Bank of England

# Optimal quantitative easing and tightening


**Staff Working Paper No. 1,063**

March 2024



**Richard Harrison**




Bank of England

Staff Working Paper No. 1,063

# Optimal quantitative easing and tightening

Richard Harrison[(1)]

## Abstract

This paper studies optimal monetary policy in a New Keynesian model with portfolio frictions that create a role for the central bank balance sheet as a policy instrument. Central bank purchases of long-term government debt ('quantitative easing') reduce average portfolio returns, thereby increasing aggregate demand and inflation. Optimal time-consistent policy prescribes large and rapid asset purchases when the policy rate hits the zero bound. Optimal balance sheet reduction ('quantitative tightening') is more gradual. A central bank that pursues a flexible inflation target can achieve similar welfare to optimal policy if quantitative tightening is calibrated appropriately.

**Key words:** Quantitative easing, quantitative tightening, 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 authors, and not necessarily those of the Bank of England or its committees. This paper is a substantially revised version of 'Optimal quantitative easing' (Harrison (2017)). I am grateful to Yunus Aksoy, William Pagel, Peter Sinclair, Ron Smith, Jacob Stevens, Dominik Thaler, Carlos Thomas, Stephen Wright and to conference attendees at the 2017 Money, Macro and Finance conference, the 2017 ECB workshop on non-standard monetary policy measures, the 2018 Oxford-Indiana Macroeconomic Policy Conference, the 2023 De Nederlandsche Bank Annual Research Conference and to seminar participants at the Bank of Canada, Bank of England, Sveriges Riksbank, Birkbeck and NIESR for helpful comments on earlier drafts. I thank in particular Thomas Wieladek and Martin Weale for sharing their results.

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

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


ISSN 1749-9135 (on-line)

# 1 Introduction

Central bank purchases of long-term government debt financed by the creation of reserves (quantitative easing, 'QE') were deployed in some economies during the Global Financial Crisis, when short-term interest rates became constrained at their lower bounds for the first time. Since then, QE has become a regular part of the monetary policy toolkit and subsequent shocks, particularly the Covid-19 pandemic, prompted further expansions of central bank balance sheets, in some cases to historically unprecedented levels (Figure 1).

**Figure 1:** Bank of England balance sheet as a percentage of nominal UK GDP, 1697-2023

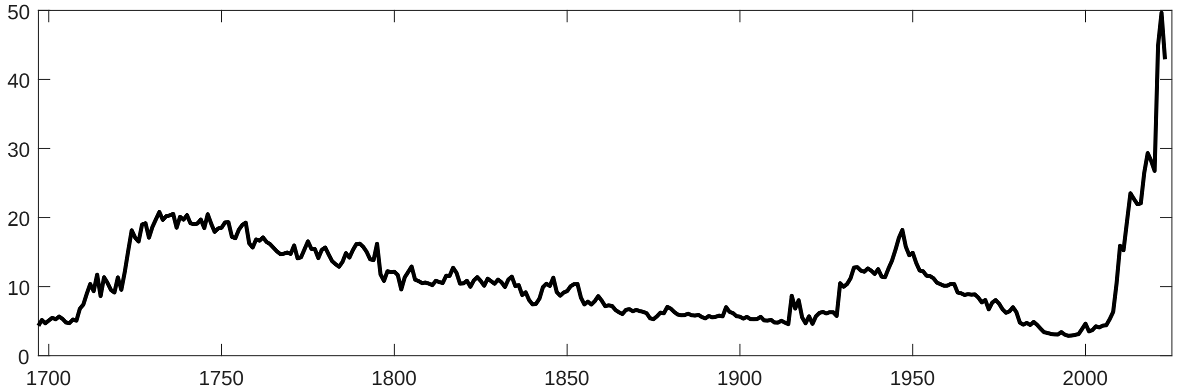


*Notes:* Data from the Bank of England "Millennium of macroeconomic data" dataset (Thomas et al., 2010). Ratio is computed as Bank of England consolidated balance sheet at the end of February of each year, divided by nominal GDP in the previous year.

How should the central bank balance sheet be optimally deployed when the short-term interest rate becomes constrained at the lower bound? Should central bank balance sheets be subsequently reduced to levels more consistent with historical norms? If so, what is the best strategy to follow? This paper explores these questions using a workhorse New Keynesian model extended to include portfolio frictions that provide a role for the central bank balance sheet as a policy instrument.

The model includes frictions that create a wedge between returns on short-term and long-term bonds. This wedge depends on the relative supplies of assets thereby capturing a "portfolio balance channel" through which many monetary policymakers believe that QE operates. The nature of the portfolio frictions is such that, in general, the policy rate and the central bank balance sheet are perfect substitutes in terms of their influence on the overall monetary policy stance, which can be summarized in terms of a 'shadow rate' of interest.

The model provides a laboratory for studying the optimal deployment of the central bank balance sheet alongside interest rate policy. The monetary policymaker sets optimal time-consistent policy to minimize a loss function derived from a quadratic approximation to the welfare of the representative household. This loss function features four components. The first two components are familiar from the workhorse New Keynesian model: infrequent price adjustment implies that there are welfare costs arising from fluctuations in inflation and the output gap. The final two components depend on the size and change in the central bank balance sheet because the portfolio frictions through which the central bank balance sheet affects real activity and inflation also reduce welfare. The weights on these components of the loss function depend on the parameters that determine the strength of the portfolio frictions.

The fact that balance sheet policies are associated with welfare-relevant costs implies an important asymmetry with the policy rate even though they are perfect substitutes in terms of their effects on the shadow rate. In particular, optimal policy equalizes the marginal portfolio costs generated by the use of balance sheet policies with the marginal benefits of better stabilizing the output gap and inflation. Therefore, relative to the workhorse New Keynesian model, optimal policy includes an additional trade-off criterion equating the costs and benefits of balance sheet policies. This trade-off criterion drives many of the key results in the paper.

Several quantitative experiments are used to study optimal asset purchases and balance sheet unwind ('quantitative tightening', QT). For these experiments the values of the parameters governing portfolio frictions are chosen to match empirical estimates of the effects of QE on long-term bond rates in the United Kingdom. That matching exercise reveals that portfolio frictions that depend both on financial intermediaries' portfolio composition and the change in the composition are required to match the estimated effects of QE on the long-term bond rate. The model is solved using projection methods, accounting for the non-linearities generated by the zero bound on the policy rate and the possibility that bounds may also apply to the size of the central bank balance sheet.

These experiments reveal that optimal balance sheet policy reduces welfare losses by more than 25% compared with the case in which the short-term policy rate is the only instrument. Asset purchases are often rapid, with large scale asset purchases commencing as soon as the short-term policy rate hits the zero bound. QT is more gradual, to mitigate the costs of changes in the portfolio composition.

These features of optimal policy behavior are consistent with asset purchase programs that have been conducted in the United Kingdom and the United States. In those economies, asset purchases were sizable and rapid, occurring when the policy rate hit the effective lower bound. The fact that the balance sheet should be unwound more slowly than the typical pace of QE is also consistent with recently adopted QT strategies in both economies.

One aspect of optimal QT in the model is less clearly consistent with recent policy actions. Quantitative tightening in the United Kingdom and United States did not begin until the short-term policy rate had been increased from the lower bound. In contrast, optimal policy in the model implies that balance sheet reduction will typically start before the short-term policy rate lifts off from the lower bound. The optimality of such a QT strategy reflects minimization a loss function based on social welfare, thereby accounting for the welfare costs of portfolio frictions. However, the monetary policy objectives of many central banks focus on the stabilization of inflation and some measure of resource utilization, without direct regard for portfolio allocation frictions.

Accordingly, the paper also explores the case in which the monetary policymaker is guided by this type of 'flexible inflation targeting' objective. To mimic real-world QT strategies, a simple QT rule determines the balance sheet away from the lower bound. This rule reduces the size of the balance sheet at a fixed rate, with the policy rate adjusting to implement the optimal monetary policy stance. If the pace of balance sheet unwind is sufficiently rapid, it is possible to achieve similar welfare losses to the fully optimal policy. To achieve this outcome requires a QT strategy that includes sales of long-term debt held on the balance sheet. Conversely, 'passive' unwind, achieved by allowing previously purchased debt to mature, leads to substantially higher welfare costs.

This paper relates to several strands of literature on unconventional monetary policy.

A large number of empirical studies have estimated the effects of quantitative easing on financial markets and the macroeconomy, most of them finding that QE has had detectable effects on output and inflation (see Bhattarai and Neely (2022) for a recent comprehensive review). One of these studies, Weale and Wieladek (2016), is used to set the values of the parameters that govern portfolio frictions in the model.

There is a growing number of models of the transmission mechanism of quantitative easing policies. In standard models without financial frictions, central bank purchases of long-dated government debt do not affect asset prices or allocations (see, for example, Wallace, 1981; Eggertsson and Woodford, 2003). Introducing frictions in financial markets can generate a role for quantitative easing as demonstrated by Andrés et al. (2004), Gertler and Karadi (2011), and Vayanos and Vila (2009, 2021) among others.[1] Using a model with bond portfolio frictions focuses the present study on the effects of central bank balance sheet policies directed at government bonds rather than private assets.

Several papers have studied QE using larger models featuring similar portfolio frictions: for example, Chen et al. (2012), Darracq Pariès and Kühl (2017), De Graeve and Theodoridis (2016), Hohberger et al. (2019), Priftis and Vogel (2016) and Mau (2022). However, all of these papers assume that agents' expectations satisfy a certainty equivalence assumption. With the exception of Darracq Pariès and Kühl (2017) and Quint and Rabanal (2017), these papers do not consider the optimal design of central bank balance sheet policies. Neither Darracq Pariès and Kühl (2017) nor Quint and Rabanal (2017) consider potential bounds on the central bank balance sheet or use a welfare-based loss function. While Mau (2022) does use a welfare-based metric the analysis does not account for occasionally binding constraints on the policy instruments.

Recent work has also focused attention on quantitative tightening and the mix of policy instruments during a normalization of monetary policy. Karadi and Nakov (2021) study optimal Ramsey policy under perfect foresight in a rich model featuring banks with occasionally binding constraints on their balance sheets. This makes it optimal for quantitative tightening to proceed slowly to ensure banks' balance sheet constraints remain slack as they recapitalize slowly following a recessionary shock. Airaudo (2022) explores the interplay between the central bank's QT strategy and the nature of fiscal policy behavior, connecting to a broader literature studying the implications of central bank balance sheet policies for monetary and fiscal interactions (see, for example, Del Negro and Sims, 2015; Benigno and Nisticò, 2020; Bhattarai et al., 2022). The scope for fiscal policy to shape the effects of central bank balance sheet policies is abstracted from in the present paper, in order to focus exclusively on the implications of the portfolio balance

[1] Recent work has also studied the role of household heterogeneity in the efficacy and design of quantitative easing and balance sheet unwind (Cui and Sterk, 2021; Sims et al., 2022; Cantore and Meichtry, 2023).

frictions.

The paper lies within a strand of literature considering optimal central bank balance sheet policies in small-scale New Keynesian models. In particular, the model implies that the effects of balance sheet policies can be captured by a 'shadow rate', as in Wu and Zhang (2019). So the policy rate and the balance sheet are perfect substitutes in terms of their influence on the shadow rate, which summarizes the overall stance of monetary policy. This feature of the model drives many of the results, but contrasts with the 'four equation' New Keynesian model of Sims et al. (2021), in which balance sheet policies also affect inflation via the pricing decisions of firms. Sims et al. (2021) study the properties of optimal policy in their model using an ad hoc loss function. The model of Benigno and Benigno (2022) features a liquidity channel, which also implies that the policy rate and central bank balance sheet are not perfect substitutes. They explore the implications of this feature for the optimal mix of policy instruments under Ramsey policy.

The present paper is most closely related to Bonciani and Oh (2021), who examine optimal policy in the Sims et al. (2021) model using the welfare-based loss function and consider both time-consistent policy and commitment policies. Their welfare-based loss function shares some features with the loss function in the present paper as it also contains terms that are determined by the central bank's balance sheet policies. However, in their model the central bank balance sheet does not feature as an endogenous state variable, simplifying some aspects of the analysis. Bonciani and Oh (2021) also examine optimal policy under simplified versions of the loss function, though focus on the importance of the relative weights placed on the output gap and inflation. In contrast to the global solution method used in the present paper, their solution and welfare analysis is based on a perfect-foresight solution (similar to Harrison, 2012) and they do not consider potential constraints on the central bank's balance sheet policies.

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 analyzes the optimal policy problem. Section 4 uses a set of quantitative experiments to explore the properties of the model under optimal policy. Section 5 examines the case in which the central bank pursues a 'flexible inflation targeting' objective and adopts a QT strategy similar to those implemented in the United States and United Kingdom. Section 6 concludes.

# 2 The model

The model is a simple extension to the workhorse three-equation New Keynesian model (Woodford, 2003; Galí, 2015). This highlights the implications of introducing a single additional friction (the portfolio balance channel) and hence the possible value of an additional monetary policy instrument, relative to a widely studied benchmark. Given widespread familiarity with the textbook model, this section focuses on the additional features and with details of the derivation provided in Appendix C.

## 2.1 Government debt and the government budget constraint

There are three assets in the economy: short and long-term nominal government bonds and central bank reserves. The long-term government bonds are modelled as zero coupon bonds with a stochastic maturity date. In particular, there is a positive probability $1-\chi$ that the bond matures in each period. Appendix B.3 shows that it is possible to write budget constraints in terms of a single bond price and a single stock of long-term bonds (see Woodford, 2001; Chen et al., 2012).[2]

The nominal government budget constraint is:

$$B_t + V_t \widetilde{D}_t = R^B_{t-1} B_{t-1} + (1 - \chi + \chi V_t)\, \widetilde{D}_{t-1} - \Omega^C_t - P_t \tau_t$$

where $B$ and $\widetilde{D}$ represent stocks of short-term and long-term debt, $\Omega^C$ denotes remittances from the central bank and $\tau$ represents net tax/transfer payments from/to households. Short-term debt (or 'bills', $B$) is in the form of one-period bonds paying a nominal gross return of $R^B$. The nominal value (price) of a long-term bond is $V$. The right hand side of the budget constraint contains the current value of outstanding issuance of long-term debt. With probability $1-\chi$, each bond matures and the government redeems at par (normalised to unity). By the law of large numbers, the fraction $1-\chi$ of previously issued bonds mature and the remaining fraction ($\chi$) have a market value of $V_t$.

The budget constraint can be expressed in terms of the market value of long-term

[2]The payoff structure of this bond is identical to a bond with a geometrically declining coupon of a particular initial size, analyzed in Harrison (2021) as a simple generalization of the bond introduced by Woodford (2001).

bonds and their one-period return:

$$B_t + D_t = R^B_{t-1} B_{t-1} + R^D_t D_{t-1} - \Omega^C_t - P_t \tau_t \tag{1}$$

where:

$$D_t \equiv V_t \widetilde{D}_t \quad ; \quad R^D_t \equiv \frac{1 - \chi + \chi V_t}{V_{t-1}}$$

## 2.2 Financial intermediation

Financial intermediaries face simple frictions that give rise to a 'portfolio balance channel' that many monetary policymakers have highlighted as a key transmission mechanism of central bank asset purchases (Joyce et al., 2012). For example, Bernanke (2010) writes:

> 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 Reserve's purchases of longer-term securities affect financial conditions by changing the quantity and mix of financial assets held by the public.

The idea that the relative supplies of assets could have effects on their relative prices has a long history, dating back to Tobin (1958, 1969) and Tobin and Brainard (1963), among others. The present paper fits into a more recent strand of literature that embeds simple financial frictions into DSGE models to capture these types of effects.[3] As discussed by Harrison (2017), a wide range of potential frictions can give rise to a relationship between relative asset supplies and relative returns and the simple specification described below is intended to capture them simply. The approach is therefore is similar in spirit to Cúrdia and Woodford (2016), who study optimal policy in a model with a "reduced-form intermediation technology" introduced with a "minimum of structure". One advantage of this type of approach is that it facilitates the comparison with the workhorse New Keynesian model, that does not include such effects.

[3]See, for example, Andrés et al. (2004), Harrison (2011), Chen et al. (2012), Ellison and Tischbirek (2014), Carlstrom et al. (2017). Harrison (2017) assumes that a similar form of portfolio friction applies directly to households, a formulation also adopted by Hohberger et al. (2019) and Cantore and Meichtry (2023).

The stylized financial intermediation sector channels household saving into government debt and consists of a population of perfectly competitive intermediaries. The assumption of perfect competition implies that it is sufficient to focus on a single representative intermediary, $I$, which collects savings ($S$) from households and invests in one-period government bonds ($B$), long-term debt ($D$) and reserves ($Z$).

The period $t$ profit, in nominal terms is given by:

$$\begin{aligned}\Omega_t^I = {} & S_t - B_t^I - D_t^I - Z_t^I + R_{t-1}Z_{t-1}^I + R_{t-1}^B B_{t-1}^I + R_t^D D_{t-1}^I - R_{t-1}^S S_{t-1} \\ & - \Theta P_t \mathcal{M}\left(\delta \frac{Z_t^I + B_t^I}{D_t^I}\right) \\ & - \Theta P_t \mathcal{A}\left(\frac{Z_t^I + B_t^I}{D_t^I} - \frac{Z_{t-1}^I + B_{t-1}^I}{D_{t-1}^I}\right)\end{aligned}$$

The intermediary collects savings, $S$ and allocates them between bonds, debt and reserves, receives interest on previous investments in these instruments and pays households interest on their previous savings. The intermediary also pays convex costs of maintaining and adjusting its relative mix of reserves, bonds and debt, denoted $\mathcal{M}$ and $\mathcal{A}$ respectively. These cost functions satisfy $\mathcal{M}(1) = \mathcal{M}'(1) = \mathcal{A}(0) = \mathcal{A}'(0) = 0$ and $\mathcal{M}''(1) = \tilde{\nu}$ and $\mathcal{A}''(0) = \tilde{\xi}$ and are scaled by the steady-state level of government debt, $\Theta$.

The financial intermediary maximizes the flow of real profits, discounted by the marginal utility of household consumption ($\Lambda$)

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Lambda_t \frac{\Omega_t^I}{P_t}$$

subject to a balance sheet constraint:

$$S_t \geq B_t^I + D_t^I + Z_t^I$$

that will bind in equilibrium.

## 2.3 Fiscal and monetary policies

Central bank balance sheet policies are a prime candidate for study from the perspective of monetary and fiscal policy interactions.[4] However, these considerations are deliberately ignored in order to focus on the monetary policy effects of central bank balance sheet policies.[5]

In particular, the central bank follows a "passive" remittance policy (Benigno and Nisticò, 2020). Remittances are determined by the central bank budget constraint:

$$\Omega_t^C = Z_t - R_{t-1}^Z Z_{t-1} - \left[V_t \widetilde{D}_t^C - (1 - \chi + \chi V_t)\, \widetilde{D}_{t-1}^C\right]$$

and the central bank balance sheet constraint:

$$V_t \widetilde{D}_t^C = Z_t$$

which states that holdings of long-term debt are financed by reserves.[6] Combining these equations reveals that remittances are determined by portfolio revaluation effects:

$$\Omega_t^C = \frac{1 - \chi + \chi V_t}{V_{t-1}} V_{t-1} \widetilde{D}_{t-1}^C - R_{t-1}^Z Z_{t-1} = \left[R_{1,t}^D - R_{t-1}^Z\right] Z_{t-1}$$

In addition, the government implements the following debt issuance policies:

$$\frac{B_t}{P_t} \equiv b_t = b > 0, \quad \forall t \tag{2}$$

$$\frac{D_t}{P_t} \equiv d_t = \delta b, \quad \forall t \tag{3}$$

[4]There is a growing literature studying this aspect. For example, Del Negro and Sims (2015) and Benigno and Nisticò (2020) study potential non-neutralities of QE arising from the separation of government and central bank intertemporal budget constraints. Bhattarai et al. (2022) analyze the case in which the stock of long-term debt is a state variable in the model, because the government budget constraint is a constraint on policy actions. This approach gives rise to the possibility that QE can be used to provide a credible signal that interest rates will remain low in the future.

[5]The central bank balance sheet is particularly stylized, abstracting from money creation and essential reserves provision, so that the only reserves supplied are those created by purchases of long-term government debt. As such, the model is not suited to the study of the optimal long-term size and composition of the central bank balance sheet, which has been considered in other research (Bailey et al., 2020; Vissing-Jorgensen, 2023).

[6]This remittance policy implies that the net worth of the central bank is zero in each period. This policy trivially satisfies the requirements for a "passive" remittance policy as defined by Benigno and Nisticò (2020), namely that the limiting value of the expected (discounted) value of central bank net worth is zero.

which ensure that – absent balance sheet operations by the central bank – financial intermediaries achieve their desired portfolio positions. The assumption of a 'balanced' debt management policy (so that relative bond supplies are always aligned with the desired holdings of intermediaries) gives the monetary policymaker maximal control over private asset holdings.

To further focus on the role of monetary policy, fiscal policy is highly simplified. There is no government spending and net taxes on households are implemented as lump-sum transfers. Conditional on debt issuance policies and asset purchases by the central bank, net transfers to households $T$ are determined by the government budget constraint (1).[7]

Since lump sum taxes adjust to ensure that the total value of debt is fixed, fiscal policy is 'passive' (in the sense of Leeper, 1991). As Benigno and Nisticò (2020) demonstrate, a combination of passive fiscal and remittance policies implies that only the consolidated public sector budget constraint matters for equilibrium allocations. As a result, the only non-neutrality from central bank balance sheet policy operates through the portfolio balance channel.

These assumptions capture key elements of institutional arrangements in many countries, including the United Kingdom and the United States. Government treasury departments (or their agents) are typically tasked with managing the maturity structure of government debt. Their mandate is often expressed in terms of achieving favorable financing conditions for the government and ensuring adequate liquidity in government debt markets (Greenwood et al., 2016a). In the context of the model, debt issuance in line with household portfolio preferences would (other things equal) minimize portfolio adjustment costs and hence the (social) costs of financing a given debt stock. The assumption that debt management policy remains unchanged when monetary policy uses balance sheet policies at the zero bound is also consistent with the institutional arrangements in the United Kingdom and other countries.[8]

[7]This implies that changes in the value of the total government debt stock are transferred to/from households (lump sum) to keep the value of debt constant (a form of balanced budget financing).

[8]In the United Kingdom, 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 (from the letter from the Chancellor of the Exchequer to the Governor of the Bank of England, 3 March 2009): `https://www.bankofengland.co.uk/-/media/boe/files/letter/2009/chancellor-letter-050309`. In the United States, the US Treasury Assistant Secretary for Financial Markets stated that "[the Federal Reserve's] decision to purchase Treasuries in the secondary market does not, and will not, impact our debt management strategy. [...] Fed monetary policy decisions are independent of that calculus" (Miller, 2010). Nevertheless, there is some debate on

Finally, it is convenient to define the asset purchase policy instrument as the fraction of the market value of long-term bonds held by the central bank, denoted by $q$:

$$q_t \equiv \frac{Q_t}{D_t}$$

## 2.4 Households

The optimization problem of the representative household is

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

where $c$ is consumption, $n$ is hours worked and $\phi_t$ denotes a preference shock.

Maximization is subject to the budget constraint

$$S_t = R_t^S S_{t-1} + W_t n_t + T_t + \Omega_t^F + \Omega_t^I - P_t c_t \tag{4}$$

The household invests savings $S$ with the financial intermediary, earning a (gross) nominal rate of return $R^S$. The right hand side of the budget constraint captures income from working $n_t$ hours at nominal wage $W_t$, net transfers/taxes $T_t$ from the government and dividends $\Omega_t^F, \Omega_t^I$ from firms and intermediaries, less spending on consumption goods $c_t$ at price $P_t$.

## 2.5 Firms

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

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

the extent to which US government debt issuance may have offset some of the effects of FOMC asset purchases (see, for example, Greenwood et al., 2016b).

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

Firms produce using a Cobb-Douglas production function of labor and a fixed capital stock (normalized to unity):

$$y_{j,t} = A n_{j,t}^{1-\alpha}$$

where $A$ is a productivity parameter.

A fixed production subsidy is assumed to ensure that the steady state is efficient. Calvo (1983) staggered pricing gives rise to a New Keynesian Phillips curve derived in Appendix C.3 and discussed below.

## 2.6 Market clearing and aggregate output

Market clearing for short-term bonds, long-term bonds and reserves implies that:

$$b_t = b \quad ; \quad \frac{Q_t}{P_t} + d_t^I = d \quad ; \quad z_t^I = z_t$$

where lower case letters denote real-valued debt stocks (for example, $b_t \equiv B_t/P_t$).

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

$$d_t^I = (1 - q_t)\,\delta b$$

Goods market clearing implies that:

$$c_t = y_t - \Theta\mathcal{M}\left(\delta\frac{z_t + b}{d_t^I}\right) - \Theta\mathcal{A}\left(\frac{z_t + b}{d_t^I} - \frac{z_{t-1} + b}{d_{t-1}^I}\right)$$

where $\Theta \equiv b + d$ is the steady-state level of (long-term and short-term) government debt.

Total output satisfies:

$$y_t = \mathcal{D}_t^{-1} A n_t^{1-\alpha}$$

and $\mathcal{D}_t \equiv \int_0^1 \left(\frac{P_{jt}}{P_t}\right)^{-\eta} \mathrm{d}j$ is a measure of price dispersion across firms.

### 2.7 Model equations

Appendix C shows that, when log-linearized around the efficient steady state, the model can be written in terms of a Phillips curve for inflation ($\hat{\pi}$) and an Euler equation for the output gap ($\hat{x}$):[9]

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t \tag{5}$$

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[ \widetilde{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - r_t^* \right] \tag{6}$$

The slope of the Phillips curve is given by $\kappa = \Gamma \Xi$ , where $\Xi \equiv \sigma^{-1} + \frac{\psi+\alpha}{1-\alpha}$ and $\Gamma \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta} \frac{1-\alpha}{1-\alpha+\eta\alpha}$. The 'natural rate of interest' is $r_t^* \equiv -\mathbb{E}_t \left( \hat{\phi}_{t+1} - \hat{\phi}_t \right)$ and the cost push shock is defined as $u_t \equiv -\frac{(1-\theta)(1-\beta\theta)}{\theta} \frac{\eta}{\eta-1} \hat{\eta}_t$. These variables follow exogenous processes given by:

$$r_t^* = \varrho_r r_{t-1}^* + \varepsilon_t^r \tag{7}$$

$$u_t = \varrho_u u_{t-1} + \varepsilon_t^u \tag{8}$$

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

Importantly, the nominal interest rate in the Euler equation (6) is a 'shadow rate' (Wu and Zhang, 2019), which accounts for the effects of the policy rate $\hat{R}_t$ and the effects of balance sheet policies, $\widetilde{q}$:

$$\widetilde{R}_t = \hat{R}_t - \widetilde{q}_t \tag{9}$$

and where

$$\widetilde{q}_t \equiv \gamma q_t - \xi q_{t-1} - \beta \xi \mathbb{E}_t q_{t+1} \tag{10}$$

is the 'effective' balance sheet and $\gamma \equiv \nu + \xi (1+\beta)$, $\nu \equiv \tilde{\nu} (1+\delta)^2$ and $\xi \equiv \tilde{\xi} \left(1+\delta^{-1}\right)^2$. The definition of the 'effective' balance sheet, (10), incorporates the effects of expected and current changes in the balance sheet (driven by the portfolio adjustment costs) and maps them into interest rate equivalent units.

Equation (9) reveals an important feature of the model, namely the strong substitutability between the balance sheet and the policy rate. The shadow rate $\widetilde{R}$ can be

[9] Here, $\hat{z}_t \equiv \ln (z_t/z)$ denotes the log-deviation of variable $z_t$ from its non-stochastic steady state, $z$. The equations are linearized (rather than log-linearized) with respect to $q$.

adjusted either by changes in the policy rate ($\hat{R}$) or the effective balance sheet ($\widetilde{q}$). The perfect substitutability between the policy instruments, together with the fact that their only influence on the economy operates through the Euler equation (6), is important for several features of optimal policy behavior.

Finally, as shown in Appendix B.3, the yield to maturity of the long-term bond is given by:

$$\hat{\mathcal{R}}_t = \chi\beta\mathbb{E}_t\hat{\mathcal{R}}_{t+1} + (1-\chi\beta)\left(\hat{R}_t - \delta^{-1}(1+\delta)\,\widetilde{q}_t\right) \tag{11}$$

# 3 Optimal Quantitative Easing & Tightening

This section analyzes the optimal use of quantitative easing (QE) and quantitative tightening (QT) alongside the short-term policy rate. Section 3.1 sets out the optimal policy problem. Section 3.2 analyzes the first order conditions for optimal policy and focuses on several special cases to develop intuition.

## 3.1 The optimal policy problem

The monetary policymaker sets both instruments to minimize a loss function based on a second-order approximation to the utility of the representative household. Appendix D shows that the loss function is given by:

$$\mathcal{L}_0 = \frac{1}{2}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left(\omega_x\hat{x}_t^2 + \omega_\pi\hat{\pi}_t^2 + \omega_q q_t^2 + \omega_{\Delta q}\left(q_t - q_{t-1}\right)^2\right) \tag{12}$$

where $\omega_x = \Xi$, $\omega_\pi = \frac{\eta}{\Gamma}$, $\omega_q = \nu\Theta$, $\omega_{\Delta q} = \xi\Theta$.

The loss function specifies that the policymaker should seek to stabilize the output gap, inflation and the size (and change in) its balance sheet. The first two terms appear in the welfare-based loss function of the standard New Keynesian model (Woodford, 2003). The third and fourth terms appear because of the portfolio friction that gives balance sheet policies traction.[10] This friction can be mitigated by stabilizing the relative supplies of assets and the rate at which portfolio shares change. Because the maturity structure of

[10]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.

government debt issuance is matched to households' preferred portfolio mix, deviations in the relative supplies of assets are due entirely to the size of the central bank balance sheet, $q_t$.

The policymaker minimizes the loss function (12) subject to (6), (5) and the relevant constraints on the policy instruments:

$$\hat{R}_t \geq \ln \beta \tag{13}$$

$$q_t \geq \underline{q} \tag{14}$$

$$q_t \leq \bar{q} \tag{15}$$

The baseline assumption is that there is no commitment technology that allows the policymaker to make credible time-inconsistent promises about future policy actions. Examining time-consistent policy is motivated by two considerations. The first is that, in the workhorse New Keynesian model, the zero lower bound is not particularly costly if the policymaker is able to make commitments about future interest rate policy (Eggertsson and Woodford, 2003; Adam and Billi, 2006).[11] This implies that the welfare benefits of adding balance sheet policies to the monetary policy toolkit are very small under commitment, as demonstrated in Section 4.4. The second reason is that many central bankers have expressed doubts over their ability to credibly commit to future policy actions (Nakata, 2015).[12] Nevertheless, the optimal commitment solution represents a useful benchmark case and it is used in Section 5 below to assess the scale of welfare improvements for alternative delegation schemes.

In the time-consistent 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 $t+i, i \geq 1$. The equilibrium is a Markov perfect policy in which optimal decisions are a function only of the payoff relevant state variables $(\{u_t, r_t^*, q_{t-1}\})$. The policymaker therefore understands that future policymakers will choose allocations according to time-

[11] Levin et al. (2010) note that if aggregate demand is very sensitive to real interest rates ($\sigma$ is large), 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 using a piecewise-linear solution approach.

[12] This evidence is consistent with the observation that, in the United States and United Kingdom, QE was used as a policy tool before explicit forward guidance that incorporated an intention to overshoot the inflation target. Moreover, even when that types of forward guidance was deployed, there was debate over the extent to which it represented a commitment by policymakers (see, for example, Plosser, 2012).

invariant Markovian policy functions and therefore that its current policy decisions affect future outcomes through their impact on the endogenous state variable ($q$).

## 3.2 Optimal policy

Appendix E shows that the first order conditions of the policy problem are given by:

$$0 = \omega_\pi \hat{\pi}_t - \lambda_t^\pi \tag{16}$$

$$0 = \omega_x \hat{x}_t + \kappa \lambda_t^\pi - \lambda_t^x \tag{17}$$

$$0 = \Theta \widetilde{q}_t + \beta\sigma\xi \mathbb{E}_t \lambda_{t+1}^x + \beta \boldsymbol{D}_t^\Pi \omega_\pi \hat{\pi}_t - \left[ \boldsymbol{D}_t^X + \sigma \boldsymbol{D}_t^\Pi + \sigma\gamma - \beta\sigma\xi \boldsymbol{D}_t^Q \right] \lambda_t^x - \lambda_t^{\bar{q}} - \lambda_t^q \tag{18}$$

$$0 = \sigma \lambda_t^x - \lambda_t^R \tag{19}$$

where $\boldsymbol{D}_t^Z$ is used to denote the derivatives of expected policy functions for variable $Z$ with respect to $q_t$:

$$\boldsymbol{D}_t^Z \equiv \frac{\partial \mathbb{E}_t \boldsymbol{Z}_{t+1}}{\partial q_t}$$

for $\boldsymbol{Z} = \{\boldsymbol{\Pi}, \boldsymbol{X}, \boldsymbol{Q}\}$. The Lagrange multipliers on the constraints (6), (5), (13), (14) and (15) are denoted by $\lambda_t^x, \lambda_t^\pi, \lambda_t^R, \lambda_t^q, \lambda_t^{\bar{q}}$ respectively. Appendix E reports the required Kuhn-Tucker conditions for the multipliers on the inequality constraints. Equation (19) reveals that the Euler equation is a binding constraint on equilibrium allocations only if the lower bound on the policy rate is binding ($\lambda_t^x > 0$ requires $\lambda_t^R > 0$).

The first order condition for the balance sheet (18) shows how the policymaker accounts for the fact that the choice of $q$ at date $t$ may have effects on welfare and future outcomes because the 'date $t+1$' policymaker will inherit the balance sheet chosen by the 'date $t$' policymaker. To build intuition for the optimal balance sheet policy, it is instructive to consider a set of special cases.

### 3.2.1 No portfolio 'adjustment' costs and unconstrained QE

A necessary condition for the central bank balance sheet to influence future outcomes in equation (18) is that $q$ is a state variable. This in turn relies on a role for portfolio 'adjustment' costs: $\tilde{\xi} > 0$. It is therefore instructive to consider the case in which these adjustment costs are zero, so that $\tilde{\xi} = \xi = \omega_{\Delta q} = 0$.

In that case, current balance sheet actions have no effect on expectations or future losses. For an interior solution for $q_t$ ($\lambda_t^{\bar{q}} = \lambda_t^{q} = 0$) the first order condition (18) can be written as:

$$\omega_q q_t = -\sigma\nu\left(\omega_x x_t + \kappa\omega_\pi \pi_t\right) \tag{20}$$

so that the choice of $q$ depends only on the current output gap and inflation.

This condition balances the marginal cost of long-term bond holdings ($\omega_q q_t$) with the marginal benefits of improved output gap and inflation stabilization via the effect of balance sheet policies through the Euler equation, (6). Note that if the zero lower bound does not bind, $\lambda_t^R = \lambda_t^x = 0$ and equations (16) and (17) imply that $\omega_x x_t + \kappa\omega_\pi\pi_t = 0$. Therefore, for balance sheet policies to be used ($q_t > 0$), the Euler equation must be an active constraint on policy choices ($\lambda_t^x > 0$). A corollary of this observation is that the policymaker does not use the balance sheet unless the zero bound is binding.

The logic of this result is simple. The policymaker has access to two instruments that affect the output gap in the same way (via the shadow rate), but one instrument generates social welfare costs since $q$ appears in the loss function (12). When the zero bound on the policy interest rate is not binding, the policymaker is able to choose the optimal output gap using the policy rate alone and can therefore avoid costs associated with using the central bank balance sheet.

### 3.2.2 Unconstrained policy instruments

The case in which constraints on policy instruments never bind is helpful for understanding optimal quantitative tightening (QT) during policy normalization. If the zero bound never binds ($\lambda_t^x = 0$, $\forall t$), equation (19) implies that the Euler equation (6) is not a binding constraint on the equilibrium, so that $\lambda_t^R = \lambda_t^x = 0, \forall t$. If the bounds on $q$ also never bind, then $\lambda_t^{\bar{q}} = \lambda_t^{q} = 0, \forall t$.

In this special case, the model under optimal policy can be written as:

$$
\begin{aligned}
\hat{x}_t &= \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[ \hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \widetilde{q}_t - r_t^* \right] \\
\hat{\pi}_t &= \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t \\
0 &= \omega_x \hat{x}_t + \kappa \omega_\pi \hat{\pi}_t \\
0 &= \Theta \widetilde{q}_t + \beta \boldsymbol{D}_t^{\Pi} \omega_\pi \pi_t
\end{aligned}
$$

This representation shows that the equilibrium is independent of the choice of $q_t$ beyond the level required to implement the optimal value of $\tilde{q}$. This in turn implies that $\boldsymbol{D}_t^{\Pi} = 0$ since, if the zero bound never binds, it is always possible to set the shadow rate using policy rate regardless of the value of $q_t$. Therefore the level of $q$ does not affect equilibrium allocations (and $\boldsymbol{D}_t^{\Pi} = 0$).

This implies that optimal balance sheet policy satisfies $\widetilde{q}_t = 0$ which requires that:

$$
q_t = \frac{\xi}{\gamma} q_{t-1} + \beta \frac{\xi}{\gamma} \mathbb{E}_t q_{t+1}
$$

and hence that

$$
q_t = \zeta q_{t-1} \tag{21}
$$

where $\zeta$ is the solution of

$$
\zeta = \frac{1 \pm \sqrt{1 - 4\beta \left( \frac{\xi}{\gamma} \right)^2}}{2\beta \frac{\xi}{\gamma}} \tag{22}
$$

with absolute magnitude less than unity.

So when constraints on the policy instruments never bind, the policy problem effectively decouples management of the output and inflation terms in the loss function from the management of portfolio adjustment costs. The former are controlled using the policy rate and the latter using balance sheet adjustments.

Finally, note that if the central bank holds no long-term debt on its balance sheet in period 0, the solution is $q_t = 0, \forall t$ which again demonstrates that the policymaker does not use balance sheet policies if the policy rate is (always) unconstrained.

### 3.2.3 The general case

The optimality condition (18) in the general case, where constraints on both policy instruments may bind, is repeated here for convenience:

$$0 = \Theta\widetilde{q}_t + \beta\sigma\xi\mathbb{E}_t\lambda^x_{t+1} + \beta\boldsymbol{D}^\Pi_t\omega_\pi\hat{\pi}_t - \left[\boldsymbol{D}^X_t + \sigma\boldsymbol{D}^\Pi_t + \sigma\gamma - \beta\sigma\xi\boldsymbol{D}^Q_t\right]\lambda^x_t - \lambda^{\bar{q}}_t - \lambda^{\underline{q}}_t$$

In general, optimal balance sheet policy depends on the multipliers associated with bounds on the policy instruments. It is therefore instructive to consider these bounds in turn.

If the bounds on $q$ never bind, the reasoning in Section 3.2.2 applies and the policymaker is unconstrained in their ability to set $\widetilde{q}$ so that $q$ ceases to be a state variable in the model. This further implies that $\boldsymbol{D}^X_t = \boldsymbol{D}^\Pi_t = 0$, $\boldsymbol{D}^Q_t = \boldsymbol{D}^Q$ and that the optimality condition can be written as

$$\Theta\widetilde{q}_t = \left[\sigma\gamma - \beta\sigma\xi\boldsymbol{D}^Q\right]\lambda^x_t - \beta\sigma\xi\mathbb{E}_t\lambda^x_{t+1}$$

which is a generalization of (20).

Again, optimal use of the balance sheet depends on the tightness of the zero lower bound, as captured by $\lambda^x_t$. Other things equal, a more binding ZLB constraint induces greater use of QE.[13] However, it is also notable that $\widetilde{q}_t$ may be negative when the ZLB is not currently binding ($\lambda^x_t = 0$), since $\mathbb{E}_t\lambda^x_{t+1} \geq 0$. An implication is that, when 'close' to the ZLB (so that $\lambda^x_t = 0$ and $\mathbb{E}_t\lambda^x_{t+1} > 0$), the evolution of the balance sheet may have a tightening effect on the stance of policy, other things equal.[14]

Turning to the constraints on $q$, the previous reasoning suggests that unless these constraints are likely to bind, the optimal use of the balance sheet is only weakly affected by the level of $q_{t-1}$. The lower bound on $q$ may constrain the policymaker's ability to set the optimal level of $\widetilde{q}$. However, $\underline{q}$ cannot be a binding constraint on the policymaker's ability to set the optimal shadow rate (since the effects of a lower level of $q$ can be delivered by a higher policy rate, $\hat{R}$).

[13]This follows from the fact that $\boldsymbol{D}^Q < 1$ and $\gamma > \xi$ so that $\left[\sigma\gamma - \beta\sigma\xi\boldsymbol{D}^Q\right] > 0$.

[14]The 'other things equal' condition is important because, away from the ZLB, the overall stance of policy (as measured by the shadow rate $\widetilde{R}$) is implemented via an unconstrained choice of the policy rate.

The upper bound, $\bar{q}$, has the potential to influence equilibrium outcomes in a similar way to the zero bound on the policy rate. When $q_t$ is constrained by the upper bound, the policymaker is unable to reduce the shadow rate by as much as would be optimal if $q_t$ was unconstrained. This gives rise to more deflationary pressure than would be observed in the absence of the upper bound. Moreover, when $q_t$ is sufficiently close to $\bar{q}$, the risk of being constrained by the upper bound in future periods is elevated. The risk of being constrained in future periods has deflationary effects via the inflation expectations channel familiar from analysis of optimal interest-rate policy subject to a zero bound in the workhorse New Keynesian model.

The practical importance of the upper bound on the central bank balance sheet depends on the extent to which optimal balance sheet policy requires large enough asset purchases for $q_t$ to be close to $\bar{q}$. That in turn depends on the parameterization and quantitative behavior of the model, which is considered in the next section.

# 4 Quantitative experiments

This section presents results from quantitative experiments using the model. To capture the non-linearity created by the zero bound on the short-term interest rate, the model is solved with projection methods using the algorithm described in Harrison (2021). Appendix F describes the solution algorithm in more detail.

Section 4.1 details the baseline parameterization of the model and Section 4.2 studies the equilibrium policy functions. A recessionary scenario is used in Section 4.3 to trace out the dynamics of QE and QT and Section 4.4 examines the unconditional distributions of key macroeconomic variables and welfare under alternative assumptions about policy behavior. Section 4.5 summarizes a range of sensitivity tests to assess the robustness of the results. Finally, Section 4.6 provides a general discussion of the results, comparing optimal policy in the model to the way that QE and QT have been implemented in the United Kingdom and United States.

## 4.1 Parameter values

Table 1 reports the baseline parameter values used in the quantitative experiments.[15]

**Table 1:** Parameter values

| | Description | Value | | Description | Value |
|---|---|---|---|---|---|
| $\sigma$ | Intertemporal substitution elasticity | 1 | $\chi$ | Debt (non)-redemption probability | 0.982 |
| $\kappa$ | Slope of Phillips curve | 0.024 | $\delta$ | Long-term to short-term bond ratio | 1.34 |
| $\beta$ | Discount factor | 0.9925 | $\Theta$ | Debt stock/output ratio | 0.81 |
| $\varrho_r$ | Autocorrelation, natural rate | 0.875 | $\nu \times 100$ | Portfolio maintenance cost | 0.38 |
| $100\sigma_r$ | Standard deviation, natural rate | 0.20 | $\xi \times 100$ | Portfolio adjustment cost | 5.97 |
| $\varrho_u$ | Autocorrelation, cost push shock | 0 | $\underline{q}$ | Lower bound on balance sheet | 0 |
| $100\sigma_u$ | Standard deviation, cost push shock | 0.15 | $\bar{q}$ | Upper bound on balance sheet | 0.7 |
| $\eta$ | Elasticity of substitution | 9 | | | |
| $\alpha$ | Capital share in production | 0.25 | | | |
| $\theta$ | Calvo probability | 0.9 | | | |
| $\psi$ | Inverse Frisch elasticity | 5 | | | |

The values for the parameters of the workhorse New Keynesian model (left panel) are chosen in line with values commonly used in the literature. In particular, the values of $\alpha$, $\eta$, $\psi$ and $\sigma$ are taken from Galí (2015, p67), who selects them on that basis. The Calvo probability is chosen so that the slope of the Phillips curve is $\kappa = 0.024$, a commonly used value in the study of New Keynesian models in the presence of the zero bound (see, for example, Adam and Billi, 2006; Eggertsson and Woodford, 2003; Bodenstein et al., 2012). This approach requires setting $\theta = 0.9$ which is high relative to values based on assumptions about the frequency of price adjustments, but is close to estimates from macroeconomic models such as Smets and Wouters (2005).

The value of $\beta$ implies a real interest rate of 3% in the non-stochastic steady state. As shown by Adam and Billi (2007), as $\beta$ 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. This value may be considered rather high, given evidence of persistent declines in the equilibrium real interest rate over time (Del Negro et al., 2019; Cesa-Bianchi et al., 2023). However, 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.[16]

Other parameters that are important in determining the incidence of the zero bound are those governing the shock processes. The persistence of the natural real interest rate

[15] The productivity parameter $A$ is chosen to normalize output to unity in the steady state.

[16] The average UK equilibrium real interest rate for 2000–2009 is estimated to be around 1% by Del Negro et al. (2019). Together with a 2% inflation target this implies a 3% average nominal neutral rate.

is set to $\varrho_r = 0.875$, slightly higher than the values often used in studies of the effects of the lower bound in the workhorse New Keynesian model (Adam and Billi, 2006; Nakov, 2008; Bodenstein et al., 2012). Reflecting the greater persistence, the standard deviation of the process is set to a relatively small value, thereby ensuring that the global solution of the model does not feature deflationary spirals. The values of the parameters governing the cost push shock are taken from the 'RBC calibration' studied by Adam and Billi (2006).

The parameters related to government debt and balance sheet policies are shown in the right panel of Table 1. The values of these parameters are chosen using data and empirical estimates for the United Kingdom (discussed in more detail in Appendix B).

The long-term bond is interpreted as a 10-year zero coupon bond for the purposes of aligning the effects of QE on long-term interests rates in the model with empirical estimates. The value of $\chi$ is therefore chosen so that the long-term bond has a duration of ten years in the non-stochastic steady state (see Appendix B.3). The ratio of long-term to short-term debt ($\delta$) is calibrated using data on outstanding UK government debt published by the Debt Management Office and discussed further in Appendix B.1. These data are also used to calibrate the steady-state ratio of government debt to output, $\Theta$.

**Figure 2:** Impulse responses to exogenous QE shock

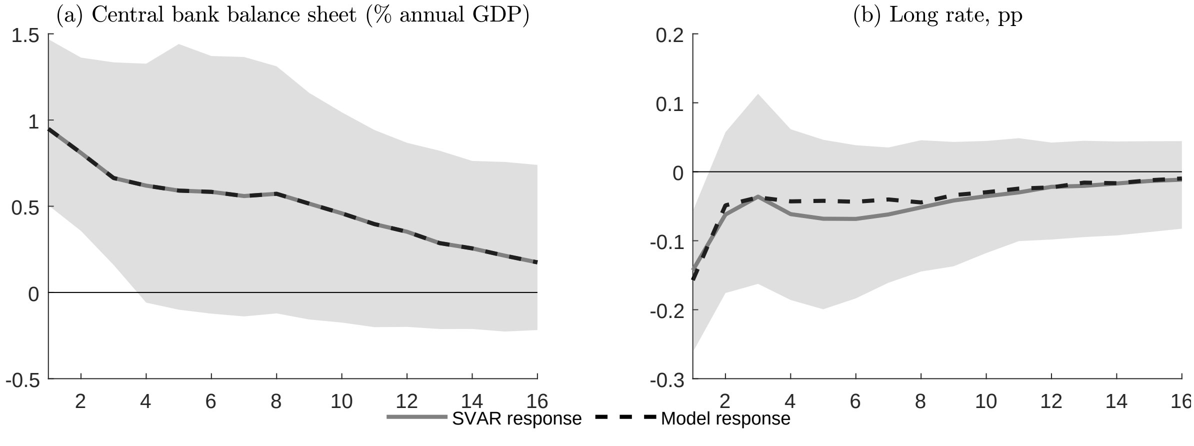


*Notes:* The solid gray line shows the impulse response to an exogenous shock to QE from the structural VAR estimated by Weale and Wieladek (2016) and the gray swathe represents the 95% posterior interval. The black dashed lines represent the response of the model to the same exogenous shock to QE. Full details of the exercise are reported in Appendix B.1.

The values of the parameters governing the portfolio frictions are chosen to match the estimated effect on long-term bond rates of a QE shock, as estimated for the UK economy by Weale and Wieladek (2016).[17] Figure 2 shows the results of matching exercise in which the values of $\nu$ and $\xi$ are chosen to minimize the distance between the empirical response of the long-term bond rate (right panel) to a shock to the path of the central bank asset stock (left panel).[18] The Weale and Wieladek (2016) long-rate response exhibits a 'kink' in the period of the shock, which requires a large elasticity of portfolio 'adjustment' costs ($\xi$) relative to the elasticity of 'maintenance' costs ($\nu$). The results also suggest that the impact of adjustment costs on bond rates are relatively short-lived, consistent with the findings of D'Amico and King (2013).[19]

Finally, the parameters $\underline{q}$ and $\bar{q}$ represent the lower and upper bounds on the central bank balance sheet. Since $q$ is the fraction of the total quantity of outstanding long-term bonds held by the central bank, $q_t$ must be greater than zero and so $\underline{q} = 0$.[20]

It must also be the case that $\bar{q} \leq 1$, since the central bank cannot purchase more than 100% of the existing stock of long-term bonds. An upper bound less than unity can be motivated by the possibility that reducing the amount of debt available to private investors may have undesirable effects on market functioning and liquidity (Logan and Blindseil, 2019; Grimaldi et al., 2021).[21] In light of these considerations, the baseline assumption is $\bar{q} = 0.7$, based on asset purchase limits reported for the United Kingdom and United States in Logan and Blindseil (2019).[22]

## 4.2 Policy functions

Figure 3 plots 'slices' of the equilibrium policy functions, assuming $u_t = 0$ and $q_{t-1} = 0$. These policy functions are compared to a case in which the central bank does not use balance sheet policies (that is $q_t = 0, \forall t$).[23] So the policy function 'slices' show how

[17] I thank Thomas Wieladek and Martin Weale for sharing their results.

[18] Appendix B.1 provides further details.

[19] The finding that adjustment costs appear to be more important than maintenance costs (since $\xi > \nu$) is consistent with the results of De Graeve and Theodoridis (2016). They estimate a flexible functional form for the mapping between maturity structure and the long-short bond spread and find that the data prefers a specification close to a first difference specification (implying $\nu \approx 0$ in the context of the present model).

[20] This rules out the possibility that the central bank can *issue* long-term liabilities that are perfect substitutes for long-term government bonds.

[21] Set against these effects is the possibility that the plentiful supply of reserves that are created as a by-product of asset purchases may enhance financial stability (Greenwood et al., 2016).

[22] Sensitivity to this assumption is considered in Appendix A.1.

[23] Formally, this case corresponds to a situation in which the policymaker acts to minimize the welfare-based loss function using only the short-term nominal interest rate as the policy instrument. This case

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: the effects of optimal balance sheet policy

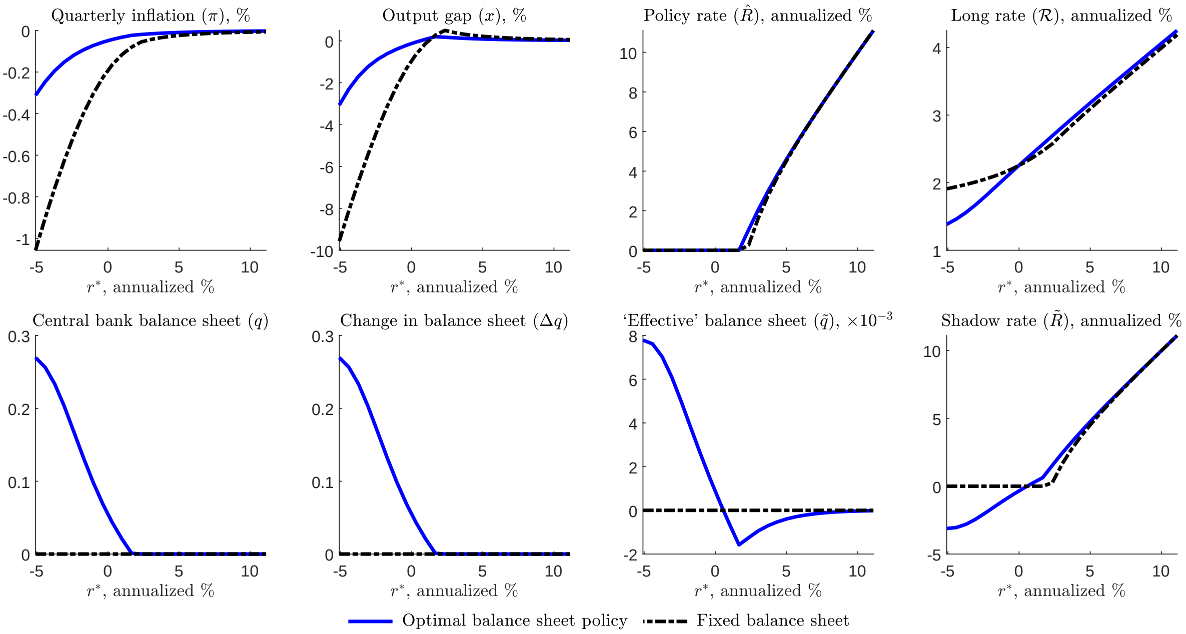


*Notes:* 'Slices' of policy functions for alternative policy assumptions. The solid blue lines are slices of the policy functions conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The dot-dash black lines are policy functions conditional on $u_t = 0$ for a version of the model in which the policymaker does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

Figure 3 demonstrates that when the balance sheet is not used (black dot-dash lines), the policy functions have the same qualitative features as those presented in Adam and Billi (2007). Low values of $r^*$ are associated with the policy rate at the zero bound, a negative output gap and negative inflation. The fact that agents understand that policy will be constrained in this way for low realizations of $r^*$ reduces inflation expectations for values of $r^*$ that are low enough that there is a substantial risk of hitting the zero bound. The downward skew in the distribution of future inflation induces the policymaker to generate a positive output gap for values of $r^*$ slightly above the value at which the policy rate is constrained by the zero bound. As described by Adam and Billi (2007), this is the optimal response to the effect of low inflation expectations on inflation.[24]

When balance sheet policy is used (solid blue lines), the recessionary consequences of low realizations of $r^*$ are mitigated, since asset purchases can be used to reduce the

therefore corresponds to the case of discretionary policy subject to the zero lower bound in the workhorse New Keynesian model, as studied by Adam and Billi (2007) and Nakov (2008), among others.

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

shadow rate $\widetilde{R}$ when the short-term rate is constrained by the zero bound. The policy functions show that the policymaker does not make substantial use of the balance sheet until the short-term policy rate is constrained by the zero bound, confirming the analytical discussion in Sections 3.2.1 and 3.2.2.

When the policy rate is constrained by the zero bound, the optimal size of the central bank balance sheet increases for lower realizations of $r^*$. A larger balance sheet reduces the long-term interest rate and provides additional monetary stimulus, hence reducing the recessionary effects of these realizations of $r^*$. Importantly, the anticipation of additional monetary easing via asset purchases when the short-term interest rate is constrained by the zero bound supports inflation expectations for 'low' values of $r^*$.

**Figure 4:** Policy function comparison: the effects of existing assets on the balance sheet

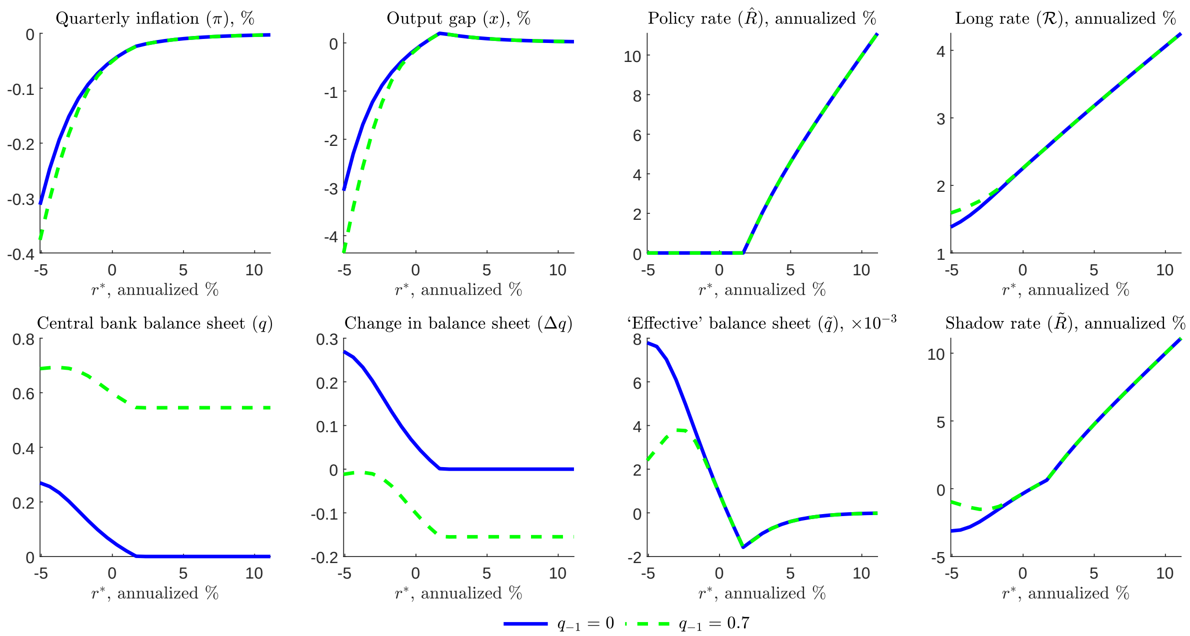


*Notes:* 'Slices' of policy functions for alternative assumptions regarding the central bank's initial balance sheet. The solid blue lines are conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The dashed green lines are conditional on $\{u_t, q_{t-1}\} = \{0, \bar{q}\}$.

To explore the implications of existing asset holdings ($q_{t-1} > 0$) on equilibrium outcomes, Figure 4 compares the policy functions for the cases in which $q_{t-1} = 0$ and $q_{t-1} = \bar{q} = 0.7$ (dashed green lines). This comparison confirms the importance of 'effective QE', $\widetilde{q}$, in determining the equilibrium allocations. For sufficiently high values of $r^*$, it is possible to implement the optimal policy stance regardless of whether the policymaker initially holds the maximal stock of long-term government debt on its balance sheet or no assets on its balance sheet. Indeed, for high realizations of $r^*$, delivering the optimal level of the 'effective' balance sheet ($\widetilde{q}$) requires a *reduction* in asset holdings, so that the change in the balance sheet ($\Delta q$) has the opposite sign to the case in which the central bank has no initial assets on its balance sheet.

However, for very low levels of $r^*$ the upper bound on $q$ constrains the effective balance sheet ($\widetilde{q}$). This prevents the shadow rate from being reduced by as much as in the case in which the central bank initially holds no assets on its balance sheet (the dashed green line lies above the solid blue line).

## 4.3 A recessionary scenario

To shed further light on optimal policy at the zero bound, Figure 5 presents simulations of a recession under alternative policy assumptions. In each case, the initial value of the cost-push shock is zero and the initial level of the natural rate of interest is $-4.3\%$ (measured as an annualized rate). The low initial level of $r^*$ is sufficient to cause a sizable recession if monetary policy is constrained. Beyond the first period, no shocks arrive and $r^*$ evolves according to (7).

**Figure 5:** Simulation of a severe recessionary scenario

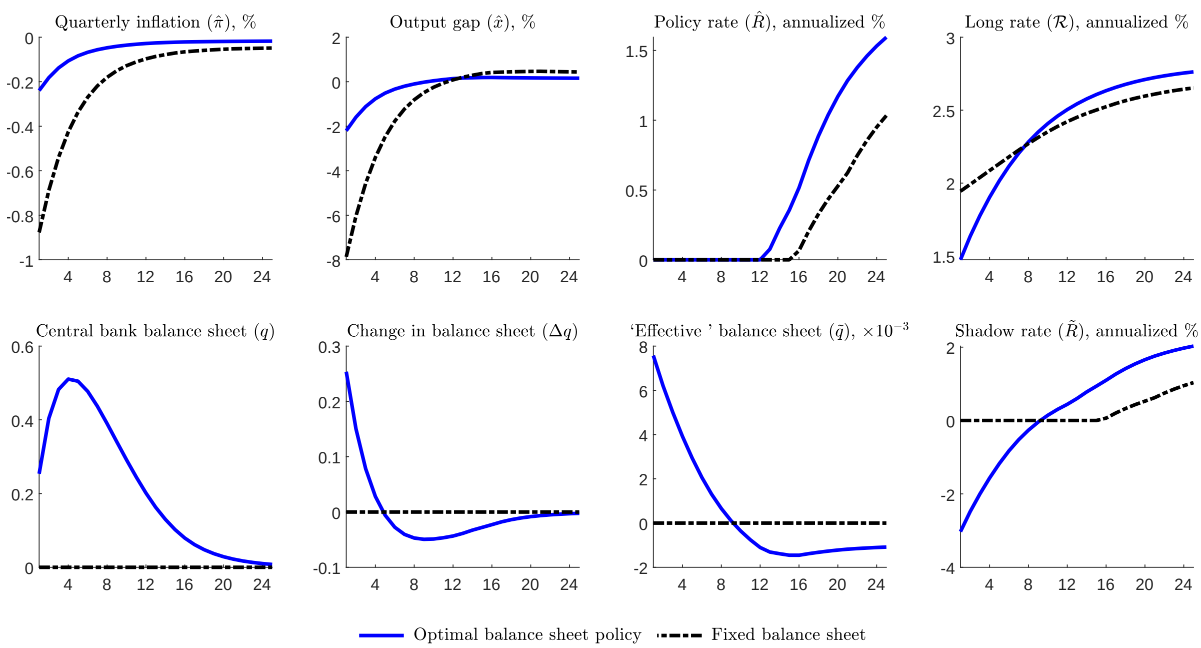


*Notes:* Each simulated path is computed using linear interpolation of the policy functions for $u_t = 0$ and $r_t^* = (r_1^*)^{\varrho_r(t-1)}$ $(t = 1, \ldots, T)$. The value of $r_1^*$ is such that the level of the equilibrium rate (in annualized units) is $-4.3\%$, substantially lower than the deterministic steady state value of 3%. The solid blue lines correspond to the case in which the initial balance sheet is $q_0 = 0$. The dash-dotted black lines show the case in which the policymaker does not use balance sheet policies (so $q_t = 0, \forall t$).

Indeed, the black dash-dot lines show the case in which the policymaker does not use balance sheet policies and tell a story familiar from analysis using the workhorse New Keynesian model. The persistent fall in $r^*$ pushes the policy rate to the zero bound and generates a large and prolonged recession. The deflationary effects of the recession are substantial and it is more than three years before policy is able to lift off from the zero lower bound.

When the policymaker optimally adjusts the central bank balance sheet (solid blue lines) the effects on inflation and the output gap are dampened substantially. This is achieved by lowering the long-term interest rate, via asset purchases, in the near term. The stimulus provided by balance sheet expansion allows the policymaker to lift off from the zero bound faster than otherwise. The bottom-left panel shows the path of $q$ (the fraction of the long-term government debt stock on the central bank balance sheet). The central bank immediately purchases around a quarter of the long-term debt stock, and makes additional asset purchases in the first few quarters. After six quarters the central bank holds more than half of the long-term debt stock.

Following this initial surge of purchases, however, the stock of assets declines gradually over time and policy moves from quantitative easing to quantitative tightening within six quarters. This is shown more clearly in the second panel on the bottom line, which plots the change in the asset stock held by the central bank ($\Delta q$). This reveals an asymmetry between the initial purchases (positive values for $\Delta q$) which are larger in magnitude than the subsequent quantitative tightening phase (the negative values of $\Delta q$ in the latter part of the simulation). The dynamic response of the balance sheet results in a sharp initial expansion in the effective balance sheet ($\widetilde{q}$) and hence an immediate and substantial reduction of the shadow rate, $\widetilde{R}$.

Importantly, the dynamics of the balance sheet demonstrate that it is optimal to start quantitative tightening before the policy rate lifts off from the zero bound. So when it becomes time to tighten policy, it is optimal to do so in a way that reduces some of the portfolio distortions generated by past balance sheet policy actions. This result is consistent with the analysis of the optimal policy conditions in Section 3.2 and discussed further in Section 4.6.

Figure 6 considers the importance of the stock of previously accumulated assets on the

**Figure 6:** Simulation of a severe recessionary scenario: the importance of the initial balance sheet

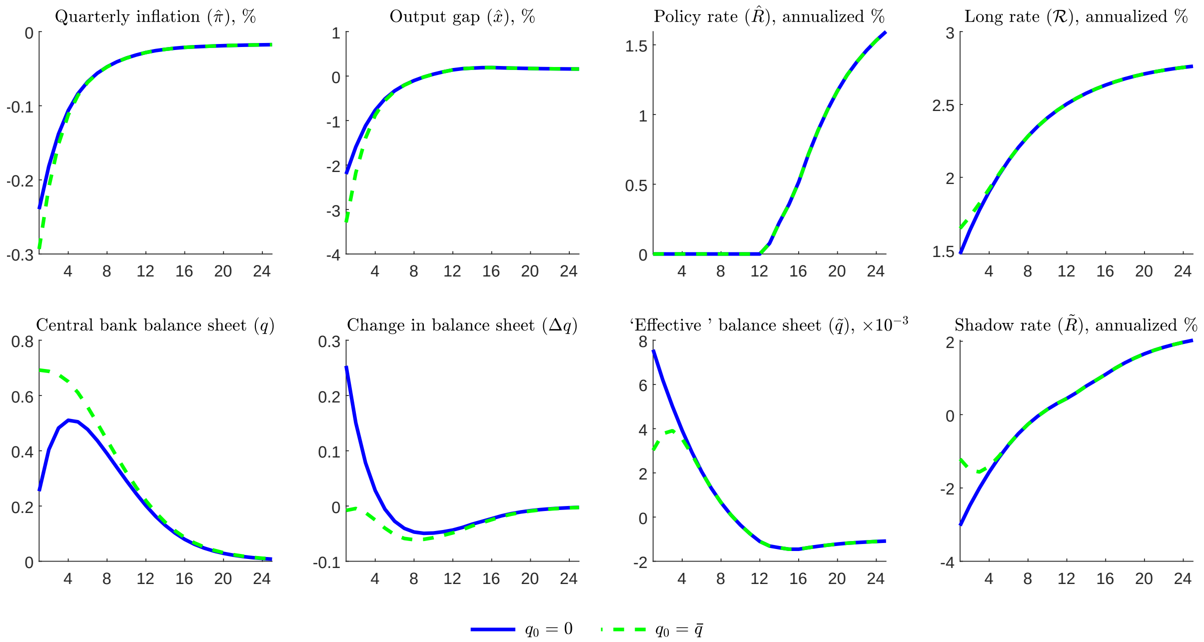


*Notes:* Each simulated path is computed using linear interpolation of the policy functions for $u_t = 0$ and $r_t^* = (r_1^*)^{\varrho_r(t-1)}$ $(t = 1, \ldots, T)$. The value of $r_1^*$ is such that the level of the equilibrium rate (in annualized units) is $-4.3\%$, substantially lower than the deterministic steady state value of 3%. The solid blue lines correspond to the case in which the initial balance sheet is $q_0 = 0$. The green dashed lines correspond to the case in which the initial balance sheet is $q_0 = 0.7$.

central bank's balance sheet ($q_0$). Consistent with the results in Section 4.2, a policymaker that experiences the same recessionary state (annualized $r^*$ of $-4.3\%$) but initially holds a maximal stock of long-term government debt on its balance sheet ($q_0 = \bar{q} = 0.7$, green dashed lines) cannot achieve as favorable outcomes for inflation and the output gap. In this case, the upper bound on the balance sheet initially binds, preventing the policymaker from implementing a sufficiently positive $\tilde{q}$. As a result, the stance of policy is tighter than the case in which $q_0 = 0$ resulting in weaker paths for the output gap and inflation. Over time, however, it becomes possible to implement the optimal stance by *reducing* the stock of purchased assets. So in the latter part of the simulation, the equilibrium outcomes are independent of the initial size of the balance sheet (the solid blue and dashed green lines coincide).

## 4.4 Macroeconomic distributions and welfare

The results in Figures 5 and 6 suggest that there may be a skew in the distribution of balance sheet policy actions so that it is more common to observe large balance sheet expansions than large contractions. That is because large scale asset purchases can be triggered by a large recessionary shock when the policy rate is constrained by the zero bound but quantitative tightening typically occurs slowly and at least partially during periods in which the short-term policy rate is unconstrained by the zero bound. Figure 7 confirms this intuition by plotting the frequency distributions of the balance sheet ($q$) and the change in the balance sheet ($\Delta q_t$) from a stochastic simulation of the model. The distribution of $\Delta q$ exhibits an upward skew.

**Figure 7:** Distributions of the balance sheet and changes in the balance sheet

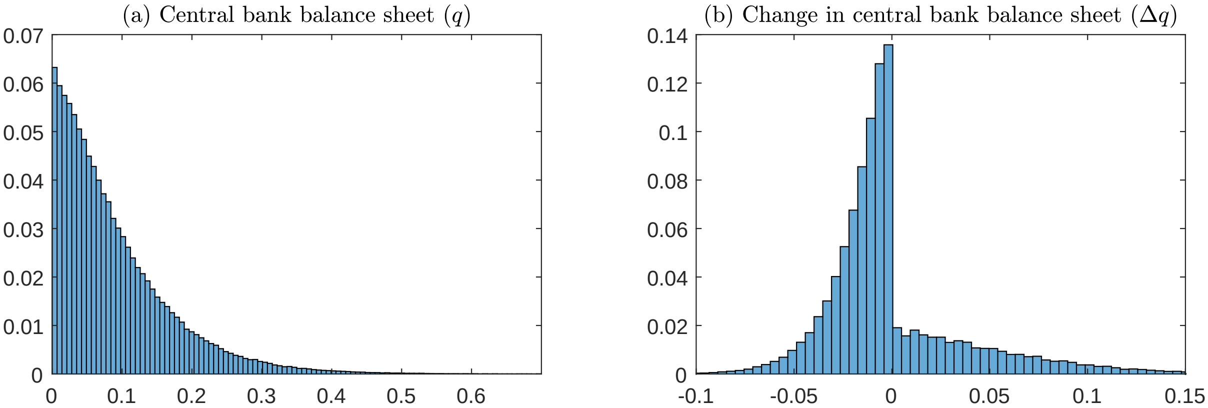


*Notes:* The histograms record the frequency distributions of outcomes for QE ($q$) the change in QE ($\Delta q_t$) from a stochastic simulation of 500,000 periods.

The preceding simulations also suggest that optimal (time-consistent) use of balance sheet policies improves welfare by allowing the policymaker to use an additional instrument to offset the effects of shocks on output and inflation. Table 2 confirms this by reporting the means of key variables for a simulation of 500,000 periods.[25] The mean of the period loss (that is, $\omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2 + \omega_q q_t^2 + \omega_{\Delta q} \left(q_t - q_{t-1}\right)^2$) is also reported.[26]

The left side of Table 2 considers time-consistent policy and shows the results from the baseline version of the model (when both $\hat{R}$ and $q$ are policy instruments) and a variant in which the balance sheet is not used and the policymaker sets $q_t = 0, \forall t$. The right side shows results for these instrument assumptions under commitment.

For time-consistent policies, optimal use of the balance sheet materially reduces the

[25] A simulation of 510,000 periods is produced and the first 10,000 periods are discarded.

[26] Results for welfare losses are often converted into consumption equivalent units (see, for example, Adam and Billi, 2007; Nakov, 2008). As in these papers, applying this conversion to the present results generates quantitatively small consumption equivalent losses. However, the *relative* sizes of the losses in consumption equivalent units are very similar to those reported in Table 2 since $\sigma = 1$. Moreover, using (12) to compare losses ignores 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.

**Table 2:** Mean outcomes and welfare under alternative policy assumptions

| | Time consistent | | Commitment | |
|---|---|---|---|---|
| | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only |
| Quarterly inflation, % | -0.02 | -0.07 | 0.00 | 0.00 |
| Output gap, % | -0.01 | -0.02 | -0.00 | -0.00 |
| Policy rate, % (annualized) | 3.06 | 2.75 | 3.02 | 3.02 |
| Long-term rate, % (annualized) | 2.82 | 2.75 | 3.01 | 3.01 |
| Balance sheet ($q$) | 0.09 | 0.00 | 0.01 | 0.00 |
| Loss ($\times$100) | 0.60 | 0.82 | 0.43 | 0.44 |
| Relative loss | 1.38 | 1.89 | 1.00 | 1.00 |
| ZLB frequency, % | 38 | 40 | 12 | 12 |

*Notes:* '$\hat{R}$ and $q$' refers to the case in which the short-term policy rate and the balance sheet are both used optimally. '$\hat{R}$ only' refers to the case in which the only instrument is the short-term policy rate. 'Relative loss' reports the ratio of the loss to the loss under optimal commitment.

downward skews in the distributions of the output gap and inflation. Relative to the case in which the balance sheet is not used, welfare losses are reduced by more than a quarter. This improved performance is associated with an average level of the balance sheet of 0.09, a higher level of the policy rate and a lower average level of the long-term rate. The relatively low average level of the balance sheet is consistent with the results in Figure 7 which suggest that encounters with the upper bound on the balance sheet are very rare.

Unsurprisingly, commitment policies deliver smaller welfare losses. The means of inflation, the output gap and interest rates are all very close to their deterministic steady state levels, indicating that the distortions generated by the lower bound on the policy rate are virtually eliminated on average. Accordingly, and in contrast to the results for time consistent policy, the benefits from the use of balance sheet policies are negligible under commitment.[27] The welfare loss when the policy rate is the only instrument ('$\hat{R}$ only') is less than 1% larger than the case in which balance sheet policies are also used ('$\hat{R}$ and $q$'). In equilibrium, policy spends *less* time constrained by the ZLB under commitment policies since they stabilize the output gap and inflation so effectively.[28]

[27] Mau (2022) finds similar results for a variant of the Sims et al. (2021) model, though abstracting from endogenously binding constraints on the short-term interest rate and central bank balance sheet.

[28] Nakov (2008) reports the same result when comparing optimal time consistent and commitment policies for the workhorse New Keynesian model.

### 4.5 Robustness

Appendix A assesses the robustness of the results with respect to alternative assumptions for key parameter values. The focus is on the parameters that are most important for the transmission and efficacy of monetary policy and hence the ability of the policymaker to stabilize macroeconomic fluctuations. This section summarizes the key results of those experiments.

Relaxing the upper bound on the central bank balance sheet to its highest feasible value ($\bar{q} = 1$) improves the policymaker's ability to stabilize the output gap and inflation, particularly when $r^*$ is very low. In such cases, the balance sheet expands by more when $\bar{q} = 1$ as the likelihood of encountering the upper bound on the balance sheet is lower. This reduces the shadow rate by more than the baseline case ($\bar{q} = 0.7$), thereby mitigating the recessionary effects of very low $r^*$ (see Appendix A.1).

Reducing the slope of the Euler equation, $\sigma$, weakens the effect of monetary policy on spending, making it more difficult to stabilize inflation and the output gap (particularly in response to cost-push shocks) and reducing the welfare gains from balance sheet policies. Increasing the slope of the Phillips curve, $\kappa$, has similar effects (see Appendix A.2).

Appendix A.3 explores the implications of alternative values of the portfolio friction parameters. The values of $\nu$ and $\xi$ are chosen to match the Weale and Wieladek (2016) estimates of the effects of a QE shock on long-term bond rates in the United States. That calibration suggests a somewhat weaker effect of balance sheet policies on the shadow rate (lower values of both $\nu$ and $\xi$). As a result, the balance sheet is used more actively when the policy rate hits the zero bound. Nevertheless, this is not sufficient to offset the weaker effects of balance sheet policies so that their overall stabilization benefits are somewhat smaller than in the baseline calibration.

Finally, the importance of the (deterministic) steady-state real interest rate is explored by varying the assumption about the household discount factor, $\beta$. A lower steady-state real interest rate (higher $\beta$) implies that the steady-state policy rate is closer to the zero lower bound. The higher likelihood of encountering the zero lower bound increases the welfare gains from balance sheet policies, which are used more actively on average (see Appendix A.4).

## 4.6 Discussion

The quantitative experiments demonstrate that it is optimal to deploy balance sheet policies when the policy rate hits the zero lower bound and that, in such circumstances, asset purchases can be large and rapid. These features of optimal policy behavior are consistent with the initiation of QE in early 2009 in the United Kingdom and the United States.[29] In both cases, the initial purchases of long-term government debt were sizable and occurred when (or very soon after) the policy rate hit the effective lower bound. In addition, as discussed in Section 3.2.2, optimal policy implies a 'de-coupling' of the optimality conditions determining the balance sheet and the policy rate away from the zero bound. This implication of optimal policy is consistent with recent FOMC and MPC quantitative tightening actions, in which the policy rate is the "primary instrument" for setting the stance of monetary policy (Monetary Policy Committee, 2021).

As discussed in Sections 3.2.2 and 3.2.3, these results are driven in part by the fact that the short-term interest rate and central bank balance sheet are perfect substitutes in terms of their effects on the stance of policy (i.e., the shadow rate). Sims and Wu (2020) present evidence of considerable substitutability between the policy rate and the central bank balance sheet. However, in the Sims et al. (2021) framework that they use, balance sheet policy actions have a direct effect via the Phillips curve so that in general both instruments will be used to stabilize the output gap and inflation, regardless of whether the policy rate is constrained by the zero bound. Bonciani and Oh (2021) study optimal policy in that model, including the implications for the optimal mix of policy instruments.[30]

Importantly, however, one aspect of optimal QT in the present model is less obviously consistent with the exit strategies that have been implemented in recent years. Quantitative tightening in both the United States and United Kingdom did not begin until the short-term policy rate had been increased from the lower bound. In contrast, optimal policy in the model implies that balance sheet reduction should typically begin before

[29] The macroeconomic backdrop to subsequent asset purchase programs in these economies was somewhat different. In particular, the exceptional circumstances of the Covid-19 pandemic prompted broad-based macroeconomic policy responses, of which quantitative easing was an important part. See Bailey et al. (2020) and Benmelech and Tzur-Ilan (2020) for discussions of these responses.

[30] The substitutability of the policy instruments in practice is a topic worthy of additional empirical research. One avenue could consider the empirical relevance of the different effects of long-term and short-term interest rates on firms' marginal costs, as in the Sims et al. (2021) model.

the short-term policy rate lifts off from the lower bound.[31]

What might explain these differences?

In the model, the absolute size of the asset stock held by the central bank is perfectly correlated with the balance sheet instrument, $q$, because the stock of long-term government debt is held fixed. This means that $q$ can only be reduced by allowing bonds to mature or selling them to financial intermediaries. However, in a more general setting, a reduction in $q$ can be achieved with an unchanged central bank balance sheet if the stock of long-term government debt is rising (Greenwood et al., 2016b).[32]

Another factor that has guided QT strategies is the relative uncertainty over the macroeconomic effects of the central bank balance sheet, compared with the effects of the short-term policy rate.[33] This uncertainty may provide an additional motivation to use the short-term policy rate as the "primary instrument" to set the overall stance of monetary policy, a result that arises in the simple model analyzed by Williams (2013). In particular, a specific uncertainty relevant for QT is the possibility that asset sales may have different effects to asset purchases. The present model abstracts from the effects of these types of uncertainty on optimal policy.

Finally, the policymaker in the model minimizes a loss function based on household welfare and therefore accounts for the welfare costs of portfolio frictions. These costs are particularly important in guiding the optimal approach to QT, as described in Section 3.2.2. However, the mandates of most central banks more closely resemble a so-called 'flexible inflation targeting' loss function which includes only the costs of output gap and inflation variability. The implications of pursuing such a mandate is considered are the next section.

[31] Harrison (2012) and Darracq Pariès and Kühl (2017) reach similar conclusions studying optimal policy under commitment (using a perfect foresight solution method), finding that the balance sheet begins to unwind at or before the date of liftoff.

[32] Harrison (2017) computes a crude proxy for $q$ the United Kingdom, showing declines driven by increases in government debt in periods during which the policy rate was held at its lower bound and the central bank balance sheet was held fixed. However, that proxy does not control for the market revaluations of the central bank asset portfolio necessary to compute $q$ and so can be regarded as indicative evidence at best.

[33] One source of uncertainty is that the factors that gave rise to large effects from some asset purchase programs may have been related to the particular state of financial stress during the period in which they were implemented. Bailey et al. (2020) explore the potential implications of such 'state contingency' for optimal balance sheet policy.

# 5 Simple QT strategies and instrument sequencing

The analysis has thus far assumed that monetary policy is set to minimize the loss function derived from household welfare. Since adjustments in the central bank balance sheet affect spending and inflation because of portfolio frictions, the loss function captures the welfare costs caused by those frictions. However, many central bank mandates are specified solely in terms of inflation and the output gap (or a similar consideration for avoiding excessive volatility in real activity and/or employment). This type of mandate is sometimes described as 'flexible inflation targeting' (Svensson, 1999, 2000, 2010).

Motivated by that observation, this section considers the case in which monetary policy is guided by a 'flexible inflation targeting' loss function, defined as:

$$\mathcal{L}_0^{FIT} = \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left( \omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2 \right) \tag{23}$$

where $\omega_x = \Xi$ and $\omega_\pi = \eta/\Gamma$. Compared with the welfare-based loss function, (12), the flexible inflation targeting loss therefore places no weight on the welfare costs from portfolio frictions.

When policy instruments are unconstrained, the optimality condition implied by minimization of (23) delivers a familiar targeting criterion:

$$\omega_x \hat{x}_t + \kappa \omega_\pi \hat{\pi}_t = 0 \tag{24}$$

as shown in Appendix G.

Appendix G shows further that an additional optimality condition, which accounts for the possibility that instrument bounds may bind in future, determines the 'instrument mix' between $\hat{R}$ and $q$ used to implement the optimal level of the shadow rate $\widetilde{R}$ required to implement (24). However, that optimality condition is entirely forward-looking and hence does not prescribe a particular sequencing between unwinding the balance sheet and liftoff from the zero lower bound.

Accordingly, this section considers a case in which the policymaker sets the shadow rate $\widetilde{R}$ to achieve the optimal flexible inflation targeting criterion (24), with the instrument mix determined by simple rules. These rules are specified to mimic the key properties of the QT strategies announced by monetary policymakers in the United Kingdom and United States (as described in Monetary Policy Committee, 2021; Federal Open Market Committee, 2022).

Assuming that it is feasible to implement (24), that condition determines the shadow rate, $\widetilde{R}_t$. The instrument mix is determined by considering the following cases.

First, the policymaker considers reducing the balance sheet according to a simple QT rule

$$q_t = \rho q_{t-1}, \quad 0 \leq \rho < 1 \tag{25}$$

and computes the policy rate $\hat{R}_t$ required to deliver the optimal shadow rate. If the required policy rate is above the ZLB, then this case determines the policy mix. This assumption implies that, away from the zero lower bound, the policy rate $\hat{R}$ is the "active instrument" in the sense that it is adjusted to deliver the shadow rate $\widetilde{R}$ required to satisfy the targeting criterion (24), while the balance sheet is adjusted in a way that is unresponsive to the state of the economy. Moreover, equation (25) implies that, away from the zero bound, the balance sheet is reduced in a predictable and stable manner. Both of these properties are consistent with the quantitative tightening strategies pursued by monetary policymakers in the United Kingdom and United States.[34]

If the first policy mix is not feasible (because it cannot be delivered without $\hat{R}_t$ violating the ZLB), the policymaker computes the policy rate required to deliver the optimal shadow rate when the balance sheet is held fixed, $q_t = q_{t-1}$. If the resulting value of the policy rate is above the ZLB, then this case determines the policy mix. Taken together with the first case, this behavior ensures that liftoff from the zero bound occurs before balance sheet unwind begins, replicating the instrument sequencing chosen by UK and US policymakers.

If neither of the preceding cases can be implemented without violating the ZLB on the policy rate, the policy rate is set equal to the zero bound and the balance sheet ($q$) is adjusted to implement the optimality condition (24). This captures the revealed preference to use *the balance sheet* as the primary instrument when the ZLB binds.[35]

[34]For example, in their July 2023 press conference, the FOMC Chair stated that "[QT and rates] are two independent things. And, you know, really, the active tool of monetary policy is rates. (Powell, 2023).

[35]Finally, if the policy rate is constrained by the zero lower bound and it is not possible to implement (24) without violating the upper bound on the balance sheet ($q_t \leq \bar{q}$), then the policy rate is set at

This specification of policy behavior gives rise to a set of alternative QT rules, each indexed by a value of $\rho \in [0, 1)$. However, two values for $\rho$ are of particular interest. The first value is $\rho = \chi$, which corresponds to the case in which the central bank does not re-invest the proceeds from maturing bonds (sometimes called 'passive unwind').[36] The second value is $\rho = \zeta$, defined in (22), which is labeled 'neutral unwind' as it corresponds to the pace of QT for which balance sheet policy has no effect on the shadow rate. As shown in Section 3.2.2, this is also (approximately) the optimal unwind pace when the economy is sufficiently far from the zero bound.

**Figure 8:** Mild recessionary scenario: alternative QT strategies and optimal balance sheet policy

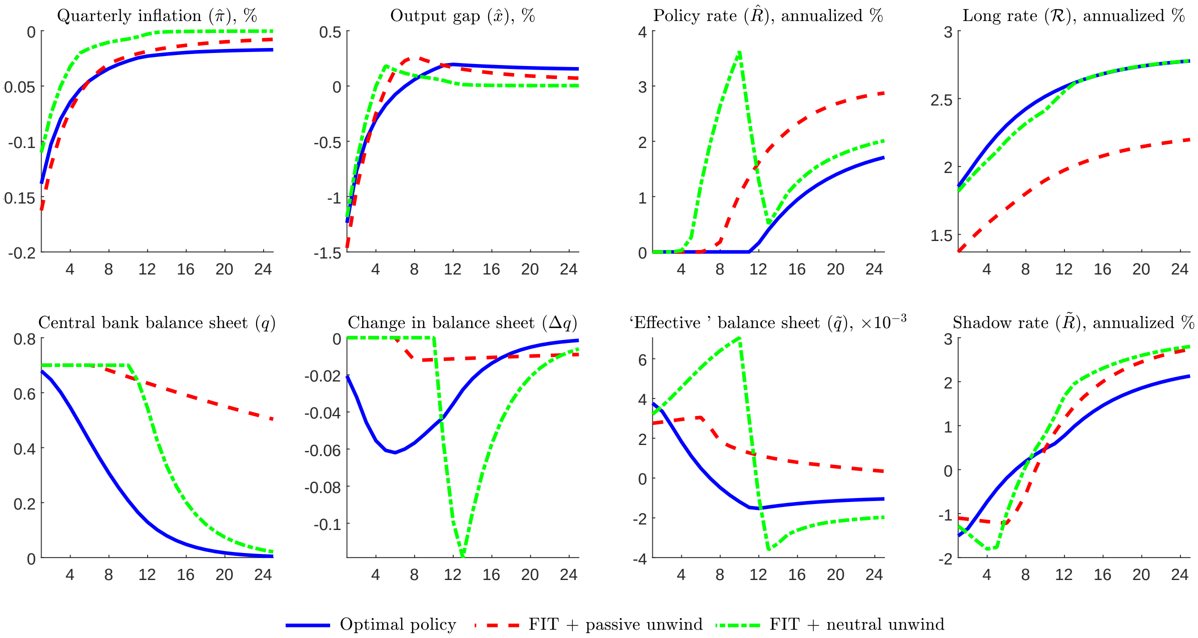


*Notes:* Each simulated path is computed using linear interpolation of the policy functions for $u_t = 0$ and $r_t^* = (r_1^*)^{\varrho_r(t-1)}$ $(t = 1, \ldots, T)$. The value of $r_1^*$ is such that the level of the equilibrium rate (in annualized units) is $-2.2\%$, much lower than the deterministic steady state value of 3%. The central bank is assumed to hold a maximal stock of assets in period 0 ($q_0 = \bar{q} = 0.7$). The solid blue lines show optimal balance sheet policy. The green dash-dotted lines correspond to a flexible inflation targeting central bank with a 'neutral unwind' QT rule ($\rho = \zeta$). The red dashed lines correspond to a flexible inflation targeting central bank with a 'passive unwind' QT rule ($\rho = \chi$).

Figure 8 shows a milder variant of the recessionary scenario considered in Section 4.3, assuming $r_1^* = -2.2\%$ (in annualized units). A milder recessionary scenario is useful in comparing differences in QT behavior, since policy will move more quickly to a tighter stance. Similarly, since QT strategies differ most when the balance sheet is large, it is

the lower bound and the balance sheet is set to its maximal level, $\bar{q}$. In this case, the output gap and inflation are then determined by the Euler equation and the Phillips curve and the optimal unconstrained trade-off (24) cannot be achieved.

[36]Recall that the long-term bonds are zero coupon bonds with stochastic redemption probability $\chi$.

assumed that the central bank initially holds the maximal stock of assets on its balance sheet ($q_0 = \bar{q} = 0.7$). Results are plotted for the 'passive' and 'neutral' QT unwind strategies, alongside the optimal policy response (solid blue line).

The bottom left panel of Figure 8 shows that, unsurprisingly, balance sheet reduction is much slower under a passive unwind strategy (red dashed line) than a neutral strategy (green dash-dotted line). Under a passive unwind strategy the shadow rate is roughly constant for around seven quarters (bottom right panel) before rising. The increase in the shadow rate is driven by a path for the policy rate that is steeper than observed under optimal policy, reflecting the actions required to maintain the desired policy stance given the very slow reduction in the balance sheet. The very persistent path for the balance sheet under a passive unwind strategy implies that the long-term interest rate remains very low for a prolonged period (top right panel).

Compared with optimal policy, the passive unwind strategy generates worse outcomes for the output gap and inflation in the near term. Agents recognize that the subsequent slow unwinding of the balance sheet will be associated with a relatively high probability that policy will be constrained if deflationary shocks arrive. In particular, there is an asymmetry between the value of building 'headroom' for future balance sheet expansions relative to headroom for future policy rate reductions because the shadow rate depends on the *change* in the central bank balance sheet. The slow pace of balance sheet reduction implied by the passive unwind strategy increases the probability that a rapid asset purchase program will be constrained by the upper bound on the balance sheet if deflationary shocks arrive. This gives rise to a deflationary bias of the form discussed in Section 4.2 whereby weaker inflation expectations raise the real interest rate, depressing activity and inflation in the near term.

In contrast, a neutral unwind strategy (green dash-dotted lines) leads to better stabilization of inflation (and similar outcomes for the output gap) compared with optimal policy (solid blue lines). The neutral unwind strategy is associated with a more variable path for the shadow rate (bottom right panel) than both optimal policy and passive unwind. The pace of balance sheet reduction is sufficient to rapidly reduce the chances of encountering future policy constraints. In the longer-term, the risk is sufficiently small that it is possible to approximate the 'divine coincidence' in which both the output gap and inflation are stabilized at zero.

The instrument mix that delivers this policy stance requires sharp movements in the policy rate. Balance sheet unwind at a relatively rapid pace ($\rho = \zeta \approx 0.8$) is only feasible several quarters after liftoff. So during the first year after liftoff, policy tightening is achieved by a substantial increase in the policy rate. However, once QT begins, the rapid pace of unwind requires a large initial change in the effective balance sheet: $\widetilde{q}$ moves rapidly from positive to negative territory. Other things equal, the initiation of the neutral QT strategy therefore has a substantial tightening effect on the shadow rate, which must be offset by a large reduction in the policy rate.[37]

The recessionary scenario therefore suggests that alternative QT strategies have different implications for the ability to stabilize inflation and the output gap in the vicinity of the lower bound and the mix of policy instruments required to implement the strategy. Figure 9(a) explores the overall implications of alternative QT strategies by plotting average welfare losses (from (12)) using stochastic simulations for alternative values of $\rho$ (red dashed line). Average losses are normalized by the loss achieved under optimal commitment. The particular cases of interest – passive unwind ($\rho = \chi$) and neutral unwind ($\rho = \zeta$) – are marked with the circle and diamond respectively.

Figure 9(a) shows that all values of $\rho$ generate smaller welfare losses than the case in which the balance sheet is not used as a policy instrument (black dot-dashed line). However, welfare losses are U-shaped, so that very fast and very slow strategies (including 'passive unwind' ($\rho = \chi$) shown by the red circle) generate relatively high losses. In contrast, the 'neutral unwind' strategy ($\rho = \zeta$, red diamond) generates welfare losses that are close to those obtained when the balance sheet is optimally deployed.

Figure 9(b) plots the frequency distributions of the balance sheet for the passive and neutral QT strategies. The shape of the frequency distribution for the neutral QT rule resembles the distribution observed under optimal policy (Figure 7(a)). The distribution of $q$ under the passive QT rule looks very different, with the mode of the distribution at the upper bound $\bar{q} = 0.7$. So one reason why the passive QT strategy delivers higher welfare losses is that it leads to a ratchet effect on the balance sheet. In other words, after

[37]The large policy rate reduction when QT begins is in part driven by the size of the initial change in $\widetilde{q}$. That, in turn, depends on the initial size of the central bank balance sheet, $q_0 = \bar{q}$. As shown in Figure 9(b), the distribution of the balance sheet under a neutral QT strategy is such that $q$ is rarely close to the upper bound, $\bar{q}$. So 'typical' policy rate adjustments at the start of QT will be much smaller than shown in Figure 8.

**Figure 9:** Normalized welfare losses and balance sheet distributions under flexible inflation targeting

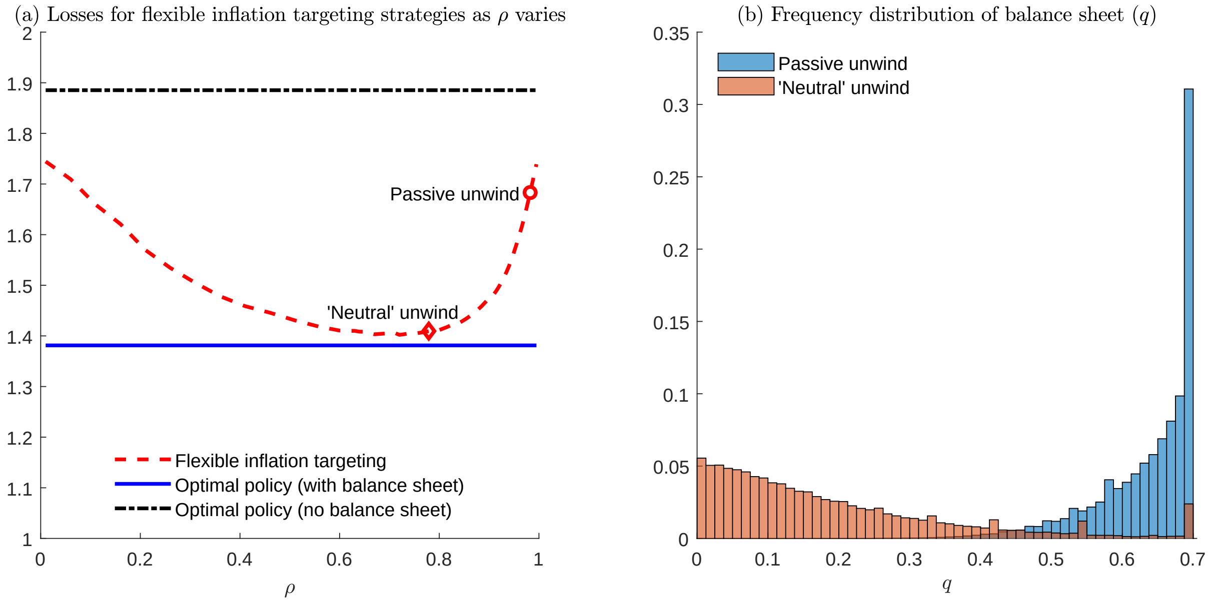


*Notes:* Welfare losses when the central bank follows a 'flexible inflation targeting' approach, minimizing $\mathcal{L}^{FIT}$ defined by (23) and following a quantitative tightening strategy defined by (25) for alternative values of $\rho \in (0, 1)$. Losses are computed as the average level of $\omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2 + \omega_q q_t^2 + \omega_{\Delta q} (q_t - q_{t-1})^2$ normalized by the average loss achieved under optimal commitment. Losses are computed using a simulation of 500,000 periods.

liftoff from the lower bound the balance sheet is not materially smaller by the time the policymaker next encounters the lower bound and needs to re-start asset purchases. Over time, this effect leads the the average size of the balance sheet to 'ratchet' upwards.[38]

# 6 Conclusion

This paper studies the optimal use of quantitative easing (QE) and quantitative tightening (QT) alongside the short-term policy rate using a workhorse New Keynesian model extended to include portfolio frictions. These frictions imply both that QE and QT can influence long-term rates (via a portfolio balance channel) and that these policies have welfare costs.

The inclusion of portfolio frictions creates an additional trade-off criterion for optimal

[38]The extent of the ratchet effect, and the welfare implications of alternative QT rules more broadly, depends on the value of $\chi$, which governs the rate at which previously issued government debt matures and hence the pace of balance sheet reduction under passive unwind. The baseline calibration uses UK government debt, which has a particularly long maturity (Harrison, 2021). For economies in which the maturity structure of government debt is substantially shorter (such as the United States) the difference between a 'passive' and 'neutral' unwind strategy may therefore be somewhat smaller than implied by Figure 9(a).

policy. The marginal cost of portfolio distortions generated by balance sheet policy should be equated to the marginal benefit of better stabilization of the output gap and inflation. Simulations of the model reveal that this criterion gives rise to optimal QE and QT behavior that is, in many respects, similar to observed policy actions in the United Kingdom and United States: asset purchases are often rapid and triggered when the policy rate hits the zero bound, while balance sheet unwind (QT) is more gradual with the short-term policy rate being adjusted in response to shocks to determine the overall monetary stance.

However, optimal policy implies that balance sheet unwind will often begin *before* the policy rate has been raised from its lower bound, while UK and US policymakers have taken the opposite approach. Simulations of a 'flexible inflation targeting' policymaker that starts QT only after liftoff from the lower bound demonstrate that it is possible to achieve similar welfare to optimal policy if the pace of QT is appropriately calibrated.

A broader lesson is that the frictions that give balance sheet policies traction should also guide their optimal deployment. The portfolio balance channel studied in the present paper gives rise to a particular optimality condition that determines balance sheet policy. However, despite the more widespread and routine use of balance sheet policies, there remains considerable uncertainty around their transmission (Bailey et al., 2020). Studying the normative implications of other frictions that may provide a role for central bank balance sheet policies is therefore a promising avenue for future research.

# References


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

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

Airaudo, F. S. (2022). Exit strategies from quantitative easing: The role of the fiscal-monetary policy mix. *mimeo*.

Alla, Z., R. A. Espinoza, and A. R. Ghosh (2016, March). Unconventional Policy Instruments in the New Keynesian Model. *IMF Working Paper No. 16/58* (16/58).

Andrés, J., D. López-Salido, and E. Nelson (2004). Tobin's imperfect asset substitution in optimizing general equilibrium. *Journal of Money, Credit and Banking 36*(4), 665–91.

Bailey, A. J., J. Bridges, R. Harrison, J. Jones, and A. Mankodi (2020). The central bank balance sheet as a policy tool: past, present and future. *Bank of England Working Paper No. 899*.

Benigno, G. and P. Benigno (2022). Managing monetary policy normalization. *FRB of New York Staff Report* (1015).

Benigno, P. and S. Nisticò (2020). Non-neutrality of open-market operations. *American Economic Journal: Macroeconomics 12*(3), 175–226.

Benmelech, E. and N. Tzur-Ilan (2020). The determinants of fiscal and monetary policies during the COVID-19 crisis. *NBER Working Paper*.

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

Bhattarai, S., G. B. Eggertsson, and B. Gafarov (2022). Time consistency and duration of government debt: A model of quantitative easing. *The Review of Economic Studies*.

Bhattarai, S. and C. J. Neely (2022). An analysis of the literature on international unconventional monetary policy. *Journal of Economic Literature 60*(2), 527–97.

Bodenstein, M., J. Hebden, and R. Nunes (2012). Imperfect credibility and the zero lower bound. *Journal of Monetary Economics 59*(2), 135–149.

Bonciani, D. and J. Oh (2021). Optimal monetary policy mix at the zero lower bound. *Bank of England Staff Working Paper No. 945*.

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

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

Cantore, C. and P. Meichtry (2023). Unwinding quantitative easing: state dependency and household heterogeneity. *Bank of England Working Paper No. 1030*.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian (2017). Targeting long rates in a model with segmented markets. *American Economic Journal: Macroeconomics 9*(1), 205–42.

Cesa-Bianchi, A., R. Harrison, and R. Sajedi (2023). Global $R^*$. *CEPR Discussion Paper DP18518*.

Chen, H., V. Cúrdia, and A. Ferrero (2012, November). The macroeconomic effects of large-scale asset purchase programmes. *Economic Journal 122*(564), F289–F315.

Cui, W. and V. Sterk (2021). Quantitative easing with heterogeneous agents. *Journal of Monetary Economics 123*, 68–90.

Cúrdia, V. and M. Woodford (2016). Credit frictions and optimal monetary policy. *Journal of Monetary Economics 84*, 30–65.

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*(2), 425–448.

Darracq Pariès, M. and M. Kühl (2017). The optimal conduct of central bank asset purchases. *Deutsche Bundesbank Working Paper No. 22/2017*.

De Graeve, F. and K. Theodoridis (2016). Forward guidance, quantitative easing, or both? *National Bank of Belgium Working Paper No. 305*.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2019). Global trends in interest rates. *Journal of International Economics 118*, 248–262.

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.

Federal Open Market Committee (2022). Principles for Reducing the Size of the Federal Reserve's Balance Sheet. *Statement released on January 26 2022, Federal Reserve Board*.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle (second edition)*. Princeton University Press.

Gertler, M. and P. Karadi (2011). A model of unconventional monetary policy. *Journal of Monetary Economics 58*(1), 17–34.

Greenwood, R., S. Hanson, J. Rudolph, and L. Summers (2016a). The optimal maturity of government debt. In D. Wessel (Ed.), *The $13 Trillion Question*, Chapter 1, pp. 1–27. Brookings Institution Press.

Greenwood, R., S. G. Hanson, J. S. Rudolph, and L. Summers (2016b). Debt Management Conflicts between the U.S. Treasury and the Federal Reserve. In D. Wessel (Ed.), *The $13 Trillion Question: How America Manages Its Debt*, Chapter 2, pp. 43–89. Brookings Institution Press.

Greenwood, R., S. G. Hanson, and J. C. Stein (2016). The Federal Reserve's Balance Sheet as a Financial-Stability Tool. *Federal Reserve Bank of Kansas Jackson Hole Symposium*.

Grimaldi, M., A. Crosta, and D. Zhang (2021). The liquidity of the government bond market-What impact does quantitative easing have? Evidence from Sweden. *Sveriges Riksbank Working Paper No. 402*.

Harrison, R. (2011). Asset purchase policies and portfolio balance effects: a DSGE analysis. In J. Chadha and S. Holly (Eds.), *Interest rates, prices and liquidity*, Chapter 5. Cambridge University Press.

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

Harrison, R. (2017). Optimal quantitative easing. *Bank of England Staff Working Paper No. 678*.

Harrison, R. (2021). Flexible inflation targeting with active fiscal policy. *Bank of England Staff Working Paper No. 928*.

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

Hohberger, S., R. Priftis, and L. Vogel (2019). The macroeconomic effects of quantitative easing in the euro area: Evidence from an estimated DSGE model. *Journal of Economic Dynamics and Control 108*.

Joyce, M. A., N. McLaren, and C. Young (2012). Quantitative easing in the United Kingdom: evidence from financial markets on QE1 and QE2. *Oxford Review of Economic Policy 28*(4), 671–701.

Karadi, P. and A. Nakov (2021). Effectiveness and addictiveness of quantitative easing. *Journal of Monetary Economics 117*, 1096–1117.

Kopecky, K. A. and R. M. Suen (2010). Finite state Markov-chain approximations to highly persistent processes. *Review of Economic Dynamics 13*(3), 701–714.

Leeper, E. M. (1991). Equilibria under 'active' and 'passive' monetary and fiscal policies. *Journal of Monetary Economics 27*(1), 129–147.

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*(1), 143–89.

Logan, L. and U. Blindseil (2019). Large central bank balance sheets and market functioning. *BIS Markets Committee Report*.

Mau, R. (2022). What is in a name? Purchases and Sales of Financial Assets as a Monetary Policy Instrument. *Journal of Money, Credit and Banking*.

Miller, M. J. (2010). Remarks before the Future Industry Association Treasury and Rates Conference New York. *October 5*.

Monetary Policy Committee (2021). Monetary Policy Report, August. *Bank of England*.

Nakata, T. (2015). Credibility of optimal forward guidance at the interest rate lower bound. *FEDS notes 08/27*.

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

Plosser, C. I. (2012). Perspectives on monetary policy. *Official Monetary and Financial Institutions Forum (OMFIF) Golden Series Lecture London, England*.

Powell, J. (2023). Transcript of Chair Powell's Press Conference, July 26, 2023. *Federal Reserve Board of Governors*.

Priftis, R. and L. Vogel (2016). The Portfolio Balance Mechanism and QE in the Euro Area. *The Manchester School 84*(S1), 84–105.

Quint, D. and P. Rabanal (2017). Should unconventional monetary policies become conventional? *IMF Working Paper No. 17/85*.

Rouwenhorst, K. (1995). Asset pricing implications of equilibrium business cylce models. In T. Cooley (Ed.), *Frontiers in business cycle research*, pp. 294–330. Princeton University Press.

Sims, E. and J. C. Wu (2020). Are QE and Conventional Monetary Policy Substitutable? *International Journal of Central Banking 16*.

Sims, E., J. C. Wu, and J. Zhang (2021). The four equation New Keynesian model. *The Review of Economics and Statistics*, 1–45.

Sims, E. R., J. C. Wu, and J. Zhang (2022). Unconventional monetary policy according to HANK. *NBER Working Paper 30329*.

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*(2), 161–183.

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

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

Svensson, L. E. (2010). Inflation targeting. In *Handbook of monetary economics*, Volume 3, pp. 1237–1302. Elsevier.

Thomas, R., S. Hills, and N. Dimsdale (2010). The UK recession in contextwhat do three centuries of data tell us? *Bank of England Quarterly Bulletin*, Q4.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *Review of Economic Studies 25*(2), 65–86.

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

Tobin, J. and W. C. Brainard (1963). Financial intermediaries and the effectiveness of monetary controls. *American Economic Review 53*(2), 383–400.

Vayanos, D. and J.-L. Vila (2009). A preferred-habitat model of the term structure of interest rates. *NBER Working Paper No. 15487* (15487).

Vayanos, D. and J.-L. Vila (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica 89*(1), 77–112.

Vissing-Jorgensen, A. (2023). Balance sheet policy above the elb. *mimeo*.

Wallace, N. (1981). A Modigliani-Miller theorem for open-market operations. *The American Economic Review 71*(3), 267–274.

Weale, M. and T. Wieladek (2016). What are the macroeconomic effects of asset purchases? *Journal of Monetary Economics 79*, 81 – 93.

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

Woodford, M. (2001). Fiscal requirements for price stability. *Journal of Money, Credit and Banking 33*(3), 669–728.

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

Wu, J. C. and J. Zhang (2019). A shadow rate New Keynesian model. *Journal of Economic Dynamics and Control 107*, 103–128.

# A Robustness

This Appendix assesses the robustness of the results with respect to alternative assumptions for key parameter values. The focus is on the parameters that are most important for the transmission and efficacy of monetary policy and hence the ability of the policymaker to stabilize macroeconomic fluctuations.

Section A.1 considers the upper bound on the central bank balance sheet and Section A.2 investigates the slopes of the Euler equation and Phillips curve. Section A.3 explores the implications of alternative values of the portfolio friction parameters and Section A.4 examines the effects of alternative assumptions about the steady-state real interest rate.

## A.1 Upper bound on the central bank balance sheet

As noted in Section 4.1 an upper bound on the balance sheet ($\bar{q}$) can be motivated by a range of factors, including the effects of a reduced availability of government debt on liquidity and market functioning. However, since these factors are not explicitly modeled, this section explores the implications of relaxing the upper bound to its highest feasible value, $\bar{q} = 1$.

**Figure A.1:** Policy function comparison: effects of the upper bound on the central bank balance sheet

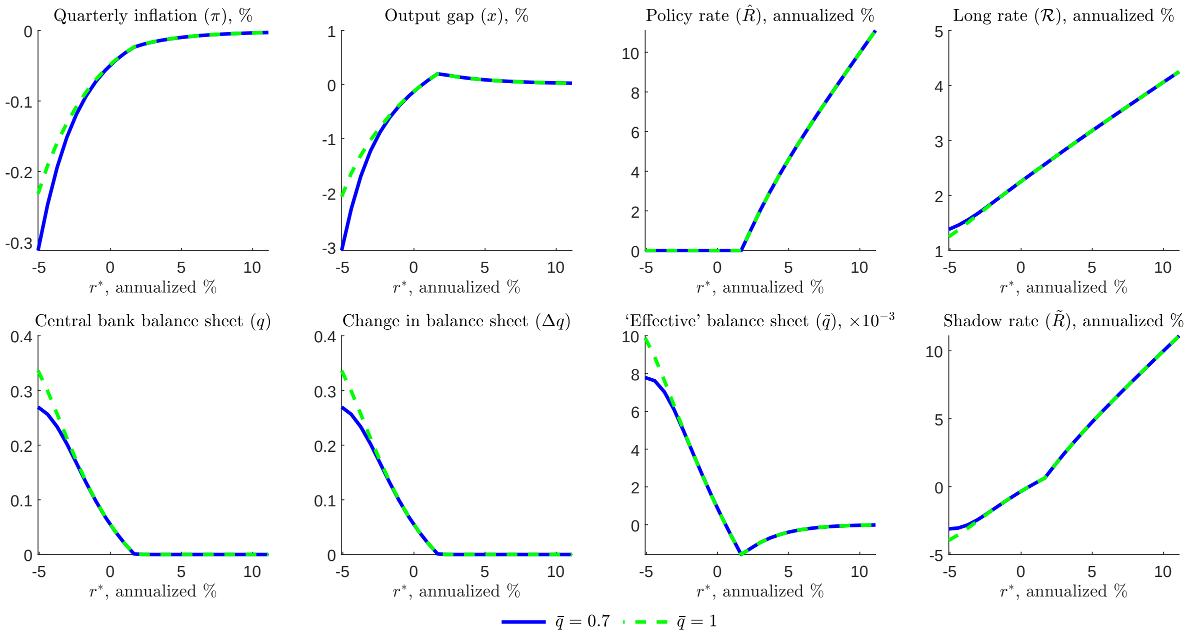


*Notes:* 'Slices' of policy functions for alternative model variants. The solid blue lines are slices of the baseline ($\bar{q} = 0.7$) policy functions conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The green dashed lines are policy functions conditional on $u_t = 0$ for a version of the model in which the upper bound on the central bank balance sheet is $\bar{q} = 1$.

Figure A.1 shows representations of the policy functions for key variables, comparing the baseline assumption of $\bar{q} = 0.7$ with the case in which the central bank is permitted to hold the entire stock of long-term government debt ($\bar{q} = 1$). Each policy function is plotted holding both the cost push shock and the existing balance sheet equal to zero ($u_t = q_{t-1} = 0$). The results show that, unsurprisingly, increasing the maximum scale of asset purchases improves the policymaker's ability to stabilize output and inflation for low realizations of $r^*$. A larger $\bar{q}$ leads to a larger expansion of the balance sheet when $r^*$ is very low. This is because agents recognize that, when $\bar{q}$ is larger, the policymaker has more 'policy space' to further expand the balance sheet in the event of future recessionary shocks. Agents therefore expect that outcomes will be better stabilized in those states. This 'stabilization effect' 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 the central bank balance sheet.

## A.2 Euler equation and Phillips curve slopes

The slope of the Euler equation, $\sigma$, 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. This case is explored by setting $\sigma = 0.5$ following Eggertsson and Woodford (2003). Conversely, a higher elasticity increases the power of monetary policy but exacerbates the effects of being constrained at the zero bound. This case is investigated by setting $\sigma = 1.5$.

**Figure A.2:** Policy function comparison: alternative assumptions for $\sigma$

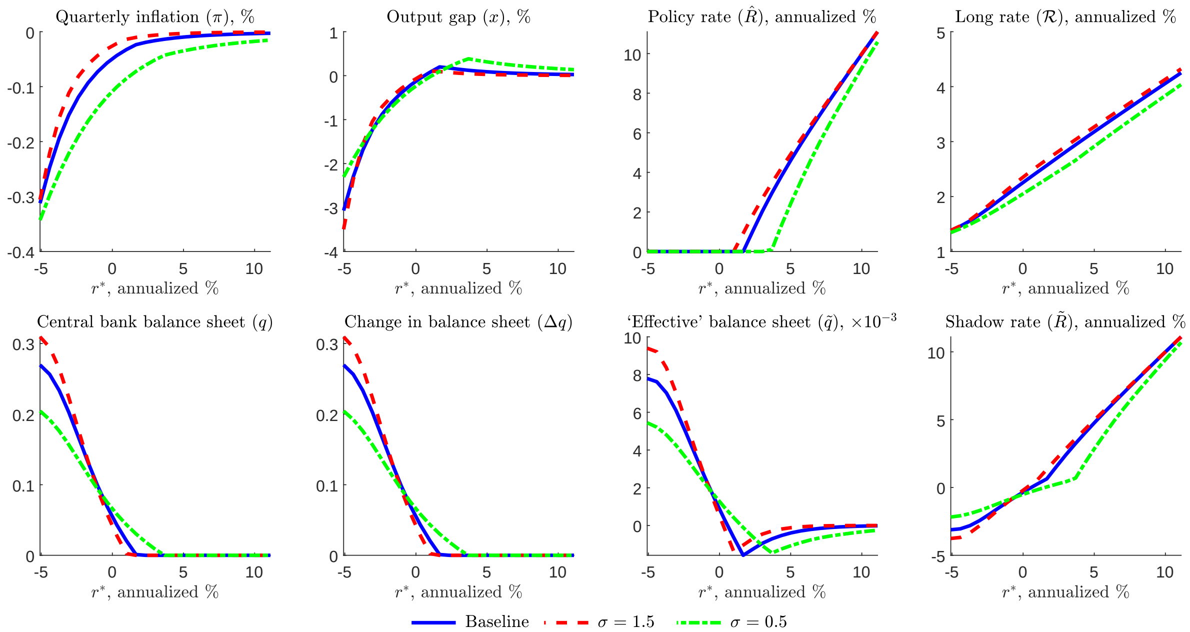


*Notes:* 'Slices' of policy functions for alternative model variants conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The solid blue lines are for the baseline model calibration, the dashed red lines show the case in which $\sigma = 1.5$ and the dash-dotted green lines show the case in which $\sigma = 0.5$.

Figure A.2 compares representations of the policy functions for alternative values of $\sigma$. When $\sigma = 0.5$, the slope of the Euler equation is smaller. This means that very low realizations of $r^*$ have less negative effects on the output gap even when the balance sheet expands by less than the baseline case. However, the policy rate becomes constrained at the lower bound at a higher level of $r^*$, increasing the (unconditional) likelihood of policy being at least partially constrained. Because of the additional deflation bias generated by this effect, the policy function for inflation lies below the corresponding policy function for the baseline case. These effects operate in reverse when $\sigma = 1.5$, though the quantitative implications for the policy functions are somewhat smaller.

These results are consistent with Table A.1, which compares distributions of key variables under alternative policy assumptions for each assumption about $\sigma$. A lower value for $\sigma$ makes

**Table A.1:** Mean outcomes and welfare for alternative values of $\sigma$

| | Baseline | | High $\sigma$ | | Low $\sigma$ | |
|---|---|---|---|---|---|---|
| | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only |
| Quarterly inflation, % | -0.02 | -0.07 | -0.01 | -0.06 | -0.06 | -0.10 |
| Output gap, % | -0.01 | -0.02 | -0.00 | -0.02 | -0.02 | -0.03 |
| Policy rate, % (annualized) | 3.06 | 2.75 | 3.06 | 2.76 | 2.98 | 2.62 |
| Long-term rate, % (annualized) | 2.82 | 2.75 | 2.91 | 2.76 | 2.59 | 2.62 |
| Balance sheet ($q$) | 0.09 | 0.00 | 0.05 | 0.00 | 0.14 | 0.00 |
| Optimal BSP gain, % | 27 | – | 36 | – | 23 | – |
| ZLB frequency, % | 38 | 40 | 24 | 33 | 55 | 62 |

*Notes:* 'Optimal BSP gain' reports the ratio of the loss with optimal balance sheet policy to the loss when the balance sheet is held fixed.

it harder to stabilize macroeconomic fluctuations using monetary policy and also reduces the gains from using the central bank balance sheet as a policy instrument.

**Figure A.3:** Policy function comparison: alternative assumptions for $\kappa$

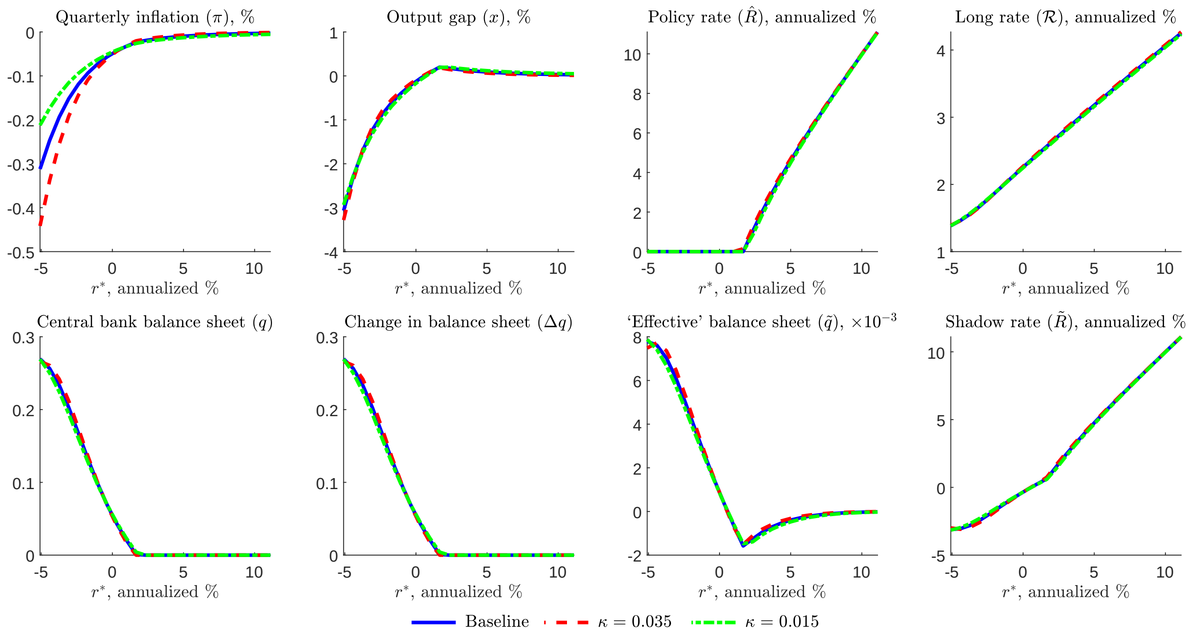


*Notes:* 'Slices' of policy functions for alternative model variants conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The solid blue lines are for the baseline model calibration, the dashed red lines show the case in which $\kappa = 0.035$ and the dash-dotted green lines show the case in which $\kappa = 0.015$.

Some of these effects also apply when the slope of the Phillips curve ($\kappa$) is varied. Other things equal, a flatter Phillips curve (lower $\kappa$) mitigates the downward drag on inflation expectations near the zero bound. This leads to improved inflation stabilization. The converse applies for a steeper Phillips curve. Figure A.3 verifies this by plotting representations of the policy functions computed using alternative values of $\kappa$. Though the policy functions for inflation show this effect clearly, the remaining policy functions are little changed. While $\kappa$ influences the effects of the output gap on inflation it does not affect the required policy actions required to achieve a particular output gap, unlike the case considered in Figure A.2.

Again, the insights from the policy functions are consistent with Table A.2, which compares distributions of key variables under alternative policy assumptions for each assumption about $\kappa$. Macroeconomic performance when balance sheet policies are used is relatively invariant to

**Table A.2:** Mean outcomes and welfare for alternative values of $\kappa$

| | Baseline | | High $\kappa$ | | Low $\kappa$ | |
|---|---|---|---|---|---|---|
| | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only |
| Quarterly inflation, % | -0.02 | -0.07 | -0.02 | -0.08 | -0.02 | -0.06 |
| Output gap, % | -0.01 | -0.02 | -0.00 | -0.01 | -0.01 | -0.03 |
| Policy rate, % (annualized) | 3.06 | 2.75 | 3.06 | 2.68 | 3.07 | 2.77 |
| Long-term rate, % (annualized) | 2.82 | 2.75 | 2.84 | 2.68 | 2.80 | 2.77 |
| Balance sheet ($q$) | 0.09 | 0.00 | 0.08 | 0.00 | 0.10 | 0.00 |
| Optimal BSP gain, % | 27 | – | 47 | – | 18 | – |
| ZLB frequency, % | 38 | 40 | 32 | 41 | 38 | 40 |

*Notes:* 'Optimal BSP gain' reports the ratio of the loss with optimal balance sheet policy to the loss when the balance sheet is held fixed.

the value of $\kappa$, whereas a higher $\kappa$ leads to worse outcomes when only the short-term policy rate is available as an instrument. As a result, the gains from balance sheet policies are larger in the case when $\kappa$ is higher.

## A.3 Portfolio friction parameters

As discussed in Section 4.1, the baseline values for the portfolio friction parameters $\nu$ and $\xi$ are chosen to match empirical estimates of the effects of asset purchases on UK long-term interest rates. The Weale and Wieladek (2016) study used for that exercise also contains estimates of the effects of asset purchases on US long-term interest rates. This section studies a version of the model in which the parameters relating to government debt and portfolio frictions $(\delta, \Theta, \nu, \xi)$ are chosen to match US data, shown in Table A.3. Appendix B.2 provides details of the data and parameterization approach.

**Table A.3:** UK and US calibrations for debt and portfolio frictions

| | Baseline (UK) | US calibration |
|---|---|---|
| $\delta$ | 1.34 | 0.20 |
| $\Theta$ | 0.81 | 1.44 |
| $\nu \times 100$ | 0.38 | 0.07 |
| $\xi \times 100$ | 5.97 | 0.75 |

The parameter values from the US calibration are somewhat different to the baseline values, in part reflecting the different properties of government debt markets. The smaller value of $\delta$ reflects the shorter maturity structure of US government debt. The value of $\Theta$ is larger, reflecting the smaller fraction of US index-linked debt issuance.[39] The US calibration also delivers much smaller values for both $\xi$ and $\nu$. This reflects the more muted response of long rates to asset purchases found by Weale and Wieladek (2016). Nevertheless, the value of $\xi$ is somewhat larger than $\nu$, as in the baseline UK calibration, reflecting the short-term 'kink' in the estimated response of long-term interest rates to an asset purchase shock in both economies.

Figure A.4 plots representations of the policy functions for the baseline model (solid blue lines) and the version calibrated for US data (red dashed lines). Since balance sheet policies are weaker in the alternative (US) calibration, the output gap and inflation are less well stabilized for very low values of $r^*$. This is the case even though the balance sheet is adjusted more forcefully in those states of the world (bottom left panels). However, these actions are not

[39]Index-linked debt is neither incorporated in the model nor included in central bank QE programs. It is therefore excluded from the debt stock data used to calibrate the model. See Appendix B.

**Figure A.4:** Policy function comparison: alternative assumptions for $\nu$ and $\xi$

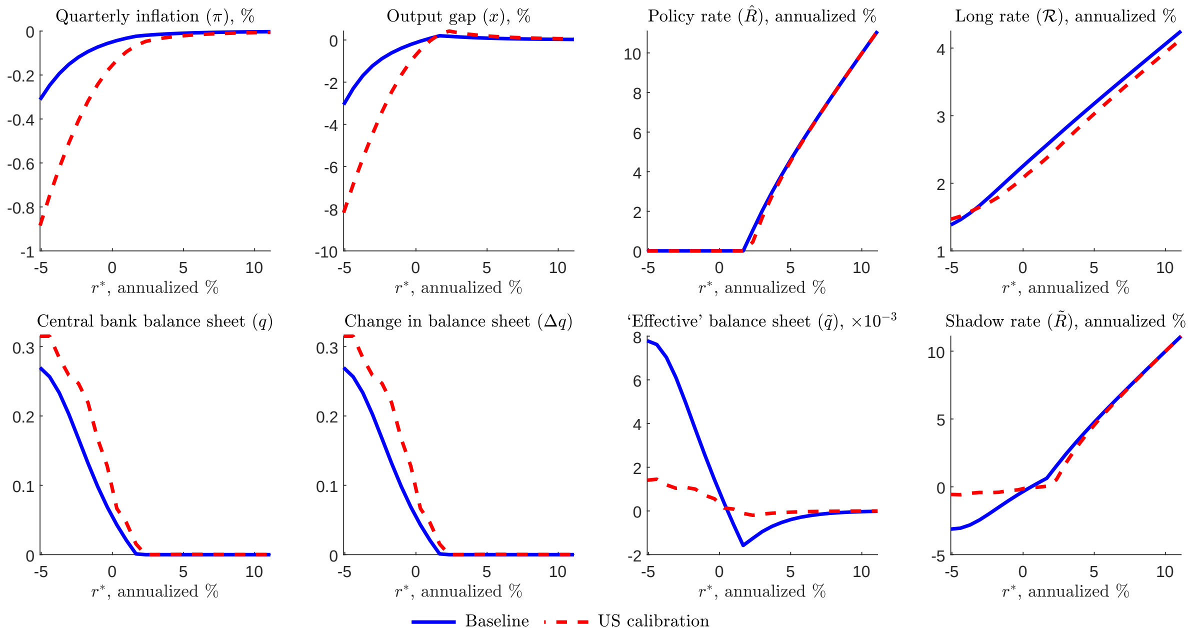


*Notes:* 'Slices' of policy functions for alternative model variants conditional on $\{u_t, q_{t-1}\} = \{0, 0\}$. The solid blue lines are for the baseline model calibration, the dashed red lines show a calibration based on US data, shown in Table A.3.

sufficient to expand the effective balance sheet ($\widetilde{q}$) by enough to depress the shadow rate to levels that provide significant support to spending (bottom right panels).

**Table A.4:** Mean outcomes and welfare for alternative values of $\nu$ and $\xi$

| | Baseline | | US calibration | |
|---|---|---|---|---|
| | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only |
| Quarterly inflation, % | -0.02 | -0.07 | -0.06 | -0.07 |
| Output gap, % | -0.01 | -0.02 | -0.02 | -0.02 |
| Policy rate, % (annualized) | 3.06 | 2.75 | 2.81 | 2.75 |
| Long-term rate, % (annualized) | 2.82 | 2.75 | 2.66 | 2.75 |
| Balance sheet ($q$) | 0.09 | 0.00 | 0.10 | 0.00 |
| Optimal BSP gain, % | 27 | – | 9 | – |
| ZLB frequency, % | 38 | 40 | 40 | 40 |

*Notes:* 'Optimal BSP gain' reports the ratio of the loss with optimal balance sheet policy to the loss when the balance sheet is held fixed.

Table A.4 compares macroeconomic outcomes and welfare under the baseline and alternative US calibration. The weaker efficacy of QE in the US calibration results in a less marked improvement in macroeconomic performance when balance sheet policies are used. As suggested by the policy function comparison above, this is the case even though the balance sheet is used more extensively on average for the US calibration.

## A.4 Steady-state real interest rate

The value of the discount factor, $\beta$, determines the level of the real interest in the non-stochastic steady state. It therefore has potentially important implications for the average amount of (conventional) monetary policy space and hence on the policymaker's ability to stabilize output and inflation in the presence of a lower bound on the policy rate. As discussed in Section 4.1, the baseline assumption for the value of $\beta$ is consistent with a nominal interest rate of around 3% in the non-stochastic steady state.

Table A.5 examines average outcomes for key variables and other population statistics for alternative values of $\beta$. The 'high $\beta$' variant implies a *lower* real interest rate (2.75%) in the deterministic steady state. The 'low $\beta$' variant is chosen so that the real interest rate is 3.25% in the deterministic steady state.

**Table A.5:** Mean outcomes and welfare for alternative values of $\beta$

| | Baseline | | High $\beta$ | | Low $\beta$ | |
|---|---|---|---|---|---|---|
| | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only | $\hat{R}$ and $q$ | $\hat{R}$ only |
| Quarterly inflation, % | -0.02 | -0.07 | -0.03 | -0.11 | -0.02 | -0.05 |
| Output gap, % | -0.01 | -0.02 | -0.01 | -0.03 | -0.01 | -0.02 |
| Policy rate, % (annualized) | 3.06 | 2.75 | 2.77 | 2.26 | 3.25 | 2.98 |
| Long-term rate, % (annualized) | 2.82 | 2.75 | 2.48 | 2.27 | 3.03 | 2.98 |
| Balance sheet ($q$) | 0.09 | 0.00 | 0.11 | 0.00 | 0.08 | 0.00 |
| Optimal BSP gain, % | 27 | – | 52 | – | 19 | – |
| ZLB frequency, % | 38 | 40 | 38 | 50 | 34 | 39 |

*Notes:* 'Optimal BSP gain' reports the ratio of the loss with optimal balance sheet policy to the loss when the balance sheet is held fixed.

The results in Table A.5 show that when balance sheet policies are not used ('$\hat{R}$ only' columns) the lower average real interest rate associated with the high $\beta$ variant leads to a substantial increase in the frequency with which the zero lower bound is binding. This is associated with larger average shortfalls of output and inflation. The converse is true for the low $\beta$ case.

However, when the balance sheet is optimally deployed ('$\hat{R}$ and $q$' columns) average inflation and output gap shortfalls are relatively insensitive to the value of $\beta$. This result is achieved via more active use of balance sheet policies (higher mean $q$) when $\beta$ is high (and conversely when $\beta$ is low). These results explain why the benefits from balance sheet policies are larger when the average real interest rate is lower: as shown in the 'optimal BSP gain' row. So optimal deployment of the central bank balance sheet may be particularly beneficial when the average level of $r^*$ is low.

# B Debt and quantitative easing parameters

This appendix provides details of how parameter values relating to debt and quantitative easing are chosen, based on UK data.

## B.1 Baseline parameterization (UK data)

The parameters relating to the scale and composition of government debt ($\Theta$ and $\delta$ respectively) are set using UK data published by the Debt Management Office.

Data on gross UK government debt at market prices is used to match the model concept. In the model, there is no difference between government debt measured on a gross and net basis. However, gross debt is the appropriate data concept since what matters in the model is the quantity of debt held by private agents. The total debt stock is simply the sum of short-term and long-term debt. Short-term debt is defined as Treasury bills plus (nominal) debt with less than three years residual maturity. Long-term debt is defined as the sum of (DMO defined) short conventional and medium conventional debt, plus half of the long conventional debt stock (which has an extremely long residual maturity of at least 15 years).[40] The ratio of these debt stocks gives $\delta$. The sum of the two (i.e., total debt) relative to nominal GDP gives $\Theta$, where the average to end-2008 is used for the calibration.

**Figure B.1:** Impulse responses to exogenous QE shock (United Kingdom)

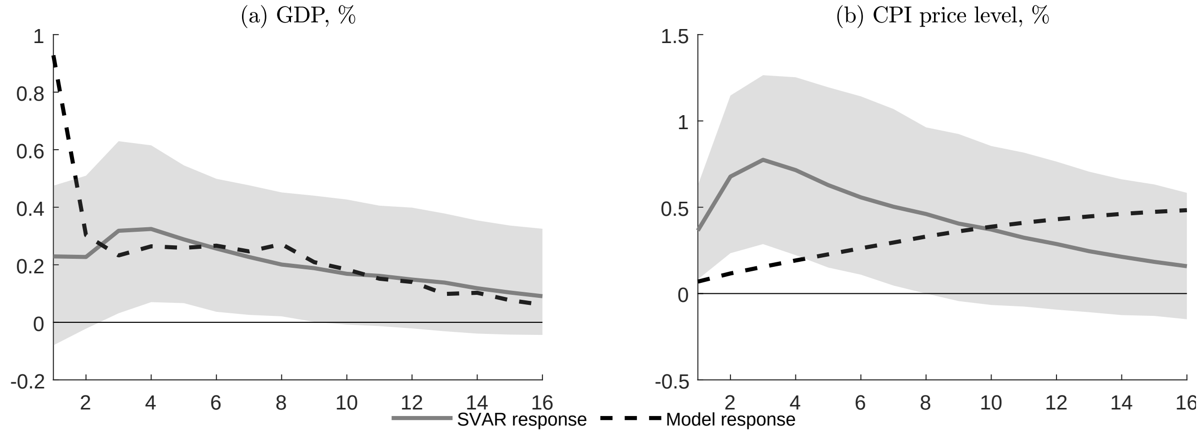


*Notes:* The solid gray line shows the impulse response to an exogenous shock to QE from the structural VAR estimated by Weale and Wieladek (2016) and the gray swathe represents the 95% posterior interval. The black dashed lines represent the response of the model to the same exogenous shock to QE.

The values of the parameters governing the effects of QE ($\nu$ and $\xi$) are set by minimizing the distance between the impulse response of the model to a QE shock with the SVAR estimates from Weale and Wieladek (2016). Responses from the Weale and Wieladek (2016) 'portfolio balance' identification scheme are used, since that corresponds to the mechanism through which QE has an effect in the model. For this experiment it is assumed that the balance sheet instrument follows a simple autoregressive process, so that $q_t = \rho_q q_{t-1} + \varepsilon_t^q$, where $\rho_q = 0.9875$ based on the estimated impulse response of the balance sheet. A sequence of fully anticipated disturbances $\varepsilon_t^q$ are chosen to exactly match the endogenous response of the balance sheet estimated by Weale and Wieladek (2016). This is implemented using the 'inversion algorithm' described in Burgess et al. (2013).

A numerical search is undertaken to find the vector $\{\nu, \xi\}$ that minimizes the distance between the impulse responses generated by the model and the SVAR median estimates. The unweighted sum of squared distances between the responses of the long-term bond rate are minimized, together with a much smaller weighting on the squared deviations of the output and inflation responses (discussed below). In computing the impulse response we assume that the short-term policy rate is held fixed for eight quarters (consistent with the identification assumption in Weale and Wieladek (2016)) again using a sequence of anticipated shocks to a simple rule for the short-term policy rate.[41]

Figure 2 in the main text shows the response of the long-term bond rate to a QE shock in the model alongside the 'target' from the SVAR. Figure B.1 shows the corresponding responses for output and inflation. While these responses are targeted in the impulse response matching

[40] The exclusion of some very long-dated debt is an adjustment to better match the maturity range of 'long term' debt with the maturity range of asset purchases typically undertaken in quantitative easing operations.

[41] The rule is $\hat{R}_t = (1 - \phi_R)\, \phi_\pi \hat{\pi}_t + \phi_R \hat{R}_{t-1} + \varepsilon_t^R$ with $\phi_\pi = 5$ and $\phi_R = 0.975$ to ensure that the policy rate does not adjust too quickly after the period during which it is held fixed, but also that inflation is stabilized aggressively in the longer run.

exercise, they receive a very small weight (0.01 compared with a unit weight on the balance sheet and long rate). Nevertheless, the model replicates the empirical responses relatively well, providing some support to the calibration of the model. The exception to this is the response of output in the first period, which lies outside the SVAR confidence interval. This reflects the fact that the effective real interest rate in the Euler equation jumps down on impact, given the sharp drop in the long-term bond rate shown in Figure 2. The lack of inertia in the IS curve implies that output rises strongly, driven by the intertemporal substitution response to the fall in the effective real interest rate.[42]

## B.2 Alternative parameterization (US data)

The US parameterization follows the same steps as described for the United Kingdom above. The steady state debt to (annual) GDP ratio is calibrated to 0.4, consistent with Greenwood et al. (2016a, Fig. 1.1) and multiplied by 0.9 to account for the fact that around 10% of US Treasuries are index-linked (and therefore not included as eligible debt for quantitative easing in the model). The share of long term debt is set to $\delta = 0.2$, following Greenwood et al. (2016a, Fig. 1.1).

The parameters $\nu$ and $\xi$ are selected by using the same minimum distance approach described in Appendix B.1.[43] Figures B.2 and B.3 show the responses of the balance sheet, long-term interest rate, output and inflation. The long-rate response has a similar kink to the UK response shown in the main text and matching this behavior requires $\xi > \nu$, also consistent with the UK calibration. The effects of the balance sheet shock on output and inflation are initially somewhat smaller in the model, though the price level effect builds over time.

**Figure B.2:** Impulse responses to exogenous QE shock (United States)

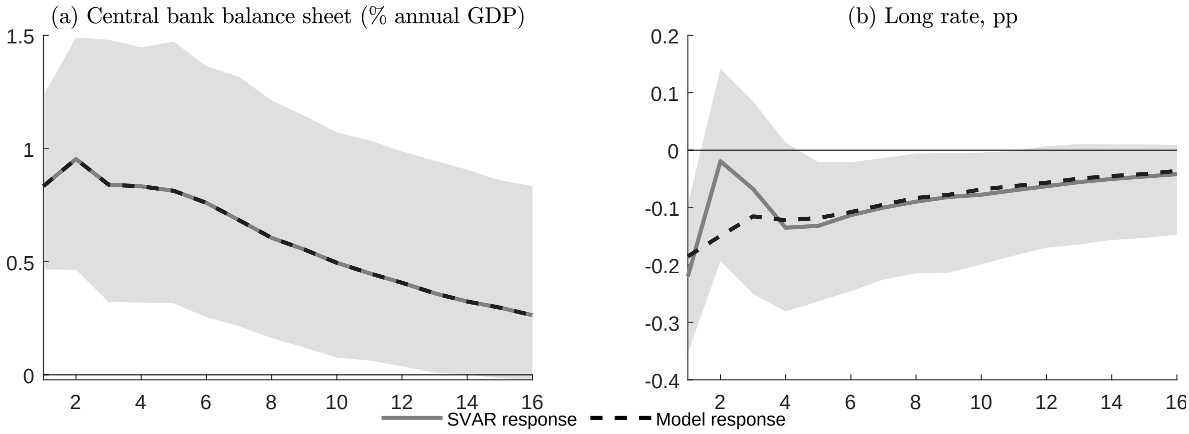


*Notes:* The solid gray line shows the impulse response to an exogenous shock to QE from the structural VAR estimated by Weale and Wieladek (2016) and the gray swathe represents the 95% posterior interval. The black dashed lines represent the response of the model to the same exogenous shock to QE.

## B.3 Rates of return

This Appendix analyzes the basic properties of the type of long-term government debt in the model. First consider the steady-state duration of the bond.

[42]Indeed, the 'impact' response of the output gap is given a weight of zero in the impulse response matching exercise.

[43]The exercise is identical except for the assumption that $\phi_\pi = 1.5$, which ensures that inflation is not excessively stabilized by an anticipated future monetary policy reaction.

**Figure B.3:** Impulse responses to exogenous QE shock (United States)

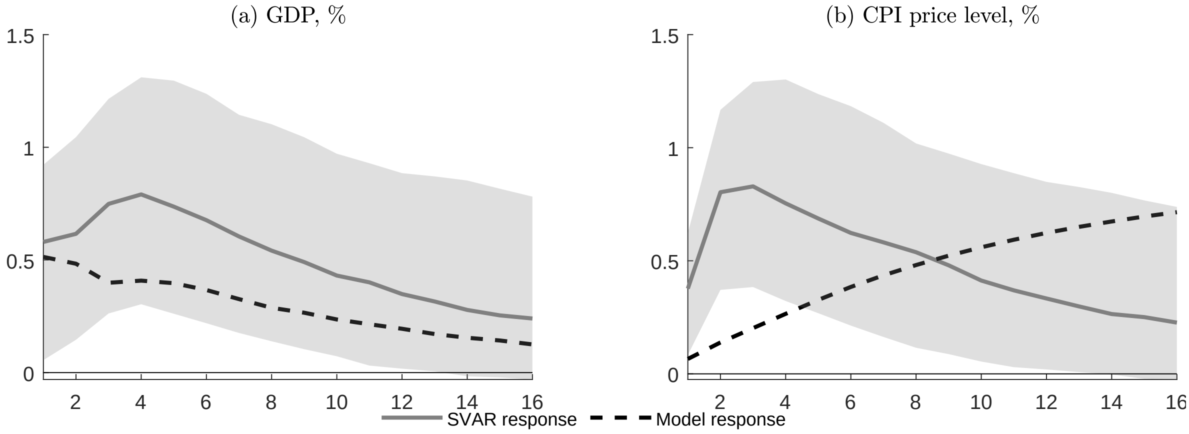


*Notes:* The solid gray line shows the impulse response to an exogenous shock to QE from the structural VAR estimated by Weale and Wieladek (2016) and the gray swathe represents the 95% posterior interval. The black dashed lines represent the response of the model to the same exogenous shock to QE.

The *expected* payoffs from the bond discounted using the steady state real rate ($\beta^{-1}$):

$$\beta(1-\chi)+\beta^2\chi(1-\chi)+\beta^3\chi^2(1-\chi)+\cdots=\beta(1-\chi)\sum_{i=0}^{\infty}(\beta\chi)^i=\frac{\beta(1-\chi)}{1-\beta\chi}$$

No arbitrage implies $R_L = R$ and $R = \beta^{-1}$ in steady state. So:

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

Duration is the ratio of payoffs weighted by payment date and the value of the bond:

$$D=\frac{W}{V}$$

where

$$W\equiv\beta(1-\chi)\times 1+\beta^2\chi(1-\chi)\times 2+\beta^3\chi^2(1-\chi)\times 3+\ldots$$

We can see that:

$$\begin{aligned}W=&\ \beta(1-\chi)+\beta^2\chi(1-\chi)+\beta^3\chi^2(1-\chi)+\ldots\\ &+\qquad\qquad \beta^2\chi(1-\chi)+\beta^3\chi^2(1-\chi)+\ldots\\ &+\qquad\qquad\qquad\qquad\quad \beta^3\chi^2(1-\chi)+\ldots\\ &=V+\beta\chi V+(\beta\chi)^2V+\ldots\\ &=\frac{V}{1-\beta\chi}\end{aligned}$$

This means that

$$D=\frac{1}{1-\beta\chi}$$

This implies that a bond that pays a zero coupon and being redeemed with a fixed probability is equivalent to one that is never redeemed but pays a geometrically declining coupon. This means that the following results, shown by Woodford (2001) and Chen et al. (2012) can be used:

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

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

Log-linearizing the first expression gives:

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

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

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

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

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

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 equalized.

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

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

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

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

or

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

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

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

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

$$\mathbb{E}_t\hat{R}_{t+1}^D = \hat{R}_t - \nu\left(1+\delta^{-1}\right)q_t - \xi\left(1+\delta^{-1}\right)\left(q_t - q_{t-1}\right) + \beta\xi\left(1+\delta^{-1}\right)\mathbb{E}_t\left(q_{t+1} - q_t\right)$$

This implies that the yield to maturity is given by:

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

# C Model derivation

## C.1 Financial intermediation

The intermediary, $I$, collects savings ($S$) from households and invests in one-period government bonds ($B$), long-term debt ($D$) and reserves ($Z$).

The period $t$ profit, in nominal terms is given by:

$$\begin{aligned}\Omega_t^I =& S_t - B_t^I - D_t^I - Z_t^I + R_{t-1}Z_{t-1}^I + R_{t-1}^B B_{t-1}^I + R_t^D D_{t-1}^I - R_{t-1}^S S_{t-1} \\ &- \left(z^I + b^I + d^I\right)P_t\mathcal{M}\left(\delta\frac{Z_t^I + B_t^I}{D_t^I}\right) - \left(z^I + b^I + d^I\right)P_t\mathcal{A}\left(\frac{Z_t^I + B_t^I}{D_t^I} - \frac{Z_{t-1}^I + B_{t-1}^I}{D_{t-1}^I}\right)\end{aligned}$$

where $z^I$, $b^I$ and $d^I$ are the real steady-state values of the intermediary's holdings of reserves, bonds and debt respectively.

The intermediary collects savings, $S$ and makes investments in bonds, debt and reserves, receives interest on previous investments in these instruments and pays households interest on their previous savings. The intermediary also pays convex costs of maintaining and adjusting its relative mix of reserves, bonds and debt, denoted $\mathcal{M}$ and $\mathcal{A}$ respectively. These functions satisfy $\mathcal{M}(1) = \mathcal{M}'(1) = \mathcal{A}(0) = \mathcal{A}'(0) = 0$ and $\mathcal{M}''(1) = \tilde{\nu}$ and $\mathcal{A}''(0) = \tilde{\xi}$.

The real profit of the intermediary is given by:

$$\begin{aligned}\omega_t^I \equiv \frac{\Omega_t^I}{P_t} =& s_t - b_t^I - d_t^I - z_t^I + \frac{R_{t-1}}{\pi_t}z_{t-1}^I + \frac{R_{t-1}^B}{\pi_t}b_{t-1}^I + \frac{R_t^D}{\pi_t}d_{t-1}^I - \frac{R_{t-1}^S}{\pi_t}s_{t-1} \\ &- \left(z^I + b^I + d^I\right)\mathcal{M}\left(\delta\frac{z_t^I + b_t^I}{d_t^I}\right) - \left(z^I + b^I + d^I\right)\mathcal{A}\left(\frac{\left(z_t^I + b_t^I\right)/d_t^I}{\left(z_{t-1}^I + b_{t-1}^I\right)/d_{t-1}^I}\right)\end{aligned}$$

The balance sheet constraint is:

$$S_t \geq B_t^I + D_t^I + Z_t^I$$

which will bind in equilibrium.

The maximization problem of the financial intermediary is

$$\max \mathbb{E}_t\sum_{k=0}^{\infty}\beta^k\Lambda_{t+k}\omega_{t+k}^I$$

where $\Lambda$ is the household's marginal utility of consumption.

The first order conditions are:

$$
\begin{aligned}
0 = {} & \Lambda_t - \beta \mathbb{E}_t \Lambda_{t+1} \frac{R_t^S}{\pi_{t+1}} \\
0 = {} & -\Lambda_t - \Lambda_t \left(z^I + b^I + d^I\right) \frac{\delta}{d_t^I} \mathcal{M}' \left(\delta \frac{z_t^I + b_t^I}{d_t^I}\right) \\
& - \Lambda_t \frac{z^I + b^I + d^I}{d_t^I} \mathcal{A}' \left(\frac{z_t^I + b_t^I}{d_t^I} - \frac{z_{t-1}^I + b_{t-1}^I}{d_{t-1}^I}\right) \\
& + \beta \mathbb{E}_t \Lambda_{t+1} \frac{R_t}{\pi_{t+1}} + \beta \mathbb{E}_t \Lambda_{t+1} \frac{z^I + b^I + d^I}{d_t^I} \mathcal{A}' \left(\frac{z_{t+1}^I + b_{t+1}^I}{d_{t+1}^I} - \frac{z_t^I + b_t^I}{d_t^I}\right) \\
0 = {} & -\Lambda_t - \Lambda_t \left(z^I + b^I + d^I\right) \frac{\delta}{d_t^I} \mathcal{M}' \left(\delta \frac{z_t^I + b_t^I}{d_t^I}\right) \\
& - \Lambda_t \frac{z^I + b^I + d^I}{d_t^I} \mathcal{A}' \left(\frac{z_t^I + b_t^I}{d_t^I} - \frac{z_{t-1}^I + b_{t-1}^I}{d_{t-1}^I}\right) \\
& + \beta \mathbb{E}_t \Lambda_{t+1} \frac{R_t^B}{\pi_{t+1}} + \beta \mathbb{E}_t \Lambda_{t+1} \frac{z^I + b^I + d^I}{d_t^I} \mathcal{A}' \left(\frac{z_{t+1}^I + b_{t+1}^I}{d_{t+1}^I} - \frac{z_t^I + b_t^I}{d_t^I}\right)
\end{aligned}
$$

and

$$
\begin{aligned}
0 = {} & -\Lambda_t + \Lambda_t \left(z^I + b^I + d^I\right) \delta \frac{z_t^I + b_t^I}{\left(d_t^I\right)^2} \mathcal{M}' \left(\delta \frac{z_t^I + b_t^I}{d_t^I}\right) \\
& + \Lambda_t \left(z^I + b^I + d^I\right) \frac{z_t^I + b_t^I}{\left(d_t^I\right)^2} \mathcal{A}' \left(\frac{z_t^I + b_t^I}{d_t^I} - \frac{z_{t-1}^I + b_{t-1}^I}{d_{t-1}^I}\right) \\
& + \beta \mathbb{E}_t \Lambda_{t+1} \frac{R_{t+1}^D}{\pi_{t+1}} - \beta \mathbb{E}_t \Lambda_{t+1} \left(z^I + b^I + d^I\right) \frac{z_t^I + b_t^I}{\left(d_t^I\right)^2} \mathcal{A}' \left(\frac{z_{t+1}^I + b_{t+1}^I}{d_{t+1}^I} - \frac{z_t^I + b_t^I}{d_t^I}\right)
\end{aligned}
$$

Several results can be used to simplify the log-linearization process. In particular, the total stocks of bonds and long-term debt are held fixed:

$$b_t = \bar{b} \quad , \quad d_t = \bar{d}$$

and the steady state level of reserves and bonds satisfies:

$$z^I + b^I = z + b = \delta^{-1} \bar{d}$$

since reserves are zero in steady state, so that $z^I = 0$.

Real debt purchases by the central bank ($d^C$) are denoted by:

$$d_t^C = q_t \bar{d}$$

Note also that market clearing requires:

$$\bar{d} = d_t^I + d_t^C \Rightarrow d_t^I = \left(1 - q_t\right) \bar{d}$$

and since $d^I = \bar{d}$, in steady state:

$$\bar{d} \hat{d}_t^I = -\bar{d} \left(q_t - 0\right) \Rightarrow \hat{d}_t^I = -q_t$$

where the expression is *linearized* with respect to $q_t$.

Finally, note that the central bank balance sheet implies that

$$z_t = q_t d_t$$

so that

$$z_t - 0 = \bar{d}\,(q_t - 0) \Rightarrow z_t = \bar{d} q_t$$

The above observations imply that linearizing $\frac{z_t^I + b_t^I}{d_t^I}$ gives:

$$\frac{1}{d^I}\bar{d}q_t + \frac{b^I}{d^I}\hat{b}_t^I - \frac{b^I}{d^I}\hat{d}_t^I = \left(1+\delta^{-1}\right)q_t$$

With these results in hand, log linearizing the first order conditions gives:

$$\begin{aligned}
0 =& \Lambda\hat{\Lambda}_t - \Lambda\mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_t^S - \hat{\pi}_{t+1}\right] \\
0 =& -\Lambda\hat{\Lambda}_t + \Lambda\mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_t - \hat{\pi}_{t+1}\right] - \Lambda\left(z^I + b^I + d^I\right)\frac{\delta}{d^I}\mathcal{M}''(1)(1+\delta)q_t \\
& - \Lambda\frac{z^I + b^I + d^I}{d^I}\mathcal{A}''(0)\left(1+\delta^{-1}\right)(q_t - q_{t-1}) \\
& + \beta\Lambda\frac{z^I + b^I + d^I}{d^I}\mathcal{A}''(0)\left(1+\delta^{-1}\right)\mathbb{E}_t(q_{t+1} - q_t) \\
0 =& -\Lambda\hat{\Lambda}_t + \Lambda\mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_t^B - \hat{\pi}_{t+1}\right] - \Lambda\left(z^I + b^I + d^I\right)\frac{\delta}{d^I}\mathcal{M}''(1)(1+\delta)q_t \\
& - \Lambda\frac{z^I + b^I + d^I}{d^I}\mathcal{A}''(0)\left(1+\delta^{-1}\right)(q_t - q_{t-1}) \\
& + \beta\Lambda\frac{z^I + b^I + d^I}{d^I}\mathcal{A}''(0)\left(1+\delta^{-1}\right)\mathbb{E}_t(q_{t+1} - q_t) \\
0 =& -\Lambda\hat{\Lambda}_t + \Lambda\mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_{t+1}^D - \hat{\pi}_{t+1}\right] + \Lambda\left(z^I + b^I + d^I\right)\delta\frac{z^I + b^I}{\left(d^I\right)^2}\mathcal{M}''(1)(1+\delta)q_t \\
& + \Lambda\left(z^I + b^I + d^I\right)\frac{z^I + b^I}{\left(d^I\right)^2}\mathcal{A}''(0)\left(1+\delta^{-1}\right)(q_t - q_{t-1}) \\
& - \beta\Lambda\left(z^I + b^I + d^I\right)\frac{z^I + b^I}{\left(d^I\right)^2}\mathcal{A}''(0)\left(1+\delta^{-1}\right)\mathbb{E}_t(q_{t+1} - q_t)
\end{aligned}$$

Note that the second and third conditions show that

$$\hat{R}_t^B = \hat{R}_t$$

so the focus in what follows will be on the first, second and fourth conditions.

Making cancellations and collecting terms gives:

$$\begin{aligned}
0 =& \hat{\Lambda}_t - \mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_t^S - \hat{\pi}_{t+1}\right] \\
0 =& -\hat{\Lambda}_t + \mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_t - \hat{\pi}_{t+1}\right] - \tilde{\nu}(1+\delta)^2 q_t - \tilde{\xi}\left(1+\delta^{-1}\right)^2(q_t - q_{t-1}) \\
& + \beta\tilde{\xi}\left(1+\delta^{-1}\right)^2\mathbb{E}_t(q_{t+1} - q_t) \\
0 =& -\hat{\Lambda}_t + \mathbb{E}_t\left[\hat{\Lambda}_{t+1} + \hat{R}_{t+1}^D - \hat{\pi}_{t+1}\right] + \tilde{\nu}\delta^{-1}(1+\delta)^2 q_t + \tilde{\xi}\delta^{-1}\left(1+\delta^{-1}\right)^2(q_t - q_{t-1}) \\
& - \beta\tilde{\xi}\delta^{-1}\left(1+\delta^{-1}\right)^2\mathbb{E}_t(q_{t+1} - q_t)
\end{aligned}$$

Using the first equation in the third and fourth gives:

$$
\begin{aligned}
\hat{R}_t^S &= \hat{R}_t - \tilde{\nu}\left(1+\delta\right)^2 q_t - \tilde{\xi}\left(1+\delta^{-1}\right)^2\left(q_t - q_{t-1}\right) + \beta\tilde{\xi}\left(1+\delta^{-1}\right)^2 \mathbb{E}_t\left(q_{t+1} - q_t\right) \\
\hat{R}_t^S &= \mathbb{E}_t\hat{R}_{t+1}^D + \tilde{\nu}\delta^{-1}\left(1+\delta\right)^2 q_t + \tilde{\xi}\delta^{-1}\left(1+\delta^{-1}\right)^2\left(q_t - q_{t-1}\right) - \beta\tilde{\xi}\delta^{-1}\left(1+\delta^{-1}\right)^2 \mathbb{E}_t\left(q_{t+1} - q_t\right)
\end{aligned}
$$

Taking a linear combination of the two equations gives:

$$
\hat{R}_t^S = \frac{1}{1+\delta}\hat{R}_t + \frac{\delta}{1+\delta}\mathbb{E}_t\hat{R}_{t+1}^D
$$

The first of the two equations can be written as:

$$
\hat{R}_t^S = \hat{R}_t - \nu q_t - \xi\left(q_t - q_{t-1}\right) + \beta\xi\mathbb{E}_t\left(q_{t+1} - q_t\right) \tag{C.1}
$$

where

$$
\begin{aligned}
\nu &\equiv \tilde{\nu}\left(1+\delta\right)^2 \\
\xi &\equiv \tilde{\xi}\left(1+\delta^{-1}\right)^2
\end{aligned}
$$

Subtracting the two equations gives:

$$
\mathbb{E}_t\hat{R}_{t+1}^D = \hat{R}_t - \tilde{\nu}\delta^{-1}\left(1+\delta\right)^3 q_t - \tilde{\xi}\left(1+\delta^{-1}\right)^3\left(q_t - q_{t-1}\right) + \beta\tilde{\xi}\left(1+\delta^{-1}\right)^3 \mathbb{E}_t\left(q_{t+1} - q_t\right)
$$

so that

$$
\mathbb{E}_t\hat{R}_{t+1}^D = \hat{R}_t - \nu\left(1+\delta^{-1}\right) q_t - \xi\left(1+\delta^{-1}\right)\left(q_t - q_{t-1}\right) + \beta\xi\left(1+\delta^{-1}\right) \mathbb{E}_t\left(q_{t+1} - q_t\right)
$$

## C.2 Households

The optimization problem considered in Section 2.4 is

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

subject to

$$
S_t = R_{t-1}^S S_{t-1} + W_t n_t + T_t + F_t - P_t c_t \tag{C.2}
$$

The first-order conditions of the optimization problem are:

$$
\phi_t c_t^{-\frac{1}{\sigma}} = \mu_t P_t \tag{C.3}
$$

$$
\phi_t n_t^{\psi} = W_t \mu_t \tag{C.4}
$$

$$
0 = -\mu_t + \beta R_t^S \mathbb{E}_t \mu_{t+1} \tag{C.5}
$$

where $\mu$ is the Lagrange multiplier on the nominal budget constraint (C.2).

Define the real Lagrange multiplier as:

$$
\Lambda_t \equiv P_t \mu_t
$$

The first order condition for savings, (C.5), can be written in terms of real-valued variables as:

$$0 = -\Lambda_t + \beta R_t^S \mathbb{E}_t \Lambda_{t+1} \pi_{t+1}^{-1} \tag{C.6}$$

Combining (C.3) and (C.6) creates an Euler equation for consumption:

$$\phi_t c_t^{-\frac{1}{\sigma}} = \beta R_t^S \mathbb{E}_t \phi_{t+1} c_{t+1}^{-\frac{1}{\sigma}} \pi_{t+1}^{-1}$$

which can be log-linearized to give:

$$\hat{c}_t = \mathbb{E}_t \hat{c}_{t+1} - \sigma \left[ \hat{R}_t^S - \mathbb{E}_t \hat{\pi}_{t+1} + \mathbb{E}_t \Delta \hat{\phi}_{t+1} \right] \tag{C.7}$$

The first order conditions for labor supply (C.4) and consumption (C.3) can be combined and log-linearized to give

$$\psi \hat{n}_t = \hat{w}_t - \sigma^{-1} \hat{c}_t \tag{C.8}$$

## C.3 Firms

The nominal profit of producer $j$ is:

$$(1+s)\, P_{j,t} y_{j,t} - \mathcal{C}_{j,t}$$

where $s$ is a subsidy paid to producers to ensure that the steady-state level of output is efficient and $\mathcal{C}_{j,t}$ is the nominal cost function.

The demand function for producer $j$ is

$$y_{j,t} = \left( \frac{P_{j,t}}{P_t} \right)^{-\eta_t} y_t$$

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

$$\max \mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k P_k^{-1} (\beta\theta)^{k-t} \left( (1+s)\, P_{j,t} y_{j,t+k} - \mathcal{C}_{j,k} \right)$$

where $\Lambda$ represents the marginal utility of consumption and $0 \leq \theta < 1$ is the probability that the producer is *not* allowed to reset its price each period.

The first order condition is

$$\mathbb{E}_t \sum_{k=t}^{\infty} \Lambda_k P_k^{-1} y_{j,t+k} (\beta\theta)^{k-t} \left( (1-\eta_t)(1+s)\, P_{j,t} + \eta_t \mathcal{C}'_{j,k} \right) = 0$$

where $\mathcal{C}'$ is the nominal marginal cost.

Log-linearizing around the nominal path under flexible prices (and no markup disturbances), in which $P_{j,k} = \frac{\eta}{\eta-1} \mathcal{C}'_k$ gives:

$$\mathbb{E}_t \sum_{k=t}^{\infty} (\beta\theta)^{k-t} \left( \hat{p}_t^* - \widehat{\mathcal{C}}'_{j,k} + \frac{\eta}{\eta-1} \hat{\eta}_k \right) = 0$$

where the price that all firms that change prices at date $t$ is denoted $\hat{p}_t^*$.

Rearranging the first order condition shows that this price satisfies

$$\hat{p}_t^* = (1-\beta\theta)\,\mathbb{E}_t \sum_{k=t}^{\infty} (\beta\theta)^{k-t} \left( \widehat{\mathcal{C}'_{j,k}} - \frac{\eta}{\eta-1}\hat{\eta}_k \right)$$

The log-linearized marginal cost of producer $j$ in period $k \geq t$ is

$$\widehat{\mathcal{C}'_{j,k}} = \hat{p}_k + \hat{w}_k + \alpha\hat{n}_{j,k}$$

which is the nominal wage less the marginal product of labor.

Average marginal costs in period $k$ satisfy

$$\widehat{\mathcal{C}'_k} = \hat{p}_k + \hat{w}_k + \alpha\hat{n}_k$$

where $\hat{n}_k$ is the log-deviation of $n_k \equiv \int_{j=0}^{1} n_{k,j}$.

This implies that

$$\begin{aligned}
\widehat{\mathcal{C}'_{j,k}} &= \hat{p}_k + \hat{w}_k + \alpha\hat{n}_{j,k} \\
&= \widehat{\mathcal{C}'_k} + \alpha\left(\hat{n}_{j,k} - \hat{n}_k\right) \\
&= \widehat{\mathcal{C}'_k} + \frac{\alpha}{1-\alpha}\left(\hat{y}_{j,k} - \hat{y}_k\right) \\
&= \widehat{\mathcal{C}'_k} - \frac{\eta\alpha}{1-\alpha}\left(\hat{p}_t^* - \hat{p}_k\right)
\end{aligned}$$

which uses first-order approximations to the aggregate production ($\hat{y}_t = (1-\alpha)\,\hat{n}_t$) and the demand function for variety $j$.

Using this relationship in the pricing equation gives

$$\hat{p}_t^* = (1-\beta\theta)\,\mathbb{E}_t \sum_{k=t}^{\infty} (\beta\theta)^{k-t} \left( \widehat{\mathcal{C}'_k} - \frac{\eta\alpha}{1-\alpha}\left(\hat{p}_t^* - \hat{p}_k\right) - \frac{\eta}{\eta-1}\hat{\eta}_k \right)$$

which can be rearranged to give

$$\hat{p}_t^* = (1-\beta\theta)\,\mathbb{E}_t \sum_{k=t}^{\infty} (\beta\theta)^{k-t} \left( \hat{p}_k - \Upsilon\hat{\vartheta}_k - \frac{\Upsilon\eta}{\eta-1}\hat{\eta}_k \right)$$

where

$$\hat{\vartheta}_t \equiv p_t - \widehat{\mathcal{C}'_t} \tag{C.9}$$

and

$$\Upsilon \equiv \frac{1-\alpha}{1-\alpha+\eta\alpha}$$

The pricing equation can be expressed recursively as:

$$\hat{p}_t^* = \beta\theta\mathbb{E}_t\hat{p}_{t+1}^* + (1-\beta\theta)\left( \hat{p}_t - \Upsilon\hat{\vartheta}_t - \frac{\Upsilon\eta}{\eta-1}\hat{\eta}_t \right) \tag{C.10}$$

The aggregate price is:

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

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 $(1-\theta)\,\theta^k$. This means that the aggregate price level can be written as

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

so that

$$\Pi_t = \left[ \theta + (1-\theta) \left( \frac{P_t^*}{P_t} \right)^{1-\eta_t} \right]^{\frac{1}{1-\eta_t}}$$

and hence

$$\hat{\pi}_t = (1-\theta)\left( \hat{p}_t^* - \hat{p}_{t-1} \right)$$

Using this expression in (C.10) implies that

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} - \frac{(1-\beta\theta)(1-\theta)}{\theta} \left( \Upsilon \hat{\vartheta}_t + \frac{\Upsilon \eta}{\eta - 1} \hat{\eta}_t \right)$$

Finally, note that the markup satisfies:

$$\begin{aligned} \hat{\vartheta}_t &= p_t - \widehat{\mathcal{C}'_t} \\ &= -\hat{w}_t - \alpha \hat{n}_t \\ &= -\left( \sigma^{-1} \hat{c}_t + \psi \hat{n}_t \right) - \alpha \hat{n}_t \\ &= -\sigma^{-1} \hat{y}_t - (\alpha + \psi)\, \hat{n}_t \\ &= -\sigma^{-1} \hat{y}_t - \frac{\alpha + \psi}{1-\alpha} \hat{y}_t \end{aligned}$$

where the second line uses the definition of the markup (price over marginal product of labor), the third line uses the household labor supply condition and the fourth line uses market clearing (and collects terms). The final line uses the aggregate production relationship.

This means that the Phillips curve can be written as:

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \frac{(1-\beta\theta)(1-\theta)}{\theta} \Upsilon \left[ \sigma^{-1} + \frac{\alpha+\psi}{1-\alpha} \right] \hat{y}_t - \frac{(1-\beta\theta)(1-\theta)}{\theta} \frac{\Upsilon \eta}{\eta - 1} \hat{\eta}_t \tag{C.11}$$

## C.4 Market clearing and the efficient allocation

Goods market clearing requires:

$$c_t = y_t - (b+d)\, \mathcal{M} \left( \delta \frac{z_t + b}{d_t^I} \right) - (b+d)\, \mathcal{A} \left( \frac{z_t + b}{d_t^I} - \frac{z_{t-1} + b}{d_{t-1}^I} \right)$$

Output for each variety $j$ satisfies:

$$y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\eta_t} y_t$$
$$y_{jt} = An_{jt}^{1-\alpha}$$

Equating the previous expressions and integrating over $j$ gives:

$$\int_0^1 An_{jt}^{1-\alpha}\mathrm{d}j = \int_0^1 \left(\frac{P_{jt}}{P_t}\right)^{-\eta_t} y_t\mathrm{d}j$$

which implies that:

$$An_t^{1-\alpha} = \mathcal{D}_t y_t$$

where

$$\mathcal{D}_t \equiv \int_0^1 \left(\frac{P_{jt}}{P_t}\right)^{-\eta_t} \mathrm{d}j \tag{C.12}$$

is a measure of price dispersion.

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

$$\hat{b}_t^h - \hat{b}_{L,t}^h = -\hat{b}_{L,t}^h = q_t \tag{C.13}$$

It is straightforward to show that in the absence of price-setting and imperfect asset substitutability frictions, the efficient level of output is constant. 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 labor. So the efficient allocations, denoted with an asterisk, can be found from the labor supply relation (C.8):

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

Imposing market clearing (with zero portfolio adjustment costs and price dispersion equal to 1) gives $c_t^* = (n_t^*)^{1-\alpha} = y_t^*$. This in turn implies that $\hat{c}_t^* = (1-\alpha)\,\hat{n}_t^* = \hat{y}_t^*$. Since the labor supply equation requires that $\hat{n}_t^* = -\left((\psi+\alpha)\,\sigma\right)^{-1}\hat{c}_t^*$, we must have $\hat{c}_t^* = \hat{n}_t^* = \hat{y}_t^* = 0$.

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 ($\hat{y}_t^* = 0$).

The Phillips curve (C.11) can be written as:

$$\hat{\pi}_t = \kappa\hat{x}_t + \beta\mathbb{E}_t\hat{\pi}_{t+1} + u_t \tag{C.14}$$

where

$$\kappa \equiv \Gamma\Xi$$

where

$$\Xi \equiv \sigma^{-1} + \frac{\psi+\alpha}{1-\alpha}$$

and

$$\Gamma \equiv \frac{(1-\theta)\,(1-\beta\theta)}{\theta}\frac{1-\alpha}{1-\alpha+\eta\alpha}$$

and the cost push shock, $u$, is defined as:

$$u_t \equiv -\frac{(1-\theta)\,(1-\beta\theta)}{\theta}\frac{\Upsilon\eta}{\eta-1}\hat{\eta}_t$$

The Euler equation for consumption (C.7) can be written as:

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

which incorporates the market clearing condition for output and uses the equation for the saving rate, $\hat{R}^S$, (C.1).

Collecting terms, this can be written as:

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

In terms of the output gap we have:

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

where

$$\gamma \equiv \nu + \xi\left(1+\beta\right)$$

and the efficient rate of interest $r^*$ satisfies

$$r_t^* \equiv -\mathbb{E}_t\Delta\hat{\phi}_{t+1} \tag{C.15}$$

# D Utility-based loss function

Ignoring constants, the period utility function is:

$$U_t = \phi_t\left[\frac{c_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{n_t^{1+\psi}}{1+\psi}\right] \equiv \phi_t\bar{U}_t$$

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

The second order approximation of utility around the efficient steady state is given by:

$$\begin{aligned} U_t - U \approx\ & U_c c\left(\frac{c_t - c}{c}\right) + U_n n\left(\frac{n_t - n}{n}\right) + \frac{1}{2}U_{cc}c^2\left(\frac{c_t - c}{c}\right)^2 \\ & + \frac{1}{2}U_{nn}n^2\left(\frac{n_t - n}{n}\right)^2 + U_c c\left(\frac{c_t - c}{c}\right)\left(\frac{\phi_t - \phi}{\phi}\right) \\ & + U_n n\left(\frac{n_t - n}{n}\right)\left(\frac{\phi_t - \phi}{\phi}\right) + t.i.p. \end{aligned}$$

where terms that are independent of policy are denoted as '$t.i.p.$'.

To proceed, it is useful to 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 $\hat{z}_t \equiv \ln z_t - \ln z$.

Given the isoelastic form of the utility functions for consumption and hours worked, the approximation can be written as:

$$
\begin{aligned}
U_t - U \approx U_c c & \left( \hat{c}_t \left(1 + \hat{\phi}_t\right) + \frac{1-\sigma^{-1}}{2} \hat{c}_t^2 \right) \\
& + U_n n \left( \hat{n}_t \left(1 + \hat{\phi}_t\right) + \frac{1+\psi}{2} \hat{n}_t^2 \right) + t.i.p.
\end{aligned}
$$

The goods market clearing condition is:

$$
c_t = y_t - (b+d)\,\mathcal{M}\left(\delta \frac{z_t + b}{d_t^I}\right) - (b+d)\,\mathcal{A}\left(\frac{z_t + b}{d_t^I} - \frac{z_{t-1} + b}{d_{t-1}^I}\right)
$$

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

$$
\begin{aligned}
\hat{c}_t + \frac{1}{2}\hat{c}_t^2 &= \hat{y}_t + \frac{1}{2}\hat{y}_t^2 - \frac{\tilde{\nu}\,(b+d)}{2}(1+\delta)^2 q_t^2 - \frac{\tilde{\xi}\,(b+d)}{2}\left(\left(1+\delta^{-1}\right)\Delta q_t\right)^2 \\
&\approx \hat{y}_t + \frac{1}{2}\hat{y}_t^2 - \frac{\nu\Theta}{2} q_t^2 - \frac{\xi\Theta}{2}(\Delta q_t)^2
\end{aligned}
$$

where $\Theta \equiv b + d$.

This implies that:

$$
\begin{aligned}
\hat{c}_t \approx & -\frac{1}{2}\hat{c}_t^2 + \hat{y}_t + \frac{1}{2}\hat{y}_t^2 - \frac{\nu\Theta}{2} q_t^2 - \frac{\xi\Theta}{2}(\Delta q_t)^2 \\
\approx & \ \hat{y}_t - \frac{\nu\Theta}{2} q_t^2 - \frac{\xi\Theta}{2}(\Delta q_t)^2 + h.o.t.
\end{aligned}
$$

where the second equality follows from using the first expression for $\hat{c}_t$ on the right hand side of the first equality and where $h.o.t.$ denotes higher order terms.

This can be used to eliminate $\hat{c}_t$ from the approximation:

$$
\begin{aligned}
U_t - U \approx U_c c & \left( \hat{y}_t \left(1 + \hat{\phi}_t\right) + \frac{1-\sigma^{-1}}{2} \hat{y}_t^2 - \frac{\nu\Theta}{2} q_t^2 - \frac{\xi\Theta}{2}(\Delta q_t)^2 \right) \\
& + U_n n \left( \hat{n}_t \left(1 + \hat{\phi}_t\right) + \frac{1+\psi}{2} \hat{n}_t^2 \right) + t.i.p.
\end{aligned}
$$

The aggregate production function implies an exact log-linear relationship:

$$
\hat{y}_t = (1-\alpha)\,\hat{n}_t - \hat{\mathcal{D}}_t
$$

which can be used to eliminate $\hat{n}_t$ from the approximation:

$$
\begin{aligned}
U_t - U \approx U_c c & \left( \hat{y}_t \left(1 + \hat{\phi}_t\right) + \frac{1-\sigma^{-1}}{2} \hat{y}_t^2 - \frac{\nu\Theta}{2} q_t^2 - \frac{\xi\Theta}{2}(\Delta q_t)^2 \right) \\
& + \frac{U_n n}{1-\alpha} \left( \hat{y}_t \left(1 + \hat{\phi}_t\right) + \hat{\mathcal{D}}_t + \frac{1+\psi}{2\,(1-\alpha)} \hat{y}_t^2 \right) + t.i.p.
\end{aligned}
$$

which uses the fact that $\hat{\mathcal{D}}_t$ is a second-order term.

The steady-state labor supply relationship is

$$
n^{\psi} = w c^{-1/\sigma} = (1-\alpha)\,\frac{y}{n} c^{-1/\sigma}
$$

which follows from the assumption that subsidies to firms are set to eliminate the distortion from monopolistic competition.

This relationship implies that in steady state (when $c = y$), $\frac{U_n n}{1-\alpha} = -U_c c$, so that the approximation to utility can be written as:

$$U_t - U \approx U_c c\left(\frac{1-\sigma^{-1}}{2}\hat{y}_t^2 - \frac{\nu\Theta}{2}q_t^2 - \frac{\xi\Theta}{2}\left(\Delta q_t\right)^2\right) - U_c c\left(\hat{\mathcal{D}}_t + \frac{1+\psi}{2\left(1-\alpha\right)}\hat{y}_t^2\right) + t.i.p.$$

so that

$$\begin{aligned}\frac{U_t - U}{U_c c} &\approx \frac{1}{2}\left(\left(1-\sigma^{-1} - \frac{1+\psi}{1-\alpha}\right)\hat{x}_t^2 - \nu\Theta q_t^2 - \xi\Theta\left(\Delta q_t\right)^2 - 2\hat{\mathcal{D}}_t\right) \\ &\approx -\frac{1}{2}\left(\Xi\hat{x}_t^2 + \nu\Theta q_t^2 + \xi\Theta\left(\Delta q_t\right)^2 + 2\hat{\mathcal{D}}_t\right)\end{aligned}$$

where

$$\Xi \equiv \sigma^{-1} + \frac{\alpha+\psi}{1-\alpha}$$

and the fact that efficient output is constant (so that $\hat{y}_t = \hat{x}_t$)has been used.

The intertemporal loss function can be defined in terms of the discounted approximated utility loss:

$$\mathcal{L}_t \equiv -\mathbb{E}_t\sum_{\tau=t}^{\infty}\beta^{\tau-t}\frac{U_\tau - U}{U_c c}$$

Using the above approximation, this implies that:

$$\mathcal{L}_t = \frac{1}{2}\mathbb{E}_t\sum_{\tau=t}^{\infty}\beta^{\tau-t}\left(\Xi\hat{x}_\tau^2 + \nu\Theta q_\tau^2 + \xi\Theta\left(\Delta q_\tau\right)^2 + 2\hat{\mathcal{D}}_\tau\right)$$

Since the price setting assumptions are standard, Lemmas 1 and 2 of Galí (2015, p118–119) can be applied, so that:

$$\mathcal{L}_t = \frac{1}{2}\mathbb{E}_t\sum_{\tau=t}^{\infty}\beta^{\tau-t}\left(\Xi\hat{x}_\tau^2 + \frac{\eta}{\Gamma}\hat{\pi}_\tau^2 + \nu\Theta q_\tau^2 + \xi\Theta\left(\Delta q_\tau\right)^2\right)$$

where

$$\Gamma \equiv \frac{\left(1-\theta\right)\left(1-\beta\theta\right)}{\theta}\frac{1-\alpha}{1-\alpha+\eta\alpha}$$

# 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 payoff relevant state variables. 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 recognizes that future policymakers will choose allocations according to time-invariant Markovian policy functions. Upper case bold letters denote these policy functions. For example, inflation at date $t+j$ is given by the function:

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

where $z_{t+j} \equiv \begin{bmatrix} u_{t+j} & r^*_{t+j} \end{bmatrix}'$ are the exogenous state variables. To simplify notation in what follows, the dependence of the policy functions on $z$ is suppressed and they are presented as dependent only on $q$.

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

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

The problem can therefore be expressed as a Lagrangean:

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

The first order conditions are:

$$0 = \omega_\pi \hat{\pi}_t + \lambda^\pi_t \tag{E.2}$$

$$0 = \omega_x \hat{x}_t - \kappa \lambda^\pi_t + \lambda^x_t \tag{E.3}$$

$$\begin{aligned}
0 = & \omega_q q_t + \omega_{\Delta q} \left( q_t - q_{t-1} \right) + \beta \frac{\partial \mathbb{E}_t \widetilde{\mathcal{L}}_{t+1}}{\partial q_t} - \beta \frac{\partial \mathbb{E}_t \mathbf{\Pi} \left( q_t \right)}{\partial q_t} \lambda^\pi_t \\
& - \left[ \frac{\partial \mathbb{E}_t \boldsymbol{X} \left( q_t \right)}{\partial q_t} + \sigma \frac{\partial \mathbb{E}_t \mathbf{\Pi} \left( q_t \right)}{\partial q_t} + \sigma \gamma - \beta \sigma \xi \frac{\partial \mathbb{E}_t \boldsymbol{Q} \left( q_t \right)}{\partial q_t} \right] \lambda^x_t - \lambda^{\bar{q}}_t - \lambda^{q}_t
\end{aligned} \tag{E.4}$$

$$0 = \sigma \lambda^x_t - \lambda^R_t \tag{E.5}$$

with contemporary slackness conditions:

$$\begin{aligned}
0 &= \lambda^R_t \left( \hat{R}_t - \ln \beta \right) \\
0 &= \lambda^{\bar{q}}_t \left( q_t - \bar{q} \right) \\
0 &= \lambda^{q}_t \left( q_t - \underline{q} \right)
\end{aligned}$$

The first order condition for quantitative easing, (E.4), indicates that the policymaker accounts for the effects of current QE decisions on the losses incurred by future policymakers.

The envelope condition implies that:

$$\frac{\partial \widetilde{\mathcal{L}}_t}{\partial q_{t-1}} = -\omega_{\Delta q} \left( q_t - q_{t-1} \right) + \sigma \xi \lambda^x_t$$

and hence that:

$$\frac{\partial \mathbb{E}_t \widetilde{\mathcal{L}}_{t+1}}{\partial q_t} = -\omega_{\Delta q} \mathbb{E}_t \left( q_{t+1} - q_t \right) + \sigma \xi \mathbb{E}_t \lambda^x_{t+1}$$

This means that (E.4) can be written as:

$$\begin{aligned} 0 = \; & \omega_q q_t + \omega_{\Delta q}\left(q_t - q_{t-1}\right) - \beta\omega_{\Delta q}\mathbb{E}_t\left(q_{t+1} - q_t\right) + \beta\sigma\xi\mathbb{E}_t\lambda^x_{t+1} - \beta\boldsymbol{D}^{\Pi}_t\lambda^{\pi}_t \\ & - \left[\boldsymbol{D}^X_t + \sigma\boldsymbol{D}^{\Pi}_t + \sigma\gamma - \beta\sigma\xi\boldsymbol{D}^Q_t\right]\lambda^x_t - \lambda^{\bar{q}}_t - \lambda^{q}_t \end{aligned} \tag{E.6}$$

where partial derivatives of expectations are expressed using the following notation:

$$\boldsymbol{D}^X_t \equiv \frac{\partial\mathbb{E}_t\boldsymbol{X}\left(q_t\right)}{\partial q_t} \tag{E.7}$$

The equilibrium can be written in terms of 'effective QE', $\widetilde{q}$, defined as

$$\widetilde{q}_t \equiv \gamma q_t - \xi q_{t-1} - \beta\xi\mathbb{E}_t q_{t+1}$$

The equilibrium is then given by:

$$\hat{x}_t = \mathbb{E}_t\hat{x}_{t+1} - \sigma\left[\hat{R}_t - \mathbb{E}_t\hat{\pi}_{t+1} - \tilde{q}_t - r^*_t\right] \tag{E.8}$$

$$\hat{\pi}_t = \beta\mathbb{E}_t\hat{\pi}_{t+1} + \kappa\hat{x}_t + u_t \tag{E.9}$$

$$0 = \omega_x\hat{x}_t + \kappa\omega_\pi\hat{\pi}_t + \lambda^x_t \tag{E.10}$$

$$0 = \Theta\widetilde{q}_t + \beta\sigma\xi\mathbb{E}_t\lambda^x_{t+1} + \beta\boldsymbol{D}^{\Pi}_t\omega_\pi\hat{\pi}_t - \left[\boldsymbol{D}^X_t + \sigma\boldsymbol{D}^{\Pi}_t + \sigma\gamma - \beta\sigma\xi\boldsymbol{D}^Q_t\right]\lambda^x_t - \lambda^{\bar{q}}_t - \lambda^{q}_t \tag{E.11}$$

where $\Theta \equiv (b + d)$.

The solutions for a number of cases corresponding to whether or not the constraints on the instruments are binding are considered in turn. Expectations are taken as given (i.e., known). 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

First assume that the ZLB does not bind in the current period, so that $\lambda^x_t = 0$. This means that the solution can be computed as follows:

1. Use the targeting criterion $\hat{x}_t = -\frac{\omega_\pi\kappa}{\omega_x}\hat{\pi}_t$ to eliminate the output gap from the Phillips curve and solve for inflation:

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

2. Use the targeting criterion to solve for the output gap: $\hat{x}_t = -\frac{\omega_\pi\kappa}{\omega_x}\hat{\pi}_t$.

3. Set $\tilde{q}_t = -\Theta^{-1}\beta\sigma\xi\mathbb{E}_t\lambda^x_{t+1} - \Theta^{-1}\beta\boldsymbol{D}^{\Pi}_t\omega_\pi\hat{\pi}_t$.

4. Compute the implied solution for $q_t = \gamma^{-1}\left(\tilde{q}_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1}\right)$ and check whether it violates the upper bound. If the upper bound is violated, set $q_t = \bar{q}$ and $\tilde{q}_t = \gamma\bar{q} - \xi q_{t-1} - \beta\xi\mathbb{E}_t q_{t+1}$.

5. Use the IS surve to compute the required level of the policy rate:

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

6. Check the the solution for the policy rate is above the lower bound.

## E.2 ZLB binds

If the procedure above delivers a solution for the policy rate that violates the ZLB, we move on to solve the system with the ZLB applied. We first assume that there is an interior solution for $q$. In this case, the equilibrium is given by

$$\begin{aligned}
\hat{R}_t &= \ln \beta \\
\hat{x}_t &= \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[ \hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \tilde{q}_t - r_t^* \right] \\
\hat{\pi}_t &= \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t \\
\omega_x \hat{x}_t &= -\kappa \omega_\pi \hat{\pi}_t - \lambda_t^x \\
0 &= \Theta \tilde{q}_t + \beta \sigma \xi \mathbb{E}_t \lambda_{t+1}^x + \beta \boldsymbol{D}_t^\Pi \omega_\pi \hat{\pi}_t - \left[ \boldsymbol{D}_t^X + \sigma \boldsymbol{D}_t^\Pi + \sigma \gamma - \beta \sigma \xi \boldsymbol{D}_t^Q \right] \lambda_t^x
\end{aligned}$$

The equilibrium allocations therefore satisfy:

$$\begin{bmatrix} 0 & 1 & -\sigma & 0 \\ 1 & -\kappa & 0 & 0 \\ \kappa \omega_\pi & \omega_x & 0 & 1 \\ -\beta \boldsymbol{D}_t^\Pi \omega_\pi & 0 & -\Theta & \boldsymbol{D}_t^X + \sigma \boldsymbol{D}_t^\Pi + \sigma \gamma - \beta \sigma \xi \boldsymbol{D}_t^Q \end{bmatrix} \begin{bmatrix} \hat{\pi}_t \\ \hat{x}_t \\ \tilde{q}_t \\ \lambda_t^x \end{bmatrix} = \begin{bmatrix} \mathbb{E}_t \hat{x}_{t+1} - \sigma \ln \beta + \sigma \mathbb{E}_t \hat{\pi}_{t+1} + \sigma r_t^* \\ \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \\ 0 \\ \beta \sigma \xi \mathbb{E}_t \lambda_{t+1}^x \end{bmatrix}$$

which can be solved by matrix inversion.

The process of computing the equilibrium when the ZLB binds is therefore as follows:

1. Compute the equilibrium allocations for inflation, the output gap, effective QE and the multiplier on the IS curve by matrix inversion.
2. Compute the implied solution for $q_t = \gamma^{-1} (\tilde{q}_t + \xi q_{t-1} + \beta \xi \mathbb{E}_t q_{t+1})$ and check whether it violates the upper bound.

## E.3 QE instrument is constrained

If the check in step 2 reveals that $q$ violates the upper bound, then both instruments are constrained and the equilibrium is computed as follows:

1. The policy instruments satisfy $\hat{R}_t = \ln \beta$ and $q_t = \bar{q}$.
2. Effective QE is $\tilde{q}_t = \gamma \bar{q} - \xi q_{t-1} - \beta \xi \mathbb{E}_t q_{t+1}$.
3. The output gap is given by $\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left[ \hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \tilde{q}_t - r_t^* \right]$.
4. Inflation is given by $\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t$.
5. The multiplier on the IS curve is given by: $\lambda_t^x = -\omega_x \hat{x}_t - \kappa \omega_\pi \hat{\pi}_t$.

# F Solution algorithm

## F.1 State space and policy functions: notation

The description of the algorithm can be simplified by introducing some notation for the key objects that will be solved for.

The vector of endogenous variables are denoted by $z$, defined implicitly above, but explicitly here:

$$z \equiv \begin{bmatrix} \hat{\pi} \\ \hat{x} \\ \hat{R} \\ \hat{V} \\ q \\ \tilde{q} \\ \lambda^x \end{bmatrix}$$

The exogenous states are denoted $s$:

$$s \equiv \begin{bmatrix} r^* \\ u \end{bmatrix}$$

and full state vector for relevant policy functions is, $\tilde{s}$:

$$\tilde{s} \equiv \begin{bmatrix} s \\ q \end{bmatrix}$$

The exogenous state is defined as a set of fixed values for the cost push shock and natural rate. Specifically, $S_r \equiv \left\{r_1^* \quad \ldots \quad r_{n_r}^*\right\}$ and $S_u \equiv \{u_1 \quad \ldots \quad u_{n_u}\}$. The transition matrices for the Markov processes are $\Omega_r$ and $\Omega_u$, computed using the Rouwenhorst (1995) approach.

The combined (exogenous) state-space is given by $S = S_u \times S_r$ with transition matrix $\Omega = \Omega_r \otimes \Omega_u$.[44] The endogenous state is $q$, which is discretized on a grid $S_q \equiv \left\{q_1 \quad \ldots \quad q_{n_q}\right\}$, with $q_i > q_{i-1}$, $i = 2, \ldots, n_q$. The endogenous state is assumed to be ordered last. So the full state space is given by $\tilde{S} = S \times S_q$.[45] Thus, $\tilde{S}$ is a $n_{\tilde{s}} \times 3$ matrix, where $n_{\tilde{s}} \equiv n_s \times n_q$. The index of the element $\left\{u_i, r_j^*, q_k\right\} \in \tilde{S}$ is $(k-1) \times n_s + (j-1) \times n_u + i$.

This implies that the combined state can be written as:

$$\tilde{S} = \begin{bmatrix} S & q_1 \mathbf{1}_{n_s} \\ \vdots & \vdots \\ S & q_{n_d} \mathbf{1}_{n_s} \end{bmatrix}$$

where $\mathbf{1}_{n_s}$ is a $n_s \times 1$ unit vector.

This representation of the state space is useful for subsequent computations since approximation of expectations requires interpolation between grid points for the endogenous state, while integrating across the exogenous state $S$. Similar methods are used for the estimation of derivatives of expectations.

The objects of interest are policy functions. These are $n_{\tilde{s}} \times n_z$ matrices. Let a generic policy

[44] The first $n_u$ elements of the state space are $\{u_1, r_1^*\}, \ldots, \{u_{n_u}, r_1^*\}$, followed by $\{u_1, r_2^*\}, \ldots, \{u_{n_u}, r_2^*\}$ and so on.

[45] Thus the first $n_s \equiv n_u \times n_r$ elements are given by the triples $\{u_1, r_1^*, q_1\}, \ldots, \left\{u_{n_u}, r_{n_r}^*, q_1\right\}$, the next $n_s$ elements are $\{u_1, r_1^*, q_2\}, \ldots, \left\{u_{n_u}, r_{n_r}^*, q_2\right\}$ and so on.

function be denoted $\mathbf{f}$.

## F.2 Expectations

It is useful to define an 'expectation' operator that integrates out exogenous state uncertainty but holds the endogenous state vector constant:

$$\mathbb{E}_S \mathbf{f} \equiv \overline{\mathbf{f}}^{\mathbf{S}} \equiv \begin{bmatrix} \Omega & 0 & \dots & 0 & 0 \\ 0 & \Omega & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & \Omega & 0 \\ 0 & 0 & \dots & 0 & \Omega \end{bmatrix} \mathbf{f}$$

so that the 'bar' is used as a summary notation for expectations and the $S$ superscript indicates that the expectation is computed with respect to the exogenous state variables only.

To compute the actual expectation requires conditioning on the solution for $q'$ at the particular point in the state space. This can be done as follows.

- Extract the relevant column of $\mathbf{f}$ that corresponds to $q$. Let this column vector be denoted $\mathbf{q}$. The elements of this vector denote the solutions $q' \in z$ for each state $1, \dots, n_{\tilde{s}}$.
- Let the elements of $\mathbf{q}$ be denoted $\mathbf{q}_k, k = 1, \dots, n_{\tilde{s}}$. Let the exogenous state corresponding to this solution be $S_{<k>}$. For each $k$, perform the following:
  - Compute the indices in $S_q$ that bracket this element.[46] This gives two indices $i_1, i_2 \in S_q$ with $1 \le i_1 < i_2 \, (= i_1 + 1) \le n_q$ .
  - Compute the weights that should apply to each of these gridpoints (using linear interpolation). This gives $\phi_2 = \frac{\mathbf{q}_k - S_q(i_1)}{S_q(i_2) - S_q(i_1)}$ and $\phi_1 = 1 - \phi_2$.
  - Compute the indices of the elements of $\tilde{S}$ corresponding to the elements in $\tilde{S}$ for which (a) $S = S_{<k>}$ and (b) $S_q = S_q(i_1)$ and $S_q = S_q(i_2)$. Denote these indices as $\tilde{i}_1$ and $\tilde{i}_2$.
  - Estimate the expectation using linear interpolation as:

    $$\overline{\mathbf{f}}_{\mathbf{k},\cdot} = \phi_1 \overline{\mathbf{f}}^{\mathbf{S}}_{\tilde{\mathbf{i}}_1,\cdot} + \phi_2 \overline{\mathbf{f}}^{\mathbf{S}}_{\tilde{\mathbf{i}}_2,\cdot}$$

    where the subscript '$j, \cdot$' denotes the $j$-th row of a matrix.

The penultimate step (finding $\tilde{i}_1$ and $\tilde{i}_2$) can be aided by a pre-computation operation. To see this, recall that for each $k \in \{1, \dots, n_{\tilde{s}}\}$, there is an exogenous state, $S_{<k>}$. The indices corresponding to different values of $q$ for the same value of $S_{<k>}$ are multiples of $n_s$ away from from $k$. This allows us to form a $n_{\tilde{s}} \times n_q$ matrix of indices – a 'lookup matrix', denoted $\Lambda$ – as follows.

For each $k \in \{1, \dots, n_{\tilde{s}}\}$

- Compute $j$, defined as the index of the grid point $q'_k$ within $S_q$. Recall that $q$ is the final (third) state.

[46] Extrapolation is conceptually identical, but for the purposes of exposition, I assume that interpolation is required.

- For $m = 1, \ldots, n_s$, form the $k$-th row of $\Lambda$ as:

$$\Lambda_{k,m} = k - (j - m)\, n_s$$

Then, in the computation of expectations, for each $k$ the indices are found by setting $\tilde{i}_1 = \Lambda_{k,i_1}$ and $\tilde{i}_2 = \Lambda_{k,i_2}$.

## F.3 Derivatives

The first order conditions depend on derivatives of expectations of the policy functions. To approximate these derivatives, a two-sided finite difference approach is used. The derivatives are computed in two steps. In the first step, two-sided finite difference derivatives of the static expectations are computed, using adjacent gridpoints for $q$. In the second step, linear interpolation is used to approximate the derivatives at the relevant values of $\mathbf{q}$.

The first step is to approximate the derivative of the static expectation function $\bar{\mathbf{f}}^{\mathbf{S}}$. We seek the finite difference approximation to the derivative of $\bar{\mathbf{f}}^{\mathbf{S}}$ for each row $m = 1, \ldots, n_{\tilde{s}}$. First note that $S_q$ is assumed to be formed of an evenly-spaced grid of values: $S_q = \begin{Bmatrix} q_1 & \ldots & q_{n_q} \end{Bmatrix}$, with $q_{i+1} = q_i + h_q$. So the difference between each grid point is $h_q$.

Consider an $m$ for which the corresponding element of $S_q$ is $q_i$ with $1 < i < n_q$, that is, an interior point. Then the 'static derivative' at point $m$ is given by:

$$\mathbf{D}^{\mathbf{S}}_{\mathbf{m},\cdot} = \frac{1}{2h_q}\left(\bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m}+\mathbf{n_s},\cdot} - \bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m}-\mathbf{n_s},\cdot}\right)$$

Now consider the endpoints. For $1 \leq m \leq n_s$, $i = 1$ and a one-sided difference is used:

$$\mathbf{D}^{\mathbf{S}}_{\mathbf{m},\cdot} = \frac{1}{h_q}\left(\bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m}+\mathbf{n_s},\cdot} - \bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m},\cdot}\right)$$

Similarly, for $n_{\tilde{s}} - n_s + 1 \leq m \leq n_{\tilde{s}}$, $i = n_q$ and the one-sided approximation is given by:

$$\mathbf{D}^{\mathbf{S}}_{\mathbf{m},\cdot} = \frac{1}{h_q}\left(\bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m},\cdot} - \bar{\mathbf{f}}^{\mathbf{S}}_{\mathbf{m}-\mathbf{n_s},\cdot}\right)$$

The second step is to form an estimate of the derivative at $\mathbf{d}$ using linear interpolation. This step is set out in the description of the algorithm below.

## F.4 Algorithm

The objective of the algorithm is to solve for the policy function $\mathbf{f}$ by iterating directly on it.

1. Initialize a guess, $\mathbf{f}^{<\mathbf{0}>}$, for the policy function $\mathbf{f}$.
2. Build the 'lookup matrix', $\Lambda$, as described above.
3. For each iteration $j = 1, \ldots$

   **Update expectations**

   (a) Update the guess for 'static' expectations. As described above, this integrates out exogenous state uncertainty but holds the endogenous state vector constant:

$$\mathbf{\bar{f}^{S<j>}} = \begin{bmatrix} \Omega & 0 & \dots & 0 & 0 \\ 0 & \Omega & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & \Omega & 0 \\ 0 & 0 & \dots & 0 & \Omega \end{bmatrix} \mathbf{f^{<j-1>}}$$

(b) Extract the vector of $q'$ values, $\mathbf{q^{<j-1>}}$ as the relevant column of $\mathbf{f^{<j-1>}}$.

(c) Compute the indices and weights of the elements of $S_q$ that bracket the values in $\mathbf{q^{<j-1>}}$. Use the lookup matrix $\Lambda$ to convert these into $n_{\tilde{s}} \times 2$ matrices of indicators and interpolation/extrapolation weights, denoted $\Upsilon$ and $\Phi$ respectively.

(d) For each $m = 1, \dots, n_{\tilde{s}}$: Compute expectations by extracting interpolation indices $[\iota_1 \;\; \iota_2] = \Upsilon_m$ and weights $[\phi_1 \;\; \phi_2] = \Phi_m$. Translate the interpolation weights into $\tilde{S}$ space by setting $\tilde{\iota}_1 = \Lambda_{m,\iota_1}$ and $\tilde{\iota}_2 = \Lambda_{m,\iota_2}$. Now set

$$\mathbf{\bar{f}_{m,\cdot}} = \phi_1 \mathbf{\bar{f}^{S<j>}_{\tilde{\iota}_1,\cdot}} + \phi_2 \mathbf{\bar{f}^{S<j>}_{\tilde{\iota}_2,\cdot}}$$

**Update the estimate of the derivative of expectations**

(e) Update the estimate of the 'static' derivatives, $\mathbf{D^S}$, as described in F.3.

(f) Compute the derivatives prevailing at $\mathbf{q}$ by linear interpolation. For each $m = 1, \dots, n_{\tilde{s}}$, set:

$$\mathbf{D_{m,\cdot}} = \phi_1 \mathbf{D^S_{\tilde{\iota}_1,\cdot}} + \phi_2 \mathbf{D^S_{\tilde{\iota}_2,\cdot}}$$

where the indices $\tilde{\iota}_1, \tilde{\iota}_2$ and weights $\phi_1, \phi_2$ are the same as in step 3d.

Update the guess for the policy function

(g) For each $m = 1, \dots, n_{\tilde{s}}$:

i. Extract latest guesses for expectations and their derivatives:

$$\mathbb{E}z' = \mathbf{\bar{f}_{m,\cdot}} \qquad \mathcal{D} = \mathbf{D_{m,\cdot}}$$

ii. Solve for outcomes conditional on the derivatives and expectations as described in Appendix E.

$$\mathbf{f^{<j>}_{m,\cdot}} = z$$

4. Check for convergence. If $\left|\mathbf{f^{<j>}} - \mathbf{f^{<j-1>}}\right| < \varepsilon$, set $\mathbf{f} = \mathbf{f^{<j>}}$ and stop, otherwise set $j = j + 1$ and return to step 3.

## F.5 Practical implementation

The Rouwenhorst (1995) Markov chain approximation of autoregressive processes performs well even for a relatively small number of nodes (Kopecky and Suen, 2010). Given the persistence properties of the neutral rate and cost-push shock processes, the number of approximation nodes are set to $n_u = 15$ and $n_r = 25$ respectively.[47] For the endogenous state variable, $q$, $n_q = 100$ nodes are used.

[47]The simulation responses shown in Figures 5, 6 and 8 use a solution computed with $n_r = 41$ nodes to give a smoother interpolation of the policy functions. However, the choice of $n_r$ does not affect the results pertaining to the baseline model (e.g., welfare comparisons), so to lighten the computational burden $n_r = 25$ is the assumption for all other cases, including robustness exercises.

A good initial guess for the policy functions is constructed by initializing the algorithm using policy functions computed using the piecewise linear algorithms in Harrison and Waldron (2021).

# G Flexible inflation targeting

Suppose that the policymaker minimizes a 'flexible inflation targeting' loss function:

$$\mathcal{L}_0^{FIT} = \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\omega_x \hat{x}_t^2 + \omega_\pi \hat{\pi}_t^2\right) \tag{G.1}$$

where $\omega_x = \Xi$ and $\omega_\pi = \eta/\Gamma$.

The Lagrangean is:

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

The policy problem specifies upper and lower bounds on $q$ (as in the baseline case) together with a bound on the shadow rate ($\widetilde{R}$) which embeds the lower bound on the policy rate and the upper bound on the balance sheet, but is also a function of lagged and expected $q$. This formulation helps to highlight the perfect substitutability of the policy rate and balance sheet in setting the shadow rate.

The first order conditions are:

$$\begin{aligned}
0 = & \ \omega_\pi \hat{\pi}_t + \lambda_t^\pi \\
0 = & \ \omega_x \hat{x}_t - \kappa\lambda_t^\pi + \lambda_t^x \\
0 = & \ \beta \frac{\partial \mathbb{E}_t \widetilde{\mathcal{L}}_{t+1}^{FIT}}{\partial q_t} - \beta \frac{\partial \mathbb{E}_t \mathbf{\Pi}\left(q_t\right)}{\partial q_t}\lambda_t^\pi - \left[\frac{\partial \mathbb{E}_t \boldsymbol{X}\left(q_t\right)}{\partial q_t} + \sigma \frac{\partial \mathbb{E}_t \mathbf{\Pi}\left(q_t\right)}{\partial q_t}\right]\lambda_t^x \\
& + \beta\xi \frac{\partial \mathbb{E}_t \boldsymbol{Q}\left(q_t\right)}{\partial q_t}\lambda_t^R - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}} \\
0 = & \ \sigma\lambda_t^x - \lambda_t^R
\end{aligned} \tag{G.2}$$

with contemporary slackness conditions:

$$\begin{aligned}
0 &= \lambda_t^R \left(\widetilde{R}_t - \ln\beta + \gamma\bar{q} - \xi q_{t-1} - \beta\xi \mathbb{E}_t \boldsymbol{Q}\left(q_t\right)\right) \\
0 &= \lambda_t^{\bar{q}}\left(q_t - \bar{q}\right) \\
0 &= \lambda_t^{\underline{q}}\left(q_t - \underline{q}\right)
\end{aligned}$$

The envelope condition implies that:

$$\frac{\partial \widetilde{\mathcal{L}}_t^{FIT}}{\partial q_{t-1}} = \xi\lambda_t^R$$

and hence that:

$$\frac{\partial \mathbb{E}_t \widetilde{\mathcal{L}}_{t+1}^{FIT}}{\partial q_t} = \xi \mathbb{E}_t \lambda_{t+1}^R$$

So we can write the first order condition for $q$ as

$$0 = \beta\sigma\xi \mathbb{E}_t \lambda_{t+1}^x + \beta \boldsymbol{D}_t^\Pi \omega_\pi \hat{\pi}_t - \left[ \boldsymbol{D}_t^X + \sigma \boldsymbol{D}_t^\Pi - \sigma\beta\xi \boldsymbol{D}_t^Q \right] \lambda_t^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}}$$

using the same notation for partial derivatives of expectations introduced earlier.

If the constraint on total policy space is not binding, so that $\lambda_t^x = 0$, then the FOCs are:

$$\begin{aligned} 0 =& \ \omega_x \hat{x}_t + \kappa \omega_\pi \hat{\pi}_t \\ 0 =& \ \beta\sigma\xi \mathbb{E}_t \lambda_{t+1}^x + \beta \boldsymbol{D}_t^\Pi \omega_\pi \hat{\pi}_t - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}} \end{aligned} \tag{G.3}$$

The first equation, together with the Phillips curve will determine $x$ and $\pi$ (conditional on expectations). The consistency of the assumption of an interior solution (i.e., that it is feasible to implement the equilibrium with $\widetilde{R}_t$ above the ZLB) can be checked by inverting the Euler equation.

Suppose that is indeed the case, so that the required value of the effective policy rate is such that $\widetilde{R}_t > \ln\beta - \gamma\bar{q} + \xi q_{t-1} - \beta\xi \mathbb{E}_t q_{t+1}$. Then the second equation pins down the policy mix (i.e., determines $q_t$, from which $\hat{R}_t$ can be computed), since $\boldsymbol{D}_t^\Pi$ is an implicit function of $q$.

Importantly, the condition that determines the policy mix, (G.3), depends only on the balance sheet only through its influence on the forward-looking term $\boldsymbol{D}_t^\Pi$. So it does not depend on the existing stock of assets held by the central bank, $q_{t-1}$, and therefore has no direct implications for the 'sequencing' of instruments used to deliver a monetary tightening following a period at the ZLB with $q_{t-1} > 0$.

# H Commitment

The commitment problem can be expressed as a Lagrangean:

$$\min_{\left\{\hat{\pi}_t, \hat{x}_t, \hat{R}_t, q_t\right\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[ \begin{array}{c} \frac{\omega_x}{2}\hat{x}_t^2 + \frac{\omega_\pi}{2}\hat{\pi}_t^2 + \frac{\omega_q}{2}q_t^2 + \frac{\omega_{\Delta q}}{2}\left(q_t - q_{t-1}\right)^2 + \lambda_t^\pi \left(\hat{\pi}_t - \kappa\hat{x}_t - \beta\mathbb{E}_t\hat{\pi}_{t+1} - u_t\right) \\ +\lambda_t^x \left(\hat{x}_t - \mathbb{E}_t\hat{x}_{t+1} + \sigma\left(\hat{R}_t - \mathbb{E}_t\hat{\pi}_{t+1} - \gamma q_t + \xi q_{t-1} + \beta\xi\mathbb{E}_t q_{t+1} - r_t^*\right)\right) \\ -\lambda_t^R \left(\hat{R}_t - \beta^{-1} + 1\right) - \lambda_t^{\bar{q}}\left(q_t - \bar{q}\right) - \lambda_t^{\underline{q}}\left(q_t - \underline{q}\right) \end{array} \right]$$

with first order conditions

$$\begin{aligned} 0 &= \omega_x \hat{x}_t - \kappa\lambda_t^\pi + \lambda_t^x - \beta^{-1}\lambda_{t-1}^x \\ 0 &= \omega_\pi \hat{\pi}_t + \lambda_t^\pi - \beta^{-1}\left(\beta\lambda_{t-1}^\pi + \sigma\lambda_{t-1}^x\right) \\ 0 &= \sigma\lambda_t^x - \lambda_t^R \\ 0 &= \omega_q q_t + \omega_{\Delta q}\left(q_t - q_{t-1}\right) - \beta\omega_{\Delta q}\mathbb{E}_t\left(q_{t+1} - q_t\right) - \sigma\gamma\lambda_t^x + \beta\sigma\xi\mathbb{E}_t\lambda_{t+1}^x + \beta^{-1}\sigma\beta\xi\lambda_{t-1}^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}} \end{aligned}$$

Note that the FOC for $q$ can be written in terms of $\widetilde{q}$:

$$0 = \Theta\tilde{q}_t - \sigma\gamma\lambda_t^x + \beta\sigma\xi\mathbb{E}_t\lambda_{t+1}^x + \beta^{-1}\sigma\beta\xi\lambda_{t-1}^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}}$$

## H.1 Solution

This subsection details how to compute equilibrium allocations taking expectations as given, since the solution algorithm detailed below will iterate between a step that solves such conditional outcomes and a step that computes expectations conditional on the implied decision rules.

### H.1.1 Equilibrium when the ZLB is not binding

In this case, we have $\lambda_t^x = \lambda_t^R = 0$ and the first order conditions can be written as:

$$
\begin{aligned}
0 &= \omega_x \hat{x}_t - \kappa\lambda_t^\pi - \beta^{-1}\lambda_{t-1}^x \\
0 &= \omega_\pi \hat{\pi}_t + \lambda_t^\pi - \beta^{-1}\left(\beta\lambda_{t-1}^\pi + \sigma\lambda_{t-1}^x\right) \\
0 &= \omega_q q_t + \omega_{\Delta q}\left(q_t - q_{t-1}\right) - \beta\omega_{\Delta q}\mathbb{E}_t\left(q_{t+1} - q_t\right) + \beta\sigma\xi\mathbb{E}_t\lambda_{t+1}^x + \beta^{-1}\sigma\beta\xi\lambda_{t-1}^x - \lambda_t^{\bar{q}} - \lambda_t^{\underline{q}}
\end{aligned}
$$

The following steps will solve for equilibrium allocations conditional on expectations:

1. Set $\lambda_t^x = 0$.
2. Combine first order conditions for output and inflation to give:

$$0 = \omega_x \hat{x}_t - \kappa\left(\lambda_{t-1}^\pi + \beta^{-1}\sigma\lambda_{t-1}^x - \omega_\pi\hat{\pi}_t\right) - \beta^{-1}\lambda_{t-1}^x$$

   which implies that

$$\hat{x}_t = \omega_x^{-1}\kappa\left(\lambda_{t-1}^\pi + \beta^{-1}\sigma\lambda_{t-1}^x - \omega_\pi\hat{\pi}_t\right) + \left(\beta\omega_x\right)^{-1}\lambda_{t-1}^x \tag{H.1}$$

3. Plug this into the Phillips curve to give

$$
\begin{aligned}
\hat{\pi}_t &= \kappa\hat{x}_t + \beta\mathbb{E}_t\hat{\pi}_{t+1} + u_t \\
&= \kappa\left[\omega_x^{-1}\kappa\left(\lambda_{t-1}^\pi + \beta^{-1}\sigma\lambda_{t-1}^x - \omega_\pi\hat{\pi}_t\right) + \left(\beta\omega_x\right)^{-1}\lambda_{t-1}^x\right] + \beta\mathbb{E}_t\hat{\pi}_{t+1} + u_t \\
&= \frac{1}{1+\kappa^2\omega_\pi\omega_x^{-1}}\left(\kappa\left[\omega_x^{-1}\kappa\left(\lambda_{t-1}^\pi + \beta^{-1}\sigma\lambda_{t-1}^x\right) + \left(\beta\omega_x\right)^{-1}\lambda_{t-1}^x\right] + \beta\mathbb{E}_t\hat{\pi}_{t+1} + u_t\right)
\end{aligned} \tag{H.2}
$$

   where all the terms on the right hand side are 'known' for the conditional solution.

4. Plug the solution for $\hat{\pi}$ from (H.6) into (H.5) to solve for the output gap.
5. Compute the optimal level of QE assuming that the solution is interior:

$$
\begin{aligned}
q_t = &\frac{\omega_{\Delta q}}{\omega_q + (1+\beta)\,\omega_{\Delta q}} q_{t-1} + \frac{\beta\omega_{\Delta q}}{\omega_q + (1+\beta)\,\omega_{\Delta q}}\mathbb{E}_t q_{t+1} \\
&- \frac{\beta\sigma\xi}{\omega_q + (1+\beta)\,\omega_{\Delta q}}\mathbb{E}_t\lambda_{t+1}^x - \frac{\sigma\xi}{\omega_q + (1+\beta)\,\omega_{\Delta q}}\lambda_{t-1}^x
\end{aligned} \tag{H.3}
$$

6. Check whether the solution to (H.3) is interior, if not, set $q_t = \underline{q}$ or $q_t = \bar{q}$ as appropriate.
7. Compute effective QE

$$\widetilde{q}_t = \gamma q_{t-1} - \xi q_{t-1} - \beta\xi\mathbb{E}_t q_{t+1}$$

8. Compute the interest rate implements this equilibrium:

$$\hat{R}_t = \sigma^{-1}\left(\mathbb{E}_t x_{t+1} - x_t\right) + \mathbb{E}_t \hat{\pi}_{t+1} + r_t^* + \widetilde{q}_t \tag{H.4}$$

If the solution from this procedure respects the ZLB and the bounds on QE, it is the (conditional) solution. If not, we follow the steps in the next subsection.

### H.1.2 ZLB binds

In this case, we can form a system to be solved for $x$, $\pi$, $\lambda^x$, $\lambda^\pi$, and $\widetilde{q}$ by using the first order conditions for $x$, $\pi$ and $q$ along with the Phillips curve and IS curve (setting $r_t = \ln\beta$ in the latter). This gives:

$$\begin{bmatrix} 0 & \omega_x & 0 & 1 & -\kappa \\ \omega_\pi & 0 & 0 & 0 & 1 \\ 0 & 0 & \Theta & -\sigma\gamma & 0 \\ 0 & 1 & -\sigma & 0 & 0 \\ 1 & -\kappa & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{\pi}_t \\ \hat{x}_t \\ \tilde{q}_t \\ \lambda_t^x \\ \lambda_t^\pi \end{bmatrix} = \begin{bmatrix} \beta^{-1}\lambda_{t-1}^x \\ \lambda_{t-1}^\pi + \beta^{-1}\sigma\lambda_{t-1}^x \\ -\beta\sigma\xi\mathbb{E}_t\lambda_{t+1}^x - \sigma\xi\lambda_{t-1}^x \\ -\sigma\left(b - \mathbb{E}_t\pi_{t+1} - r_t^*\right) \\ \mathbb{E}_t\pi_{t+1} + u_t \end{bmatrix}$$

which can be solved by matrix inversion.

### H.1.3 Computing conditional expectations

The exogenous states are denoted $s$:

$$s \equiv \begin{bmatrix} u \\ r^* \end{bmatrix}$$

and full state vector for relevant policy functions is, $\tilde{s}$:

$$\tilde{s} \equiv \begin{bmatrix} s \\ q \\ \lambda^x \\ \lambda^\pi \end{bmatrix}$$

where the endogenous state variables $q, \lambda^\pi, \lambda^x$ are the values of these variables in the previous period.

The exogenous state is defined as a set of fixed values for the cost push shock and natural rate. Specifically, $S_r \equiv \left\{r_1^* \quad \dots \quad r_{n_r}^*\right\}$ and $S_u \equiv \{u_1 \quad \dots \quad u_{n_u}\}$. The transition matrices for the Markov processes are $\Omega_r$ and $\Omega_u$. The combined (exogenous) state-space is given by $S = S_u \times S_r$ with transition matrix $\Omega = \Omega_r \otimes \Omega_u$.[48]

The endogenous states are discretized on a grid $S_\lambda = S_q \times S_{\lambda^\pi} \times S_{\lambda^x}$, where $S_q \equiv \left\{q_1 \quad \dots \quad q_{n_q}\right\}$, $S_{\lambda^x} \equiv \left\{\lambda_1^x \quad \dots \quad \lambda_{n_x}^x\right\}$ and $S_{\lambda^\pi} \equiv \left\{\lambda_1^\pi \quad \dots \quad \lambda_{n_\pi}^\pi\right\}$, with $q_i > q_{i-1}$, $i = 2, \dots, n_q$, $\lambda_i^x > \lambda_{i-1}^x$, $i = 2, \dots, n_x$ and $\lambda_i^\pi > \lambda_{i-1}^\pi$, $i = 2, \dots, n_\pi$.

The endogenous states are assumed to be ordered last. So the full state space is given by $\tilde{S} = S \times S_\lambda$. Thus, $\tilde{S}$ is a $n_{\tilde{s}} \times 5$ matrix, where $n_{\tilde{s}} \equiv n_s \times n_\lambda$.[49] The index of the element $\left\{u_i, r_j^*, q_k, \lambda_l^x, \lambda_m^\pi\right\} \in \tilde{S}$ is $(m-1)\times(n_s \times n_q \times n_x) + (l-1)\times(n_s \times n_q) + (k-1)\times n_s + (j-1)\times n_u + i$.

[48]The first $n_u$ elements of the state space are $\{u_1, r_1^*\}, \dots, \{u_{n_u}, r_1^*\}$, followed by $\{u_1, r_2^*\}, \dots, \{u_{n_u}, r_2^*\}$ and so on.

[49]Thus the first $n_s \equiv n_u \times n_r$ elements are given by $\{u_1, r_1^*, q_1, \lambda_1^\pi, \lambda_1^x\}, \dots, \left\{u_{n_u}, r_{n_r}^*, q_1, \lambda_1^\pi, \lambda_1^x\right\}$, the next $n_s$ elements are $\{u_1, r_1^*, q_2, \lambda_1^\pi, \lambda_1^x\}, \dots, \left\{u_{n_u}, r_{n_r}^*, q_2, \lambda_1^\pi, \lambda_1^x\right\}$ and so on.

## H.2 Solution when balance sheet policy is not used

When the central bank does not use the balance sheet instrument ($q$) the model collapses to the standard New Keynesian model and the FOCs are the usual ones:

$$
\begin{aligned}
0 &= \omega_x \hat{x}_t - \kappa \lambda_t^\pi + \lambda_t^x - \beta^{-1} \lambda_{t-1}^x \\
0 &= \omega_\pi \hat{\pi}_t + \lambda_t^\pi - \lambda_{t-1}^\pi - \beta^{-1} \sigma \lambda_{t-1}^x \\
0 &= \sigma \lambda_t^x - \lambda_t^R
\end{aligned}
$$

If the ZLB never binds, we have:

$$
\begin{aligned}
0 &= \omega_x \hat{x}_t - \kappa \lambda_t^\pi \\
0 &= \omega_\pi \hat{\pi}_t + \lambda_t^\pi - \lambda_{t-1}^\pi
\end{aligned}
$$

The following steps allow us to compute allocations when the ZLB does/does not bind (taking expectations as given, as in previous sections).

### H.2.1 ZLB not binding

1. Set $\lambda_t^x = 0$.
2. Combine first order conditions for output and inflation to give:

$$0 = \omega_x \hat{x}_t - \kappa \left(\lambda_{t-1}^\pi + \beta^{-1} \sigma \lambda_{t-1}^x - \omega_\pi \hat{\pi}_t\right) - \beta^{-1} \lambda_{t-1}^x$$

which implies that

$$\hat{x}_t = \omega_x^{-1} \kappa \left(\lambda_{t-1}^\pi + \beta^{-1} \sigma \lambda_{t-1}^x - \omega_\pi \hat{\pi}_t\right) + \left(\beta \omega_x\right)^{-1} \lambda_{t-1}^x \tag{H.5}$$

3. Plug this into the Phillips curve to give

$$
\begin{aligned}
\hat{\pi}_t &= \kappa \hat{x}_t + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \\
&= \kappa \left[\omega_x^{-1} \kappa \left(\lambda_{t-1}^\pi + \beta^{-1} \sigma \lambda_{t-1}^x - \omega_\pi \hat{\pi}_t\right) + \left(\beta \omega_x\right)^{-1} \lambda_{t-1}^x\right] + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t \\
&= \frac{1}{1 + \kappa^2 \omega_\pi \omega_x^{-1}} \left(\kappa \left[\omega_x^{-1} \kappa \left(\lambda_{t-1}^\pi + \beta^{-1} \sigma \lambda_{t-1}^x\right) + \left(\beta \omega_x\right)^{-1} \lambda_{t-1}^x\right] + \beta \mathbb{E}_t \hat{\pi}_{t+1} + u_t\right)
\end{aligned} \tag{H.6}
$$

where all the terms on the right hand side are 'known' for the conditional solution.

4. Plug the solution for $\hat{\pi}$ from (H.6) into (H.5) to solve for the output gap.
5. Compute the interest rate that would implement this equilibrium:

$$\hat{R}_t = \sigma^{-1} \left(\mathbb{E}_t x_{t+1} - x_t\right) + \mathbb{E}_t \hat{\pi}_{t+1} + r_t^* \tag{H.7}$$

If the solution from this procedure respects the ZLB, it is the (conditional) solution. If not, we follow the steps in the next subsection.

### H.2.2 ZLB binds

1. Set $\hat{R}_t = \ln \beta$.

2. Solve for the output gap using the IS curve:

$$\hat{x}_t = \mathbb{E}_t \hat{x}_{t+1} - \sigma \left( \ln \beta - \mathbb{E}_t \hat{\pi}_{t+1} - r_t^* \right)$$

3. Plug this solution into the Phillips curve to solve for inflation:

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

4. Solve for $\lambda_t^\pi$ using the FOC for inflation:

$$\lambda_t^\pi = \lambda_{t-1}^\pi + \beta^{-1} \sigma \lambda_{t-1}^x - \omega_\pi \hat{\pi}_t$$

5. Use solutions for the output gap, and $\lambda_t^\pi$ to solve for $\lambda_t^x$:

$$\lambda_t^x = -\omega_x \hat{x}_t + \kappa \lambda_t^\pi + \beta^{-1} \lambda_{t-1}^x$$

## H.3 Solution algorithm and implementation

The solution algorithm for optimal commitment is a simple variation on the algorithm presented in Appendix F with two modifications. First, the step in which the derivatives of policy functions are approximated is omitted (as it is unnecessary). Second, the steps used to interpolate with respect to the endogenous state variables (when computing expectations) are modified to incorporate multivariate rather than univariate interpolation.

As for the baseline model, the initial guess for the policy functions is computed using the piecewise linear algorithms in Harrison and Waldron (2021). That toolkit is also used to stochastically simulate a piecewise linear approximation to the model and the results are used to calibrate the endpoints of the grids for the co-states $\lambda^x$ and $\lambda^\pi$. The baseline assumptions of $n_u = 15$ and $n_r = 25$ for the time-consistent solution are also used in this case. For the endogenous state variables, the assumptions are $n_q = 10$, $n_x = 20$, $n_\pi = 30$.