The Fiscal Legacy of QE: Central Bank Losses and Monetary Policy

Charles Labrousse* and Yann Perdereau†

Abstract

How do central bank losses affect output and inflation, and how should monetary policy respond? When rates rise, QE portfolios shift capital losses from households to the central bank, a debt-financed transfer that raises demand, especially if losses are deferred. Repayment through labor taxes is inflationary, and the monetary response can reverse the stimulus. In a US-calibrated HANK model, a loss of 3.8% of GDP causes a small expansion, then a persistent contraction costing 1.4% of GDP in present value. Deferral mostly shifts this cost in time. Fully offsetting inflation raises it to 2.2%; optimal policy cuts it to 0.4%.

JEL classification – E52, E58, E62, E63, E65, D31
Keywords – central bank losses; quantitative easing; fiscal-monetary interactions; distortionary taxation; HANK

We are deeply grateful to Axelle Ferriere, Julien Matheron and Gilles Saint-Paul for their help since the beginning of this project. We thank for their precious advice Adrien Auclert, Olivier Blanchard, Tobias Broer, Edouard Challe, Eustache Elina, Luigi Iovino, Donggyu Lee, Pascal Meichtry, Morten Ravn, Ricardo Reis, José-Víctor Ríos-Rull, Ciaran Rogers, Vincent Sterk, Oreste Tristani, Gianluca Violante, and participants at several seminars and conferences.

*Banque de France.
†Banca d'Italia.

1 Introduction

Central banks financed their purchases of long-term government bonds under quantitative easing (QE) largely with interest-bearing reserves. When rates rose in 2022, this duration exposure turned into large losses. The Federal Reserve's portfolio lost 4.6% of GDP in market value in 2022, and by the end of 2025 the UK Treasury had paid about £100 billion, 3.3% of GDP, to cover the losses of the Bank of England. These losses are often dismissed as an accounting matter: central banks can operate with negative equity, and losses only shift resources within the public sector. Yet someone must eventually pay for them. Lower remittances and higher taxes transmit the losses to households, at a time set by institutional arrangements. The UK Treasury covers losses as they occur, whereas the Federal Reserve and several euro-area central banks suspend remittances until their own income repays the loss, postponing the fiscal adjustment by years. How do these losses affect output, inflation and household welfare, and how should monetary policy respond?

We show that households first gain from these losses and later pay for them: under standard fiscal and monetary rules, the economy first expands, then contracts. When rates rise, the central bank bears a capital loss that households would otherwise have borne on the long-term bonds it holds. Relative to an economy without QE, the losses are therefore a transfer to households, financed by public liabilities that taxpayers will later repay. With non-Ricardian households, this transfer raises demand. Carrying the losses at the central bank strengthens the stimulus: fiscal rules respond to Treasury debt, not to central bank equity, so deferral postpones taxation further. Repayment, however, goes through distortionary labor taxes, which lower the net wage and raise inflation, even before they are levied. The monetary response to this inflation can outweigh the demand stimulus and turn the losses into a persistent contraction. When the losses are repaid using lump-sum taxes instead, output rises at all horizons. Delayed taxation thus makes central bank losses expansionary absent a monetary response, while labor-tax financing makes them inflationary and, through the monetary response, contractionary. Incidence adds a distributional cost: the transfer accrues mostly to wealthy households, while repayment falls broadly on labor income. The cost of central bank losses therefore depends on household behavior, loss-absorption arrangements and the monetary policy response.

We develop this argument in three steps. First, we extend the tractable overlapping-generations model of Angeletos, Lian and Wolf (2024) with a central bank balance sheet. The model separates the transfer created by central bank losses from the additional stimulus created by deferring their repayment, and shows how labor-tax repayment and the monetary response can reverse this stimulus. Second, we document how losses

were passed to public finances in the United States, the United Kingdom, Germany, France and Italy, from immediate Treasury indemnity at the Bank of England to the deferred asset of the Federal Reserve and the loss carryforward of the Bundesbank. Third, we quantify these mechanisms in a heterogeneous-agent New Keynesian (HANK) model with uninsurable income risk, borrowing constraints, nominal long-term debt and distortionary labor taxes. We calibrate it to the United States, with cost-push and monetary shocks that reproduce the paths of inflation and the policy rate over 2022–2025. Comparing economies with and without an inherited QE portfolio isolates the consequences of public duration exposure when rates rise.

Our main findings are the following. First, the analytical model decomposes the effect of central bank losses into a transfer effect, which does not depend on the loss-absorption regime, and a deferral effect, which arises because fiscal rules ignore central bank equity (Proposition 1). When monetary policy keeps inflation unchanged, output instead falls with the labor taxes that repay the loss, and deferral changes the timing of this fall but barely its present value (Proposition 2).

Second, the HANK model quantifies these forces for the United States. The calibrated loss amounts to 3.8% of annual GDP. Under the Federal Reserve's arrangements and holding real rates fixed, the losses raise output by 0.25% on impact. Under the Taylor rule, the monetary response to their inflationary effect overturns this stimulus: output rises by only 0.06% on impact, then falls persistently, to -0.09% after about ten years. In present value, the output cost is 1.4% of annual GDP, and the price level ends 1.25% higher. Deferral softens the initial adjustment but moves the contraction to the years in which the loss is repaid.

Third, household heterogeneity shapes both the aggregate response and its incidence. The transfer accrues mostly to wealthy households, while repayment falls on labor income: net of taxes, the bottom seven deciles of liquid wealth lose and the top three gain. This redistribution toward low-MPC households dampens the stimulus, but incomplete markets make even wealthy households spend part of their gain. Relative to a representative-agent economy, heterogeneity makes the losses less contractionary and more inflationary: under the Taylor rule, output falls by 0.11% on impact in the RANK, but rises in the HANK.

Finally, the cost of central bank losses depends more on monetary policy than on the loss-absorption regime. Keeping inflation on its no-QE path forces employment and output to absorb the adjustment to higher labor taxes, raising the present-value output cost from 1.4% to 2.2% of annual GDP. Optimal policy under commitment instead balances inflation and output stabilization. With equal weights on both objectives, it reduces the output cost of CB losses from 1.4% to 0.4% of GDP relative to the

Taylor rule. The welfare cost falls from 0.5–0.9% of annual GDP under the Taylor rule, depending on the loss-absorption regime, to approximately 0.2% across all regimes. This improvement also reflects smaller capital losses: optimal policy raises long-term rates less. Almost the same policy would be chosen without QE, so smaller losses are mainly a by-product of better macroeconomic stabilization.

Related literature

Our paper connects two strands of literature: central bank losses and their fiscal consequences, and monetary–fiscal transmission with heterogeneous households.

Central bank losses and fiscal consequences. Wallace (1981) establishes a neutrality benchmark for open-market operations. Bassetto and Messer (2013) show that paying interest on reserves exposes public finances to interest-rate risk, and Del Negro and Sims (2015); Hall and Reis (2015); Reis (2015) study central bank solvency and fiscal support. Benigno and Nistico (2020) characterize departures from neutrality when fiscal support is insufficient, the central bank seeks to limit losses, or transfers to the private sector are not offset by taxes. We instead assume full fiscal support: non-neutrality arises because the loss-absorption rule determines when losses reach the Treasury, and hence the timing of the distortionary taxes that repay them. Cecchetti and Hilscher (2025) document how accounting and indemnification arrangements shape the fiscal consequences of central bank losses; we derive and quantify their macroeconomic effects. d'Avernas et al. (2026) quantify the expected distortionary tax cost of the fiscal risk created by QE; we trace the adjustment after losses are realized and its interaction with monetary policy.

Monetary–fiscal transmission with heterogeneous households. Our analytical model extends Angeletos, Lian and Wolf (2024) with a separate central bank balance sheet, so that loss-absorption rules, rather than fiscal choices, delay fiscal adjustment, and with labor-tax repayment, which is inflationary. Auclert, Rognlie and Straub (2024) show that the timing of household consumption responses shapes the dynamic effects of fiscal policy; in our model, the loss-absorption rule determines the tax path that drives these responses. Kaplan, Moll and Violante (2018) emphasize the fiscal response to monetary policy, and Auclert (2019); Sterk and Tenreyro (2018) highlight redistribution as a transmission channel. We combine both: the fiscal response is the repayment of central bank losses, and it falls on different households than the gains. Legacy QE shields low-MPC bondholders from capital losses, while repayment falls on labor income. This incidence dampens the demand stimulus, but the contraction comes from labor-tax financing, not from incidence. Cui and Sterk (2021) study how QE stimulates demand by raising household liquidity; we take the portfolio as given and study the repayment

of its losses. Finally, we compute optimal monetary policy under commitment as in McKay and Wolf (2023), along the nonlinear transition, and show that it should respond to the distortions created by central bank losses rather than to the losses themselves.

The remainder of the paper is organized as follows. Section 2 develops the analytical model. Section 3 documents central bank losses. Section 4 presents the quantitative model, and Section 5 reports the results and studies alternative monetary policy responses. Section 6 concludes.

2 A simple model of central bank losses

This section uses a minimal model to show why central bank (CB) losses affect the economy. We call QE the stock of long-term bonds held by the CB and funded by reserves: what matters is the size and duration of this stock, not the purchases that built it.

After an unexpected rate increase, long-term bonds lose value. Without QE, households bear this loss immediately. With QE, the CB bears it, and taxpayers repay it later. Relative to an economy without QE, QE is therefore a transfer to households financed by a public liability (Lemma 1). With non-Ricardian households, this transfer raises demand. Proposition 1 splits this effect in two. The *transfer effect* does not depend on how the loss is absorbed. The *deferral effect* arises when the CB carries the loss instead of being recapitalized: the fiscal rule responds to Treasury debt, not to CB equity, so taxes are delayed further. Deferral is thus a fiscal stimulus decided by the CB.

The taxes that repay the loss fall on labor income, so they also lower natural output and raise inflation. Monetary policy determines whether these effects show up in output, inflation or the real rate. If the CB keeps inflation unchanged, output falls when taxes rise, and deferral delays this fall without much changing its present value (Proposition 2). Under a Taylor rule, whether deferral raises or lowers output on impact depends on the size of the tax distortion.

2.1 Environment

Time and monetary policy. Time is discrete. Wages are sticky and prices are flexible. The CB sets the one-period real rate: R_t denotes the gross real return between $t-1$ and t, with steady-state value $R \equiv 1+r$. All assets are real, so inflation does not revalue them.

Assets. Households hold reserves, through pass-through banks, and government bonds. Reserves pay the policy rate. The Treasury issues one-period bonds and long-term bonds. A long-term bond issued at t pays δ^k at date $t+1+k$, for all $k \geq 0$, where $\delta \in (0,1)$ governs maturity. At a constant rate R, its price is $P = 1/(R-\delta)$. Along any perfect-foresight path of rates, all assets earn the same return R_t (Appendix A.1). Stocks are measured at market value, at the end of the period.

Households. Households have perpetual youth, as in Blanchard (1985). Population has mass one. Each period, a household survives with probability $\omega \in (0,1]$ and is replaced by a newborn with no financial wealth. A household born at date s maximizes

$$\sum_{t \geq s} (\beta\omega)^{t-s} \left[\ln c_{s,t} - \nu \frac{N_t^{1+\zeta}}{1+\zeta} \right]$$

subject to

$$c_{s,t} + a_{s,t} = \frac{R_t}{\omega} a_{s,t-1} + (1-\tau_t) Y_t - \sigma \tag{1}$$

where R_t/ω reflects actuarially fair annuities and σ is a lump-sum payment representing constant seigniorage (Appendix A.9). Aggregating across cohorts gives household wealth

$$A_t = R_t A_{t-1} + Y_t - \mathcal{T}_t - \sigma - C_t \tag{2}$$

where $\mathcal{T}_t \equiv \tau_t Y_t$ is tax revenue, and aggregate consumption (Appendix A.2)

$$C_{t+1} = \beta R_{t+1} C_t - \psi_{t+1} A_t \tag{3}$$

where $\psi_{t+1} \equiv (1-\omega)(1-\beta\omega) R_{t+1}/\omega$, with steady-state value ψ. The term $-\psi_{t+1} A_t$ is the only departure from the representative-agent Euler equation: newborns, who hold no financial wealth, consume less than the households they replace. With $\omega = 1$, $\psi_{t+1} = 0$ and households are Ricardian.

Labor market. As in Section 4.3, unions set nominal wages subject to a quadratic adjustment cost, and set hours N_t, which are the same for all households. Competitive firms set prices flexibly and produce $Y_t = N_t$ at a real wage of one. Labor income is therefore Y_t, and price inflation equals wage inflation (Appendix A.3).

Treasury. The Treasury repays its debt with new debt, tax revenue $\mathcal{T}_t$ and remittances S_t from the CB:

$$B_t = R_t B_{t-1} - \mathcal{T}_t - S_t \tag{4}$$

where B_t is the market value of Treasury debt. A recapitalization of the CB is a negative remittance. Tax revenue responds to Treasury debt:

$$\mathcal{T}_t = \mathcal{T} + \phi\,(B_{t-1} - B) \tag{5}$$

where $\phi \equiv R - \mu$ and $\mu \in [0, 1)$ measures the persistence of debt. Steady-state revenue $\mathcal{T} = r\,B - \sigma$ covers interest on the steady-state debt B net of remittances. The tax rate adjusts to deliver $\mathcal{T}_t$. As in actual fiscal rules, taxes respond to Treasury debt, not to CB equity (Section 3). The Treasury finances the deficits created by QE with one-period bonds.

Central bank. The CB holds bonds, funded by reserves and equity E_t. Since bonds and reserves earn the same return, its income net of interest on reserves is $r_t\,E_{t-1} + \sigma$, with $r_t \equiv R_t - 1$, and

$$E_t = R_t\,E_{t-1} + \sigma - S_t \tag{6}$$

In normal times, the CB remits its income and its equity is zero.

Equilibrium. Households hold all reserves and the bonds not held by the CB, which equal reserves plus equity. Asset-market clearing therefore gives

$$A_t = B_t - E_t \tag{7}$$

Household wealth is Treasury debt minus CB equity. A loss carried by the CB therefore raises household wealth without raising Treasury debt, and hence without triggering the fiscal rule. This is the source of the deferral effect. The goods market clears, $Y_t = C_t$, and one of the two conditions is redundant by Walras' law (Appendix A.4).

Steady state. Inflation and CB equity are zero, and $A = B$. By (3), output is

$$Y = \frac{\psi}{\beta R - 1}\,B \tag{8}$$

The CB sets R, and output clears the goods market. With $B > 0$ and $\omega < 1$, this requires $\beta R > 1$: individual consumption grows with age, and newborns, who consume less, keep aggregate consumption constant.

Neutrality. A purchase of bonds paid with reserves at market price raises the assets and the liabilities of the CB by the same amount. It leaves E_t, and hence A_t, unchanged. Since all assets earn the same return, it also leaves CB income unchanged, as in Wallace (1981). This holds along any perfect-foresight path of real rates, and hence for any path of the portfolio. QE matters only through unexpected capital gains and losses, which change E_t.

2.2 An unexpected rate increase

The shock. The economy is in the steady state until the end of $t = -1$. At the beginning of $t = 0$, agents learn that the real rate follows a path $\{R_t\}_{t\geq 1}$ that returns to R, and have perfect foresight afterwards. Returns between $t = -1$ and $t = 0$ on reserves and one-period bonds were set at $t = -1$, so $R_0 = R$, and long-term coupons are fixed. The shock therefore acts at $t = 0$ only through the price of long-term bonds: (2), (4) and (6) hold at $t = 0$ with the realized return on each portfolio. From $t = 1$ on, all assets earn R_t. All results are to a first order in the size of the shock.

The experiment. We compare two economies with the same Treasury debt and the same shock, which differ only in the duration of the CB portfolio at the end of $t = -1$. In both, the CB has issued reserves Q, has zero equity, and holds bonds worth Q: long-term bonds in the *QE economy*, one-period bonds in the *no-QE economy*, where households hold the long-term bonds instead. The experiment isolates who bears duration risk, not the size of the balance sheet (Appendix A.4). By neutrality, the two economies are identical until $t = -1$, whatever the history of purchases. Treasury issuance does not respond to QE. For any variable X, let $\hat{X}_t \equiv X_t^{QE} - X_t^{\text{no QE}}$.

The loss. Under the path of rates of the no-QE economy, the CB in the QE economy bears a capital loss (Appendix A.1)

$$L = Q \sum_{t\geq 1} R^{-t}\, \delta^t\, (R_t - R) \tag{9}$$

Each rate increase is weighted by the value of the payments that remain at that date. With $R_t - R = \rho_r^{t-1}\Delta r$, $L = Q\,\delta\,\Delta r/(R - \rho_r\delta)$. For instance, $Q = 30\%$ of GDP, a duration of 5 years ($\delta = 0.81$ at $R = 1.01$) and a rate increase of 2 percentage points with persistence $\rho_r = 0.9$ give $L \approx 1.7\%$ of GDP.

Monetary policy. In the no-QE economy, the real rate follows $\{R_t\}$. QE changes demand and, through taxes, supply. How the CB responds in the QE economy determines whether these effects show up in output, inflation or rates. Under *fixed real rates*, the CB does not respond: the QE economy faces the same path of rates, as under the "neutral" monetary policy of Angeletos, Lian and Wolf (2024), and output is determined by demand. Under *equal inflation*, the CB fully offsets the effects of losses on inflation: the real rate adjusts and output is determined by supply. A *Taylor rule* lies between these polar closures. Sections 2.3 and 2.4 use fixed real rates, under which the path of rates matters only through L. Section 2.5 turns to the other closures.

2.3 QE after a rate surprise: a debt-financed transfer

The Treasury may take over part of the loss, by recapitalizing the CB at the end of $t = 0$ or through lower remittances later. Let

$$\ell_t \equiv -\hat{E}_t$$

denote the part of the loss carried by the CB at the end of t, after any recapitalization. By (7), $\hat{B}_t = \hat{A}_t - \ell_t$ is the part carried by the Treasury. The path of ℓ_t defines the loss-absorption regime. We consider two polar regimes. Under *Treasury support*, the Treasury recapitalizes the CB at the end of $t = 0$, so $\ell_t = 0$. Under *deferral*, the CB carries negative equity and suspends its remittances until the loss is repaid, as with the deferred asset discussed by Hall and Reis (2015), so $\ell_t > 0$ until repayment. With franchise value $F \equiv \sigma/r$, the present value of seigniorage, the loss is repaid in finite time if $L < F$, which we assume (Appendix A.8).

Lemma 1 (QE as a debt-financed transfer)**.** *Under fixed real rates, households in the QE economy hold additional public liabilities $\hat{A}_0 = L$ at the end of $t = 0$, whatever the regime. For $t \geq 1$, additional taxes are $\hat{\mathcal{T}}_t = \phi\,(\hat{A}_{t-1} - \ell_{t-1})$ and*

$$\hat{A}_t = \underbrace{\mu^t L}_{transfer} + \underbrace{\phi \sum_{j=0}^{t-1} \mu^{t-1-j}\, \ell_j}_{deferral} \tag{10}$$

Lemma 1 is an accounting result: it uses only the public-sector budget constraints, asset-market clearing and the fiscal rule, not household behavior (proof in Appendix A.5). QE shifts the capital loss L from households to the public sector at $t = 0$. The QE economy is therefore the no-QE economy plus a transfer L to households, financed by public liabilities. How these liabilities are split between the Treasury and the CB does not change $\hat{A}_0$, but it changes the timing of taxes, because the fiscal rule responds only to Treasury debt. Under Treasury support, $\ell_t = 0$ and $\hat{A}_t = \mu^t L$. When the CB carries part of the loss, taxes are postponed and households hold public liabilities for longer: this is the deferral term.

2.4 The effect on demand

Let $\gamma \equiv 1/(\beta R) \in (0, 1)$. Under fixed real rates, solving forward the difference between the Euler equations (3) of the two economies, with $Y_t = C_t$, gives to a first order (Appendix A.5)

$$\hat{Y}_t = \psi \sum_{k \geq 0} \gamma^{k+1}\, \hat{A}_{t+k} \tag{11}$$

Output depends on how much public liability households hold, and for how long, before taxes repay it. Substituting (10) gives the main result.

Proposition 1 (Transfer and deferral effects)**.** *Under fixed real rates, consider any bounded path of CB losses* $\{\ell_t\}$*, and let* $\Lambda \equiv \sum_{j\geq 0} \gamma^{j+1}\ell_j$ *denote the discounted losses. Output at* $t = 0$ *is*

$$\hat{Y}_0 = \mathcal{M}\left(\underbrace{L}_{transfer} + \underbrace{\phi\,\Lambda}_{deferral} \right) \tag{12}$$

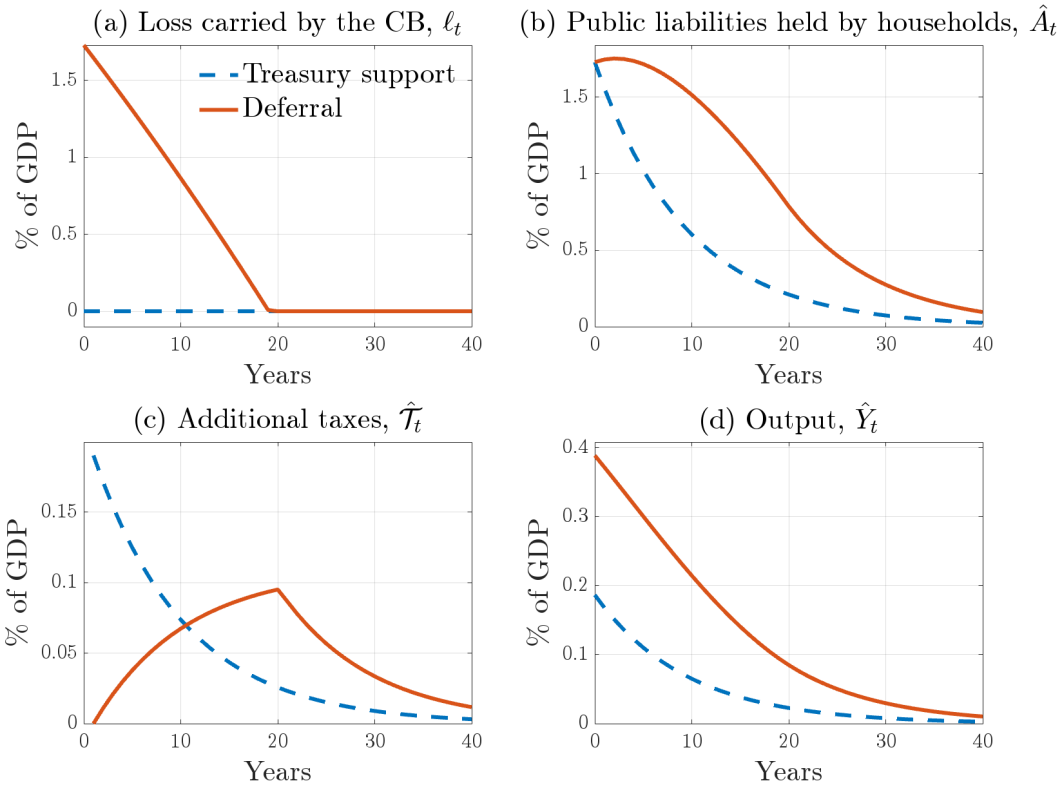
where $\mathcal{M} \equiv \psi/(\beta R - \mu)$*. Both effects are zero if households are Ricardian.*

The transfer effect does not depend on the regime. Under Treasury support, $\Lambda = 0$ and only the transfer effect remains. The deferral effect is proportional to the discounted taxes that the CB defers: relative to a rule on consolidated debt, $\phi\,\ell_j$ is the tax not levied at $t = j+1$. The multiplier $\mathcal{M}$ is the effect on output of a unit transfer financed by debt repaid at rate ϕ. It increases with ψ, since more of the taxes then fall on future generations, and with μ, since taxes are then raised more slowly.

Only the timing of taxes matters. In every regime, the present value of additional taxes is L: QE changes when taxes are paid, not how much. Holding this present value constant, delaying taxes raises output (Appendix A.6): the public sector discounts taxes at R, whereas households alive today discount them at R/ω, so a delay shifts taxes to future generations. If the loss were taxed at $t = 0$, output would not change. The deferral effect is therefore a temporary relaxation of the fiscal rule, decided by the CB rather than by the government. It exists only because CB losses escape the fiscal rule: if taxes responded to consolidated debt $B_t - E_t$, then $\hat{A}_t = \mu^t L$ whatever the regime, and $\hat{Y}_0 = \mathcal{M}L$.

Illustration. Figure 1 compares the two regimes for the same loss. Under Treasury support, taxes jump at $t = 1$ and decline geometrically. Under deferral, the Treasury first loses remittances, its debt rises gradually, and taxes peak when the loss is repaid. Households hold public liabilities for longer (panel b), and the output effect at $t = 0$ doubles, from 0.2% to 0.4% of GDP (panel d). These magnitudes are illustrative: they depend on ω, a reduced form for marginal propensities to consume. They are also large because the real rate does not respond: higher demand raises income, which raises demand again, and this feedback grows with the delay (Appendix A.6). Section 2.5 shows that monetary policy dampens, and can reverse, this effect.

Figure 1: Treasury support versus deferral



Notes: QE minus no-QE economy, % of GDP, annual. Fixed real rates. Rate increase of 2 percentage points with persistence $\rho_r = 0.9$. $r = 1\%$, $\omega = \mu = 0.9$, $Q = 30\%$, duration 5 years, franchise value $F = 10\%$, $B/Y = 60\%$. $\beta \approx 0.997$ is set by (8). The loss is $L \approx 1.7\%$ of GDP and is repaid after 20 years under deferral.

2.5 Distortionary taxes and monetary policy

The taxes that repay the loss are levied on labor income, so they also act on supply. A higher tax rate lowers the after-tax value of work and raises wage inflation, like a higher wage markup. To a first order, the difference in inflation between the two economies follows the wage Phillips curve, written in terms of revenue since the fiscal rule sets revenue (Appendix A.3):

$$\hat{\pi}_t = \tilde{\kappa}\left(\hat{Y}_t + \Theta\,\hat{\mathcal{T}}_t\right) + \beta\,\hat{\pi}_{t+1} \tag{13}$$

where $\tilde{\kappa} > 0$ and

$$\Theta \equiv \frac{1}{(1+\zeta)(1-\tau) - \tau}$$

We assume $(1+\zeta)(1-\tau) > \tau$, so that the economy is on the increasing side of the Laffer curve and $\Theta > 0$. Natural output falls by Θ per unit of additional revenue: $\hat{Y}_t^n = -\Theta\hat{\mathcal{T}}_t$. This is the *fiscal cost-push*. Lump-sum taxes correspond to $\Theta = 0$.

Fixed real rates. Demand depends on taxes only through revenue, so (11) and Proposition 1 hold unchanged. The cost-push shows up only in inflation. Deferral raises inflation, both through demand and because it delays the cost-push (Appendix A.7).

Equal inflation. If inflation does not move, (13) requires output to equal natural output, which falls when taxes rise.

Proposition 2 (Equal inflation)**.** *Under equal inflation, output is proportional to additional taxes:*

$$\hat{Y}_t = -\Theta\,\hat{\mathcal{T}}_t \qquad t \geq 0$$

Its present value is

$$\sum_{t\geq 1} R^{-t}\,\hat{Y}_t = -\Theta\,(L+\Omega) \tag{14}$$

where Ω, defined in (39), is the interest cost of the tightening.

Output does not move at $t = 0$, since taxes respond to debt with a lag. The demand effects of Proposition 1 now move the real rate instead of output: the CB keeps rates higher while households hold the transfer. Deferral keeps the transfer in households' hands for longer, so the tightening lasts longer. Deferral therefore changes when output falls, not how much. By (14), the regime affects the present value of output only through Ω, the interest cost of the tightening on the public liabilities whose payments were not set at $t = 0$. With lump-sum taxes, QE leaves output unchanged and only moves real rates: the tax distortion, not the loss itself, lowers output.

The two polar closures reverse the roles of demand and supply. Under fixed real rates, demand moves output and taxes move inflation. Under equal inflation, demand moves rates and taxes move output.

Taylor rule. A Taylor rule, $\hat{R}_{t+1} = \phi_\pi\,\hat{\pi}_t$ with $\phi_\pi > 0$, lies between the two polar closures: output lies between demand-determined output, under fixed real rates, and natural output, under equal inflation. Two forces then shape the effect of deferral on impact. First, the CB leans more against the deferral effect than against the transfer effect: the liabilities added by deferral build up over time, so the demand and inflation they create peak later, and households anticipate the tightening. Second, with distortionary taxes, deferral delays the cost-push, and hence the tightening that the CB undertakes in anticipation of it. The first force makes deferral less expansionary on impact, the second more: which force dominates is a quantitative question that we

address in the HANK model.

2.6 From the analytical model to HANK

QE is not expansionary per se. After an unexpected rate cut, $L < 0$: future taxes fall, demand is lower than without QE, and the deferral effect is negative if the CB retains its gains. After a rate increase, QE can also lower output on impact under a Taylor rule when taxes are distortionary.

Given monetary policy and MPCs, the analytical model isolates four objects that govern the effects of CB losses: the initial capital loss L, the path of losses carried by the CB ℓ_t, whether the fiscal rule recognizes CB equity, and the fiscal cost of raising taxes Θ. It highlights two opposing forces. Delaying fiscal backing raises household net wealth and stimulates demand, while financing the loss with distortionary labor taxes lowers natural output and creates inflationary pressure.

The HANK model of Section 4 relaxes three simplifications. First, MPCs differ: the transfer accrues to holders of liquid wealth, who have low MPCs, whereas the taxes that repay it fall on labor income, which drives the consumption of constrained households. Second, assets are nominal, so an inflation surprise at $t = 0$ partly offsets the transfer. Third, seigniorage rises with rates, which speeds up repayment under deferral (Appendix A.9). QE still shifts the capital loss L from households to the public sector at $t = 0$, but the timing of taxes is no longer a sufficient statistic: the incidence of the loss and of taxes also matters. The HANK model therefore determines the sign of the deferral effect on impact.

3 Central bank losses in practice

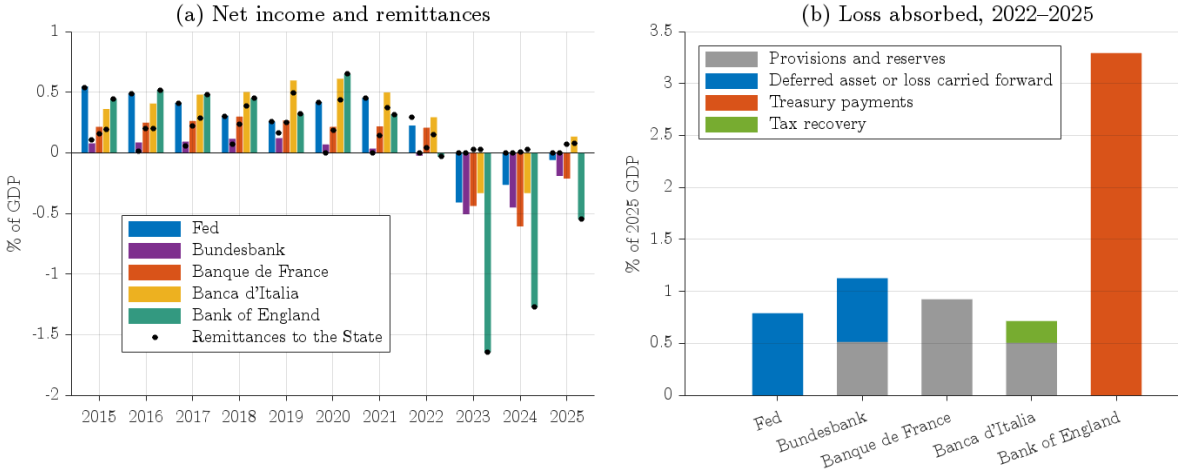
This section documents the three objects that govern the effect of central bank losses in Section 2: the size of the capital loss L, the path of losses carried by the CB ℓ_t, and whether fiscal rules respond to CB equity. Appendix B gives definitions, legal sources and data.

3.1 Size and recognition of losses

The Fed and the Eurosystem hold their QE portfolios at amortised cost. A rise in rates therefore creates a market-value loss at once, but it is recognised only gradually, as negative net interest income: assets keep earning their yield at purchase, while reserves pay the policy rate. The Bank of England's Asset Purchase Facility (APF) marks its gilts to market, but its transfers to the Treasury follow realised income, so its losses

also reach the public finances gradually. Figure 2.a shows net income before provisions and remittances. It turned negative in 2022 for the Bundesbank, in 2023 for the Fed, the Banque de France and the Bank of England, and in 2023–2024 only for the Banca d'Italia, whose portfolio of higher-yield domestic bonds limited the fall.

Figure 2: Net income, remittances and cumulative loss



Notes: Panel a: net income before provisions, reserves and taxes; remittances include corporate taxes. Panel b: losses absorbed in 2022–2025, in % of 2025 GDP. Definitions and sources in Appendix B.

3.2 Who carries the loss

Three arrangements absorb the recognised loss (Table 1, Figure 2.b).

Table 1: Absorption of central bank losses, 2022–2025

	% of 2025 GDP				
	Provisions	Deferred	Treasury	Remittances	Model regime
Fed	–	0.79	–	None	Deferral
Bundesbank	0.51	0.61	–	None	Deferral
Banque de France	0.92	–	–	Tax only, 2025	Deferral
Banca d'Italia	0.50	–	0.21	Continued	Provisions
Bank of England	–	–	3.29	Reversed	Treasury support

Notes: Accounting losses absorbed in 2022–2025 by released provisions and reserves, by a deferred asset or loss carried forward at end-2025, and by the Treasury. Remittances: payments to the State in 2023–2025. Definitions and sources in Appendix B.

Deferral. The Fed records a deferred asset and suspends remittances until future income repays it, as in Hall and Reis (2015). The deferred asset reached $243.5 billion

at the end of 2025, about 0.8% of GDP. The ECB, the Bundesbank and the Banque de France carry their losses forward and offset them against future profits; the accumulated loss of the Bundesbank reached €27.8 billion at the end of 2025. Remittances stop in both cases. This is the deferral regime of Section 2: the central bank carries the loss, $\ell_t > 0$, until its own income repays it.

Provisions. Eurosystem NCBs first released general risk provisions, built from income retained in earlier years. The Banca d'Italia used €11.4 billion of provisions in 2023–2024 and kept paying the State. Provisions front-load the fiscal cost: remittances were lower before the shock, not after it. In the model, this corresponds to positive central bank equity before the shock, which we abstract from.

Indemnity. The Bank of England's APF is indemnified by HM Treasury, which received its income until 2022 and has covered its losses as they occur since then. This is the Treasury support regime of Section 2. Treasury payments reached £100 billion by the end of 2025, 3.3% of GDP, more than elsewhere because the APF portfolio was large, about 35% of GDP at its peak, and active gilt sales brought losses forward.

3.3 Fiscal rules ignore central bank equity

The deferral effect requires that taxes respond to Treasury debt, not to CB equity. In the euro area, the debt in the excessive deficit procedure is general government gross debt at nominal value, which excludes NCBs. In the United States, remittances are budget revenue, and the deferred asset does not enter debt held by the public. A loss carried by the CB therefore lowers fiscal revenue only when remittances would have been paid, without raising measured debt. In the United Kingdom, indemnity payments are government spending, so the loss enters the budget as it is recognised.

4 Quantitative model

Section 2 isolates the timing of taxes. This section adds the features that determine incidence and magnitudes: heterogeneous marginal propensities to consume, a Taylor rule, and a demand for cash, which determines seigniorage. We keep the notation of Section 2. Time is quarterly. Variables are in real terms, except the nominal rate i_t and bond prices, and π_t denotes gross inflation. There is no aggregate uncertainty: shocks are unexpected, and agents have perfect foresight afterwards (Section 5.1).

4.1 Assets and returns

Reserves and one-period bills pay the nominal rate i_t. Their real return between $t-1$ and t is

$$R_t \equiv \frac{1+i_{t-1}}{\pi_t} \tag{15}$$

The Treasury also issues long-term bonds, which pay a nominal coupon δ^k at date $t+1+k$. Their nominal price is

$$P_t = \frac{1+\delta P_{t+1}}{1+i_t} \tag{16}$$

and their realized real return is

$$R_t^\ell \equiv \frac{1+\delta P_t}{P_{t-1}\,\pi_t} \tag{17}$$

Appendix C.1 derives these returns and the capital loss. A portfolio with a share s of long-term bonds at the end of $t-1$ earns $R_t + s\,(R_t^\ell - R_t)$. Absent surprises, (16) implies $R_t^\ell = R_t$, so the composition of portfolios does not affect returns. In the steady state, $P = 1/(1+i-\delta)$.

4.2 Households

A continuum of households $i \in [0,1]$ consume and save in liquid assets a, subject to a borrowing constraint. Liquid assets are a claim on a mutual fund that holds cash, reserves, bills, and the long-term bonds not held by the central bank. Households face idiosyncratic productivity shocks z that follow an exogenous stationary Markov process with transition probabilities $P_z(z_{t+1}|z_t)$. Hours N_t are set by unions and are the same for all households (Section 4.3). Each household i solves

$$\max_{\{c_{it},\,a_{it}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\eta}}{1-\eta} - \nu \frac{N_t^{1+\zeta}}{1+\zeta} \right]$$

subject to

$$\underbrace{c_{it} + a_{it}}_{\text{Uses}} = \underbrace{R_t^a\, a_{it-1}}_{\text{Return on wealth}} + \underbrace{(1-\tau_t)\, w_t z_{it} N_t}_{\text{Net labor income}} + \underbrace{T}_{\text{Transfers}} - \underbrace{\sigma_t}_{\text{Seigniorage}}$$

$$a_{it} \geq 0$$

The parameter η is the inverse of the elasticity of intertemporal substitution, ν determines the disutility of labor, and ζ is the inverse Frisch elasticity. Since hours are chosen by unions, households take N_t as given. Labor income $w_t z_{it} N_t$ is taxed at the rate τ_t. Households receive a constant lump-sum transfer T from the Treasury. Transfers are

the same in all the economies we compare, so the fiscal adjustment goes through the tax rate only (Section 4.4).

The fund holds the cash used for transactions on behalf of households (Section 4.5). Cash pays no nominal interest. The fund charges households the forgone interest as a lump sum σ_t, which is the seigniorage of the CB, as in Section 2. We thus assume that the cost of holding cash is shared equally across households. The fund therefore earns the same return on all its short-term assets, and its real return is

$$R_t^a = (1 - s_{t-1}^h)\, R_t + s_{t-1}^h\, R_t^\ell \tag{18}$$

where s_{t-1}^h is the share of long-term bonds in household wealth at the end of $t-1$ (Section 4.6). Households therefore bear the capital gains and losses on the long-term bonds they hold, in proportion to their wealth.

The borrowing constraint binds for households with low wealth and low productivity. These households consume their income, and their marginal propensity to consume is one. We denote by C_t and A_t the aggregates of c_{it} and a_{it}.[1]

4.3 Labor market and firms

Unions. As in Auclert, Rognlie and Straub (2024), a continuum of unions supply differentiated labor services, with elasticity of substitution ϵ_w, and ration hours equally across households. Unions set nominal wages to maximize the average utility of households, subject to a quadratic adjustment cost on wage inflation π_t^w, relative to the inflation target π. This cost enters utility and does not affect household decisions. This gives the wage Phillips curve (Appendix C.2)

$$\frac{\pi_t^w}{\pi}\left(\frac{\pi_t^w}{\pi} - 1\right) = \frac{\epsilon_w}{\theta_w}\, N_t\, \Xi_t + u_t + \beta\, \frac{\pi_{t+1}^w}{\pi}\left(\frac{\pi_{t+1}^w}{\pi} - 1\right) \tag{19}$$

where θ_w governs wage rigidity, u_t is a cost-push shock, and

$$\Xi_t \equiv \nu N_t^\zeta - \frac{1}{\mu^w}\,(1 - \tau_t)\, w_t \int z_{it}\, c_{it}^{-\eta}\, di \tag{20}$$

is the labor wedge: the marginal disutility of work minus the marked-down value of after-tax labor income, with $\mu^w \equiv \epsilon_w/(\epsilon_w - 1)$ the desired wage markup. Marginal utility is weighted by productivity, since hours are rationed equally and labor income

[1] We model cash demand at the aggregate level and assume that the fund holds it, so that households bear the forgone interest as a lump-sum payment, as in Section 2. Seigniorage matters for our mechanism only through its aggregate level, which determines how fast the CB repays a loss under deferral. Introducing cash in the household problem would make cash a store of value for constrained households and lower their marginal propensity to consume, at odds with the evidence.

is proportional to z_{it}. When demand raises hours above the level unions would choose, $\Xi_t > 0$ and wage inflation rises: hours are demand-determined in the short run. The shock u_t is in units of quarterly wage inflation.

The labor tax is distortionary. In the steady state, a higher τ lowers hours. In the short run, a higher τ_t raises the labor wedge and wage inflation, like the cost-push shock, and monetary policy determines how much of it passes to output.

Firms. Competitive firms produce with $Y_t = N_t$ and set prices flexibly. The real wage is then constant, $w_t = 1$, price inflation equals wage inflation, $\pi_t = \pi_t^w$, and profits are zero.

4.4 Treasury

The Treasury spends a constant G, pays transfers T, levies labor taxes $\mathcal{T}_t \equiv \tau_t w_t N_t$ and receives remittances S_t. In the steady state, long-term bonds are a share s^T of Treasury debt. The Treasury keeps the real face value of its long-term debt at its steady-state level $s^T B$, and finances any deficit with one-period bills B_t^S, as in Section 2. The market value of its debt is

$$B_t = B_t^S + s^T B \frac{P_t}{P} \tag{21}$$

where P is the steady-state price of long-term bonds, and its budget constraint is

$$B_t = R_t B_{t-1}^S + R_t^\ell s^T B \frac{P_{t-1}}{P} + G + T - \mathcal{T}_t - S_t \tag{22}$$

Fiscal rule. Tax revenue follows the rule of Section 2, applied to the face value of Treasury debt, which values long-term bonds at their steady-state price:

$$B_t^F \equiv B_t^S + s^T B \tag{23}$$

$$\mathcal{T}_t = \mathcal{T} + \phi \left(B_{t-1}^F - B \right) \tag{24}$$

where $\mathcal{T}$ and B are steady-state values, $\phi \equiv R - \mu$, and $\mu \in [0, 1)$ measures the persistence of debt at constant rates. The tax rate τ_t adjusts to deliver $\mathcal{T}_t$. As in actual fiscal rules, taxes respond to Treasury debt, not to CB equity, and to its face value, not to its market value.[2] A rise in long-term yields therefore lowers the market value of existing debt but leaves its face value, and hence taxes, unchanged: only deficits, financed with bills, move B_t^F.

[2] In the euro area, the debt that enters the excessive deficit procedure is total gross debt at nominal value outstanding at the end of the year.

4.5 Central bank

Monetary policy. The CB sets the nominal rate according to a Taylor rule

$$i_t = \rho_i \, i_{t-1} + (1 - \rho_i)\big[i + \phi_\pi(\pi_t - \pi)\big] + \varepsilon_t \tag{25}$$

with ε_t a monetary policy shock.

Cash. Cash demand is

$$M_t = \chi \, C_t \, e^{-\xi \, i_t^a} \tag{26}$$

where $i_t^a \equiv (1 + i_t)^4 - 1$ is the annualized nominal rate and $\xi \geq 0$ is the interest semi-elasticity of cash demand. Cash demand increases with consumption and decreases with the nominal rate. The CB supplies cash elastically at the policy rate, in exchange for bills.

Balance sheet. The CB holds long-term bonds worth Q, inherited from past asset purchases and funded by reserves $V_t = Q$, and bills $H_t = M_t + E_t$. Its balance sheet is

$$Q + H_t = V_t + M_t + E_t \tag{27}$$

The CB holds the portfolio to maturity, without reinvestment. Since $R_t^\ell = R_t$ for $t \geq 1$, this assumption matters only for the timing of recognized losses under historical cost. Between $t-1$ and t, long-term bonds earn R_t^ℓ, bills and reserves earn R_t, and cash pays no interest, so its real value is M_{t-1}/π_t. Equity before remittances is therefore

$$\Psi_t = R_t \, E_{t-1} + \sigma_t + Q \left(R_t^\ell - R_t\right) \tag{28}$$

where

$$\sigma_t \equiv \Big(R_t - \frac{1}{\pi_t}\Big) M_{t-1} = \frac{i_{t-1}}{\pi_t} \, M_{t-1} \tag{29}$$

is seigniorage: the return on the bills that back cash, net of the real cost of cash. Households pay it as a lump sum, since the fund holds cash on their behalf. The last term of (28) is the excess return on long-term bonds, which is zero absent surprises. Equity evolves as $E_t = \Psi_t - S_t$.

Seigniorage determines how fast the CB can repay a loss on its own. When ξ is small, cash demand is inelastic and seigniorage rises with the policy rate: the rate increase that creates the loss also raises the income that repays it. When ξ is large, cash demand falls with the rate, and so does seigniorage.

Loss-absorption regimes. We consider two polar regimes:

1. **Treasury support:** $S_t = \Psi_t$ and $E_t = 0$. The Treasury absorbs any loss immediately and recapitalizes the CB.
2. **Deferral:** $S_t = \max\{0, \Psi_t\}$ and $E_t = \min\{0, \Psi_t\}$. The CB keeps negative equity and retains its income until the loss is repaid, as under the Fed's deferred asset, see Hall and Reis (2015).

Without losses, both regimes remit $S_t = \sigma_t$.

4.6 Equilibrium

Markets clear:

$$\begin{cases} Y_t = C_t + G & \text{(Goods)} \\ A_t = B_t - E_t & \text{(Liquid assets)} \end{cases} \tag{30}$$

The asset-market condition is the second line of 30: household wealth, which includes cash, equals Treasury debt minus CB equity. Households hold the long-term bonds not held by the CB. At the end of $t = -1$, their share in household wealth is

$$s^h = \frac{s^T B - Q}{B} \tag{31}$$

Since $R_t^\ell = R_t$ for $t \geq 1$, only this share matters. By Walras' law, one market-clearing condition is redundant. An equilibrium is a path of prices $\{\pi_t, i_t, P_t\}$, hours N_t, policies $\{\tau_t, S_t\}$ and household decisions such that households, firms and unions optimize, cash demand is met, the budget constraints of the Treasury and the CB hold, and markets clear.

In the steady state, the real rate clears the asset market, $A = B$, and β is calibrated to match the target real rate. Absent surprises, $R_t^\ell = R_t$: the size Q of the QE portfolio affects neither the steady state nor any perfect-foresight path, and the neutrality result holds.

4.7 Calibration

The model is quarterly and calibrated to the United States at the end of 2021, before the Federal Reserve starts raising rates. Output is normalized to $Y = N = 1$, one quarter of GDP. The discount factor β, the disutility of labor ν and the level of cash demand χ are set internally to match steady-state targets. The others are set externally. Table 2 reports the values.

Households. Utility is logarithmic in consumption, $\eta = 1$. The inverse Frisch elasticity is $\zeta = 2$. Log productivity follows an AR(1) process. We use the quarterly parameters of McKay, Nakamura and Steinsson (2016): the persistence is $\rho_z = 0.966$ and the innovation variance is $\sigma_z^2 = 0.017$. We discretize the process with seven states using the method of Rouwenhorst (1995).

The discount factor clears the asset market, $A = B$, at a targeted annual real rate of 1%. The disutility of labor sets the labor wedge (20) to zero at $Y = 1$. In the steady state, 18% of households are at the borrowing constraint, and the average quarterly marginal propensity to consume out of a \$500 transfer is 0.18.

Labor market. The elasticity of substitution between labor services is $\epsilon_w = 10$, a wage markup of 11%. It is a normalization for the equilibrium path: ν adjusts to deliver $N = 1$ and θ_w adjusts to deliver the slope of the wage Phillips curve. It matters only for welfare, through the steady-state distortion and the cost of wage adjustment. We set the slope of the Phillips curve with respect to output, $\kappa_Y \equiv \kappa_w (\zeta + Y/C)$, to 0.01, as in Angeletos, Lian and Wolf (2024), where $\kappa_w \equiv (\epsilon_w/\theta_w)\, \nu N^{1+\zeta}$ is the slope with respect to the labor wedge. This pins down θ_w.

Treasury. Treasury debt is 100% of annual GDP, close to federal debt held by the public at the end of 2021. Government spending and transfers match their 2019 shares of GDP, before the pandemic: government consumption and gross investment are 17.6% of GDP, and government social benefits to persons are 14.3% of GDP. Since households are infinitely lived, transfers include old-age programs. The tax rate balances the steady-state budget of the Treasury, which gives $\tau = 32.7\%$. The speed of fiscal adjustment follows Angeletos, Lian and Wolf (2024). Their preferred value implies that taxes rise by 2.6 cents per quarter for each additional dollar of debt at the beginning of the quarter. In our timing, $\phi = R \times 0.026 = 0.026$ and $\mu = 0.976$ per quarter, a debt half-life of 7.3 years. Long-term bonds are 83% of Treasury debt, one minus the share of bills in marketable debt at the end of 2021.

Central bank and cash. The inflation target is 2% a year, the Taylor coefficient is $\phi_\pi = 1.5$, and the persistence is $\rho_i = 0.85$. In Appendix, we also consider a Taylor rule responding to output gap. The steady-state nominal rate is 3.0% a year. The semi-elasticity of cash demand is $\xi = 1.8$, the estimate of Ireland (2009). The level of cash demand matches currency in circulation, 9% of annual GDP at the end of 2021. Given this target, χ has a closed form: $\chi = M/C\, e^{\xi i^a}$. Seigniorage is then 0.27% of annual GDP, and the franchise value $F = \sigma/r$ is 27% of annual GDP.

Long-term bonds and QE. The maturity parameter δ matches the average duration of the domestic Federal Reserve System Open Market Account (SOMA) at the end of 2021: 5.8 years (par-weighted valuation). Since this duration measures the sensitivity of prices to yields, it corresponds in the model to the modified duration $-\partial \ln P/\partial i = 1/(1+i-\delta)$, which equals 23.2 quarters for $\delta = 0.964$. Treasury long-term bonds have the same maturity. The reference QE portfolio is the domestic SOMA portfolio at the end of 2021, Treasuries and agency securities, worth 34.9% of annual GDP at par value. Section 5.1 calibrates the cost-push and monetary shocks (u_t, ε_t).

Table 2: Calibration

Parameter	Description	Value	Target or source
Households			
β	Discount factor	0.993	Annual real rate of 1%
η	Inverse EIS	1	Log utility
ζ	Inverse Frisch elasticity	2	Frisch elasticity of 0.5, Chetty et al. (2011)
ν	Disutility of labor	0.704	$Y = 1$
ρ_z	Persistence of productivity	0.966	McKay, Nakamura and Steinsson (2016)
σ_z	Std. of productivity innovations	0.130	McKay, Nakamura and Steinsson (2016)
Labor market			
ϵ_w	Substitutability of labor services	10	Normalization; welfare only.
κ_Y	Slope of the Phillips curve (output)	0.01	Angeletos, Lian and Wolf (2024)
κ_w	Slope of the wage Phillips curve (wedge)	0.003	Implied by κ_Y
Treasury			
B	Treasury debt	4	100% of annual GDP
G	Government spending	0.176	Consumption and investment, 2019
T	Transfers	0.143	Social benefits to persons, 2019
τ	Labor tax rate	0.327	Steady-state budget
μ	Persistence of debt	0.976	Angeletos, Lian and Wolf (2024)
ϕ	Tax response to debt	0.026	Angeletos, Lian and Wolf (2024)
s^T	Long-term share of Treasury debt	0.83	Bill share of marketable debt, 2021
δ	Maturity of long-term bonds	0.964	SOMA duration, 5.8 years
Central bank			
π	Inflation target	1.005	2% annual
ϕ_π	Taylor coefficient	1.5	Standard
ρ_i	Interest-rate smoothing	0.85	Standard
χ	Level of cash demand	0.461	M/Y: 9% of annual GDP
ξ	Semi-elasticity of cash demand	1.8	Ireland (2009)
Q	Reference QE holdings	1.40	SOMA, 34.9% of annual GDP

5 Results

5.1 Experiment

We study a rise in rates similar to that of 2022–2025 and compare two economies that differ in who holds long-term public debt when rates rise. In one, the CB holds a long-duration portfolio inherited from past QE, $Q > 0$. In the other, the private sector holds it, $Q = 0$. Treasury issuance is the same in both, s^T being unchanged, so the comparison isolates the duration exposure of the CB, not the effects of the purchases themselves. The rate surprise shifts part of the capital loss from households to the CB. We vary how the CB recognizes and absorbs this loss, and decompose the effect of CB losses into a transfer and two forms of deferral. Unlike in Section 2, the effects reflect incidence as well as timing: Section 5.4 separates these channels.

Shock and calibration. The economy starts in the steady state. At $t = 0$, two unexpected sequences of shocks hit: cost-push shocks $\{u_t\}_{t\geq 0}$ in the wage Phillips curve (19), and monetary shocks $\{\varepsilon_t\}_{t\geq 0}$ in the Taylor rule (25). Both paths are revealed at $t = 0$, and agents have perfect foresight thereafter.

We calibrate the shocks in the QE economy, at the reference Q and under the rules of the Fed: historical-cost accounting and deferral of losses (see below). The CB starts with zero equity, like the Fed, whose surplus is capped by law. The impact quarter is 2022Q1, when the Fed ended net purchases and started raising rates. We choose $\{u_t, \varepsilon_t\}$ so that the model reproduces exactly two series over 2022Q1–2025Q4: quarterly core PCE inflation, and the change in the effective federal funds rate relative to 2021Q4. From 2026Q1 on, the shocks are zero. Figure 7 in the Appendix shows the fit. The counterfactual no-QE economy receives the same shocks. Section 5.2 compares the balance sheet of the CB with the data.

The capital loss. Portfolios held at the end of $t = -1$ were chosen expecting constant rates. The shocks lower P_0, so $R_0^\ell < R_0$, and $R_t^\ell = R_t$ afterwards: all revaluations occur at $t = 0$. Relative to bills, the CB loses

$$L \equiv Q\left(R_0 - R_0^\ell\right) = Q\,\frac{\delta\,(P - P_0)}{P\,\pi_0} \tag{32}$$

Households lose less. Their fund earns $R_0^a = R_0 - s^h\left(R_0 - R_0^\ell\right)$ and CB holdings lowers the steady-state long-term share $s^h = (s^T B - Q)/B$. The loss of the CB is therefore a transfer to households, in proportion to their liquid wealth. The Treasury gains on its long-term debt, but the face value of its debt (23), and hence taxes, barely move.

Loss recognition and absorption. Remittances follow accounting income, not economic income. Let X_t be the accounting income of the QE portfolio in excess of bills. It replaces $Q(R_t^\ell - R_t)$ in (28). Under *mark-to-market accounting* (M), $X_t = Q(R_t^\ell - R_t)$, and the whole loss is recognized at $t = 0$. Under *historical cost* (H), the rule of the Fed and the Eurosystem, the portfolio earns its yield at purchase, i, and its funding costs the policy rate:

$$X_t = \frac{(i - i_{t-1})\,\delta^t\,Q}{p_t} \tag{33}$$

where $\delta^t Q$ is the nominal book value of the portfolio at the end of $t-1$ and $p_t \equiv \prod_{k=0}^{t} \pi_k$ is the price level. The loss is then recognized gradually, as negative net interest income, but its present value is still L. Book equity E_t exceeds market equity $E_t^m \equiv E_t - \mathcal{L}_t$ by the unrealized loss

$$\mathcal{L}_t = \frac{\delta^{t+1}\,Q\,(1 - P_t/P)}{p_t} \tag{34}$$

with $\mathcal{L}_0 = L$, whereas $\mathcal{L}_t = 0$ under mark-to-market. Remittances follow book equity, but household wealth reflects market values. The asset-market condition in (30) therefore becomes $A_t = B_t - E_t^m$.

We combine these rules with the loss-absorption regimes of Section 4.5, Treasury *support* (S) and *deferral* (D), and consider three regimes. Under MS (Mark-to-market M, Treasury Support S), the Treasury absorbs the loss at $t = 0$. Under HS (Historical cost H, Treasury Support S), it absorbs the loss as it is recognized, close to the Bank of England up to asset sales. Under HD (Historical cost H, Deferral D), the CB repays the loss out of its future income, like the Fed and the Eurosystem. The loss carried by the CB at market value, $\ell_t \equiv -\hat{E}_t^m$, is zero under MS and equals $\mathcal{L}_t$ under HS. Under HD, it also includes the deferred loss $-E_t$. Section 3 documents these arrangements: the Fed and, after their provisions were exhausted, the Bundesbank and the Banque de France follow HD; the Bank of England is close to HS. The effect of QE under HD decomposes as

$$\underbrace{\text{HD} - \text{no QE}}_{\text{total}} = \underbrace{\text{MS} - \text{no QE}}_{\text{transfer}} + \underbrace{\text{HS} - \text{MS}}_{\text{accounting deferral}} + \underbrace{\text{HD} - \text{HS}}_{\text{remittance deferral}} \tag{35}$$

The transfer effect is that of Proposition 1: the loss is passed to the Treasury at once. Accounting deferral arises because remittances ignore unrealized losses. Remittance deferral arises because the fiscal rule ignores CB equity.

5.2 Central bank losses: size and timing

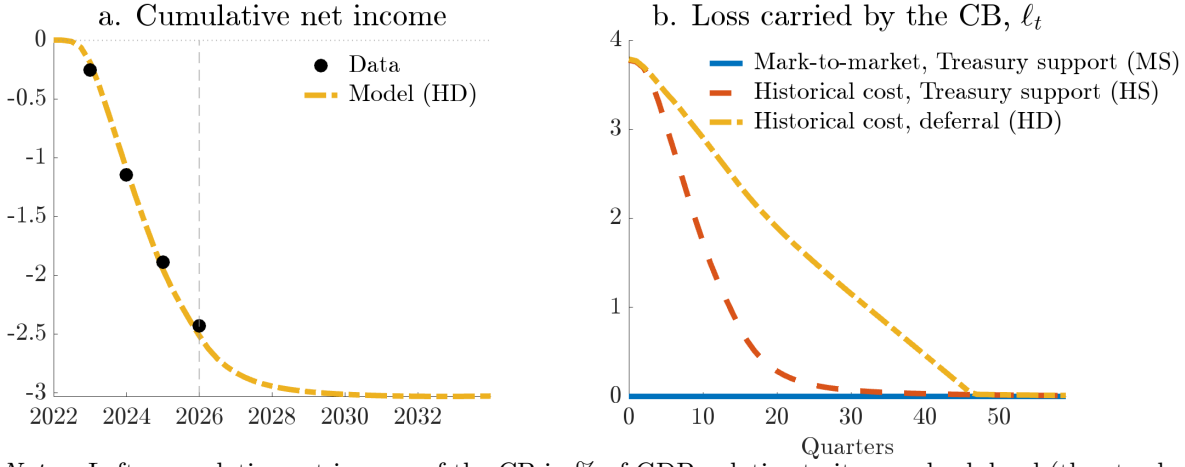
We start with the balance sheet of the CB in the QE economy, at the reference Q. Figure 3 shows cumulative net income and the loss carried by the CB under MS, HS

and HD.

Size and recognition. The shocks lower the market value of the QE portfolio by $L = 3.8\%$ of GDP, 10.9% of the portfolio, in all three regimes. The regime changes when the loss is recognized and who carries it, not its size: in all three regimes, cumulative remittances converge to $-L$ in present value. Under MS, the whole loss is recognized at $t = 0$, and the Treasury recapitalizes the CB by L. Under historical cost, $X_t < 0$ as long as the policy rate exceeds the book yield i, by (33): the loss is recognized as negative net interest income over the period of high rates. In the model, 63% of L, in present value, is recognized within three years.

Deferral. Under HS, the Treasury covers recognized losses as they occur, and remittances are negative for 15 quarters. Under HD, remittances stop, and the CB repays the loss out of its own income: while equity is negative, $E_t = R_t E_{t-1} + \sigma_t + X_t$. Seigniorage, which rises by 120% at the peak of rates, absorbs part of the loss as it is recognized. The deferred loss peaks at 1.6% of GDP at $t = 17$, almost half of L, and remittances resume after 47 quarters. The discounted loss carried by the CB, $\Lambda \equiv \sum_{t \geq 0} R^{-t} \ell_t$, is the discounted area under the curves of the right panel of Figure 3. It is 11.0 quarters of L under HS and 21.6 under HD: remittance deferral doubles the time during which the CB carries the loss.

Figure 3: Central bank losses



Notes: Left: cumulative net income of the CB in % of GDP, relative to its pre-shock level (the steady state for the model, 2021 for data). Right: loss carried by the CB at market value, ℓ_t relative to the no-QE economy, in % of annual GDP.

Comparison with the data. Table 3 compares these paths with the Fed. The model accounts for most of the capital loss, 3.8% of GDP against 4.6%, and matches the fall

in net income relative to its pre-shock level (Figure 3, left panel). It overstates the peak deferred asset, 1.6% of GDP against 0.8%. The gap comes mostly from pre-shock net income, 0.27% of GDP in the model against 0.48% for the Fed. In the steady state, the QE portfolio earns no excess return over its funding cost, so net income is seigniorage only, whereas the Fed also earned a positive carry on its long-term bonds. The same fall in income thus pushes net income further below zero in the model.

Table 3: Central bank losses: model and data (Federal Reserve)

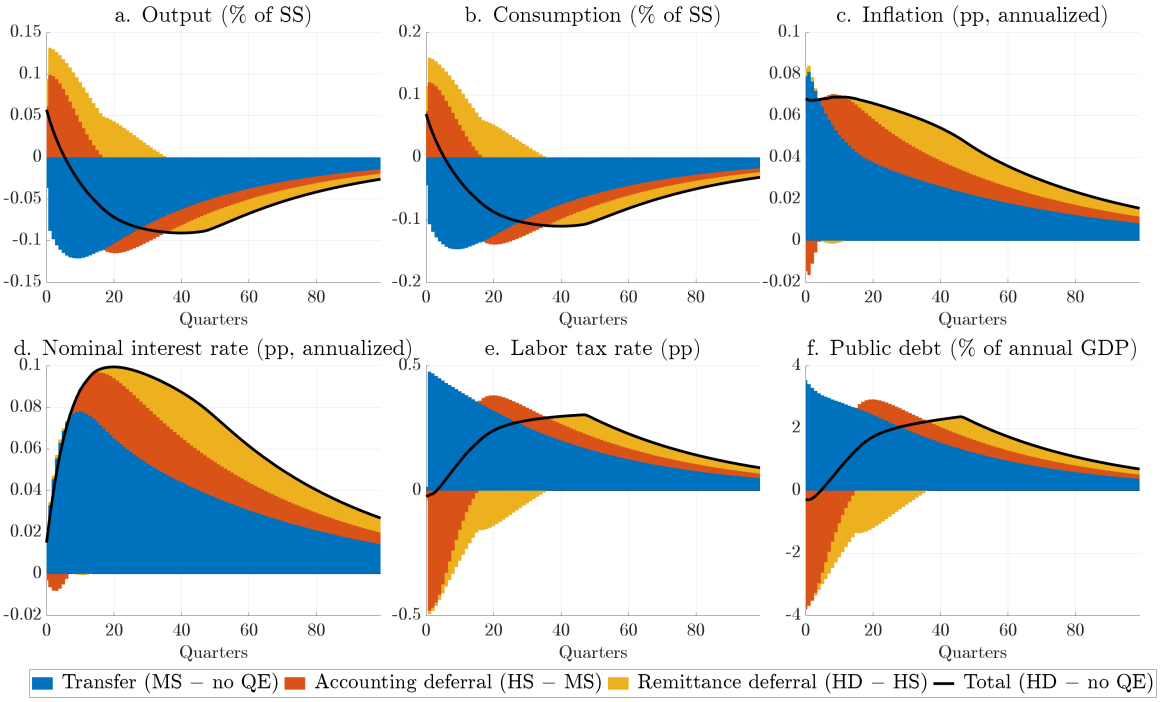
	Data	Model
% of GDP		
Capital loss, 2022	4.6	3.8
Deferred asset, end of 2022	0.06	0.06
Pre-shock net income	0.48	0.27
Cumulative net income, 2023–2025	−0.73	−1.59
relative to pre-shock	−2.17	−2.40
Peak deferred asset	0.79	1.64
Quarters from impact		
Peak deferred asset	15	17

Data: pre-shock net income is 2021 net income. Fed audited financial statements, H.4.1 and New York Fed (2023).

5.3 Transfer and deferral effects

This section presents our main quantitative result: the effect of the rate hike in an economy with a legacy QE portfolio, relative to an economy without QE. Figure 4 shows the difference between the QE and no-QE economies, at the reference portfolio size Q. The black line is the total effect, under the rules of the Fed (HD). The bars decompose it into the three terms of equation (35): the transfer (MS minus no QE), accounting deferral (HS minus MS), and remittance deferral (HD minus HS). Each term includes the endogenous fiscal and monetary responses.

Figure 4: Effects of CB losses after the rate hike: total effect and decomposition



Total effect. Under the rules of the Fed, the CB loss first raises output, then lowers it persistently. Output rises by 0.06% on impact, falls below its no-QE path after 5 quarters, reaches its bottom, −0.09%, around quarter 40, and is still 0.07% lower after 60 quarters. The impact gain is small relative to this persistent fall: in present value, discounted at the steady-state real rate, $\sum_{t\geq 0} R^{-t}\hat{Y}_t$ is $-0.37\,L$, that is −1.4% of annual GDP. The loss is also inflationary: inflation is about 0.07 pp higher for 15 quarters and stays above its no-QE path for decades, so the price level ends 1.25% higher. The policy rate is higher by up to 10 bp, for more than two decades. The labor tax is slightly lower on impact and rises until quarter 47, when remittances resume.

Decomposition. The transfer drives the medium-run contraction and most of the inflation. The two deferrals make the impact expansionary and shift the contraction and the inflation to later periods.

Transfer. Under MS, the Treasury recapitalizes the CB by L at $t = 0$. Public debt rises by L, and the fiscal rule raises the labor tax by 0.5 pp. Households hold less duration and therefore bear L less of the capital loss. This transfer from the Treasury to households is nevertheless contractionary: output falls by up to 0.12% and consumption by up to 0.15%, after about 10 quarters. The higher labor tax raises wage costs: inflation rises by 0.08 pp on impact, and the Taylor rule gradually raises the

policy rate.

Accounting deferral. Under HS, the Treasury covers the loss as it is recognized, over 15 quarters. Unrealized losses reduce CB equity at market value, but do not enter the debt to which the fiscal rule responds. Debt and taxes therefore rise gradually rather than on impact, and output rises by enough to offset its impact fall under MS. This is the deferral effect of Proposition 1. Once the loss is recognized, debt exceeds its level under MS, because the Treasury has started to repay later. The effect then reverses: output is lower from quarter 16 on. Inflation and the policy rate are lower for a few quarters only, because forward-looking wage setters anticipate the higher future taxes.

Remittance deferral. Under HD, the Treasury receives no remittances while CB equity is negative. The CB repays the loss out of its own income, and remittances resume after 47 quarters. The fiscal adjustment is postponed further: relative to HS, debt and the labor tax are lower for about 35 quarters. Output is higher for 35 quarters and lower thereafter. Inflation and the policy rate barely move at first, and are higher from quarter 12 on.

Implications for central banks. The three regimes map into the arrangements of the main central banks. The Fed defers its losses (HD). The Bank of England is indemnified by the Treasury, which covers its losses as they are recognized (HS); its active sales bring recognition forward, toward MS. The Eurosystem also defers losses, but first absorbs them with provisions built from retained income, that is, from lower past remittances. This front-loads part of the fiscal cost and shortens the period of negative equity. Our results imply that these differences matter mostly for the timing of the effects of legacy losses: deferral postpones them, but does not avoid them. Under HD, the two deferrals turn the impact fall in output into a rise, but move the contraction and the inflation to the period in which the loss is repaid. Inflation then stays higher for longer, after the rate hike is over, which matters for the stance of monetary policy (Section 5.5).

The next Section explains why the transfer is contractionary although households gain, and why these effects differ from those of Section 2.

5.4 Channels: demand, fiscal supply, and the monetary response

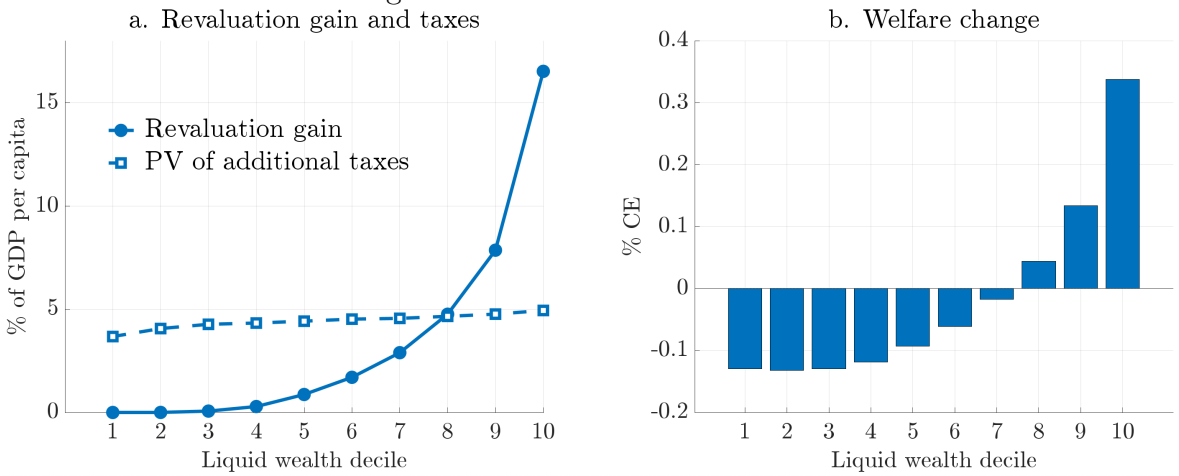
Why is the transfer contractionary although households gain? Three forces drive the effects of the CB loss. At given real rates, who receives the revaluation gain and who pays the taxes, and when, determines aggregate demand. The labor taxes that finance the loss raise wage costs. Monetary policy then responds to the resulting inflation.

Three experiments separate these forces. First, we solve the QE economies under fixed real rates: they face the ex ante real rates $\{R_t\}$ of the no-QE economy, as in Section 2. The difference with the baseline, in which both economies follow the Taylor rule, is the effect of the monetary response, including its feedback on public finances. Second, we compare the HANK with two benchmarks under both closures (Appendix D.2, which also reports decomposition (35) under fixed real rates). In the RANK, the representative household receives the whole gain and pays all the taxes, so neither timing nor incidence moves demand. In the TANK, savers are Ricardian and hold all the wealth, while hand-to-mouth households pay part of the taxes: the TANK removes the wealth effect and keeps incidence. Third, we repay the loss with lump-sum rather than labor taxes, which isolates the fiscal cost-push.

At given real rates, delayed taxation raises demand. Under HD, output rises by 0.25% on impact and stays above its no-QE path, and all three components are positive (Figure 8). Households gain on their portfolios at $t = 0$, while the fiscal rule spreads the taxes over decades. With incomplete markets, households are not Ricardian: they spend wealth faster than a representative household and discount future taxes at a rate above R. Even in the top decile of liquid wealth, the annual MPC is 0.055, against 0.01 in the RANK. Passing the loss to the CB thus acts as a debt-financed transfer to households: the transfer alone raises output by 0.09% on impact, and the two deferrals, which postpone the taxes further, add 0.16%. In the RANK, output does not move, as in the Ricardian limit of Proposition 1.

Incidence across households dampens it. Incidence works in the opposite direction. Figure 5 shows it in the baseline (HD, Taylor rule), by decile of pre-shock liquid wealth. Panel a compares, per household, the revaluation gain at $t = 0$ with the present value of additional net taxes, in % of annual GDP per capita. Panel b shows the welfare change, in permanent consumption equivalent.

Figure 5: Incidence of the CB loss



The gain is proportional to liquid wealth: the top decile receives 16% of GDP per capita, almost half of the total gain, and the bottom three deciles almost nothing. Taxes fall on labor income and vary much less, from 3.6% of GDP per capita in the first decile to 5.0% in the tenth. Net of taxes, the bottom seven deciles lose and the top three gain. Welfare falls by about 0.13% in the bottom four deciles and rises by 0.34% in the top one. The losers include the households with the highest MPCs, 1.00 and 0.92 per year in the first two deciles, while the winners have MPCs of at most 0.10. In isolation, this lowers demand: in the TANK, output falls by up to 0.05% at given real rates. In the HANK, the wealth effect dominates.

Fiscal supply. At given real rates, the loss raises inflation by 0.40 pp on impact in the HANK. Part of this increase comes from demand. The rest is the fiscal cost-push of Section 2.5: the higher labor tax lowers the after-tax return to work and raises wage costs at a given level of activity. Even in the RANK, where demand does not move, inflation rises by 0.25 pp: Ricardian equivalence removes the demand stimulus, but not the distortion created by labor-tax financing. Because wage setting is forward-looking, anticipated taxes raise inflation immediately: deferral postpones the taxes, but not all of their inflationary effects.

The monetary response. Under the Taylor rule, the policy rate rises gradually, by up to 10 bp in the HANK. Output then rises by only 0.06% on impact, against 0.25% at given real rates, and falls below its no-QE path after 5 quarters (Figure 4). When the Taylor rule also responds to the output gap, the contraction is milder and inflation higher, but the qualitative pattern is unchanged (Appendix D.4). In the RANK, where

no demand stimulus offsets the tightening, output falls by 0.11% on impact and 0.15% at the trough, against a rise of 0.06% and a trough of −0.09% in the HANK. The tightening keeps inflation in the RANK within 0.02 pp of its no-QE path, whereas inflation in the HANK stays 0.07 pp higher for 15 quarters. Heterogeneity thus does not create the contraction, but adds a demand stimulus, which makes the loss less contractionary and more inflationary. It also reveals distributional costs that the RANK hides: households with little wealth pay additional taxes but receive little of the gain.

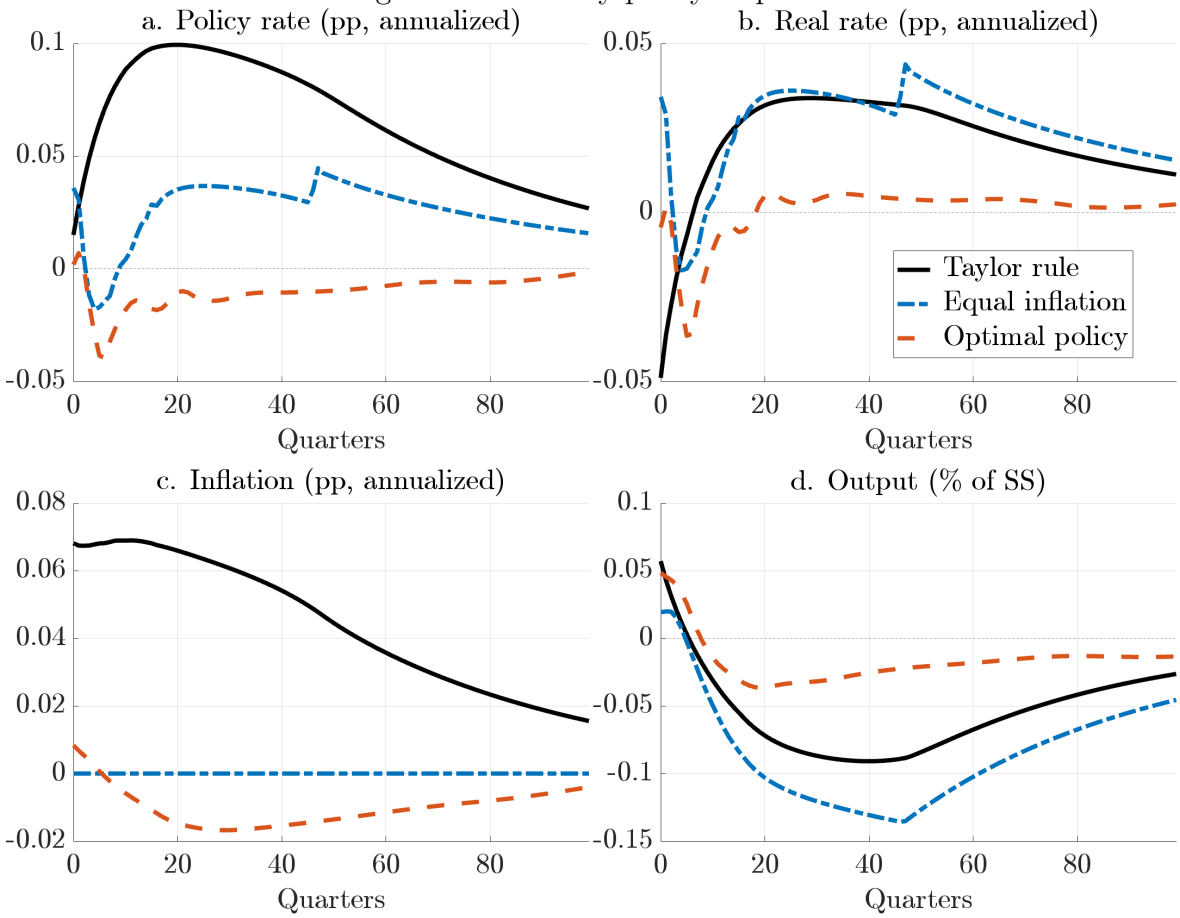
Distortionary taxation drives the contraction. When the loss is repaid with lump-sum taxes, the labor tax does not move, and under the Taylor rule CB losses raise output at all horizons. In present value under HD, output rises by 0.9% of annual GDP, against a fall of 1.4% in the baseline (Appendix D.3). Incidence does not explain the difference: the gain is as concentrated as in the baseline, and lump-sum taxes fall equally on all households, so repayment is even more regressive.

In sum, delayed taxation makes CB losses expansionary at given real rates, and labor-tax financing makes them inflationary and, through the monetary response, contractionary. Because the contraction runs through the monetary response, its size depends on how the CB reacts, which Section 5.5 examines.

5.5 Monetary policy and the cost of CB losses

Under the Taylor rule, inflation stays above its no-QE path for decades, and the price level ends 1.25% higher (Section 5.3). Yet the cost of CB losses depends on how the CB responds to them. A CB concerned with price stability may tighten more to prevent the inflation that its losses generate. It may also take into account that a higher rate path raises its own loss. We consider two alternatives to the rule. *Equal inflation* keeps inflation on its no-QE path at all dates: it measures the cost of fully neutralizing the inflationary effect of losses, as in Proposition 2. *Optimal policy* under commitment trades off inflation and output, and internalizes the effect of the rate path on the CB loss. Figure 6 compares both with the Taylor rule under HD.

Figure 6: Monetary policy responses



Notes: HD, difference with the no-QE economy under the same policy. Taylor rule: baseline of Section 5.3. Equal inflation: the policy rate keeps inflation on its no-QE path. Optimal policy: rate path minimizing $\mathbb{L}$ with $\lambda = 1$.

Equal inflation. We solve the QE economies under the inflation path of the no-QE economy and let the policy rate adjust. With inflation unchanged, the wage Phillips curve (19) keeps wage costs on their no-QE path, so a higher labor tax must be offset by lower employment: output tracks the labor tax, as in Proposition 2. Output barely moves on impact, declines as the tax rises, and reaches its trough, -0.13%, when remittances resume. In present value, it falls by about 2.2% of annual GDP, against 1.4% under the rule. This additional cost, 0.8% of annual GDP, is the price of a price level 1.25% lower. It is almost the same under MS, HS and HD: deferral delays the output loss without changing its size.

The policy rate absorbs the demand effects of Section 5.4. As long as households hold the transfer, the CB keeps the rate higher so that the transfer does not raise demand. The rate rises on impact, when the revaluation gain accrues, and stays 0.03 to 0.04 pp above its no-QE path for decades. Because deferral leaves the transfer with

households for longer, the tightening is flatter under HD than under MS, but twice as large in cumulative terms.

Why lower inflation comes with a lower policy rate. Equal inflation sets a lower policy rate than the rule, yet delivers lower inflation (Figure 6). This is because households respond to the real rate, and under the rule most of the increase in the policy rate only compensates for higher expected inflation: with $\phi_\pi = 1.5$, one point of inflation eventually raises the policy rate by 1.5 points, but the real rate by only 0.5. Equal inflation needs no such compensation, so it delivers a real rate at least as high as under the rule with a policy rate about 0.07 pp lower (panel b). The difference is largest on impact: the smoothed rule raises the policy rate slowly while inflation jumps, so the real rate falls just when households receive the transfer, whereas under equal inflation it rises. The lower policy rate reflects lower expected inflation, not a looser stance, and output falls more.

Optimal policy. The CB chooses the path of the policy rate at $t = 0$, under commitment, to minimize

$$\mathbb{L} = \sum_{t \geq 0} \beta^t \left(\pi_t^2 + \lambda y_t^2 \right) \tag{36}$$

where π_t is annualized inflation and y_t the output gap, both in deviation from the steady state. We set $\lambda = 1$, as in a dual mandate with equal weights. Appendix D.5 reports results for $\lambda = 16\kappa_Y/\epsilon_w \approx 0.016$, the weight implied by the standard second-order approximation of welfare Woodford (2003), rescaled for annualized inflation. As in McKay and Wolf (2023), we generate rate paths with news shocks to the Taylor rule. Since the deferral regime is inactive at the steady state, we compute the responses to these shocks around the nonlinear path and iterate.[3] The loss L is now endogenous to the rate path. We measure the cost of CB losses by the increase in the minimized $\mathbb{L}$ relative to the no-QE economy, also under optimal policy.

Optimal policy keeps inflation close to, and slightly below, its no-QE path. Unlike equal inflation, it does not lean against the transfer: the policy rate stays slightly below its no-QE path, the real rate close to it, and the output loss is about a third of that under the rule. Through commitment, the CB steers inflation expectations rather than raising the current rate.

Relative to the rule, optimal policy lowers inflation and raises output at the same time, because it shrinks the shock rather than moving along the Phillips curve. Since it raises rates less in response to the 2022–2025 shocks, long rates rise less and the CB

[3] The instrument excludes the impact quarter, and we add a penalty $\lambda_\Delta \sum_{t \geq 0} \beta^t (\Delta i_t)^2$, with $\lambda_\Delta = 0.1$, to remove nearly flat directions of the problem.

loses less: L falls from 3.8% to 0.9% of annual GDP (Table 4). Smaller losses require smaller increases in the labor tax, which weakens the fiscal cost-push. Balance-sheet considerations play almost no role in this: under HD, L is less than 3% smaller than what the QE portfolio would lose under the optimal rate path of the no-QE economy. The CB responds to the transfer and to the later tax increase, not to the loss itself.

Table 4: Cost of CB losses: Taylor rule and optimal policy

	Taylor rule				Optimal policy			
	$\Delta\mathbb{L}$	CE	ΔY	L	$\Delta\mathbb{L}$	CE	ΔY	L
MS	35.9	−0.88	−1.44	3.76	2.08	−0.18	−0.53	0.85
HS	18.2	−0.68	−1.48	3.78	2.68	−0.19	−0.42	0.87
HD	6.2	−0.49	−1.41	3.79	2.52	−0.18	−0.43	0.88

Notes: $\lambda = 1$. Differences with the no-QE economy under the same policy. $\Delta\mathbb{L}$: increase in the loss (36). CE: change in utilitarian welfare in consumption equivalent, present value, % of annual GDP. ΔY: present value of the change in output, % of annual GDP. L: CB loss, % of annual GDP.

Optimal policy therefore removes most of the cost of CB losses. In every regime, the increase in $\mathbb{L}$ is less than half as large as under the rule. Utilitarian welfare gives the same ranking, although optimal policy does not maximize it. In consumption equivalent, CB losses cost 0.5 to 0.9% of annual GDP under the rule, and deferral almost halves this cost. Under optimal policy, they cost about 0.2% of annual GDP in every regime. With the welfare-based weight, optimal policy brings the increase in $\mathbb{L}$ close to zero, but only by shifting the cost from inflation to output, and it lowers utilitarian welfare relative to $\lambda = 1$ (Appendix D.5).

Policy implication. The cost of CB losses depends more on the monetary policy response than on who bears the loss. Switching between loss-absorption regimes changes this cost by a factor of six under the Taylor rule, but by about 30% under optimal policy. In welfare terms, the regime changes the cost by a factor of almost two under the rule, and not at all under optimal policy. A CB would gain little from tilting its policy to limit its losses, and an optimizing CB barely does so.

6 Conclusion

When rates rise, a central bank holding a legacy QE portfolio bears a capital loss that households would otherwise have borne. Relative to an economy without this portfolio, the loss is a transfer to households, financed by public liabilities that taxpayers repay later. This paper shows that the timing of taxes makes these losses expansionary, while

their financing through labor taxes makes them inflationary and, once monetary policy responds, contractionary.

The analytical model separates a transfer effect, which does not depend on the loss-absorption regime, from a deferral effect, which arises because fiscal rules respond to Treasury debt, not to central bank equity. Loss-absorption arrangements thus relax fiscal rules temporarily, without any fiscal decision. In a HANK model calibrated to the United States, a loss of 3.8% of GDP raises output slightly on impact, then lowers it by 1.4% of annual GDP in present value, and raises the price level by 1.25%. These effects are small in any given quarter, but persist long after the rate hike is over. The contraction comes from distortionary financing, not from incidence: when the loss is repaid through lump-sum transfers, output rises at all horizons.

Our results inform three policy debates. First, loss-absorption arrangements primarily determine when output costs arise. They also affect welfare: under a Taylor rule, deferral nearly halves the welfare cost relative to immediate recapitalization, while active asset sales bring losses forward and increase their cost. Second, the financing instrument matters more than the timing of repayment: the contraction disappears when repayment does not distort labor supply. Third, monetary policy reaction is central. Fully offsetting the inflationary effect of losses raises their output cost by more than half. Optimal policy instead responds less to the 2022–2025 shocks, reducing central bank losses by roughly three quarters and their additional output cost by two thirds.

Our analysis opens several avenues for future research. We study the fiscal legacy of QE, taking the portfolio as given, not the purchases that built it. Combining both would allow a full assessment of QE, weighing the stimulus of purchases against the cost of the losses they may later generate. We also abstract from aggregate risk, and hence from the risk-premium channel of balance-sheet policies and from the ex ante fiscal risk of duration exposure. Finally, the framework could be applied to the euro area, where provisions front-load part of the fiscal cost and a single monetary policy interacts with national fiscal rules, or used to evaluate proposals to reduce the remuneration of reserves, which would shift the financing of losses to banks.

References

Angeletos, George-Marios, Lian, Chen and Wolf, Christian (2024). "Can deficits finance themselves?" In: *Econometrica* 92.5, pp. 1351–1390 (cit. on pp. 2, 4, 8, 21, 22).

Auclert, Adrien (2019). "Monetary policy and the redistribution channel". In: *American Economic Review* 109.6, pp. 2333–2367 (cit. on p. 4).

Auclert, Adrien, Rognlie, Matthew and Straub, Ludwig (2024). "The intertemporal Keynesian cross". In: *Journal of Political Economy* 132.12, pp. 4068–4121 (cit. on pp. 4, 17, 48).

Bassetto, Marco and Messer, Todd (2013). "Fiscal Consequences of Paying Interest on Reserves". In: *Fiscal Studies* 34.4, pp. 413–436. DOI: 10.1111/j.1475-5890.2013.12014.x (cit. on p. 4).

Benigno, Pierpaolo and Nistico, Salvatore (2020). "Non-neutrality of open-market operations". In: *American Economic Journal: Macroeconomics* 12.3, pp. 175–226 (cit. on p. 4).

Benigno, Pierpaolo and Woodford, Michael (2005). "Inflation Stabilization and Welfare: The Case of a Distorted Steady State". In: *Journal of the European Economic Association* 3.6, pp. 1185–1236 (cit. on p. 55).

Blanchard, Olivier (1985). "Debt, deficits, and finite horizons". In: *Journal of Political Economy* 93.2, pp. 223–247 (cit. on p. 6).

Cecchetti, Stephen G and Hilscher, Jens (2025). "Fiscal consequences of central bank losses". In: *Central Bank Capital in Turbulent Times: The Risk Management Dimension of Novel Monetary Policy Instruments*. Springer Nature Switzerland, Cham, pp. 95–124 (cit. on p. 4).

Chetty, Raj, Guren, Adam, Manoli, Day and Weber, Andrea (2011). "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins". In: *American Economic Review* 101.3, pp. 471–475 (cit. on p. 22).

Cui, Wei and Sterk, Vincent (2021). "Quantitative easing with heterogeneous agents". In: *Journal of Monetary Economics* 123, pp. 68–90 (cit. on p. 4).

d'Avernas, Adrien, Hubert de Fraisse, Antoine, Ning, Liming and Vandeweyer, Quentin (2026). "The fiscal cost of quantitative easing". In: *Working Paper* (cit. on p. 4).

Del Negro, Marco and Sims, Christopher A (2015). "When does a central bank's balance sheet require fiscal support?" In: *Journal of Monetary Economics* 73, pp. 1–19 (cit. on p. 4).

Hall, Robert and Reis, Ricardo (2015). *Maintaining central-bank financial stability under new-style central banking*. Tech. rep. National Bureau of Economic Research (cit. on pp. 4, 9, 14, 20).

Ireland, Peter N (2009). "On the welfare cost of inflation and the recent behavior of money demand". In: *American Economic Review* 99.3, pp. 1040–1052 (cit. on pp. 21, 22).

Kaplan, Greg, Moll, Benjamin and Violante, Giovanni (2018). "Monetary Policy According to HANK". In: *American Economic Review* 108 (3), pp. 697–743 (cit. on p. 4).

McKay, Alisdair, Nakamura, Emi and Steinsson, Jón (2016). "The power of forward guidance revisited". In: *American Economic Review* 106.10, pp. 3133–3158 (cit. on pp. 21, 22).

McKay, Alisdair and Wolf, Christian K. (2023). "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" In: *Econometrica* 91.5, pp. 1695–1725 (cit. on pp. 5, 33).

New York Fed (2023). *Open Market Operations during 2022*. Report prepared for the Federal Open Market Committee. Federal Reserve Bank of New York, Markets Group. URL: https://www.newyorkfed.org/medialibrary/media/markets/omo/omo2022-pdf.pdf (cit. on p. 26).

Reis, Ricardo (2015). *Different types of central bank insolvency and the central role of seignorage*. Tech. rep. National Bureau of Economic Research (cit. on p. 4).

Rouwenhorst, K Geert (1995). "Asset pricing implications of equilibrium business cycle models". In: *Frontiers of Business Cycle Research* 1, pp. 294–330 (cit. on p. 21).

Sterk, Vincent and Tenreyro, Silvana (2018). "The transmission of monetary policy through redistributions and durable purchases". In: *Journal of Monetary Economics* 99, pp. 124–137 (cit. on p. 4).

Wallace, Neil (1981). "A Modigliani-Miller theorem for open-market operations". In: *The American Economic Review* 71.3, pp. 267–274 (cit. on pp. 4, 7).

Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press (cit. on p. 33).

Appendix

Table of contents

A Appendix to the analytical Section 2 **39**

A.1 Long-term bonds and the capital loss . . . 39
A.2 Aggregation . . . 40
A.3 Phillips curve . . . 41
A.4 Wealth identity and the counterfactual . . . 41
A.5 Proofs . . . 42
A.6 Output and the timing of taxes . . . 43
A.7 Inflation under fixed real rates . . . 43
A.8 Loss-absorption regimes . . . 44
A.9 Seigniorage and cash . . . 45

B Appendix to the data Section 3 **45**

B.1 Definitions . . . 45
B.2 Central banks . . . 46
B.3 GDP . . . 47
B.4 Fiscal rules . . . 47

C Appendix to the quantitative Section 4 **47**

C.1 Nominal long-term bonds . . . 47
C.2 Wage Phillips curve . . . 48

D Additional results for Section 5 **49**

D.1 Experiment and calibration of the shocks . . . 49
D.2 Fixed real rates and representative-agent benchmarks . . . 49
D.3 Robustness: lump-sum transfers . . . 52
D.4 Robustness: Output gap in Taylor rule . . . 54
D.5 Optimal policy with the welfare-based weight . . . 54

A Appendix to the analytical Section 2

A.1 Long-term bonds and the capital loss

Prices and returns. At a constant rate R, the price of a long-term bond is

$$P = \sum_{k\geq 0} \frac{\delta^k}{R^{k+1}} = \frac{1}{R-\delta}$$

After paying its coupon, a bond is equivalent to δ new bonds, so its gross return is $(1+\delta P)/P = R$, as for a one-period bond. Its modified duration is $1/(R-\delta)$.

Budget constraint in market value. Let x_t be the position of a household in new-bond equivalents at the beginning of t, which is also the coupon it receives at t. If it buys new bonds at t, its budget constraint is

$$c_t + P_t\,(x_{t+1} - \delta x_t) = y_t + x_t$$

Let $b_t \equiv P_t\,x_{t+1}$ be the market value of its bonds at the end of t, and $R_t^\ell \equiv (1+\delta P_t)/P_{t-1}$ their realized return. The budget constraint becomes

$$c_t + b_t = y_t + R_t^\ell\, b_{t-1}$$

This identity holds whether or not prices were anticipated. When they are, $R_t^\ell = R_t$: the maturity of a portfolio does not affect its return.

Price under a rate path. Let $q_{t,t+k} \equiv \prod_{j=1}^{k} R_{t+j}^{-1}$ and $dR_t \equiv R_t - R$. At the end of $t=0$, $P_0 = \sum_{k\geq 0}\delta^k q_{0,k+1}$, and the same argument gives a return R_t from $t=1$ on. To a first order, $dq_{0,t} = -R^{-t-1}\sum_{j=1}^{t} dR_j$. Since $\sum_{k\geq t-1}\delta^k R^{-k-1} = P\,\delta^{t-1}R^{1-t}$,

$$dP_0 = -P\sum_{t\geq 1} R^{-t}\,\delta^{t-1}\,dR_t$$

A holding worth D therefore loses $D\sum_{t\geq 1} R^{-t}\delta^{t-1}dR_t$.

The loss. At the end of $t=-1$, the CB in the QE economy holds $n = Q/P$ long-term bonds. At $t=0$, it receives the coupon n and holds the equivalent of $n\delta$ bonds at price P_0. Without the shock, this position would be worth $n(1+\delta P) = R\,Q$, which is what the one-period bonds of the no-QE economy pay. The loss is therefore

$$L = n\,\delta\,(P - P_0) = Q\,\delta\Big(1 - \frac{P_0}{P}\Big)$$

which gives (9) to a first order. The weight is δ^t rather than δ^{t-1} because the coupon paid at $t=0$ is not affected. At the calibration of Figure 1, the exact loss is 1.6% of GDP, against 1.7% to a first order.

A.2 Aggregation

Individual consumption. Hours are set by unions and utility is separable, so labor does not affect consumption choices. The Euler equation is $c_{s,t+1} = \beta R_{t+1}\, c_{s,t}$. Hence $\sum_{k\geq 0} \omega^k q_{t,t+k} c_{s,t+k} = c_{s,t}/(1-\beta\omega)$, and the lifetime budget constraint gives

$$c_{s,t} = (1-\beta\omega)\Big(\frac{R_t}{\omega}\, a_{s,t-1} + h_t\Big)$$

where human wealth, common to all households, is

$$h_t \equiv \sum_{k\geq 0} \omega^k q_{t,t+k}\big[(1-\tau_{t+k})Y_{t+k} - \sigma\big]$$

At $t=0$, R_0 is replaced by the realized return on the portfolio.

Aggregate budget constraint. Survivors, a fraction ω of each cohort, hold R_t/ω times their previous wealth, and newborns hold nothing. Summing individual budget constraints gives (2).

Aggregate Euler equation. Aggregating consumption gives $C_{t+1} = (1-\beta\omega)(R_{t+1}A_t + h_{t+1})$. Survivors consume $\omega\beta R_{t+1}C_t$ in total, and each newborn consumes $(1-\beta\omega)h_{t+1} = C_{t+1} - (1-\beta\omega)R_{t+1}A_t$. Hence

$$C_{t+1} = \omega\beta R_{t+1}C_t + (1-\omega)\big[C_{t+1} - (1-\beta\omega)R_{t+1}A_t\big]$$

which rearranges to (3). It uses only the return between t and $t+1$, so it holds from $t=0$ on.

Linearization. Consider two economies whose paths of rates differ by $\hat{R}_t$. To a first order, the coefficient on $\hat{R}_{t+1}$ in the difference of (3) is $\beta Y - \psi B/R = Y/R$, since $\psi B = (\beta R - 1)Y$ by (8). Hence

$$\hat{Y}_{t+1} = \beta R\, \hat{Y}_t - \psi\, \hat{A}_t + \frac{Y}{R}\, \hat{R}_{t+1} \tag{37}$$

Calibration. Since ψ depends on β, (8) pins down β given R, ω and B/Y. It also implies $\beta R = 1 + \psi B/Y$, so $\beta < 1$ if and only if $\psi B/Y < r$, which we assume. At the calibration of Figure 1, $\psi B/Y \approx 0.007 < r$.

A.3 Phillips curve

Unions. As in Section 4.3, unions set nominal wages subject to a quadratic adjustment cost with parameter θ_w, face an elasticity of substitution ϵ_w between labor varieties, and discount at β. Let $\mu^w \equiv \epsilon_w/(\epsilon_w - 1)$. This gives the wage Phillips curve (19), with zero inflation target and labor wedge

$$\Xi_t \equiv \nu N_t^{\zeta} - \frac{1}{\mu^w}(1-\tau_t)\, w_t\, C_t^{-1}$$

where $w_t = 1$, $\pi_t = \pi_t^w$ and $N_t = Y_t = C_t$. We set ν so that the wedge is zero in the steady state:

$$\nu Y^{1+\zeta} = \frac{1-\tau}{\mu^w}$$

Linearization. Around the steady state, the left-hand side of (19) is $\pi_t - 1$ and $N_t \Xi_t$ is $Y\, d\Xi_t$. Using $(1-\tau)/(\mu^w Y) = \nu Y^{\zeta}$,

$$d\Xi_t = \nu Y^{\zeta}\left[(1+\zeta)\frac{dY_t}{Y} + \frac{d\tau_t}{1-\tau}\right]$$

With $\kappa \equiv (\epsilon_w/\theta_w)\,\nu Y^{1+\zeta}$, the difference between the two economies satisfies

$$\hat{\pi}_t = \kappa\left[(1+\zeta)\frac{\hat{Y}_t}{Y} + \frac{\hat{\tau}_t}{1-\tau}\right] + \beta\,\hat{\pi}_{t+1} \tag{38}$$

Revenue form. Substituting $\hat{\tau}_t = (\hat{\mathcal{T}}_t - \tau\hat{Y}_t)/Y$ into (38),

$$(1+\zeta)\frac{\hat{Y}_t}{Y} + \frac{\hat{\tau}_t}{1-\tau} = \frac{(1+\zeta)(1-\tau)-\tau}{(1-\tau)\,Y}\left(\hat{Y}_t + \Theta\,\hat{\mathcal{T}}_t\right)$$

which gives (13), with $\tilde{\kappa} \equiv \kappa\,[(1+\zeta)(1-\tau)-\tau]/[(1-\tau)Y]$. With flexible wages, the bracket in (38) is zero, so natural output satisfies $\hat{Y}_t = -\Theta\,\hat{\mathcal{T}}_t$.

A.4 Wealth identity and the counterfactual

Walras' law. By (4) and (6), remittances cancel and

$$B_t - E_t = R_t\,(B_{t-1} - E_{t-1}) - \mathcal{T}_t - \sigma$$

which is (2) if and only if $Y_t = C_t$. At $t = 0$, the same holds with realized returns, since every bond is held either by households or by the CB.

Counterfactual. In an alternative no-QE economy, the CB holds nothing, and households hold one-period bonds instead of reserves. Since reserves and one-period bonds earn the same return, including between $t = -1$ and $t = 0$, both counterfactuals give the same allocation.

A.5 Proofs

Proof of Lemma 1. The loss moves from households to the CB without changing Treasury debt. Before any recapitalization, remittances are identical, so $\hat{E}_0 = -L$, $\hat{A}_0 = L$ and $\hat{B}_0 = 0$. A recapitalization raises B_0 and E_0 by the same amount and leaves $\hat{A}_0$ unchanged. Taxes at $t = 0$ depend on B_{-1}, so $\hat{\mathcal{T}}_0 = 0$. For $t \geq 1$, the two economies face the same R_t, so (2) with $Y_t = C_t$ gives $\hat{A}_t = R_t \hat{A}_{t-1} - \hat{\mathcal{T}}_t$. By (5) and (7), $\hat{\mathcal{T}}_t = \phi \hat{B}_{t-1} = \phi (\hat{A}_{t-1} - \ell_{t-1})$. Hence

$$\hat{A}_t = (R_t - \phi) \hat{A}_{t-1} + \phi \ell_{t-1}$$

Since $(R_t - R) \hat{A}_{t-1}$ is second order, $\hat{A}_t = \mu \hat{A}_{t-1} + \phi \ell_{t-1}$. Iterating from $\hat{A}_0 = L$ gives (10).

Output. Under fixed real rates, (37) becomes $\hat{Y}_{t+1} = \beta R \hat{Y}_t - \psi \hat{A}_t$. Since $\beta R > 1$, this equation has an unstable root, and (11) is its unique bounded solution. The sum converges since $\gamma < 1$ and $\hat{A}_t$ is bounded.

Proof of Proposition 1. Substitute (10) into (11) at $t = 0$. The transfer term gives

$$\psi \sum_{k \geq 0} \gamma^{k+1} \mu^k L = \frac{\psi \gamma}{1 - \gamma \mu} L = \mathcal{M} L$$

For the deferral term, exchanging the sums ($k = j + 1 + m$) gives

$$\psi \phi \sum_{k \geq 1} \gamma^{k+1} \sum_{j=0}^{k-1} \mu^{k-1-j} \ell_j = \frac{\psi \gamma}{1 - \gamma \mu} \phi \sum_{j \geq 0} \gamma^{j+1} \ell_j = \mathcal{M} \phi \Lambda$$

using $\gamma / (1 - \gamma \mu) = 1 / (\beta R - \mu)$.

Proof of Proposition 2. Here $\hat{R}_t$ is the difference in rates between the two economies.

Output. With $\hat{\pi}_t = 0$ for all t, (13) gives $\hat{Y}_t = -\Theta \hat{\mathcal{T}}_t$.

Present value. We compute $\hat{A}_0$ in two ways. From $t = 1$ on, all assets earn R_t in each economy, so iterating $A_t = R_t A_{t-1} - \mathcal{T}_t - \sigma$ forward gives $A_0 = \sum_{t \geq 1} q_{0,t} (\mathcal{T}_t + \sigma)$. Linearizing the difference, with $\mathcal{T} + \sigma = rB$ and $\sum_{t \geq j} R^{-t} = R^{1-j} / r$,

$$\hat{A}_0 = \sum_{t \geq 1} R^{-t} \hat{\mathcal{T}}_t - B \sum_{t \geq 1} R^{-t} \hat{R}_t$$

Let D denote the value at the end of $t = 0$ of the long-term bonds held by households in the QE economy. These households hold the transfer L and bear the revaluation of their long-term bonds caused by the rate differential (Appendix A.1):

$$\hat{A}_0 = L - D \sum_{t \geq 1} R^{-t} \delta^{t-1} \hat{R}_t$$

Equating the two expressions gives

$$\sum_{t\geq 1} R^{-t}\,\hat{\mathcal{T}}_t = L + \Omega$$

where

$$\Omega \equiv \sum_{t\geq 1} R^{-t}\,\hat{R}_t\left(B - \delta^{t-1}D\right) \tag{39}$$

Combined with the first part, this gives (14). The weight $B - \delta^{t-1}D$ is the stock of public liabilities held by households whose payments were not set at $t = 0$. Long-term bonds held since $t = 0$ are excluded, because their loss of value at $t = 0$ offsets the higher returns they earn later.

Real rate. By (37),

$$\hat{R}_{t+1} = \frac{R}{Y}\big(\hat{Y}_{t+1} - \beta R\,\hat{Y}_t + \psi\,\hat{A}_t\big)$$

Through the term $\psi\,\hat{A}_t$, the CB keeps rates higher while households hold public liabilities.

A.6 Output and the timing of taxes

Result. Under fixed real rates, consider any path of additional taxes $\{\hat{\mathcal{T}}_t\}_{t\geq 0}$, including at $t = 0$, such that $R^{-t}\hat{A}_t \to 0$, and let $\tilde{\mathcal{T}}_t \equiv R^{-t}\hat{\mathcal{T}}_t$. To a first order, $\sum_{t\geq 0}\tilde{\mathcal{T}}_t = L$ and

$$\hat{Y}_0 = \frac{\psi}{R(1-\beta)}\sum_{t\geq 0}\left(\beta^{-t} - 1\right)\tilde{\mathcal{T}}_t \tag{40}$$

Since $\beta < 1$ (Appendix A.2), the weight $\beta^{-t} - 1$ is zero at $t = 0$ and increases with t: delaying taxes of given present value raises output.

Proof. To a first order, $\hat{A}_0 = L - \hat{\mathcal{T}}_0$ and $\hat{A}_t = R\,\hat{A}_{t-1} - \hat{\mathcal{T}}_t$. Iterating gives $\hat{A}_t = R^t\big(L - \sum_{s\leq t}\tilde{\mathcal{T}}_s\big)$. Since $R^{-t}\hat{A}_t \to 0$, $\sum_{s\geq 0}\tilde{\mathcal{T}}_s = L$ and $\hat{A}_t = R^t\sum_{s>t}\tilde{\mathcal{T}}_s$. Substituting into (11) at $t = 0$ and exchanging the sums, with $\gamma R = \beta^{-1}$,

$$\hat{Y}_0 = \psi\sum_{s\geq 1}\tilde{\mathcal{T}}_s\sum_{t=0}^{s-1}\gamma^{t+1}R^t = \frac{\psi}{R(1-\beta)}\sum_{s\geq 0}\left(\beta^{-s} - 1\right)\tilde{\mathcal{T}}_s$$

A.7 Inflation under fixed real rates

Result. Under fixed real rates, to a first order, inflation at $t = 0$ is

$$\hat{\pi}_0 = \underbrace{\tilde{\kappa}\,\Theta\,L}_{\text{cost-push}} + \tilde{\kappa}\sum_{t\geq 0}\beta^t\Big[\underbrace{\hat{Y}_t}_{\text{demand}} + \underbrace{\Theta\,(\beta R - 1)\,\hat{A}_t}_{\text{timing}}\Big] \tag{41}$$

The cost-push does not depend on the regime, since the taxes that repay the loss have present value L. Relative to Treasury support, a regime with $\ell_t \geq 0$ for all t raises $\hat{A}_t$ by (10), and hence $\hat{Y}_t$ by (11): deferral raises the demand term. The timing term reflects that wage setters discount future taxes at β, whereas the public sector discounts them at R. Since $\beta R > 1$ (Section 2.1), deferral raises it.

Proof. Under fixed real rates, output is given by (11) for any Θ. Iterating (13) forward gives

$$\hat{\pi}_0 = \tilde{\kappa} \sum_{k \geq 0} \beta^k \big(\hat{Y}_k + \Theta \, \hat{\mathcal{T}}_k \big)$$

Since $\beta^k \hat{\mathcal{T}}_k = (\beta R)^k \tilde{\mathcal{T}}_k$, a summation by parts using Appendix A.6 gives

$$\sum_{k \geq 0} (\beta R)^k \tilde{\mathcal{T}}_k = L + \sum_{t \geq 0} \big[(\beta R)^{t+1} - (\beta R)^t \big] R^{-t} \hat{A}_t = L + (\beta R - 1) \sum_{t \geq 0} \beta^t \hat{A}_t$$

where the boundary term $\beta^t \hat{A}_t$ vanishes. This gives (41).

A.8 Loss-absorption regimes

By (6), $\ell_0 = L + \hat{S}_0$ and, for $t \geq 1$, $\ell_t = R_t \, \ell_{t-1} + \hat{S}_t$.

Treasury support. $\hat{S}_0 = -L$ and $\hat{S}_t = 0$ for $t \geq 1$, so $\ell_t = 0$ for all t.

Deferral. $\hat{S}_0 = 0$ and, for $t \geq 1$,

$$\hat{S}_t = -\min\{\sigma, R_t \, \ell_{t-1}\}$$

While the loss is positive, $\ell_t = R_t \, \ell_{t-1} - \sigma$. To a first order, with $F \equiv \sigma / r$,

$$\ell_t = \max\big\{ F - R^t (F - L), 0 \big\} \tag{42}$$

If $L < F$, the loss is repaid at

$$t^* = \left\lceil \frac{\ln\big(F/(F-L)\big)}{\ln R} \right\rceil$$

which gives $t^* = 20$ at the calibration of Figure 1. If $L \geq F$, the loss is never repaid without Treasury support, and Proposition 1 does not apply.

A.9 Seigniorage and cash

With cash M, which pays no interest, the CB invests M in bonds and earns $\sigma = rM$ in the steady state: the interest forgone by cash holders. The franchise value is $F = M$. In the HANK model, seigniorage is $i_{t-1}M_{t-1}/\pi_t$, which equals rM at zero inflation: it does not require inflation. Including cash in household wealth leaves (7) unchanged, since cash adds the same amount to household wealth and to the bonds held by the CB. Holding cash costs households σ per period, the lump-sum payment in (1).

We hold σ constant after the shock. With $\sigma_t = r_t M$, the tightening that creates the loss also raises the income that repays it: under deferral, $\ell_t = R_t\,\ell_{t-1} - r_t M$, where $(r_t - r)M$ is first order. Repayment is then faster and the deferral effect smaller. Interest-elastic cash demand works in the opposite direction. The HANK model includes both effects.

B Appendix to the data Section 3

B.1 Definitions

All series are annual, in nominal terms, and in percent of nominal GDP of the same year, or of 2025 GDP for cumulative amounts. Bank of England data are by fiscal year, from March to February, assigned to the calendar year in which it starts.

Net income is the result before transfers to or from general risk provisions, before the use of reserves and before taxes. It measures the income generated by the balance sheet, independently of how the central bank absorbs or distributes it. For the Fed, it is net income before remittances. For the Bank of England, it is the net cash transfer of the Asset Purchase Facility (APF) to HM Treasury, which is both its income at historical cost and its remittance.

Remittances are all payments to the State for a given financial year: profit transfers or dividends and, where the central bank pays them, corporate taxes. They are attributed to the financial year that generated them. For the Fed, they exclude transfers of surplus above the statutory cap.

Absorption (Table 1 and Figure 2.b) splits the loss recognised in 2022–2025 into provisions and reserves released, the deferred asset or loss carried forward at the end of 2025, and payments by the Treasury, under an indemnity or through a tax recovery. These are accounting losses, which differ from the market-value loss L.

B.2 Central banks

Federal Reserve. Net income and remittances for 2015–2023 are the preliminary figures of the Board's January press releases; audited figures differ by less than 2%. For 2024 and 2025, we use the operating losses of the audited combined financial statements. Remittances in 2022 are the earnings transferred before their suspension in September. The Federal Reserve Act requires the Reserve Banks to remit their excess earnings to the Treasury; when earnings do not cover expenses, they record a deferred asset and suspend remittances until future earnings repay it. The deferred asset is the negative of the item "Earnings remittances due to the U.S. Treasury" in the H.4.1 release (FRED series RESPPLLOPNWW). It reached $16.6 billion at the end of 2022 and $243.5 billion at the end of 2025.

Deutsche Bundesbank. Data come from the profit and loss account of each Annual Report. Net income is the result before the item "transfers to or from the provision for general risks" and before withdrawals from reserves. Remittances are the profit transferred to the Federal Government. The Bundesbank raised its provision from 2016 to 2021, mainly against interest rate risk from asset purchases, and has transferred no profit since 2020. It released €1.0 billion of provisions in 2022, and the remaining €19.2 billion with €2.4 billion of reserves in 2023. It has carried its losses forward since 2024: the accumulated loss reached €27.8 billion at the end of 2025.

Banque de France. Data come from the presentation of the annual results. Net income is the ordinary result before tax and before movements of the fund for general risks (FRG) and of the gold revaluation reserve (RRRODE). Remittances are the corporate tax, the dividend and transfers related to the over-funding of the pension reserve. Releases from the FRG and the RRRODE covered the 2023 loss in full and most of the 2024 loss; the remaining €7.7 billion was carried forward. It was cleared in 2025 by an exceptional capital gain of €11 billion on the sale of gold held in New York, which we exclude from net income in Figure 2.a.

Banca d'Italia. Data come from the annual financial statements (*Il bilancio di esercizio*). Net income is the gross result (*risultato lordo*) before allocations to the general risk fund and before taxes. Remittances are the share of net profit allocated to the State plus taxes due for the year, when positive. Participants receive €340 million per year, 4.5% of capital. The Banca d'Italia used €5.6 billion of its general risk fund in 2023 and €5.8 billion in 2024, and recognised a tax recovery on its gross losses, €2.3 billion and €2.4 billion, which reduces its future taxes. Net profit remained positive, and the State received €0.6 billion in each year.

Bank of England. The APF is a subsidiary of the Bank, fully indemnified by HM Treasury: gains are owed to the Treasury and losses are borne by it. Since November

2012, its net cash is transferred to or from the Treasury every quarter. Data come from the APF Annual Reports and Quarterly Reports. Cumulative transfers to the Treasury peaked at £123.9 billion in September 2022 and fell to £23.8 billion at the end of 2025: the Treasury paid £100.1 billion to the APF over this period, including losses realised on gilt sales. The net cash transfer differs from net interest income because it includes losses realised at redemption on gilts bought above par and, in 2022/23, a profit of £3.8 billion on the temporary purchases of autumn 2022.

European Central Bank. The ECB released its provision for financial risks, €6.6 billion, in 2023 and has carried forward losses of €1.3 billion in 2023, €7.9 billion in 2024 and €1.3 billion in 2025. Under Article 33 of the Statute of the ESCB, it distributes its profits to the national central banks, which are therefore affected only through forgone distributions. We do not add ECB losses to those of the national central banks.

B.3 GDP

Nominal GDP is from the Bureau of Economic Analysis for the United States (FRED series GDPA) and from Eurostat for Germany, France and Italy (sum of quarterly GDP at current prices). For the United Kingdom, it is from the IMF International Financial Statistics for 2015–2024 and from the Office for National Statistics for 2025.

B.4 Fiscal rules

In the euro area, the debt that enters the excessive deficit procedure is the gross debt of the general government at nominal value. National central banks are outside general government, so their equity does not enter this debt. In the United States, Fed remittances are budget revenue when paid, and the deferred asset does not enter debt held by the public. In the United Kingdom, payments by HM Treasury to the APF under the indemnity raise central government borrowing as they are made.

C Appendix to the quantitative Section 4

C.1 Nominal long-term bonds

The long-term bonds are those of Appendix A.1, now nominal. Their nominal price P_t satisfies (16), and real returns are deflated by gross inflation: R_t for bills, as in (15), and R_t^ℓ for long-term bonds, as in (17). The budget constraint of Appendix A.1 holds with these real returns. In the steady state, $1 + i = (1 + \delta P)/P$. Since $i_{-1} = i$ and

$P_{-1} = P$,

$$R_0 - R_0^\ell = \frac{\delta\,(P - P_0)}{P\,\pi_0}$$

which gives (32). A surprise in π_0 revalues all nominal assets, bills, reserves and cash alike, whatever the composition of portfolios. Only the duration loss depends on who holds the long-term bonds, and hence on QE.

C.2 Wage Phillips curve

Union k sets the nominal wage $W_{k,t}$ of its labor service, and faces the demand

$$N_{k,t} = \left(\frac{W_{k,t}}{W_t}\right)^{-\epsilon_w} N_t.$$

As in Auclert, Rognlie and Straub (2024), the union maximizes the average utility of households, taking aggregates as given. Each household supplies $N_{k,t}$ hours to union k. It values the after-tax income $(1-\tau_t)\,z_{it}W_{k,t}N_{k,t}/P_t$ at its marginal utility $c_{it}^{-\eta}$, and these hours raise its disutility of work at the margin νN_t^ζ. Wage inflation $\pi^w_{k,t} \equiv W_{k,t}/W_{k,t-1}$ has a quadratic utility cost. The union solves

$$\max_{\{W_{k,t}\}} \sum_{t\geq 0} \beta^t \left[(1-\tau_t)\,\frac{W_{k,t}}{P_t}\,N_{k,t}\,MU_t - \nu N_t^\zeta N_{k,t} - \frac{\theta_w}{2}\left(\frac{\pi^w_{k,t}}{\pi} - 1\right)^2\right],$$

where $MU_t \equiv \int z_{it}\,c_{it}^{-\eta}\,di$ is the marginal utility of consumption, weighted by productivity because hours are rationed equally and labor income is proportional to z_{it}.

The wage $W_{k,t}$ affects $N_{k,t}$, with $\partial N_{k,t}/\partial W_{k,t} = -\epsilon_w N_{k,t}/W_{k,t}$, and the adjustment costs at t and $t+1$. The first-order condition, multiplied by $W_{k,t}$, is

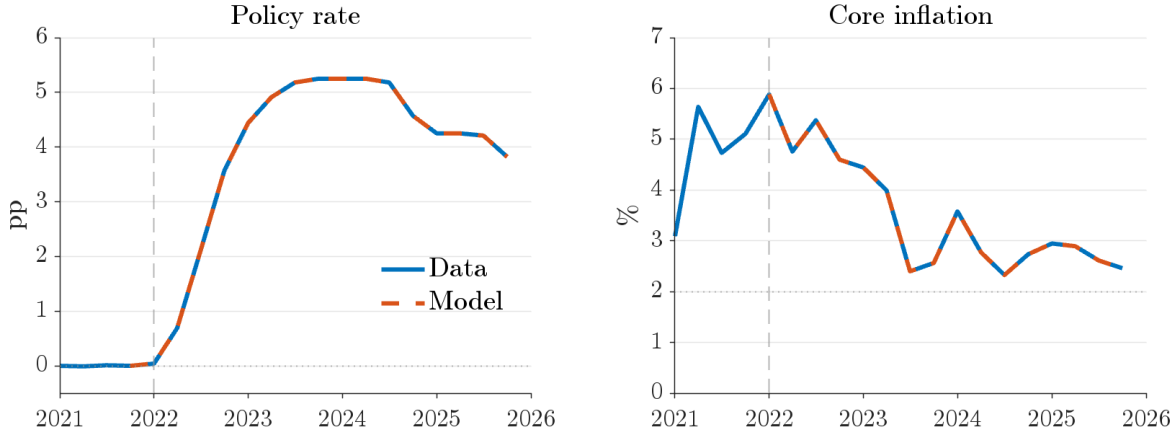
$$(1-\epsilon_w)(1-\tau_t)\,\frac{W_{k,t}}{P_t}\,N_{k,t}\,MU_t + \epsilon_w\,\nu N_t^\zeta N_{k,t} - \theta_w\,\frac{\pi^w_{k,t}}{\pi}\left(\frac{\pi^w_{k,t}}{\pi} - 1\right) + \beta\,\theta_w\,\frac{\pi^w_{k,t+1}}{\pi}\left(\frac{\pi^w_{k,t+1}}{\pi} - 1\right) = 0.$$

In a symmetric equilibrium, $W_{k,t} = W_t$ and $N_{k,t} = N_t$, and $w_t = W_t/P_t$. The first two terms are then equal to $\epsilon_w N_t\,\Xi_t$, with Ξ_t defined in (20). Dividing by θ_w gives (19). Since the adjustment cost enters utility, it does not appear in the household budget constraint.

D Additional results for Section 5

D.1 Experiment and calibration of the shocks

Figure 7: Fit of the calibrated shocks

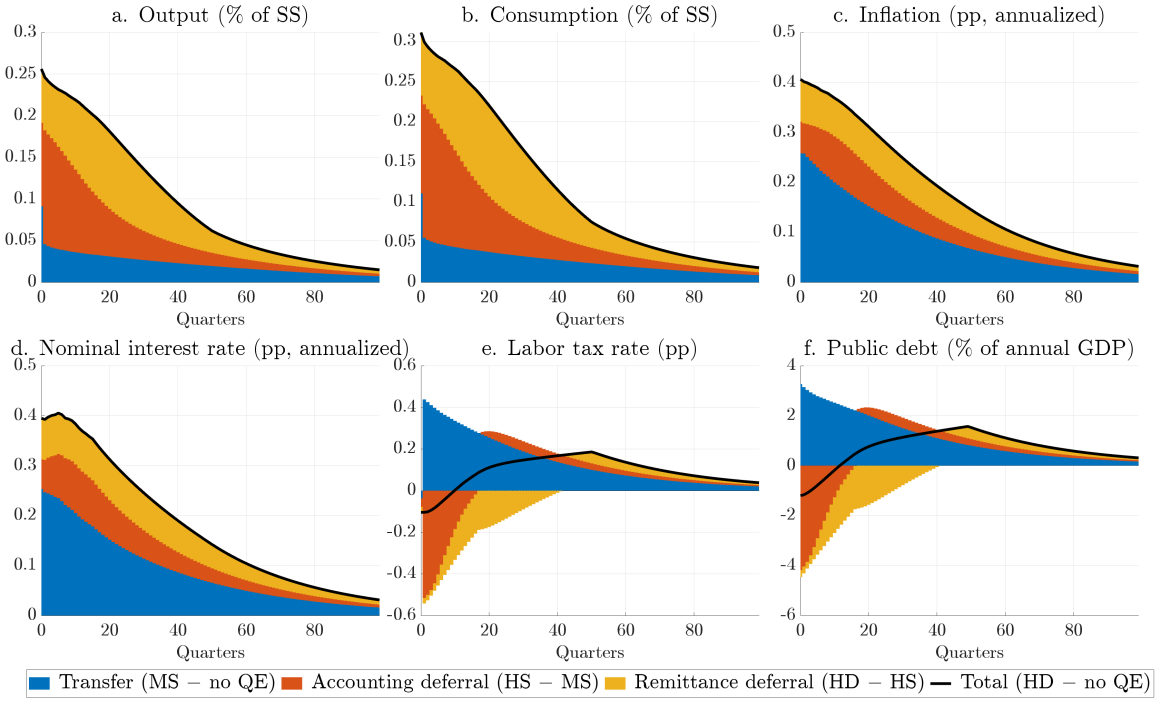


Notes: Data: FRED, quarterly averages. Policy rate: change in the effective federal funds rate (FEDFUNDS) relative to 2021Q4. Core inflation: core PCE price index (PCEPILFE), quarter-on-quarter, annualized. Vertical line: impact quarter, 2022Q1.

D.2 Fixed real rates and representative-agent benchmarks

Figure 8 shows decomposition (35) when the QE economies face the real rates of the no-QE economy.

Figure 8: Effects of CB losses at fixed real rates



TANK and RANK. In the TANK, a share $\lambda^{\mathrm{H}} = 17.8\%$ of households, the share of HANK households at the borrowing constraint, consume their after-tax income. The others hold all liquid wealth and follow an Euler equation. The RANK sets $\lambda^{\mathrm{H}} = 0$. Everything else is as in the HANK: firms, the CB, the Treasury, the fiscal and Taylor rules, the slope of the wage Phillips curve and the shocks. We set $\beta = 1/R$ and recalibrate ν so that $N = 1$. Since the wage Phillips curve is flat, inflation in the no-QE economy is driven by the cost-push shock and is almost the same in the three models. Figures 9 and 10 show the total effect of QE, HD minus no QE, under fixed real rates and under the Taylor rule. Public debt is in % of annual GDP, and rates are annualized.

Figure 9: Total effect of the CB loss at fixed real rates

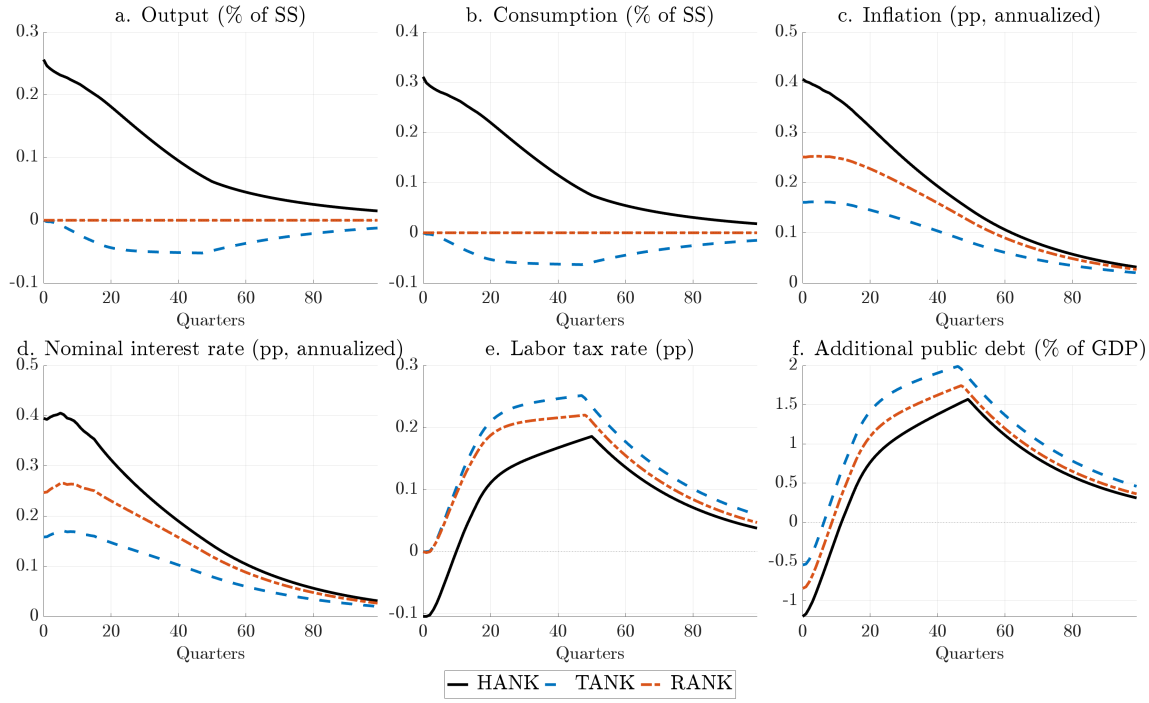
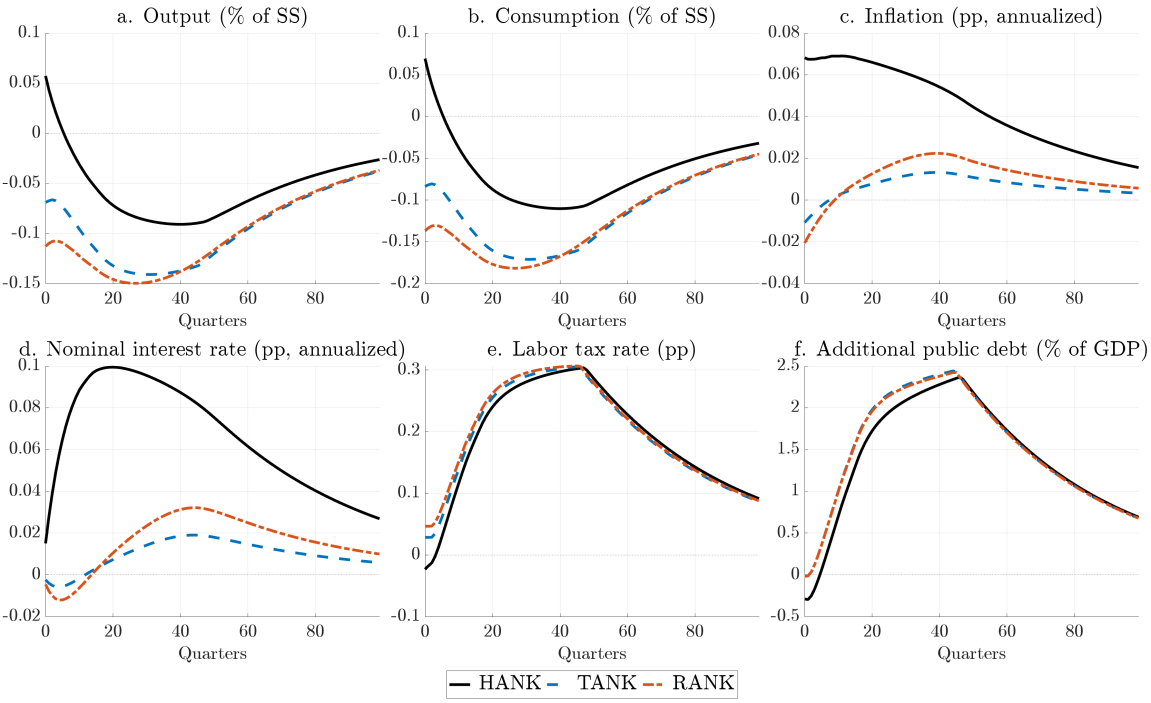


Figure 10: Total effect of the CB loss under the Taylor rule



Results. At fixed real rates, output does not move in the RANK, where the revaluation gain exactly offsets the present value of taxes. It falls by up to 0.05% in the TANK, and rises by 0.25% on impact in the HANK. Under the Taylor rule, the loss is contractionary in the medium run in the three models, but less so in the HANK. Output falls by 0.11% on impact and 0.15% at the trough in the RANK, and by 0.07% and 0.14% in the TANK. In the HANK, it rises by 0.06% on impact and falls by 0.09% at the trough. Under the Taylor rule, the paths of the labor tax and public debt are almost the same across models, so the differences come from the household block.

At fixed real rates, the TANK does not lie between the RANK and the HANK. Savers receive the whole gain L but pay only a share $1 - \lambda^{\mathrm{H}}$ of the taxes, and spend their net gain $\lambda^{\mathrm{H}} L$ as an annuity, whereas hand-to-mouth households cut consumption when taxes are levied. The TANK thus isolates incidence, which lowers demand, while the HANK adds a wealth effect, which raises it. Lower demand in the TANK also lowers inflation, which erodes less public debt at $t = 0$ and raises the labor tax.

D.3 Robustness: lump-sum transfers

Figure 11 reports decomposition (35) when the fiscal adjustment goes through lump-sum transfers instead of the labor tax. The tax rate stays at its steady-state value, $\tau_t = \tau$, and transfers adjust so that revenue net of transfers follows the same rule as in the baseline:

$$T_t = T + (\mathcal{T}_t - \mathcal{T}) - \phi \left(B^F_{t-1} - B \right) \tag{43}$$

where $\mathcal{T}_t = \tau w_t N_t$ is labor-tax revenue. Only the instrument changes. The fiscal cost-push of Section 2.5 disappears, $\Theta = 0$. The burden of repayment is also shared equally across households, instead of in proportion to labor income, which makes its incidence more regressive. The shocks $\{u_t, \varepsilon_t\}$ are those of the baseline and are not recalibrated, so the QE economy under HD no longer exactly reproduces the data.

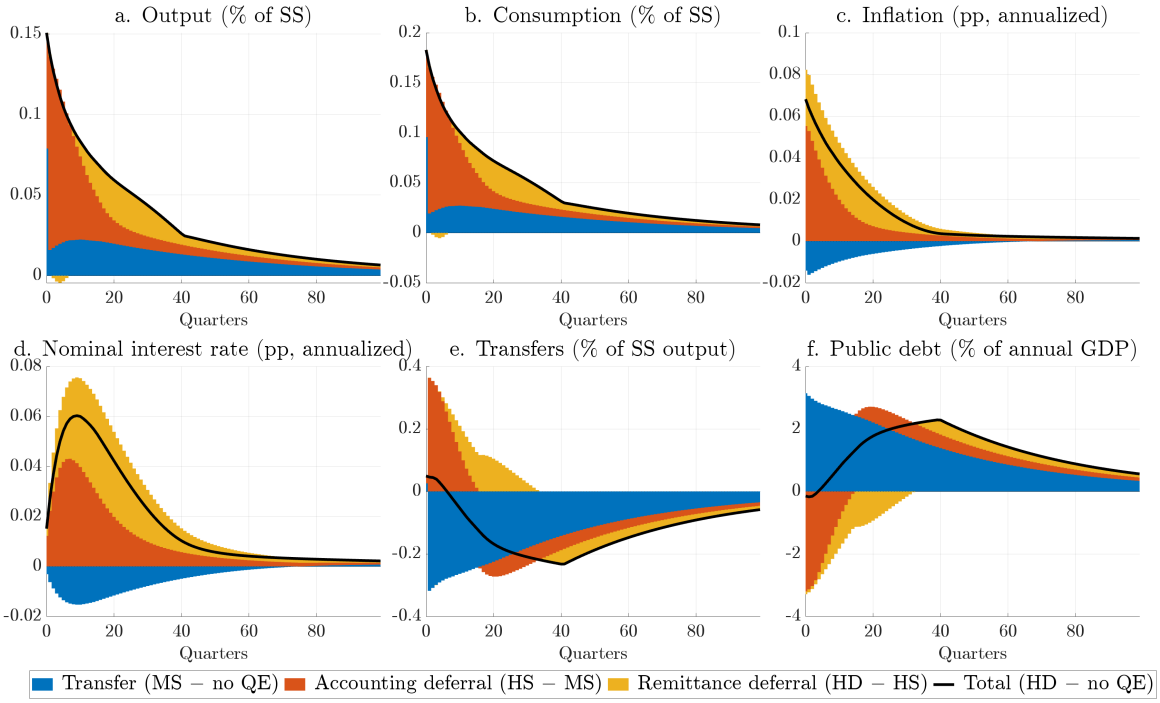
Total effect. With lump-sum transfers, CB losses are expansionary at all horizons. Under HD, output rises by 0.15% on impact, against 0.06% in the baseline. It stays above its no-QE path throughout, whereas in the baseline it falls to -0.09% around quarter 40. Inflation rises by 0.07 pp on impact, as in the baseline, but returns to its no-QE path after about 40 quarters instead of staying higher for decades. The policy rate rises by at most 6 bp, against 10 bp. Public debt follows a path close to the baseline: it rises until remittances resume, to 2.3% of annual GDP, and transfers then fall by up to 0.23% of quarterly output. In present value, discounted at the steady-state real rate, output rises by $0.28\,L$, or 0.9% of annual GDP, against a fall of $0.4\,L$ in the baseline. The price level ends 0.30% higher, against 1.25%. The loss is smaller, $L = 3.2\%$ of

annual GDP against 3.8%, because the policy rate rises less, and the CB repays it after 40 quarters, against 47.

Decomposition. The transfer component changes sign. Under MS, it raises output by 0.08% on impact and by about 0.02% over the following decades, against a fall of up to 0.12% in the baseline. It also slightly lowers inflation, by 0.015 pp on impact, so the policy rate falls. Lower transfers raise the marginal utility of consumption, especially for low-income households, which lowers the labor wedge (20) and wage inflation. Both deferrals remain expansionary on impact and add 0.07% to output. They are inflationary from impact on, since postponing taxes no longer postpones a cost-push.

Interpretation. Under a Taylor rule, output lies between demand-determined output and natural output. With $\Theta = 0$, natural output does not move, and the demand effects of Section 5.4 raise output. The contraction in the baseline therefore comes from the distortionary financing of the loss, not from the loss itself. It does not come from its incidence either: lump-sum financing is more regressive, yet output rises.

Figure 11: Effects of CB losses with lump-sum transfers: total effect and decomposition



Notes: Difference with the no-QE economy and decomposition (35), Taylor rule, baseline shocks. The labor tax rate is fixed at its steady-state value and lump-sum transfers adjust. Transfers in % of quarterly steady-state output, public debt in % of annual GDP, rates annualized.

D.4 Robustness: Output gap in Taylor rule

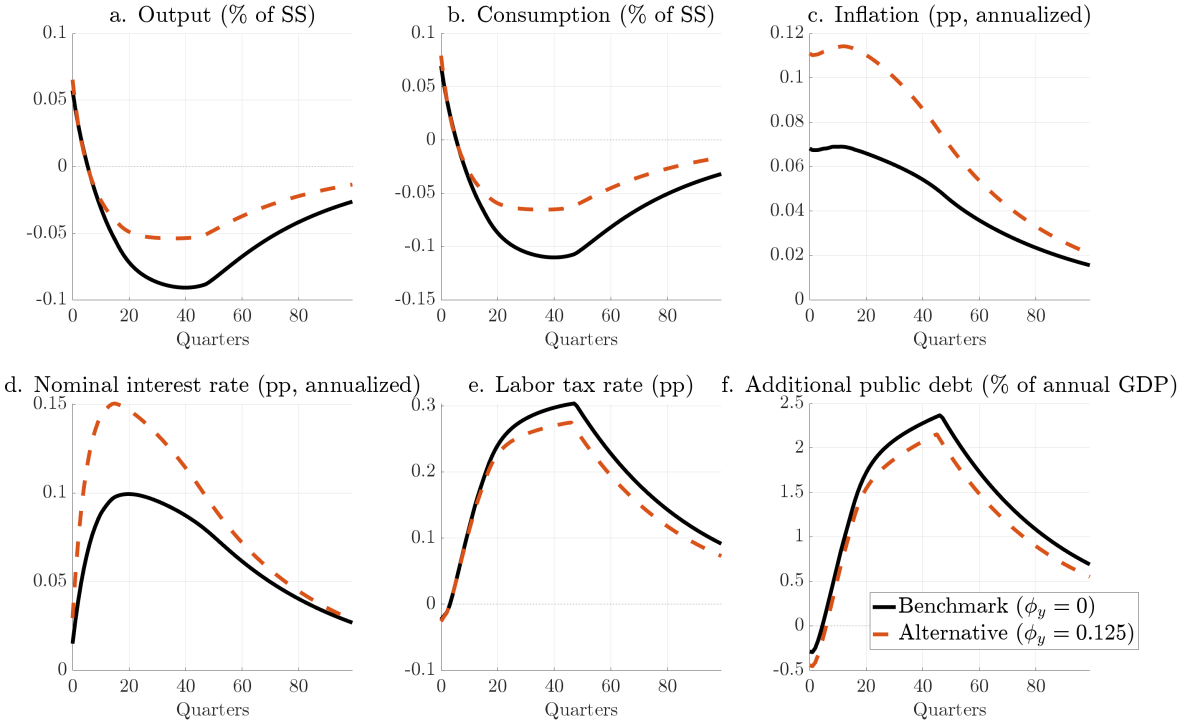
Figure 12 shows the total effect under HD with the Taylor rule

$$i_t = \rho_i\, i_{t-1} + (1-\rho_i)\big[i + \phi_\pi(\pi_t - \pi) + \phi_y(Y_t - Y)\big] + \varepsilon_t$$

In the *benchmark* scenario of the main text, $\phi_y = 0$. Here, we consider an *alternative* calibration with $\phi_y = 0.5/4 = 0.125$, since the model is quarterly. We recalibrate $\{u_t, \varepsilon_t\}$ to match the same data under this rule.

For given inflation, the output-gap term lowers the policy rate. The resulting rise in inflation raises it through ϕ_π: on net, the nominal rate is higher than in the benchmark, while the real rate is lower. The recession is therefore milder: at the peak, output and consumption fall by about 40% less. Higher inflation and output also reduce the increase in public debt and in the labor tax rate.

Figure 12: Total effect of CB losses with different Taylor rules



D.5 Optimal policy with the welfare-based weight

The weight. With sticky wages and flexible prices, a second-order approximation of household welfare around an efficient steady state gives

$$\mathcal{W} \approx -\frac{\epsilon_w}{2\kappa_Y} \sum_{t\geq 0} \beta^t \left(\pi_t^2 + \frac{\kappa_Y}{\epsilon_w} y_t^2\right) \tag{44}$$

where π_t is quarterly wage inflation, equal to price inflation in our model since $w = 1$. The relative weight on output is the slope of the wage Phillips curve with respect to output divided by the elasticity of substitution between labor services. With annualized inflation, $\pi_t^2 = (\pi_t^a)^2/16$, so the weight in (36) is $\lambda = 16\kappa_Y/\epsilon_w \approx 0.016$. This weight is only approximate in our model, for two reasons. The steady state is distorted by the labor tax and the wage markup, which adds linear terms to the approximation, cf. Benigno and Woodford (2005). Utilitarian welfare also depends on consumption dispersion across households, which CB losses affect through the transfer and the tax.

Results. Table 5 reports the results. Under the Taylor rule, CB losses raise $\mathcal{L}$ by 3.2 to 3.7, almost entirely through inflation. Under optimal policy, the increase is close to zero with historical-cost accounting and 0.7 under MS, where it comes from inflation. The cost has not disappeared. It has moved to output, which falls almost as much as under the rule. Optimal policy also halves L.

Utilitarian welfare gives a different picture of the regimes. In consumption equivalent, CB losses cost about 0.35% of annual GDP in present value in every regime, MS included, against 0.5 to 0.9% under the rule. This cost is twice as large as with $\lambda = 1$ (Table 4). In our model, the weight implied by the textbook approximation therefore does not deliver higher welfare, which is why we use $\lambda = 1$ as the baseline.

Because output enters the loss with a small weight, the problem is nearly flat in directions that move output but not inflation. Output paths are weakly identified, and we do not show them. The values in Table 5 are nonetheless stable across starting points.

Table 5: Cost of CB losses with the welfare-based weight

	Taylor rule	Optimal policy			
	$\Delta\mathbb{L}$	$\Delta\mathbb{L}$	CE	ΔY	L
MS	3.70	0.71	−0.35	−0.79	1.93
HS	3.33	0.05	−0.35	−1.11	1.98
HD	3.25	0.04	−0.37	−1.15	1.98

Notes: $\lambda = 0.016$. Differences with the no-QE economy under the same policy. $\Delta\mathbb{L}$: increase in the loss (9). CE: change in utilitarian welfare in consumption equivalent, present value, % of annual GDP. ΔY: present value of the change in output, % of annual GDP. L: CB loss, % of annual GDP. Under the Taylor rule, CE, ΔY and L are as in Table 4.