\left[\Delta\left(\hat{b}^h_t - \hat{b}^h_{L,t}\right) - \beta\mathbb{E}_t\Delta\left(\hat{b}^h_{t+1} - \hat{b}^h_{L,t+1}\right)\right]\end{aligned} \tag{A.10}$$

The final equation can be used to write the Euler equation in terms of returns on short-term and long-term bonds. First note that (A.10) can be rearranged to give:

$$\begin{aligned}\hat{R}_t - \mathbb{E}_t R^1_{L,t+1} =& \tilde{\nu}\left(1+\delta\right)^2\left[\hat{b}^h_t - \hat{b}^h_{L,t}\right] \\ & +\tilde{\xi}\delta^{-1}\left(1+\delta\right)^2\left[\Delta\left(\hat{b}^h_t - \hat{b}^h_{L,t}\right) - \beta\mathbb{E}_t\Delta\left(\hat{b}^h_{t+1} - \hat{b}^h_{L,t+1}\right)\right]\end{aligned}$$

The right hand side of this expression appears on the right hand side of (A.8), multiplied by $\sigma\delta\left(1+\delta\right)^{-1}$. This implies that the Euler equation can be written as:

$$\begin{aligned}\hat{c}_t =& \mathbb{E}_t\hat{c}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t\hat{\pi}_{t+1}\right] + \sigma\delta\left(1+\delta\right)^{-1}\left[\hat{R}_t - \mathbb{E}_t R^1_{L,t+1}\right] \\ =& \mathbb{E}_t\hat{c}_{t+1} - \sigma\left[\frac{1}{1+\delta}\hat{R}_t + \frac{\delta}{1+\delta}\mathbb{E}_t R^1_{L,t+1} - \mathbb{E}_t\hat{\pi}_{t+1}\right]\end{aligned}$$

A.2 Firms

As noted in the text, the objective function for a producer j resetting its price at date t is:

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

The first order condition is

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

or

$$\mathbb{E}_t\sum_{k=t}^{\infty}\Lambda_k\left(\beta\alpha\right)^{k-t}\left(\left(1-\eta_t\right)\frac{\left(1+s\right)p_{j,t}}{\Pi_{t,k}} + \eta_t\frac{w_k}{A}\right)\left(\frac{p_{j,k}}{\Pi_{t,k}}\right)^{-\eta_t}c_k = 0 \tag{A.11}$$

which defined the price set by firm j relative to the aggregate price level as:

$$p_{j,t} \equiv \frac{P_{j,t}}{P_t}$$

and defines the relative inflation factor as

$$\begin{aligned}\Pi_{t,k} &\equiv \frac{P_k}{P_t} = \Pi_k \times \Pi_{k+1} \times ... \times \Pi_{t+1} \text{ for } k \geq t+1 \\ &\equiv 1 \text{ for } k = t\end{aligned}$$

Since all firms are identical in terms of their information and production constraints, all firms that are able to change prices at date t will choose the same price, denoted p_t^*. Thus

$$\mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k (\beta\alpha)^{k-t} \left((1-\eta_t) \frac{(1+s)\, p_t^*}{\Pi_{t,k}} + \eta_t \frac{w_k}{A} \right) \left(\frac{p_t^*}{\Pi_{t,k}} \right)^{-\eta_t} c_k = 0$$

The aggregate price is:

$$\begin{aligned} P_t &= \left[\int_0^1 P_{j,t}^{1-\eta_t} dj \right]^{\frac{1}{1-\eta_t}} \\ &= \left[\sum_{k=0}^{\infty} (1-\alpha)\, \alpha^k \left(P_{t-k}^* \right)^{1-\eta_t} \right]^{\frac{1}{1-\eta_t}} \end{aligned}$$

where the equality follows from grouping the firms into cohorts according to the date 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 $(1-\alpha)\,\alpha^k$. This means that the aggregate price level can be written as

$$P_t = \left[\alpha \left(P_{t-1} \right)^{1-\eta_t} + (1-\alpha) \left(P_t^* \right)^{1-\eta_t} \right]^{\frac{1}{1-\eta_t}}$$

so that

$$1 = \alpha \left(\frac{1}{\pi_t} \right)^{1-\eta_t} + (1-\alpha) \left(p_t^* \right)^{1-\eta_t} \tag{A.12}$$

Log-linearising the pricing equation gives

$$\mathbb{E}_t \sum_{k=t}^{\infty} (\beta\alpha)^{s-t} \left[\hat{p}_t^* - \hat{\Pi}_{t,k} - \hat{w}_k + \frac{\eta}{\eta-1} \hat{\eta}_k \right] = 0$$

which can be rearranged to give:

$$\hat{p}_t^* = (1-\beta\alpha) \left(\hat{w}_t - \frac{\eta}{\eta-1} \hat{\eta}_t \right) + \beta\alpha \mathbb{E}_t \hat{\pi}_{t+1} + \beta\alpha \mathbb{E}_t \hat{p}_{t+1}^*$$

by using the law of iterated conditional expectations. Linearising the expression for the aggregate price level (A.12) implies that:

$$\hat{p}_t^* = \frac{\alpha}{1-\alpha} \hat{\pi}_t$$

Using this information in the log-linearised pricing equation gives:

$$\hat{\pi}_t = \frac{(1-\beta\alpha)(1-\alpha)}{\alpha} \left(\hat{w}_t - \frac{\eta}{\eta-1} \hat{\eta}_t \right) + \beta \mathbb{E}_t \hat{\pi}_{t+1} \tag{A.13}$$

A.3 Market clearing and the efficient allocation

Goods market clearing requires:

$$\mathcal{D}_t c_t = y_t - \frac{\tilde{\nu}\left(b^h + b_L^h\right)}{2}\left[\delta \frac{b_t^h}{b_{L,t}^h} - 1\right]^2 - \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{2}\left[\frac{b_t^h}{b_{t-1}^h}\frac{b_{L,t-1}^h}{b_{L,t}^h} - 1\right]^2$$

where $\mathcal{D}_t$ is a measure of price dispersion (defined in Appendix B).

As noted in the main text, market clearing in government bond markets implies

$$\hat{b}_t^h - \hat{b}_{L,t}^h = -\hat{b}_{L,t}^h = q_t \tag{A.14}$$

It is straightforward to show that in the absence of price-setting and imperfect asset substitutability frictions, the efficient level of output is proportional to ϕ. To see this, note that in a flexible price equilibrium with no distortion from monopolistic competition, the real wage will equal the marginal product of labour, which is constant and equal to A. So the efficient allocations, denoted with an asterisk, can be found from the labour supply relation (A.9):

$$\psi \hat{n}_t^* + \hat{\phi}_t = -\sigma^{-1}\hat{c}_t^*$$

where $\hat{w}_t^* = 0$ because the real wage is constant. Imposing market clearing, $c_t^* = n_t^* = y_t^*$ implies that potential output is given by:

$$\hat{y}_t^* = -\left(\psi + \sigma^{-1}\right)^{-1}\hat{\phi}_t$$

so that

$$\hat{\phi}_t = -\left(\psi + \sigma^{-1}\right)\hat{y}_t^*$$

A.4 The 'gap' representation

The Phillips curve and Euler equation can be written in terms of the output gap, defined as the deviation between output and the efficient level of output.

Substituting the labour supply equation (A.9) into the Phillips curve (A.13) gives:

$$\begin{aligned}
\hat{\pi}_t =& \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\left(\psi\hat{n}_t + \hat{\phi}_t + \sigma^{-1}\hat{c}_t\right) + \beta\mathbb{E}_t\hat{\pi}_{t+1} - \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\frac{\eta}{\eta-1}\hat{\eta}_t \\
=& \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\left(\left(\psi + \sigma^{-1}\right)\hat{y}_t + \hat{\phi}_t\right) + \beta\mathbb{E}_t\hat{\pi}_{t+1} - \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\frac{\eta}{\eta-1}\hat{\eta}_t \\
=& \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\left(\psi + \sigma^{-1}\right)\hat{x}_t + \beta\mathbb{E}_t\hat{\pi}_{t+1} + u_t
\end{aligned}$$

where the second line uses market clearing and the third line uses the definition of the output gap $\hat{y}_t - \hat{y}_t^* \equiv \hat{x}_t$ and defines the cost push shock, u, as:

$$u_t \equiv -\frac{(1-\alpha)(1-\beta\alpha)}{\alpha}\frac{\eta}{\eta-1}\hat{\eta}_t$$

The Phillips curve can therefore be written as:

$$\hat{\pi}_t = \kappa \hat{x}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \tag{A.15}$$

where

$$\kappa \equiv \frac{(1-\beta\alpha)(1-\alpha)}{\alpha}\left(\psi + \sigma^{-1}\right)$$

The Euler equation for consumption (A.8) can be written as:

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1}\right] + \sigma\tilde{\nu}\delta(1+\delta)q_t + \sigma\tilde{\xi}(1+\delta)\left[\Delta q_t - \beta\mathbb{E}_t \Delta q_{t+1}\right]$$

which incorporates the market clearing conditions for output and government bonds.

Collecting terms, this can be written as:

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - (\nu\delta + \xi(1+\beta))q_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1}\right]$$

where

$$\nu \equiv \tilde{\nu}(1+\delta)$$
$$\xi \equiv \tilde{\xi}(1+\delta)$$

In terms of the output gap we have:

$$\begin{aligned}\hat{y}_t - \hat{y}_t^* + \hat{y}_t^* =& \mathbb{E}_t\left(\hat{y}_{t+1} - y_{t+1}^* + y_{t+1}^*\right) \\ &- \sigma\left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - (\nu\delta + \xi(1+\beta))q_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1}\right]\end{aligned}$$

or

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \gamma q_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1} - \sigma^{-1}\left(\mathbb{E}_t y_{t+1}^* - y_t^*\right)\right]$$

where

$$\gamma \equiv (\nu\delta + \xi(1+\beta))$$

The efficient rate of interest r^* satisfies

$$r_t^* = \sigma^{-1}\left(\mathbb{E}_t y_{t+1}^* - y_t^*\right)$$

so that the Euler equation can be written as:

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \gamma q_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1} - r_t^*\right] \tag{A.16}$$

B Utility-based loss function

Ignoring constants, the period utility function is:

$$U_t = \frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{\phi_t n_t^{1+\psi}}{1+\psi}$$

In what follows markup shocks are ignored (by setting $\eta_t = \eta, \forall t$) to simplify notation. Since cost push shocks are independent of policy this does not affect the derivation.

To derive the loss function, first note that the percentage deviation of any variable z_t from steady state can itself be approximated to second order as:

$$\frac{z_t - z}{z} \approx \hat{z}_t + \frac{1}{2}\hat{z}_t^2$$

where $\hat{z}_t \equiv \ln z_t - \ln z$.

Approximating the utility from consumption to second order gives:

$$\frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} \approx c^{1-\frac{1}{\sigma}}\left(\frac{c_t - c}{c}\right) - \frac{1}{2\sigma}c^{1-\frac{1}{\sigma}}\left(\frac{c_t - c}{c}\right)^2 \tag{B.1}$$

and using the second order approximation for the percentage changes in consumption implies that:

$$\frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} \approx c^{1-\frac{1}{\sigma}}\left(\hat{c}_t + \frac{1}{2}\left(1-\sigma^{-1}\right)\hat{c}_t^2\right)$$

The sub-utility function for labour supply is:

$$\begin{aligned}\frac{\phi_t n_t^{1+\psi}}{1+\psi} \approx& \frac{n^{1+\psi}}{1+\psi} + n^{1+\psi}\frac{n_t - n}{n} + \frac{\psi n^{1+\psi}}{2}\left(\frac{n_t - n}{n}\right)^2 + \frac{n^{1+\psi}}{1+\psi}\frac{\phi_t - \phi}{\phi} \\ &+ n^{1+\psi}\frac{n_t - n}{n}\frac{\phi_t - \phi}{\phi} \\ \approx& n^{1+\psi}\frac{n_t - n}{n} + \frac{\psi n^{1+\psi}}{2}\left(\frac{n_t - n}{n}\right)^2 + n^{1+\psi}\frac{n_t - n}{n}\frac{\phi_t - \phi}{\phi} + t.i.p.\end{aligned}$$

Using the mapping from percentage changes to log-deviations, to second order, implies that:

$$\frac{\phi_t n_t^{1+\psi}}{1+\psi} \approx n^{1+\psi}\left[\hat{n}_t + \frac{(1+\psi)}{2}\hat{n}_t^2 + \hat{n}_t\hat{\phi}_t\right] + h.o.t.$$

where $h.o.t.$ are higher order terms. Using the fact that output is proportional to hours worked ($y_t = An_t$) gives:

$$\frac{\phi_t n_t^{1+\psi}}{1+\psi} \approx n^{1+\psi}\left[\hat{y}_t + \frac{(1+\psi)}{2}\hat{y}_t^2 + \hat{y}_t\hat{\phi}_t\right] + h.o.t.$$

The second-order approximation to the utility function is therefore

$$U_t \approx c^{1-\frac{1}{\sigma}}\left(\hat{c}_t + \frac{1}{2}\left(1-\sigma^{-1}\right)\hat{c}_t^2\right) - n^{1+\psi}\left[\hat{y}_t + \frac{(1+\psi)}{2}\hat{y}_t^2 + \hat{y}_t\hat{\phi}_t\right]$$

The steady-state labour supply relationship is

$$n^{\psi} = wc^{-1/\sigma} = Ac^{-1/\sigma}$$

which follows from the assumption that subsidies to firms are set to eliminate the distortion from monopolistic competition. Steady-state market clearing is

$$c = y = An$$

since steady-state dispersion is $\mathcal{D} = 1$.

This implies that

$$n^{1+\psi} = c^{1-1/\sigma}$$

so that the utility function can be written as

$$U_t \approx c^{1-\frac{1}{\sigma}} \left[\hat{c}_t + \frac{1}{2}\left(1-\sigma^{-1}\right)\hat{c}_t^2 - \hat{y}_t - \frac{(1+\psi)}{2}\hat{y}_t^2 - \hat{y}_t\hat{\phi}_t \right]$$

The goods market clearing condition is:

$$\mathcal{D}_t c_t = y_t - \frac{\tilde{\nu}\left(b^h + b_L^h\right)}{2}\left[\delta\frac{b_t^h}{b_{L,t}^h} - 1\right]^2 - \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{2}\left[\frac{b_t^h}{b_{t-1}^h}\frac{b_{L,t-1}^h}{b_{L,t}^h} - 1\right]^2$$

A first order approximation to the market clearing condition is:

$$\hat{c}_t = \hat{y}_t$$

which implies that:

$$\hat{c}_t^2 = \hat{y}_t^2$$

A second order approximation to the goods market clearing condition is:

$$\hat{\mathcal{D}}_t + \hat{c}_t + \frac{1}{2}\hat{c}_t^2 = \hat{y}_t + \frac{1}{2}\hat{y}_t^2 - \frac{\tilde{\nu}\left(b^h + b_L^h\right)}{2}q_t^2 - \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{2}\left(\Delta q_t\right)^2$$

where $\hat{b}_t^h - \hat{b}_{L,t}^h = q_t$ has been imposed.

Using these results implies that:

$$U_t \approx c^{1-\frac{1}{\sigma}} \left[-\hat{\mathcal{D}}_t - \frac{\psi + \sigma^{-1}}{2}\hat{y}_t^2 - \hat{y}_t\hat{\phi}_t - \frac{\tilde{\nu}\left(b^h + b_L^h\right)}{2}q_t^2 - \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{2}\left(\Delta q_t\right)^2 \right]$$

We know that $\hat{\phi}_t = -\left(\psi + \sigma^{-1}\right)\hat{y}_t^*$ which means that we can write the loss function as:

$$U_t \approx c^{1-\frac{1}{\sigma}} \left[-\hat{\mathcal{D}}_t - \frac{\psi + \sigma^{-1}}{2}\hat{y}_t^2 + \hat{y}_t\left(\psi + \sigma^{-1}\right)\hat{y}_t^* - \frac{\tilde{\nu}\left(b^h + b_L^h\right)}{2}q_t^2 - \frac{\tilde{\xi}\left(b^h + b_L^h\right)}{2}\left(\Delta q_t\right)^2 \right]$$

Note that we can write:

$$\begin{aligned}
-\frac{\psi + \sigma^{-1}}{2}\hat{y}_t^2 + \hat{y}_t\left(\psi + \sigma^{-1}\right)\hat{y}_t^* &= -\frac{\psi + \sigma^{-1}}{2}\left(\hat{y}_t^2 - 2\hat{y}_t\hat{y}_t^*\right) \\
&= -\frac{\psi + \sigma^{-1}}{2}\left(\hat{y}_t^2 - 2\hat{y}_t\hat{y}_t^* + \left(\hat{y}_t^*\right)^2\right) + \frac{\psi + \sigma^{-1}}{2}\left(y_t^*\right)^2 \\
&= -\frac{\psi + \sigma^{-1}}{2}\left(\hat{y} - \hat{y}_t^*\right)^2 + \frac{\psi + \sigma^{-1}}{2}\left(y_t^*\right)^2 \\
&= -\frac{\psi + \sigma^{-1}}{2}\hat{x}_t^2 + t.i.p.
\end{aligned}$$

Define the discounted loss function to be minimised as:

$$
\begin{aligned}
\mathcal{L} &= -2c^{\frac{1}{\sigma}-1}\sum_{t=0}^{\infty}\beta^t U_t \\
&= \sum_{t=0}^{\infty}\beta^t\left[2\hat{\mathcal{D}}_t+\left(\psi+\sigma^{-1}\right)\hat{x}_t^2+\tilde{\nu}\left(b^h+b_L^h\right)q_t^2+\tilde{\xi}\left(b^h+b_L^h\right)\left(\Delta q_t\right)^2\right]
\end{aligned}
$$

Recall that the price dispersion term is

$$
\mathcal{D}_t=\int_0^1\left(\frac{P_t\left(i\right)}{P_t}\right)^{-\eta}di
$$

which in equilibrium is given by

$$
\mathcal{D}_t=\alpha\mathcal{D}_{t-1}\pi_t^{\eta}+\left(1-\alpha\right)\left(p_t^*\right)^{-\eta}
$$

Using the price index (A.12), the optimal price can be written as

$$
p_t^*=\left[\frac{1-\alpha\pi_t^{\eta-1}}{1-\alpha}\right]^{\frac{1}{1-\eta}}
$$

so the price dispersion is

$$
\mathcal{D}_t=\alpha\mathcal{D}_{t-1}\pi_t^{\eta}+\left(1-\alpha\right)\left[\frac{1-\alpha\pi_t^{\eta-1}}{1-\alpha}\right]^{\frac{\eta}{\eta-1}}
$$

Taking a second-order Taylor expansion gives

$$
\begin{aligned}
\hat{\mathcal{D}}_t &\approx \alpha\left(\hat{\mathcal{D}}_{t-1}+\eta\hat{\pi}_t\right)+\left(1-\alpha\right)\left[\frac{-\alpha\eta\hat{\pi}_t}{1-\alpha}\right] \\
&\quad+\frac{\alpha\eta\left(\eta-1\right)}{2}\hat{\pi}_t^2+\frac{1}{2}\left[\frac{\alpha^2\eta}{1-\alpha}-\alpha\eta\left(\eta-2\right)\right]\hat{\pi}_t^2 \\
&\approx \alpha\hat{\mathcal{D}}_{t-1}+\frac{\alpha\eta}{2\left(1-\alpha\right)}\hat{\pi}_t^2
\end{aligned}
$$

Noting that

$$
\begin{aligned}
\sum_{t=0}^{\infty}\beta^t\hat{\mathcal{D}}_t &= \alpha\sum_{t=0}^{\infty}\beta^t\hat{\mathcal{D}}_{t-1}+\sum_{t=0}^{\infty}\beta^t\frac{\alpha\eta}{2\left(1-\alpha\right)}\hat{\pi}_t^2 \\
&= \alpha\hat{\mathcal{D}}_{-1}+\alpha\beta\sum_{t=1}^{\infty}\beta^{t-1}\hat{\mathcal{D}}_{t-1}+\sum_{t=0}^{\infty}\beta^t\frac{\alpha\eta}{2\left(1-\alpha\right)}\hat{\pi}_t^2 \\
&= \alpha\hat{\mathcal{D}}_{-1}+\alpha\beta\sum_{t=0}^{\infty}\beta^t\hat{\mathcal{D}}_t+\sum_{t=0}^{\infty}\beta^t\frac{\alpha\eta}{2\left(1-\alpha\right)}\hat{\pi}_t^2
\end{aligned}
$$

reveals that

$$\sum_{t=0}^{\infty}\beta^t\hat{\mathcal{D}}_t=\frac{\alpha}{1-\alpha\beta}\hat{\mathcal{D}}_{-1}+\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\frac{\alpha\eta}{(1-\alpha\beta)(1-\alpha)}\hat{\pi}_t^2$$

Using this information in the definition of the loss function gives

$$\mathcal{L} \;=\; \sum_{t=0}^{\infty}\beta^t\left[\frac{\alpha\eta}{(1-\alpha\beta)(1-\alpha)}\hat{\pi}_t^2+\left(\psi+\sigma^{-1}\right)\hat{x}_t^2+\tilde{\nu}\left(b^h+b_L^h\right)q_t^2+\tilde{\xi}\left(b^h+b_L^h\right)(\Delta q_t)^2\right]$$

because the term in $\hat{\mathcal{D}}_{-1}$ is independent of policy and can be ignored.

C Rates of return and calibration of stock and flow effects

The following results, shown by Woodford (2001) and Chen et al. (2012) are useful:

$$\text{Yield to maturity} \;\equiv \mathcal{R}_t=V_t^{-1}+\chi \tag{C.1}$$

$$\text{Duration} \;\equiv D_t=\frac{\mathcal{R}_t}{\mathcal{R}_t-\chi} \tag{C.2}$$

Log-linearising the first expression gives:

$$\mathcal{R}\hat{\mathcal{R}}_t=-\frac{1}{V}\hat{V}_t$$

By defintion, the one-period return is also linked to the price of the long-term bond. Log-linearising that relationship gives:

$$R\hat{R}_{L,t}^1=-\frac{1+\chi V}{V}\hat{V}_{t-1}+\chi\hat{V}_t \implies \hat{R}_{L,t}^1=-\hat{V}_{t-1}+\frac{\chi}{R_L^1}\hat{V}_t$$

In a zero inflation steady state, with bond issuance in line with household preferences, returns on short-term and long-term bonds are equalised at $R=R_L^1=\beta^{-1}$. Hence:

$$\hat{R}_{L,t}^1=-\hat{V}_{t-1}+\chi\beta\hat{V}_t$$

Steady-state one-period returns can be used to pin down steady-state V

$$\beta^{-1}=\frac{1+\chi V}{V} \implies V\left(\beta^{-1}-\chi\right)=1 \implies V=\frac{1}{\beta^{-1}-\chi}=\frac{\beta}{1-\beta\chi}$$

In steady state, the yield to maturity is:

$$\mathcal{R}=V^{-1}+\chi=\frac{1-\beta\chi}{\beta}+\chi=\beta^{-1}$$

which implies yield to maturity and one period returns are equalised.

So the yield to maturity can be related to the price of the bond by:

$$\hat{\mathcal{R}}_t = -\beta \frac{1-\beta\chi}{\beta} \hat{V}_t = -\left(1-\beta\chi\right) \hat{V}_t \tag{C.3}$$

This expression can also be used to compute the yield to maturity from model outcomes. Note first that the expected one-period return satisfies:

$$\mathbb{E}_t \hat{R}^1_{L,t+1} = -\hat{V}_t + \chi\beta \mathbb{E}_t \hat{V}_{t+1}$$

or

$$\hat{V}_t = -\mathbb{E}_t \hat{R}^1_{L,t+1} + \chi\beta \mathbb{E}_t \hat{V}_{t+1}$$

which can be written in terms of the yield to maturity:

$$\hat{\mathcal{R}}_t = \left(1-\chi\beta\right) \mathbb{E}_t \hat{R}^1_{L,t+1} + \chi\beta \mathbb{E}_t \hat{\mathcal{R}}_{t+1}$$

Recall that arbitrage between short-term and long-term bonds implies:

$$\begin{aligned}\mathbb{E}_t R^1_{L,t+1} =& \hat{R}_t - \tilde{\nu}\left(1+\delta\right)^2 \left[\hat{b}^h_t - \hat{b}^h_{L,t}\right] \\ &- \tilde{\xi}\delta^{-1}\left(1+\delta\right)^2 \left[\Delta\left(\hat{b}^h_t - \hat{b}^h_{L,t}\right) - \beta\mathbb{E}_t\Delta\left(\hat{b}^h_{t+1} - \hat{b}^h_{L,t+1}\right)\right]\end{aligned}$$

Imposing bond market clearing and the parameter definitions $\nu \equiv \tilde{\nu}\left(1+\delta\right)$ and $\xi \equiv \tilde{\xi}\left(1+\delta\right)$ gives:

$$\begin{aligned}\mathbb{E}_t R^1_{L,t+1} =& \hat{R}_t - \nu\left(1+\delta\right) q_t - \xi\delta^{-1}\left(1+\delta\right)\left[\Delta q_t - \beta\mathbb{E}_t\Delta q_{t+1}\right] \\ =& \hat{R}_t - \nu\left(1+\delta\right) q_t - \xi\delta^{-1}\left(1+\delta\right)\left[q_t - q_{t-1} - \beta\mathbb{E}_t q_{t+1} + \beta q_t\right] \\ =& \hat{R}_t - \delta^{-1}\left(1+\delta\right)\gamma q_t + \xi\delta^{-1}\left(1+\delta\right) q_{t-1} + \beta\xi\delta^{-1}\left(1+\delta\right)\mathbb{E}_t q_{t+1}\end{aligned}$$

where $\gamma \equiv \nu\delta + \xi\left(1+\beta\right)$ as before.

This implies that the yield to maturity is given by:

$$\begin{aligned}\hat{\mathcal{R}}_t =& \chi\beta\mathbb{E}_t\hat{\mathcal{R}}_{t+1} \\ &+ \left(1-\chi\beta\right)\left(\hat{R}_t - \delta^{-1}\left(1+\delta\right)\gamma q_t + \xi\delta^{-1}\left(1+\delta\right) q_{t-1} + \beta\xi\delta^{-1}\left(1+\delta\right)\mathbb{E}_t q_{t+1}\right)\end{aligned}$$

I use these relationships to generate model-consistent measures of the responses of bond yields to QE auctions found by D'Amico and King (2013). Specifically, D'Amico and King (2013, p441) note that \$1bn of asset purchases generates a price increase of around 0.02% for the targeted assets. They argue that this translates into a yield effect of around 0.3 basis points for a representative ten year security.

Equation (C.3) generates a similar result. Given the model calibration, we have

$$\begin{aligned}\hat{\mathcal{R}}_t - \mathbb{E}_{t-1}\hat{\mathcal{R}}_t =& -\left(1-\beta\chi\right)\left(\hat{V}_t - \mathbb{E}_{t-1}\hat{V}_t\right) \\ =& -\left(1-0.9918\times 0.975\right) 0.0002 \\ =& -0.000006599\end{aligned}$$

which when multiplied by 400 to convert into an annualised rate of return gives -0.00264, which is approximately 0.3 basis points.

Average QE auctions were around $\$5$bn, so the target bond yield change is $-0.00264 \times 5 = -0.013$. Repeating the same calculation for the price changes implied by the point estimate of the elasticity of price to QE purchases plus and minus one standard deviation gives the target range used in Figure 2.

D Profits and losses on the central bank's asset portfolio

I assume that the central bank finances asset purchases by issuing interest-bearing reserves. Reserves earn the same (risk free) nominal interest rate as short-term bonds. They are therefore perfect substitutes for short-term bonds and (in equilibrium) households will willingly hold whatever supply of reserves is created by the central bank. I assume that any profits/losses on the central bank's portfolio are transfered to/from the government. At the start of period t the central bank's balance sheet is assumed to have a simple structure. The central bank holds $V_{t-1}\widetilde{Q}_{t-1}$ of previously purchased long-term bonds on the asset side, which is matched by Z_{t-1} of central bank reserves on the liabilities side.

The revaluation effect (or capital gain) on the central bank's existing portfolio is defined as:

$$K_t \equiv [1 + \chi V_t - V_{t-1}] \widetilde{Q}_{t-1} - [R_{t-1} - 1] Z_{t-1}$$

which is the change in the value of the assets minus the change in the cost of the liabilities. The former includes the coupon payment on the long-term bond holdings and the latter includes the risk free interest rate payment on previously issued reserves.

The revaluation effect can be written as:

$$\begin{aligned} K_t &= \left[R^1_{L,t} - 1\right] Q_{t-1} - [R_{t-1} - 1] Z_{t-1} \\ &= \left[R^1_{L,t} - R_{t-1}\right] Q_{t-1} \end{aligned}$$

where the first line uses the definition of the value of assets $\left(Q_{t-1} \equiv V_{t-1}\widetilde{Q}_{t-1}\right)$ and the one-period return on long bonds $\left(R^1_{L,t} \equiv V^{-1}_{t-1}\left[1 + \chi V_t\right]\right)$ and the second line uses the fact that the central bank balance sheet satisfies $Z_{t-1} = V_{t-1}\widetilde{Q}_{t-1}$ at the end of period $t-1$.

Since steady-state output is normalised to unity, K_t can be interpreted as a ratio to steady-state output. Given the assumed debt issuance policy, the revaluation effect can be written in real terms as:

$$\mathcal{K}_t \equiv \frac{K_t}{P_t} = \left[R^1_{L,t} - R_{t-1}\right] q_{t-1}\delta b \approx \frac{\delta\left(b + b_L\right)}{1+\delta} \left[\hat{R}^1_{L,t} - \hat{R}_{t-1}\right] q_{t-1}$$

Using the relationships derived in Appendix C, the ex post one-period return on long-term bonds can be written:

$$\hat{R}^1_{L,t} = -\hat{V}_{t-1} + \chi\beta\hat{V}_t = (1 - \chi\beta)^{-1} \hat{\mathcal{R}}_{t-1} - \chi\beta\left(1 - \chi\beta\right)^{-1} \hat{\mathcal{R}}_t$$

E The optimal policy problem

The policymaker sets policy under discretion, with no ability to commit to future policy plans. I seek a Markov perfect policy in which optimal decisions are a function only of the relevant state variables in the model. The policymaker at date t is treated as a Stackelberg leader with respect to both private agents and policymakers in dates $t+i, i \geq 1$.

Under this interpretation, the policymaker understands that future policymakers will choose allocations according to time-invariant Markovian policy functions. I use upper case bold letters to denote these policy functions. For example, inflation at date $t+j$ is given by the function:

$$\hat{\pi}_{t+j} = \mathbf{\Pi}\left(q_{t+j-1}; z_{t+j}\right) \quad , \ j \geq 1 \tag{E.1}$$

where $z_{t+j} \equiv [u_{t+j} \ \ r_{t+j}]'$ are the exogenous state variables. To simplify notation, I present the policy functions as dependent only on q in what follows.

The loss function that the policymaker minimises is therefore given by:

$$\begin{aligned}
\tilde{\mathcal{L}}_t &= \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left(\frac{\omega_x}{2} \hat{x}_{t+i}^2 + \frac{\omega_\pi}{2} \left(\hat{\pi}_{t+i} - \pi^*\right)^2 + \frac{\omega_q}{2} \left(q_{t+i} - q^*\right)^2 + \frac{\omega_{\Delta q}}{2} \left(q_{t+i} - q_{t+i-1}\right)^2 \right) \\
&= \frac{\omega_x}{2} \hat{x}_t^2 + \frac{\omega_\pi}{2} \left(\hat{\pi}_t - \pi^*\right)^2 + \frac{\omega_q}{2} \left(q_t - q^*\right)^2 + \frac{\omega_{\Delta q}}{2} \left(q_t - q_{t-1}\right)^2 + \beta \mathbb{E}_t \tilde{\mathcal{L}}_{t+1}
\end{aligned}$$

where I consider the variant analysed in Section 6 because it nests the loss function derived in Appendix B when $\pi^* = q^* = 0$.

The problem can therefore be expressed as a Lagrangean:

$$\begin{aligned}
\min_{\left\{\hat{\pi}_t, \hat{x}_t, \hat{R}_t, q_t\right\}} & \frac{\omega_x}{2} \hat{x}_t^2 + \frac{\omega_\pi}{2} \left(\hat{\pi}_t - \pi^*\right)^2 + \frac{\omega_q}{2} \left(q_t - q^*\right)^2 + \frac{\omega_{\Delta q}}{2} \left(q_t - q_{t-1}\right)^2 + \beta \mathbb{E}_t \tilde{\mathcal{L}}_{t+1} \\
& - \lambda_t^\pi \left(\hat{\pi}_t - \kappa \hat{x}_t - \beta \mathbb{E}_t \mathbf{\Pi}\left(q_t\right) - u_t\right) \\
& - \lambda_t^x \left(\hat{x}_t - \mathbb{E}_t \boldsymbol{X}\left(q_t\right) + \sigma \left(\hat{R}_t - \mathbb{E}_t \mathbf{\Pi}\left(q_t\right) - \gamma q_t + \xi q_{t-1} + \beta \xi \mathbb{E}_t \boldsymbol{Q}\left(q_t\right) - r_t^*\right)\right) \\
& - \lambda_t^R \left(\hat{R}_t - \beta^{-1} + 1\right) - \lambda_t^{\bar{q}} \left(q_t - \bar{q}\right) - \lambda_t^{\underline{q}} \left(q_t - \underline{q}\right)
\end{aligned} \tag{E.2}$$

The first order conditions are:

$$0 = \omega_\pi \left(\hat{\pi}_t - \pi^*\right) - \lambda_t^\pi \tag{E.3}$$

$$0 = \omega_x \hat{x}_t + \kappa \lambda_t^\pi - \lambda_t^x \tag{E.4}$$

$$\begin{aligned}
0 = \ & \omega_q \left(q_t - q^*\right) + \omega_{\Delta q} \left(q_t - q_{t-1}\right) + \beta \frac{\partial \mathbb{E}_t \tilde{\mathcal{L}}_{t+1}}{\partial q_t} + \beta \frac{\partial \mathbb{E}_t \mathbf{\Pi}\left(q_t\right)}{\partial q_t} \lambda_t^\pi \\
& + \left[\frac{\partial \mathbb{E}_t \boldsymbol{X}\left(q_t\right)}{\partial q_t} + \sigma \frac{\partial \mathbb{E}_t \mathbf{\Pi}\left(q_t\right)}{\partial q_t} + \sigma \gamma - \beta \sigma \xi \frac{\partial \mathbb{E}_t \boldsymbol{Q}\left(q_t\right)}{\partial q_t} \right] \lambda_t^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}}
\end{aligned} \tag{E.5}$$

$$0 = -\sigma \lambda_t^x - \lambda_t^R \tag{E.6}$$

The first order condition for quantitative easing, (E.5), indicates that the policymaker accounts for the effects of current QE decisions on the losses incurred by future policymakers.

I now consider the solution for a number of cases corresponding to whether or not the constraints on the instruments are binding. Expectations are taken as given (ie known). As described in Appendix E.3, the solution procedure uses the previous guess of the policy functions to compute expectations and then refines the policy function guess conditional on those expectations, iterating in this way until the policy functions converge.

E.1 Interior optimum for the policy instruments

Note that we can write the Euler equation as

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[\tilde{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \tilde{r}_t^* \right] \tag{E.7}$$

where

$$\tilde{R}_t \equiv \hat{R}_t - \gamma q_t \tag{E.8}$$

denotes 'effective' policy rate and

$$\tilde{r}_t^* \equiv r_t^* - \xi q_{t-1} - \beta \xi \mathbb{E}_t q_{t+1} \tag{E.9}$$

is the 'effective' efficient real interest rate.

This variant of the model is isomorphic to the standard New Keynesian model, conditional on past QE and expected future QE. When the zero bound on the short term interest rate $\hat{R}$ does not bind, we have

$$\lambda_t^R = \lambda_t^x = 0$$

In this case, the optimal effective policy rate can be computed using the following steps.

1. When $\lambda_t^x = 0$ the first order conditions imply a targeting criterion:

$$\hat{x}_t = -\frac{\omega_\pi \kappa}{\omega_x} (\hat{\pi}_t - \pi^*) \tag{E.10}$$

2. Using (E.10) to eliminate the output gap from the Phillips curve implies a solution for inflation:

$$\hat{\pi}_t = \left(1 + \frac{\omega_\pi \kappa^2}{\omega_x}\right)^{-1} \left[\kappa \frac{\omega_\pi \kappa}{\omega_x} \pi^* + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \right]$$

3. A solution for the output gap can be computed by plugging the solution for inflation derived in Step 2 into the targeting criterion (E.10).

4. With these solutions in hand, the optimal effective policy rate can be computed from the Euler equation as:

$$\tilde{R}_t = \sigma^{-1} (\mathbb{E}_t \hat{x}_{t+1} - \hat{x}_t) + \mathbb{E}_t \hat{\pi}_{t+1} + \tilde{r}_t^*$$

5. The next step is to determine whether the optimal effective policy rate can be delivered as an interior optimum for the policy instruments. Under the assumption that the solution for q is an interior solution ($\underline{q}_t \leq q_t \leq \bar{q}_t$), it is the case that $\lambda_t^{\bar{q}} = \lambda_t^{\underline{q}} = 0$ and so the first order condition for q can be solved to give:[66]

$$\begin{aligned} q_t = & \frac{\omega_q}{\omega_q + \omega_{\Delta q}} q^* + \frac{\omega_{\Delta q}}{\omega_q + \omega_{\Delta q}} q_{t-1} \\ & - \frac{\beta}{\omega_q + \omega_{\Delta q}} \left[\frac{\partial \mathbb{E}_t \mathbf{\Pi}\left(q_t\right)}{\partial q_t} \omega_\pi \left(\hat{\pi}_t - \pi^*\right) + \frac{\partial \mathbb{E}_t \tilde{\mathcal{L}}_{t+1}}{\partial q_t} \right] \end{aligned} \tag{E.11}$$

6. If the solution for q_t from equation (E.11) is indeed an interior solution, the optimal policy rate can be computed as $\hat{R}_t = \tilde{R}_t + \gamma q_t$. If this value of $\hat{R}_t$ is greater than the zero bound, the solution computed from these steps represents the equilibrium.

E.2 Bounded instruments

The steps presented in Appendix E.1 may fail to deliver a valid equilibrium for two reasons: the implied level of quantitative easing may violate the upper and lower bounds on q; or the implied level of the policy rate required to deliver the desired effective policy rate may violate the zero bound.

Suppose first that Step 5 in Appendix E.1 delivers a solution for q_t which violates the bounds. In this case, the solution for q_t is set to the relevant bound value.[67] If the optimal policy rate $\hat{R}_t = \tilde{R}_t + \gamma q_t$ computed using this value for q_t is greater than the zero bound, then this represents the equilibrium.

In the event that the value of $\hat{R}_t$ computed in Step 6 in Appendix E.1 is below the zero bound, the system is solved as follows. I first assume that, even though the zero bound on the policy instrument is binding, there is an interior solution for QE (so that

[66] Recall also that at this stage in the solution process, it is assumed that the zero bound on the policy rate does not bind, so that $\lambda_t^x = 0$.

[67] If the solution to (E.11) is less than $\underline{q}$, then set $q_t = \underline{q}$. If the solution is greater than $\bar{q}$, set $q_t = \bar{q}$.

$\lambda_t^{\bar{q}} = \lambda_t^{\underline{q}} = 0$).[68] In this case, the equilibrium solves the following system:

$$\begin{bmatrix} \omega_\pi & 0 & 0 & 0 & 0 & -1 \\ 0 & \omega_x & 0 & 0 & -1 & \kappa \\ 0 & 0 & 0 & \omega_q + \omega_{\Delta q} & \mathcal{D}_{\hat{x}} + \sigma \mathcal{D}_{\hat{\pi}} + \sigma\gamma - \beta\sigma\xi\mathcal{D}_q & \beta\mathcal{D}_{\hat{\pi}} \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & \sigma & -\sigma\gamma & 0 & 0 \\ 1 & -\kappa & 0 & 0 & 0 & 0 \end{bmatrix} \times \begin{bmatrix} \hat{\pi}_t \\ \hat{x}_t \\ \hat{R}_t \\ q_t \\ \lambda_t^x \\ \lambda_t^\pi \end{bmatrix}$$
$$= \begin{bmatrix} \omega_\pi \pi^* \\ 0 \\ \omega_q q^* + \omega_{\Delta q} q_{t-1} - \beta \frac{\partial \mathbb{E}_t \tilde{\mathcal{L}}_{t+1}}{\partial q_t} \\ 1 - \beta^{-1} \\ \mathbb{E}_t \hat{x}_{t+1} + \sigma \left[\mathbb{E}_t \hat{\pi}_{t+1} + \tilde{r}_t^*\right] \\ \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \end{bmatrix}$$

where the following notation for derivatives of expectations of variable z_{t+1} with respect to q_t is used to simplify the expressions:

$$\mathcal{D}_z \equiv \frac{\partial \mathbb{E}_t z_{t+1}}{\partial q_t}$$

This system can be solved by matrix inversion. If the solution for q_t is an interior solution, then the equilibrium allocations have been found. Otherwise the solution is found by setting q_t to the relevant bound value and solving the following recursive solutions.

1. The output gap is

$$\hat{x}_t = \begin{cases} \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[1 - \beta^{-1} - \gamma\bar{q} - \mathbb{E}_t \hat{\pi}_{t+1} - \tilde{r}_t^*\right] & \text{if } q_t = \bar{q} \text{ binds} \\ \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[1 - \beta^{-1} - \gamma\underline{q} - \mathbb{E}_t \hat{\pi}_{t+1} - \tilde{r}_t^*\right] & \text{if } q_t = \underline{q} \text{ binds} \end{cases}$$

2. Inflation is:

$$\hat{\pi}_t = \kappa \hat{x}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t$$

E.3 Solution algorithm details

The preceding subsections detailed the optimal policy problem and elements of the solution. To solve for policy functions $\boldsymbol{X}, \boldsymbol{\Pi}, \boldsymbol{Q}, \boldsymbol{R}$ a simple policy function iteration scheme is used. The policy functions are defined over a grid for the state vector $s \equiv \{u, r^*, q_{-1}\}$ formed as a tensor product of three linearly spaced vectors. The vector for q_{-1} is defined on the range $[\underline{q}, \bar{q}]$ with 101 nodes and the grids for u_t and r_t are specified across ± 4 standard deviations with 25 and 101 nodes respectively. The state space is therefore $\mathcal{S} \equiv \mathcal{S}_u \times \mathcal{S}_{r^*} \times \mathcal{S}_{q_{-1}}$ with typical element s.

The iteration scheme to update the estimates of the policy functions is as follows, where the superscript j refers to iteration:

[68] The reasoning is the following. The solution computed in the previous subsection assumed that there was an interior solution. Therefore the solution for q_t was derived conditional on $\lambda_t^x = 0$. If the zero bound on $\hat{R}_t$ is binding, the multiplier λ_t^x is non-zero and hence the solution for q_t implied by the general first order condition (E.5) may admit an interior solution.

1. Form $\boldsymbol{P}^j \equiv \left\{\boldsymbol{X}^j, \boldsymbol{\Pi}^j, \boldsymbol{Q}^j, \boldsymbol{R}^j\right\}$. To do so:

   (a) For each $s \in \mathcal{S}$ compute solutions for $\{\hat{x}, \hat{\pi}, q, \hat{r}\}$ by using the procedure set out in Appendix E.1 and (if necessary) the steps set out in Appendix E.2.

   (b) These solutions are conditional on expected outcomes, which are computed using Gauss-Hermite quadrature for the shocks to exogenous states (ε^u and ε^r) using five nodes for each shock.

   (c) Expectations are computed by integrating over the estimated policy functions from the previous iteration: $\left\{\boldsymbol{X}^{j-1}, \boldsymbol{\Pi}^{j-1}, \boldsymbol{Q}^{j-1}, \boldsymbol{R}^{j-1}\right\}$.

   (d) Derivatives of the loss function are computed using a finite difference method on $\mathcal{S}_{q_{-1}}$. Linear interpolation is used to estimate the value of the derivative conditional on the value of the current stock of QE (i.e., q, which is the relevant state variable for expected outcomes) using the *previous* estimate of the policy function $\boldsymbol{Q}^{j-1}$.

2. Update the estimate of the loss function using

$$\tilde{\mathcal{L}}_t = \frac{\omega_x}{2}\hat{x}_t^2 + \frac{\omega_\pi}{2}\left(\hat{\pi}_t - \pi^*\right)^2 + \frac{\omega_q}{2}\left(q_t - q^*\right)^2 + \frac{\omega_{\Delta q}}{2}\left(q_t - q_{t-1}\right)^2 + \beta \mathbb{E}_t \tilde{\mathcal{L}}_{t+1}$$

   (a) The loss function is computed for each $s \in \mathcal{S}$. Expectations are computed using the quadrature scheme described above.

   (b) Convergence of the solution for the loss function is very slow, so the solution algorithm performs policy function iteration a number of times (up to 10 depending on whether successive iterations are within solution tolerance) for each iteration on the policy functions.

3. If $\min\left|\text{vec}\left(\boldsymbol{P}^j - \boldsymbol{P}^{j-1}\right)\right| < \epsilon$ then stop, otherwise return to Step 1. I set $\epsilon = 10^{-6}$.

One thing to note about this solution approach is that it does not involve solving a fixed point problem in step 1(a). For example, the first order condition for q for an interior solution, (E.11), implies that q is a function of the derivative of expected inflation with respect to q. So the right hand side of the equation for q is itself a function of q. One approach would be to solve for q as the fixed point of the equation for each point in $\mathcal{S}$. However, such an approach is computationally intensive, so instead I use derivative of expected inflation with respect to q evaluated using the previous iteration of the q policy function (that is, $\boldsymbol{Q}^{j-1}$). As the policy functions converge, $\boldsymbol{Q}^j \to \boldsymbol{Q}^{j-1}$, so the derivative is computed at the fixed point value of q. Approximations for the policy functions for the yield to maturity are computed using a simple variant of the policy function iteration approach.

F Equilibrium without the zero bound

In the case in which there is no zero bound on the short term interest rate, the model is linear and it is possible to derive analytical expressions for the endogenous variables in terms of the state variables r^* and u.

I consider the case in which the initial stock of QE inherited by the monetary policy-maker is $q_{-1} = q^*$. In this case, the subsequent choices of QE satisfy $q_t = q^*, \forall t$. This follows from inspection of the first order condition (E.5). First note that an interior solution for QE implies $\lambda_t^{\bar{q}} = \lambda_t^{\underline{q}} = 0$ and the absence of a zero bound on $\hat{R}_t$ implies that $\lambda_t^R = 0$. The remaining condition for the conjectured policy $q_t = q^*$ to be optimal is that $\frac{\partial \mathbb{E}_t \hat{\pi}_{t+1}}{\partial q_t} = 0$. This is indeed the case if $\hat{R}_t$ can always be freely chosen. In that case, equilibrium outcomes for the output gap and inflation are uniquely pinned down by the effective policy rate $\tilde{R}_t$ which in turn can be set to any required value by an appropriate choice of $\hat{R}_t$. This implies that q_{t-1} ceases to be a meaningful state variable in the model, because monetary conditions can be chosen independently of the level of quantitative easing.

If $q_t = q^*$, then the optimal allocations for inflation and the output gap satisfy:

$$\hat{x}_t = -\frac{\kappa\omega_\pi}{\omega_x}\left(\pi_t - \pi^*\right)$$

which can be substituted into the Phillips curve (13) to give:

$$\pi_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} - \frac{\kappa^2\omega_\pi}{\omega_x}\left(\pi_t - \pi^*\right) + u_t$$

or

$$\begin{aligned}
\pi_t &= \frac{\beta\omega_x}{\omega_x + \kappa^2\omega_\pi}\mathbb{E}_t\hat{\pi}_{t+1} + \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi}u_t \\
&= \sum_{i=0}^{\infty}\left(\frac{\beta\omega_x}{\omega_x + \kappa^2\omega_\pi}\right)^i\left(\frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi}u_{t+i}\right) \\
&= \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi}\sum_{i=0}^{\infty}\left(\frac{\beta\omega_x\rho_u}{\omega_x + \kappa^2\omega_\pi}\right)^i u_t \\
&= \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t
\end{aligned}$$

where we make use of the fact that:

$$\sum_{i=0}^{\infty}\left(\frac{\beta\omega_x}{\omega_x + \kappa^2\omega_\pi}\right)^i = \frac{1}{1 - \frac{\beta\omega_x}{\omega_x + \kappa^2\omega_\pi}} = \frac{\omega_x + \kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}$$

and similarly

$$\sum_{i=0}^{\infty}\left(\frac{\beta\rho_u\omega_x}{\omega_x + \kappa^2\omega_\pi}\right)^i = \frac{1}{1 - \frac{\beta\rho_u\omega_x}{\omega_x + \kappa^2\omega_\pi}} = \frac{\omega_x + \kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}$$

The targeting criterion implies that the output gap is:

$$\begin{aligned}
\hat{x}_t &= -\frac{\kappa\omega_\pi}{\omega_x}\left(\frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t - \pi^*\right) \\
&= -\frac{\kappa\omega_\pi}{\omega_x}\left(\frac{\kappa^2\omega_\pi - \omega_x - \kappa^2\omega_\pi + \beta\omega_x}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \frac{\omega_x}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t\right) \\
&= \frac{(1-\beta)\,\kappa\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - \frac{\kappa\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t
\end{aligned}$$

To solve for the nominal interest rate, note that:

$$\mathbb{E}_t \hat{x}_{t+1} = \frac{(1-\beta)\,\kappa\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - \frac{\kappa\omega_\pi\rho_u}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t$$
$$\mathbb{E}_t \pi_{t+1} = \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \frac{\omega_x\rho_u}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t$$

Using this in the IS curve (12) gives:

$$-\frac{\kappa\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t = -\frac{\kappa\omega_\pi\rho_u}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t - \sigma\left(\hat{R}_t - \nu\delta q^* - \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - \frac{\omega_x\rho_u}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t - r^*\right)$$

So

$$\sigma\left(\hat{R}_t - \nu\delta q^* - \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - \frac{\omega_x\rho_u}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t - r^*\right) = \frac{\kappa\omega_\pi\,(1-\rho_u)}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t$$

Or

$$\sigma\left(\hat{R}_t - \nu\delta q^* - \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - r^*\right) = \frac{\sigma\omega_x\rho_u + \kappa\omega_\pi\,(1-\rho_u)}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u}u_t$$

Which implies that the nominal interest rate satisfies:

$$\hat{R}_t = r_t^* + \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \nu\delta q^* + \frac{\sigma\omega_x\rho_u + \kappa\omega_\pi\,(1-\rho_u)}{\sigma\,(\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u)}u_t \tag{F.1}$$

With $q_t = q^*, \forall t$, the yield to maturity is given by:

$$\begin{aligned}
\hat{\mathcal{R}}_t &= (1-\chi\beta)\left(\hat{R}_t - \nu\,(1+\delta)\,q^*\right) + \chi\beta\mathbb{E}_t\hat{\mathcal{R}}_{t+1} \\
&= (1-\chi\beta)\,\mathbb{E}_t\sum_{j=0}^{\infty}(\chi\beta)^j\left(\hat{R}_t - \nu\,(1+\delta)\,q^*\right) \\
&= -\nu\,(1+\delta)\,q^* + (1-\chi\beta)\,\mathbb{E}_t\sum_{j=0}^{\infty}(\chi\beta)^j\left[\begin{array}{l} r_{t+j}^* + \frac{\kappa^2\omega_\pi}{\omega_x+\kappa^2\omega_\pi-\beta\omega_x}\pi^* + \nu\delta q^* \\ + \frac{\sigma\omega_x\rho_u+\kappa\omega_\pi(1-\rho_u)}{\sigma(\omega_x+\kappa^2\omega_\pi-\beta\omega_x\rho_u)}u_{t+j} \end{array}\right] \\
&= -\nu\,(1+\delta)\,q^* + \frac{1-\chi\beta}{1-\chi\beta\rho_r}r_t^* + \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* + \nu\delta q^* \\
&\quad + \frac{1-\chi\beta}{1-\chi\beta\rho_u}\frac{\sigma\omega_x\rho_u + \kappa\omega_\pi\,(1-\rho_u)}{\sigma\,(\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u)}u_t
\end{aligned}$$

Collecting terms gives:

$$\hat{\mathcal{R}}_t = \frac{1-\chi\beta}{1-\chi\beta\rho_r}r_t^* + \frac{\kappa^2\omega_\pi}{\omega_x + \kappa^2\omega_\pi - \beta\omega_x}\pi^* - \nu q^* + \frac{1-\chi\beta}{1-\chi\beta\rho_u}\frac{\sigma\omega_x\rho_u + \kappa\omega_\pi\,(1-\rho_u)}{\sigma\,(\omega_x + \kappa^2\omega_\pi - \beta\omega_x\rho_u)}u_t$$