Bank of England

A welfare analysis of the central bank balance sheet

Staff Working Paper No. 1,202
August 2026

William Pagel

Bank of England

Staff Working Paper No. 1,202

A welfare analysis of the central bank balance sheet

William Pagel[1]

Abstract

What is the socially optimal long-run size of the central bank balance sheet, once interest rates are away from the effective lower bound and the balance sheet is no longer needed for monetary stimulus? I introduce a central bank into a model in which banks are liquidity mismatched and prone to sudden 'bank runs'. Supplying central bank reserves reduces the likelihood and severity of runs, but comes at a cost: constraints on the set of securities the central bank can hold mean a larger balance sheet crowds out private investment and misallocates capital. I calibrate the model to pre-financial crisis conditions and empirical estimates of the non-linear reserve demand curve, and compute optimal policy. Under full information, it is optimal to expand reserve supply to the point where reserve demand is satiated, but no further. However, because under-supplying reserves is more costly than over-supplying them, robustness to parameter uncertainty calls for an additional buffer in reserve supply. But even for high degrees of robustness, the buffer needed is no larger than two and a half percentage points of bank assets. The policy prescription remains restrained: robustness moves the balance sheet modestly beyond satiation, but fails to justify an open-ended provision of abundant reserves.

Key words: Reserve supply, central bank reserves, optimal policy, bank runs.

JEL classification: E44, E52, E58, G21.

---

(1) University of Oxford and Bank of England. Email: william.pagel@bankofengland.co.uk

The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees. I would like to thank Fabrice Collard, Andrea Ferrero, Giovanni Lombardo, Pablo Ottonello, Ricardo Reis, Francesco Zanetti and participants at the Bank of England Agenda for Research Conference, the European Economic Association Congress 2026, and ASSA 2026 for their helpful comments on this work.

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

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

ISSN 1749-9135 (on-line)

1 Introduction

During both the great financial crisis and the Covid-19 pandemic, central banks around the world purchased large volumes of assets from the private sector, financed by the creation of new central bank reserves. These purchases drove reserve supply far beyond the amount demanded by the banking sector for its daily liquidity needs (Afonso et al. 2024). The expansion in central bank balance sheets was a by-product of the unconventional monetary policy measures taken when interest rates were constrained by the effective lower bound. But as short-term rates began to rise in 2022, central banks no longer needed to maintain large balance sheets to set the stance of monetary policy. Consequently, they began to scale back their asset holdings (through a mixture of sales and allowing assets to mature) and, with them, the level of reserve supply (Figure 1). The open question now facing policymakers is where the process of balance sheet normalisation should stop.

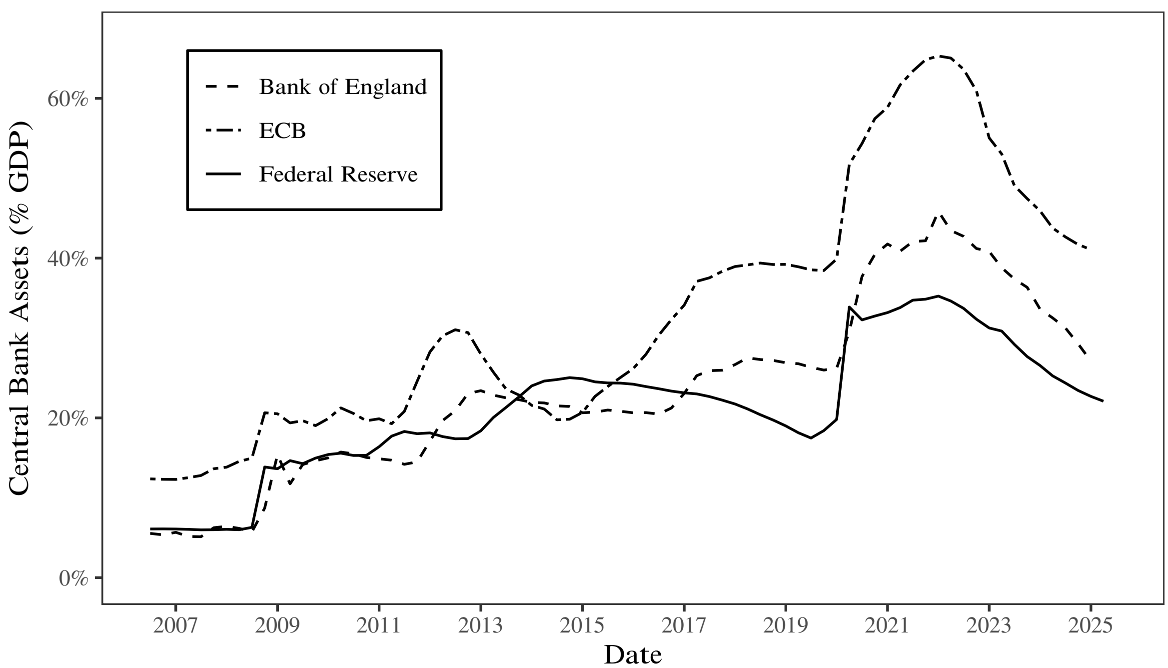

Figure 1: *Central Bank Assets as a Percentage of GDP.* Source: FRED; ECB; Bank of England.

Proponents of large central bank balance sheets argue that the expansive provision of liquidity improves financial stability, alleviating an inefficiently low supply of ultra-liquid assets which only the central bank can provide (Greenwood, Hanson, and Stein 2018; Arce et al. 2020; Eren, Jackson, and Lombardo 2024; Sack and Gagnon 2019; Karadi, Thaler, and

Tristani 2024). Critics argue that distortions created by a large central bank footprint in private markets can outweigh these benefits, at least for some levels of reserve supply (Borio 2023; Vissing-Jørgensen 2023; Bindseil 2016; Acharya and Rajan 2024). This paper builds a unified, general equilibrium welfare framework in which both the liquidity benefits of reserves and the cost of supplying them arise endogenously.

The modelling environment centres on a banking sector exposed to endogenous liquidity panics (as in Gertler and Kiyotaki (2015) and Gertler, Kiyotaki, and Prestipino (2020)[1]). In this setting, a central bank can improve welfare by supplying reserves — a perfectly liquid, risk-free asset not otherwise available to the financial sector. Increasing reserve supply reduces the likelihood and severity of financial crises by ensuring the banking sector is better equipped to meet rapid liquidity outflows. However, this liquidity insurance comes at a cost. Central banks must hold assets to back their reserve liabilities. But constraints on the central bank's asset mandate mean it must concentrate its portfolio in a narrow set of securities, artificially depressing their yields and distorting the allocation of capital across sectors. In the long run, this capital misallocation reduces output. The cost of expanding the central bank balance sheet therefore emerges endogenously, whereas existing work has largely treated it in reduced form.[2] The optimal level of reserve supply trades off liquidity benefits against the costs of capital misallocation.

The paper makes three contributions to this debate.

1. *Under full information, ample reserve supply is optimal*

The model characterises the socially optimal long-run reserve supply. I match the model to the empirical evidence on the non-linear reserve demand curve, to quantify the costs and benefits of reserve supply and map the empirical evidence into a normative assessment. In this model, under full information, the welfare-maximising balance sheet supplies reserves up to the point where reserve demand is close to satiation. That implies an "ample" supply regime, close to the kink in the reserve demand curve, where the marginal liquidity value of additional reserves falls close to zero. Here, reserves are roughly 15% of banking sector assets, comparable to the level of reserve supply in 2023-24. Beyond this point, further expansion of reserve supply is welfare-reducing even if the marginal cost of doing so is close to zero.

[1] These models capture banking instability through both financial accelerator effects and liquidity mismatch in the spirit of D. Diamond and Dybvig (1983).

[2] See for instance Gertler, Kiyotaki, and Prestipino (2016); Afonso et al. (2023) and Eren, Jackson, and Lombardo (2024).

2. Robustness to uncertainty requires only a modest increase in reserve supply

Relaxing the assumption of full information, I consider a policymaker who acknowledges that the data cannot pin down the model parameters with certainty. I assess the robustness of different reserve policies across a set of plausible parametrisations (in the spirit of the robust policy literature (Hansen and Sargent 2008)), and estimate the cost of achieving robustness. I find that under-supplying reserves is on average more costly than over-supplying them.[3] As a result, robustness implies supplying more reserves, not fewer. Rules that insure against deeper tails of adverse risks prescribe larger buffers over the optimal policy under full information. However, achieving a high degree of robustness requires only a relatively small increase in reserve supply: no more than around two and a half percentage points of bank assets. Even the most risk-averse rules therefore prescribe restraint: a small buffer above the ample optimum, but not a return to open-ended, abundant reserve supply.

3. The cost of reserve supply depends on the composition of central bank assets

I develop a new theory of the cost of expanding the central bank balance sheet, centred on capital misallocation. I show that the cost of reserve supply depends on the concentration of the central bank's asset portfolio used to back reserves; not on its size alone. For a given stock of assets, broadening the set of securities the central bank is permitted to hold lowers the dispersion in marginal products across asset classes, and reduces the output loss from capital misallocation. Policies which seek to widen eligibility for the central bank balance sheet, and reduce the distortions created by the central bank's footprint in private markets, are therefore welfare improving.

Related literature. The normative exercise at the heart of this paper relates most closely to a growing literature on the optimal size of the central bank balance sheet. Vissing-Jørgensen (2023) characterises the optimum from the perspective of maximising the total surplus created by the availability of liquid, safe ('convenient') assets. That framework emphasises that the convenience provided by reserves should be offset against the convenience of the assets central banks must hold to supply them. I build on this work in two ways. First, both sides of the trade-off arise endogenously in my framework, such that both are quantified directly in terms of household welfare: the convenience of reserves is micro-founded in the reduction of the risk of systemic liquidity crises which destroy consumption, while the social cost of expanding the central bank arises from capital misallocation which compresses output.

[3] Liquidity crunches such as the September 2019 repo turmoil support the notion that misjudging where reserve scarcity begins can be very costly (Copeland, Duffie, and Yang 2025).

The cost channel in this paper is consistent with Vissing-Jørgensen's argument that large central bank balance sheets drive an inefficient scarcity in the net supply of the assets they hold. However, in my framework, that scarcity is not costly *per se*. Rather, it pushes private markets to lend at sub-optimally low rates, distorting the economy away from the efficient allocation of capital, consistent with recent empirical evidence (Acharya et al. 2025).[4] Second, with a structural model, I can extend the analysis to consider robust supply policies: in addition to the benchmark which takes the reserve demand curve as known, I characterise the supply a central bank should choose when it is uncertain about the true parameters of the economy. I show that, depending on the policymaker's attitude to risk, robustness can imply policy rules which do not necessarily maximise total convenience, expanding reserve supply beyond the point of satiating reserve demand.

In more quantitative macroeconomic models, Eren, Jackson, and Lombardo (2024) study the size and composition of the balance sheet that allocates duration risk optimally in a model with financial frictions, but abstract from the details of reserve demand: reserve demand enters in reduced form, through the effect of reserve holdings on the incentive constraint, rather than through a non-linear demand schedule with a satiation point. In the absence of costs related to the size of the balance sheet, the optimum in their framework is a corner: the central bank would optimally hold all eligible assets. Arce et al. (2020) study how reserves alleviate interbank frictions and find that steady-state welfare is monotonically increasing in balance sheet size, even when the central bank must fund itself with distortionary taxes. Karadi, Thaler, and Tristani (2024) argue for supplying reserves beyond banks' aggregate liquidity needs, but suggest limits to optimal reserve supply are driven by off-model costs such as political constraints and exposure to financial losses. By contrast, in my framework, decreasing marginal benefits and increasing marginal costs arise endogenously. Both forces are quantified within a single welfare framework. As a result, the optimal policy has an interior solution, at which the central bank holds a non-zero asset portfolio but does not hold the maximum eligible amount, without assuming the central bank faces any cost the private sector does not.[5]

[4] W. Diamond, Jiang, and Ma (2024) estimate a liability-side cost: in a structural model of the US banking system, reserves raise banks' marginal cost of lending, so that reserves injected by QE crowd out lending to firms. Their channel and mine, though different in mechanism, both imply a marginal cost of reserve supply that does not vanish once demand is satiated.

[5] Alexandrov and Brunnermeier (2026) study balance sheet and interest rate policy jointly, but from a different perspective: the balance sheet serves to allocate interest-rate risk between the central bank and the private sector, and deliberately plays no liquidity role. Here, by contrast, the long-run level of reserve supply sets the socially efficient level of liquidity risk, disciplined by evidence on estimated demand for liquidity.

In my framework, the optimal level of reserve supply can be estimated by fitting the model to the empirical reserve demand curve (Afonso et al. 2024; Lopez-Salido and Vissing-Jørgensen 2025; Lagos and Navarro 2026) and related empirical targets. I take the estimates in Afonso et al. (2024) as a key empirical target. The model reproduces the non-linear features of the empirical reserve demand curve: the demand curve is steeper and more volatile when reserves are scarce; flatter and more stable when supply is greater. Matching the model to this data uses the estimated demand curve to generate a quantitative statement about optimal policy.

The policy robustness exercise in this paper relates to a large literature on the supply of reserves under uncertainty. This topic has been studied since Poole (1968), who emphasises how payment-flow uncertainty generates a precautionary demand for reserves. Most closely related to my paper, Afonso et al. (2023) characterise the optimal supply of reserves when the parameters of the reserve demand curve are uncertain. In their analytically tractable framework, uncertainty about the location and shape of the reserve demand curve raises optimal supply relative to an exogenous target level of supply in the policymaker's loss function. This logic underpins the Federal Reserve's reserve supply framework, which sets supply with a buffer against unpredictable shifts in reserve supply and demand (Ihrig, Senyuz, and Weinbach 2020). This paper reaches a similar qualitative prescription in a general equilibrium environment. Rather than assuming a loss function with exogenous targets and penalty weights, my approach micro-founds the policymaker's target level of reserves under full information and explicitly models the (possibly asymmetric) costs of deviating from it. Using this framework, I consider a setting where the policymaker is uncertain not only about the state of reserve demand, but about all the model's parameters governing the costs and the benefits of reserve supply. I find that robust policy rules expand reserve supply above the full information benchmark, in line with existing literature. I build on this result by quantifying both how large a buffer is needed to ensure robustness, and how costly such a buffer would be, in terms of household welfare. The results suggest the quantitative impact of uncertainty on optimal policy is small: even a modest buffer over the full information case provides a high degree of robustness to parameter uncertainty.

Finally, I note two areas that are beyond the scope of my paper. First, I do not consider the effect of supplying reserves on the implementation of monetary policy.[6] Equally, my

[6] Bianchi and Bigio (2022) is the now canonical framework for analysing how reserve supply interacts with the implementation and transmission of monetary policy. For policy-focussed summaries of this topic and its interaction with different balance sheet regimes, see Borio (2023) and Schnabel (2023).

framework abstracts from the political economy implications of central bank balance sheet risk and its implications for central bank independence from the fiscal authority.[7]

Outline. Section 2 briefly sets out the banking sector, household problem, and the risk of financial panics, focussing on how these factors provide a micro-foundation for the convenience benefits of reserves. Section 3 introduces the production side of the economy and develops the theory for how the central bank can generate capital misallocation: the cost of reserve supply. Section 4 characterises the optimal size of the steady-state balance sheet analytically. Section 5 fits the model to data on reserve demand and presents quantitative welfare results. Section 6 investigates reserve policy which is robust to policymaker uncertainty about the true parameters of the model. Section 7 discusses policy implications, and concludes.

2 Households, banks and the central bank

The household and financial blocks of the model are broadly as in Gertler and Kiyotaki (2015), with the addition of a central bank that issues reserves. Importantly, agents internalise the liquidity benefits of reserves, generating a reserve demand curve that reproduces salient features of the empirical reserve demand literature.

A continuum of identical households of measure 1 populate the economy. Households can save by holding financial assets directly ("capital") or bank deposits. Banks raise funds by issuing deposits, in order to invest in capital and central bank reserves. Banks are subject to a moral hazard friction and the possibility of 'bank runs'. The central bank issues reserves, and backs these liabilities by purchasing capital.

2.1 Capital

The total capital stock is fixed (and normalised to one) and does not depreciate. Households, banks and the central bank can all hold capital, so that

$$K_t^h + K_t^b + K_t^{cb} = 1. \tag{1}$$

A bank that intermediates capital in period t earns $Z_{t+1}K_t^b$ in the following period. Households are less efficient than banks at intermediating capital and must pay an additional 'management

[7] Hall and Reis (2015) and Del Negro and Sims (2015) consider the risk that central banks become insolvent as a result of losses on their expanded asset portfolios, and the implications of such risks for inflation, monetary policy, and the relationship between the central bank and the fiscal authority.

cost' $F(K_t^h) = \alpha^h (K_t^h)^2/2$ upfront, denominated in units of the consumption good.

2.2 Households

Households maximize

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s u(C_{t+s}),$$

where $\beta \in (0,1)$ is the individual discount factor and $u(C) = \ln(C)$, subject to the budget constraint

$$C_t + D_t + Q_t K_t^h + F(K_t^h) \leq (Q_t + Z_t) K_{t-1}^h + W^h + R_t^d D_{t-1} + \Pi_t + \tau_t.$$

The price of the overall consumption basket C_t is the numeraire, D_t are deposits that pay the gross return R_t^d, Q_t is the market price of capital K_t, Z_t its aggregate dividend, W^h an endowment, and Π_t is the sum of dividends from exiting bankers (Π_t^b) and producing firms (Π_t^f). τ_t is a lump-sum rebate of central bank profits or losses to households.

For banks, deposits are one-period debt. With no bank runs, depositors receive the gross return $\bar{R}_{t+1}$. In the event of a run, banks are fully liquidated and all depositors face the same haircut, receiving $x_{t+1}\bar{R}_{t+1}$, where $x_{t+1} \in [0,1]$ is the recovery rate. The gross deposit rate that households receive is therefore

$$R_{t+1}^d = \begin{cases} \bar{R}_{t+1} \text{ if no bank run;} \\ x_{t+1}\bar{R}_{t+1} \text{ if bank run.} \end{cases}$$

The first-order condition for deposits is

$$1 = \bar{R}_{t+1} \mathbb{E}_t \left[(1-p_t)\Lambda_{t,t+1} + p_t \Lambda_{t,t+1}^* x_{t+1} \right], \tag{2}$$

where X_t^* denotes the value of variable X_t in the equilibrium with a bank run. Similarly, the first-order condition for capital is

$$1 = \mathbb{E}_t \left[(1-p_t)\Lambda_{t,t+1} \left(\frac{Q_{t+1} + Z_{t+1}}{Q_t + \alpha^h K_t^h} \right) + p_t \Lambda_{t,t+1}^* \left(\frac{Q_{t+1}^* + Z_{t+1}}{Q_t + \alpha^h K_t^h} \right) \right]. \tag{3}$$

where $\Lambda_{t,t+1} = \beta \frac{C_t}{C_{t+1}}$ is the stochastic discount factor implied by log utility.

For reference, the unconditional risk-free rate, accounting for the uncertainty from bank runs, is given by:

$$R^f_{t+1} = \Big[(1-p_t)\mathbb{E}_t\Lambda_{t,t+1} + p_t\mathbb{E}_t\Lambda^*_{t,t+1}\Big]^{-1} \tag{4}$$

2.3 Banks

A fraction $f \in (0,1)$ of members of the household become bankers in every period and receive some funds to start their operations. Over time, a bank remains in the market with probability σ and exits with the complementary probability, so that, on average, a bank operates for $(1-\sigma)^{-1}$ periods. When a bank exits, its net worth goes back to the household. The fraction of banks that exit each period is replaced by an equal mass so that the size of the banking sector remains constant.

The balance sheet of a bank that continues its operations at time t is

$$Q_t k^b_t + m_t = d_t + n_t,$$

where net worth is

$$n_t = (Q_t + Z_t)k^b_{t-1} + R^m_t m_{t-1} - R^d_t d_{t-1}.$$

Using the balance sheet constraint to substitute for deposits in the previous expression gives the law of motion of net worth

$$n_t = (R^k_t - R^d_t)Q_{t-1}k^b_{t-1} + (R^m_t - R^d_t)m_{t-1} + R^d_t n_{t-1}, \tag{5}$$

where the return on capital is $R^k_t \equiv \frac{Q_t+Z_t}{Q_{t-1}}$. The first term of (5) captures the spread that banks earn from intermediating capital relative to the cost of funds. The second term measures the spread on reserves.

Banking activities suffer a moral hazard problem. In every period, bankers can abscond with a fraction θ of the value of their capital. As a result, the value function of a bank is

$$V_t = \max_{k^b_t, d_t, m_t} (1-p_t)\mathbb{E}_t\left\{\Lambda_{t,t+1}\left[(1-\sigma)N_{t+1} + \sigma V_{t+1}\right]\right\} \tag{6}$$

subject to (5) and the incentive compatibility constraint

$$V_t \geq \theta Q_t k^b_t. \tag{7}$$

The bank's optimisation problem can be expressed as a choice of leverage ratio $\phi_t \equiv \frac{Q_t k_t^b}{n_t}$ and reserves to net worth, $\gamma_t \equiv \frac{m_t}{n_t}$ so that

$$\max_{\phi_t,\gamma_t} V_t/n_t \equiv \psi_t = (1-p_t)\mathbb{E}_t\left\{\Omega_{t,t+1}[(R^k_{t+1}-\bar{R}_{t+1})\phi_t + (R^m_{t+1}-\bar{R}_{t+1})\gamma_t + \bar{R}_{t+1}]\right\} \tag{8}$$

where $\Omega_{t,t+1} \equiv \Lambda_{t,t+1}(1-\sigma+\sigma V_{t+1}/N_{t+1})$ is the modified discount factor that takes into account the probability that banks may exit.

Under the assumption that the incentive compatibility constraint binds in every period ($V_t = \theta Q_t k_t^b\ \forall t$), the solution to the banks' problem pins down an optimal leverage ratio

$$\phi_t = \psi_t/\theta. \tag{9}$$

2.3.1 The convenience yield of reserves

When choosing the optimal level of reserve holdings, bankers internalise the marginal effects of reserve holdings on their probability of experiencing a run ($\frac{\partial p_t}{\partial \gamma_t} = p'_\gamma$), and the funding cost through their effect on the depositors' pricing of deposits, $\bar{R}'_\gamma$.[8]

Bankers' first order condition for reserve holdings ($\frac{\partial \psi_t}{\partial \gamma_t} = 0$), determines the equilibrium spread between the remuneration rate on reserves and their marginal funding cost

$$\mathbb{E}\Big[\Omega_{t,t+1}(R^m_{t+1}-\bar{R}_{t+1})\Big]$$

$$= \frac{\overbrace{p'_\gamma\,\mathbb{E}\Big[\Omega_{t,t+1}\Big((R^k_{t+1}-\bar{R}_{t+1})\phi_t+\bar{R}_{t+1}\Big)\Big]}^{\text{liquidity insurance}} + \overbrace{(1-p_t)\,\mathbb{E}\Big[\Omega_{t,t+1}\,\bar{R}'_\gamma(d_t/n_t)\Big]}^{\text{funding cost effect}}}{\underbrace{1-(p_t+p'_\gamma\gamma_t)}_{\substack{\text{survival-probability}\\ \text{adjustment}}}} \tag{10}$$

When the spread on the left-hand side of (10) is negative, banks are willing to hold reserves at below the marginal funding cost in order to capture liquidity benefits. In this case, reserves earn a 'convenience yield'. If the bank is fully insured, or further reserves have no effect on their degree of liquidity risk, this yield collapses to zero. Reserve demand is then satiated, and further supply expansion has no marginal price effect. Micro-founding the

[8] The functional form of p_t is defined in Section 2.4. In equilibrium, p_t will be a weakly decreasing function of the bank's expected recovery rate, and therefore, depend on the composition of the bank's assets. In the symmetric equilibrium studied, these bank-level conditions coincide such that p_t represents the systemic run probability.

convenience yield in this way reveals that it comprises two components.

First, *liquidity insurance*: by reducing the probability that the bank experiences a run, higher reserve holdings also increase the present discounted value of future returns on capital, captured by the first term in the numerator.

Second, *funding cost effects*: higher reserve holdings increase the volume of deposits the bank must issue to fund them, feeding through to the bank's funding costs. Depositors monitor these variables and price them in their first order condition on deposits. The impact on equilibrium deposit rates is captured in the second term in the numerator.

The total convenience yield is the total effect of the insurance and funding-cost channels, scaled by a discount-factor adjustment in the denominator that accounts for how reserve holdings affect the bank's survival probability. When the net impact of liquidity insurance and funding cost is sufficiently large, banks accept a negative return on reserves, forgoing higher returns on capital to guarantee their liquidity position.

2.3.2 Aggregation

Summing over both the net interest income of the surviving banks and the endowment of newly entered banks in the absence of bank runs, aggregate net worth evolves according to

$$N_t = \sigma \left[(R^k_t - \bar{R}_t) Q_{t-1} K^b_{t-1} + \left(R^m_t - \bar{R}_t \right) m_{t-1} + \bar{R}_t N_{t-1} \right] + w^b_t, \tag{11}$$

where w^b_t is the aggregate initial transfer that bankers receive to start their operations in period t, whereas aggregate dividends paid back to households are

$$\Pi^b_t = (1-\sigma) \left[(R^k_t - \bar{R}_t) Q_{t-1} K^b_{t-1} + \left(R^m_t - \bar{R}_t \right) m_{t-1} + \bar{R}_t N_{t-1} \right]. \tag{12}$$

In the symmetric equilibrium, aggregate bank capital and reserves are given by $K^b_t = (\phi_t N_t)/Q_t$ and $m_t = \gamma_t N_t$ respectively, directly imposing reserve market clearing. Finally, aggregate deposits are implied by the balance sheet identity $D_t = (\phi_t + \gamma_t - 1) N_t$.

2.4 Bank runs

Runs are systemic (that is, a run affects all banks). When a run takes place, the banking sector goes out of business (so that $K^{b*}_t = N^*_t = D^*_t = 0$) and households must absorb their assets.

As a result, the liquidation price of capital in the event of a run must be consistent with the household first-order condition when $K_t^{h*} = 1 - K_t^{cb}$

$$Q_t^* = \mathbb{E}_t \left(\Lambda_{t,t+1}(Q_{t+1} + Z_{t+1}) \right) - \alpha^h K_t^{h*} \tag{13}$$

Because households are less efficient at intermediating capital than banks, the liquidation price is strictly lower than the normal times value of capital, whenever banks hold capital in equilibrium. This drives a liquidity mismatch in banks' balance sheets: bank assets cannot be liquidated at their full market price.

In contrast to capital, central bank reserves are perfectly liquid. Reserves can be liquidated by the bank at full value and transferred directly to the household as 1 period risk-free bonds in the event of a run. I assume households hold these risk-free bonds until banks re-enter the market, when they return to deposits (whose returns weakly dominate those of reserves).

Following a run at t, aggregate bank net worth follows

$$\begin{aligned} N_t^* &= 0 \\ N_{t+1} &= w^b + \sigma w^b \end{aligned} \tag{14}$$

After this period, net worth continues to evolve according to the law of motion defined in (11).

2.4.1 The probability of bank runs

The recovery rate for depositors in a bank run, x_t, is the ratio of the liquidation value of the bank's assets to the value of its liabilities in the run state,

$$x_t = \min \left[1, \frac{R_t^{k*} Q_{t-1} K_{t-1}^b + R_t^m m_{t-1}}{\bar{R}_t D_{t-1}} \right].$$

This rate has an upper bound of one as deposits never return more than their promised rate, even when the bank makes excess returns.

Like Gertler, Kiyotaki, and Prestipino (2020) and previous works, the probability of a run is assumed to be a decreasing function of the recovery rate, linking the risk of bank runs endogenously to liquidity fundamentals. However, I introduce a non-fundamental shock to depositor expectations, in order to model uncertainty about the aggregate fragility of the banking sector to capture precautionary liquidity motives. Specifically, I assume household expectations for the recovery rate are noisy. The expected recovery rate is the sum of the 'true' recovery rate and a shock ξ_t, but remains bounded at one, so that

$$\mathbb{E}_t(x_{t+1}) = \mathbb{E}_t \min\left[1, \frac{R^{k*}_{t+1} Q_t K^b_t + R^m_{t+1} m_t}{\bar{R}_{t+1} D_t} + \xi_{t+1}\right]. \tag{15}$$

A bank run equilibrium exists if and only if the household believes that the bank would not be able to repay the full value of its obligation, $\bar{R}_t D_{t-1}$, should all other depositors run on the bank. That condition is satisfied when the recovery rate is less than 1, implying that a run would wipe out the banks' entire net worth. Depositors rationally run on the banks if and only if their expectation is for all other depositors to do the same. Therefore, the bank run equilibrium exists as a sunspot whenever $\mathbb{E}(x_{t+1}) < 1$.

I assume the probability of the sunspot crystallising is

$$p_t = \delta_0 \left[1 - \mathbb{E}_t(x_{t+1})\right]^{\delta_1}, \tag{16}$$

where δ_0 and δ_1 are parameters governing the level and sensitivity of the probability of a run to liquidity risk, respectively. This functional form ensures that the risk of runs is continuously differentiable in $\mathbb{E}(x_{t+1})$ and that $p_t > 0$ only if the run equilibrium is possible.[9]

2.5 Central bank

The central bank issues reserves, m_t, by purchasing capital. The central bank balance sheet is therefore given by the identity:

$$Q_t K^{cb}_t = m_t \tag{17}$$

where Q_t is the price of capital. The central bank earns a return on its capital holdings, and remunerates reserves at rate R^m_t.

Any profits or losses from the central bank are rebated to the household in the form of a lump-sum transfer, τ_t.

[9] In principle this functional form needs to be capped such that $p \leq 1$. In practice, for the states visited by the model, this cap never binds since the probability remains low.

3 The cost of supplying reserves

Central banks face tight political and legal constraints on the assets they are permitted to hold (Vissing-Jørgensen 2023). For instance, outside times of crisis, the Fed cannot in general purchase securities which are not government guaranteed.[10] While these asset mandates differ across jurisdictions, they all typically limit central banks to holding a narrow range of securities, relative to the universe of assets held by the private sector. For the purposes of this model, the precise nature of these mandates will not matter: the crucial point is that central banks cannot purchase all possible securities.

This section sets out the production block of the economy, intended to be the simplest possible environment which captures how these constrained asset mandates can drive capital misallocation in the aggregate when the central bank expands its balance sheet. This block of the model is static so for the remainder of this section, I drop time subscripts to simplify notation.

There is a unit mass of firms which produce the consumption good, split into two groups: an Eligible sector E with mass χ and an Ineligible sector N with mass $1-\chi$. The central bank can only purchase assets in the Eligible sector.

Each firm's production technology exhibits decreasing returns to scale, such that distortions in capital allocation can drive dispersion in the marginal product of capital and reduce aggregate TFP (Hsieh and Klenow 2009). The firm in sector i uses loans, k_i, to produce output y_i:

$$y_i = Ak_i^{\alpha}$$

All firms have identical productivity, given by A and capital share α, such that sector E can differ from sector N only in scale. The total output of the economy is given by the sum of outputs across all sectors, such that outputs from different sectors are perfect substitutes. Firms choose k_i to maximise profits, taking z_i and A as given, implying the standard first order condition:

$$z_i = A\alpha k_i^{\alpha-1} \tag{18}$$

so that the yield z_i on the working-capital claim equals the marginal product of capital in sector i.

[10]Federal Reserve Act, Sections 14 and 13(3)

3.1 Capital allocation

The aggregate supply of capital across sectors will depend on the combined holdings of the private sector and the central bank.

The entire central bank asset portfolio is allocated to the sector eligible for central bank purchases, implying a holding of K^{cb}/χ in each firm in E. In contrast, private capital holdings are allocated by fund managers who hold a portfolio of private capital services across firms from both sectors. Define private capital as the sum of household and bank capital holdings:

$$K^p = K^b + K^h$$

I assume that fund managers have inelastic asset demand such that they have a downward sloping demand curve for securities from any given firm. Imperfectly elastic demand in financial markets has been shown to be an important empirical factor in explaining variation in asset prices (Koijen and Yogo 2019; Gabaix and Koijen 2023), and has been studied at length in the literature on quantitative easing. It has been explained with reference to a wide range of factors, including the presence of institutional investment mandates, regulatory requirements, risk aversion, market segmentation, balance sheet constraints, preference for diversification or other non-pecuniary properties of various assets.[11] To represent these factors in a tractable manner, I follow the empirical literature on inelastic asset demand and assume that fund managers allocate private capital across production sectors according to a logit demand system. Investor j allocates to sector i to maximise utility $U_{ij} = z_i + \varepsilon_{ij}$, where ε_{ij} are i.i.d. Type-I Extreme Value with scale parameter $1/\lambda$. Aggregating over investors, the portfolio shares of private capital allocated to each sector are:

$$w_E = \frac{\chi e^{\lambda z_E}}{\chi e^{\lambda z_E} + (1-\chi)e^{\lambda z_N}}, \qquad w_N = 1 - w_E \tag{19}$$

so that private capital allocations per firm are $k_E^p = w_E K^p/\chi$ and $k_N^p = w_N K^p/(1-\chi)$. This is equivalent to the ratio condition:

$$\frac{k_E^p}{k_N^p} = e^{\lambda(z_E - z_N)} \tag{20}$$

[11] See for instance Greenwood and Vayanos (2010), Krishnamurthy and Vissing-Jørgensen (2012), Chen, Cúrdia, and Ferrero (2012), Del Negro et al. (2017), Vayanos and Vila (2021), Gabaix and Koijen (2023), Abadi (2025).

As $\lambda \to 0$, investor demand becomes perfectly inelastic and capital is allocated in proportion to sector sizes regardless of yields. As $\lambda \to \infty$, investors perfectly arbitrage yield differences and capital flows entirely to the higher-yielding sector.

3.2 Equilibrium

An equilibrium consists of vectors of capital allocations (k_E, k_N) and returns (z_E, z_N) such that:

1. The total supply of private capital satisfies:

$$K^p = K^b + K^h = \chi k_E^p + (1-\chi)k_N^p,$$

2. For each sector i, the capital market clears

$$k_i = \begin{cases} k_E^p + K^{cb}/\chi & \text{if } i = E \\ k_N^p & \text{if } i = N \end{cases}$$

3. The first-order conditions for the firm's problem (18) and for the fund manager's portfolio choice (20) are satisfied.

3.3 The efficiency cost of the central bank asset holdings

The unique first-best allocation of capital maximises aggregate output subject to the resource constraint. Efficiency requires that the marginal product of capital is equalised across all sectors. Given symmetric productivities and technologies, this condition implies that all firms operate at the same scale. With a unit mass of firms and unit aggregate capital, the efficient allocation is simply $k_i = k^\star = 1\, \forall i \in [0,1]$ with a corresponding common yield $z^\star = \alpha A$ satisfying the firm FOC.

In light of the constraints on the securities the central bank can hold, it is then straightforward to show that:

Proposition 3.1. *When $\lambda > 0$ and finite, the decentralised economy achieves the first-best allocation if and only if $K^{cb} = 0$.*

Proof. See Appendix A.1.

Proposition 3.1 shows under general conditions that expanding the central bank's footprint in private markets comes at a cost: when the private sector fails to fully arbitrage the distortion created by the central bank, a central bank with constraints on the set of securities it can hold will generate some degree of aggregate capital misallocation, consistent with empirical evidence from the US: Acharya et al. (2025) identify capital misallocation resulting from Fed asset purchases by comparing similar assets with differential exposure to the programme.

Given the central role of constraints on the central bank's asset mandate, a natural question is how the cost of the central bank changes as its eligibility set expands.

Proposition 3.2. *For a fixed total portfolio $\bar{K}^{cb}$ and finite $\lambda > 0$, expanding the central bank's eligibility set – redistributing holdings from eligible to ineligible assets – strictly reduces the aggregate output loss from capital misallocation.*

Proof. See Appendix A.2.

The result says that a central bank can reduce the cost of its balance sheet by spreading its holdings more broadly across asset markets. The distortion in this model comes from concentration: when the central bank can only buy one class of assets, it crowds private investors out of that market and into others, pushing yields apart and misallocating capital. Widening the set of assets the central bank can buy dilutes this concentration.

In the limit, a central bank that holds a representative market portfolio would generate no distortion at all in this model, because it displaces private capital uniformly and gives investors no reason to rebalance. Clearly this result does not consider other factors that may motivate the constraints on asset mandates – in practice, it is unlikely that central banks are best placed to hold certain types of assets. Nonetheless, absent these other factors, the cost of reserve supply can be mitigated by supplying reserves against a broader asset portfolio.

3.4 Aggregation

Total resources are the sum of each sector's output and the endowments for households and bankers

$$Y_t = \chi y_{E,t} + (1-\chi) y_{N,t} + W^h + w^b. \tag{21}$$

Total firm profits are:

$$\Pi_t^f = \chi(y_{E,t} - z_{E,t} k_{E,t}) + (1-\chi)(y_{N,t} - z_{N,t} k_{N,t}). \tag{22}$$

All output is used for management costs or consumption by households:

$$Y_t = C_t + F(K_t^h). \tag{23}$$

The aggregate dividend on capital sets the fund manager's profits to zero such that:

$$Z_t = \frac{\chi z_{E,t} k_{E,t}^p + (1-\chi) z_{N,t} k_{N,t}^p}{K_t^p} \tag{24}$$

4 The optimal long-run balance sheet problem

This section characterises the planner's problem analytically to expose the trade-off faced by the policymaker and then sorts possible optima into one of three regimes: *scarce*, *ample*, or *abundant*. Borrowed from the literature on the empirical reserve demand curve, this convention classifies different reserve supply regimes by the *slope* of the reserve demand curve in equilibrium. Supply is *abundant* when there are sufficient reserves to 'satiate' aggregate reserve demand. Here, the convenience yield on reserves is low, and insensitive to further increases in supply. In equilibrium, the demand curve is flat. Conversely, reserve supply is *scarce* in the supply region where the demand curve is downward sloping. In this region, the convenience yield on reserves responds to changes in supply. Finally, an *ample* regime falls between the two, at the minimum level of reserves at which demand is 'satiated': the kink in the demand curve beyond which the curve is flat.[12] For analytical convenience, I write the problem in terms of central bank assets $K^{cb} = m/Q$, rather than reserves m, and hold the belief shock ξ at its mean value.

The planner chooses the steady-state balance sheet level of capital $K^{cb} \in [0, \chi]$, to maximise unconditional household welfare, taking the competitive equilibrium of Sections 2 and 3 as given. First write household welfare as a function of the central bank balance sheet, $\mathcal{W}(K^{cb}) = \mathcal{W}_n(K^{cb})$. This value function satisfies the recursion

$$\mathcal{W}_n = u(C_n) + \beta[(1-p)\mathcal{W}_n + p\mathcal{W}_r], \tag{25}$$

where $\mathcal{W}_n$ is the normal times value and $\mathcal{W}_r$ is the value of entering a run state: the value function therefore reflects aggregate uncertainty with respect to whether or not the economy will enter a 'bank run' state. When $p > 0$, agents account for the possibility of endogenous

[12]See Afonso et al. (2025) for a graphical representation.

dynamics arising as a result of runs: after a bank run, bank net worth jumps to zero before the economy transitions back to steady state as new banks enter and accumulate net worth. The welfare function therefore depends on the full saddle path for all endogenous variables in the event of a run, through $\mathcal{W}_r$.

Differentiating the value function and using the equilibrium conditions, the planner's first-order condition (derived in Appendix B) yields a variation of the Friedman rule (Friedman 1969): set the marginal cost of reserve supply equal to its marginal benefit.[13] This condition holds when:

$$\underbrace{u'(C_n)\left[(z_N - z_E)\,\Phi\right]}_{\text{MC: misallocation}} = \underbrace{\beta\left[(\mathcal{W}_n - \mathcal{W}_r)(-p_K)\right]}_{\text{(i) lower run probability}} + \underbrace{\beta\left[p\,\frac{d\mathcal{W}_r}{dK^{cb}}\right]}_{\text{(ii) milder runs}} + \underbrace{\left[\mathrm{MB}_K\right]}_{\substack{\text{(iii) capital-reallocation}\\ \text{incidence}}}, \tag{26}$$

where $p_K \equiv dp/dK^{cb} \leq 0$ and $\Phi(K^{cb}) \in (0,1]$ records how much of a unit increase of central-bank asset purchase increases eligible firms' capital in equilibrium, after the private sector has rebalanced.[14] The marginal cost (LHS) measures the marginal welfare cost of capital misallocation which occurs as a result of the central bank's footprint in asset markets. It is strictly positive whenever $K^{cb} > 0$ and $\lambda < \infty$, continuous, and increasing in K^{cb}.

The marginal benefit comprises three terms. Terms (i) and (ii) both capture liquidity benefits. By raising the recovery rate $\mathbb{E}(x)$, reserves reduce the run probability p and, when runs do hit, reduce their severity. Importantly, however, these liquidity benefit terms collapse at satiation, when reserve supply drives the probability of runs to zero so that $p = p_K = 0$. At this point, a kink in the demand curve arises, and the marginal liquidity benefits beyond this point are zero.

Term (iii) measures the effect on holdings in the market for capital. Central-bank purchases change how much capital households manage directly, on which they pay the convex cost $F(K^h)$. The welfare effect is an incidence term which depends on how strongly banks substitute toward reserves as their liquidity position improves (See Appendix B for detail). In the numerical analysis, MB_K is a small net cost but variation of the net effect across reserves is an order of magnitude smaller than that of the liquidity benefit. As a result, this term shifts the MB curve's level but not its shape.

For simplicity, the characterisation above holds the belief shock ξ at its mean value,

[13] The model nests the Friedman rule as the $\lambda \to \infty$ limit: with zero marginal cost of supplying liquidity, full satiation is optimal.

[14] By construction $\Phi \to 0$ as $\lambda \to \infty$ (perfect arbitrage); a closed-form expression is given in Appendix B.

creating a sharp fall in marginal benefits where the steady state probability of a run hits zero. In practice, with shocks to depositor beliefs, this drop is smoothed. Even where the recovery rate is high enough to rule out runs at the mean shock value, the possibility of adverse draws of ξ keeps the probability of a run positive, and the effect of reserve supply on agents' value functions persists. As a result, the liquidity benefits decay sharply through a neighbourhood of the satiation point rather than vanishing. For these reasons, it is more natural to consider satiation as a range of supply levels in which there is a rapid fall in the slope of the demand curve, rather than as a particular point.

Balance sheet regimes. Let $\bar{K}$ denote the smallest K^{cb} at which reserve demand is satiated, and the demand curve is close to flat. The optimal regime is determined by where the marginal-cost curve crosses marginal benefit, relative to the kink at $\bar{K}$. Figure 2 illustrates this point schematically. If the crossing lies far below $\bar{K}$, reserves are *scarce* at the optimum (Panel A). For this to occur, the marginal cost of supplying reserves must be large enough to outstrip their liquidity benefits even when those benefits remain relatively high at the margin. If the crossing of MB and MC lies in the neighbourhood of $\bar{K}$, where marginal benefit falls steeply as the insurance channels are exhausted, reserves are *ample*: the run probability approaches zero from above, and the convenience yield is near to zero (Panel B). Because marginal benefit falls steeply around the kink, a wide range of possible misallocation costs delivers optima in this region. If marginal benefit still exceeds marginal cost at $\bar{K}$, the planner expands past satiation until marginal benefit equals marginal cost, and reserves are *abundant*. The balance sheet is larger than satiation requires (Panel C).

From an analytical perspective, any of the three reserve regimes can be optimal. In the following section, I fit the model to the data to assess which regime obtains in the empirically relevant case.

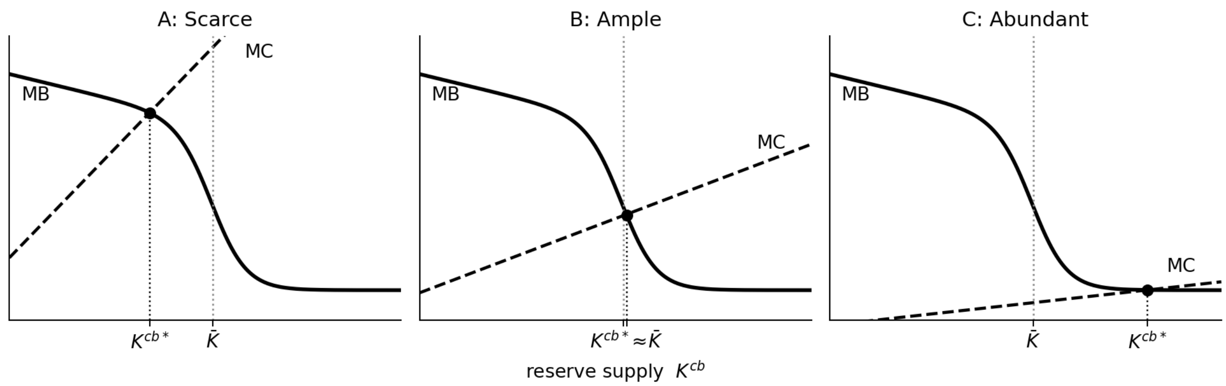

Figure 2: The taxonomy of reserve supply regimes: an illustrative depiction of conditions for different reserve regimes to be optimal. The slope of the MB schedule falls sharply as p falls to zero. The optimal level of reserve supply satisfies MB = MC. In each panel the optimal size of the central bank balance sheet, and the size required to satiate reserve demand are denoted with black and grey dotted lines respectively.

5 Numerical analysis

I solve for the global solution to the model in a setting where agents form rational expectations over the exogenous shock process ξ and over the endogenous probability of bank runs occurring. The full non-linear solution is needed to capture non-linearities around liquidity panics, as well as the effect on precautionary behaviour and risk premia from agents pricing the risk of such events in normal times. I outline the key steps of the numerical procedure in Appendix C.

5.1 Parameters

First, I fix a small number of parameters at conventional values. The quarterly discount factor β is set to 0.99 and the capital share of producers α to 0.33. I set the demand elasticity λ to 21.8, the mid-point of the range of interest-rate elasticities of demand for US government debt in the demand system reported in Eren, Schrimpf, and Xia (2023). The share of productive assets eligible for central-bank purchases, χ, is set to 0.29, which matches the average ratio of a set of eligible assets (US Treasuries, government-sponsored enterprise debt, and agency-owned mortgage-backed securities) to total credit-market debt outstanding in the years before the financial crisis.

The remaining parameters are estimated. I collect the seven free structural parameters in the vector $\Theta = (\theta, \alpha^h, W^h, w^b, \sigma, \delta_0, \delta_1)$. For any Θ, A is chosen to normalise the price of capital to 1, $q(\Theta, A) = 1$, fixing the units of the model.

Because bank runs are endogenous, the calibration cannot proceed moment-by-moment. Every parameter affects the run probability through banks' balance sheets, and the run probability feeds back into every price and quantity in the model. All seven parameters must therefore be disciplined by all of the targets jointly, which calls for a more involved strategy: a single criterion over the full moment set, minimised globally.

I estimate Θ by minimising

$$J(\Theta) \;=\; \underbrace{\sum_{k\in\mathcal{G}}\left(\frac{m_k(\Theta)-\tau_k}{\sigma_k}\right)^2}_{\text{point moments}} + \underbrace{\sum_{k\in\mathcal{B}}\left(\frac{\Big[\ell_k-m_k(\Theta)\Big]_+ + \Big[m_k(\Theta)-u_k\Big]_+}{(u_k-\ell_k)/2}\right)^2}_{\text{interval moments}} + \underbrace{\sum_{j=1}^{3}\left(\frac{\rho_j(\Theta)-r_j}{\sigma_{\text{RDE}}}\right)^2}_{\text{demand-curve anchors}},$$

where $[x]_+ = \max(x, 0)$. Each gap is scaled by how precisely its target is known: a measured standard error where available, and otherwise a plausible interval with a penalty only outside it.[15] The point and interval moments describe the scarce reserve regime that prevailed before 2008, evaluated where central bank liabilities represent 3% of banking-sector assets[16]; the demand-curve targets are estimates of the slope of the reserve-demand curve from January 2010 to March 2025, covering a period of significant variation in the size of the Fed's balance sheet.

I target four point moments. The convenience yield on reserves is targeted at 25 basis points with a standard error of 9, matching the spread (and HAC standard error) between the effective federal funds rate and the three-month Treasury bill from 2002 to 2008. The annual frequency of banking crises is targeted at 4.0% (following Gertler, Kiyotaki, and Prestipino (2020), Reinhart and Rogoff (2009), Jordà, Schularick, and Taylor (2017)) with a standard error of 0.68%, from the Jordà, Schularick, and Taylor (2017) panel with year-clustered standard errors. The rise in credit spreads during a run is targeted at 250 basis points, guided by the increase in the excess bond premium during the 2008 financial crisis (Gilchrist and Zakrajšek 2012). Given uncertainty over the spread measure and the window over which spreads rise, I assume a wide 95% range of 150 to 350 (a standard error of 50). Bank leverage is targeted at 10 with a standard error of 1, so that the 95% range spans 8 to 12.[17]

[15] Written this way J is a diagonal GMM criterion, $m'\Omega^{-1}m$ with $\Omega = \text{diag}(\sigma_k^2)$

[16] Before the crisis, the Fed relied on a different operating framework in which reserves were not remunerated but reserve balances were actively managed. Consequently, there is no direct empirical counterpart for remunerated reserve supply for the pre-crisis period of the model. 3% represents a scarce liquidity supply but acknowledges that banks nonetheless had access to some liquid assets and central bank liabilities such as cash.

[17] The target follows standard calibrations (Gertler, Kiyotaki, and Prestipino 2020), and the standard error

One moment enters as an interval: the bank's spread on assets. The interval brackets are the unit cost of financial intermediation estimated by Philippon (2015), whose annual series averages 1.87% and has remained within 1.46–2.35%. Values within this interval contribute nothing to the criterion inside the interval and a quadratic penalty outside it, scaled by the half-width.

Three demand-curve anchors come from the reserve-demand elasticity estimates of Afonso et al. (2024), summarised in three binned means and weighted by a common standard error of 0.46 implied by the published confidence bands. The three anchors suffice to span the level, slope, and curvature of the demand schedule without over-weighting the demand curve data (which is highly serially correlated).

The parameters of the belief shock, σ_ξ and ρ_ξ, are set to match the observed relationship between the volatility in convenience yields and level of reserve supply, shown in the bottom panel of Figure 3.

Table 1 shows the calibrated parameter values, which are all in line with typical values in the literature. Households have log utility, implying a degree of risk aversion that is low relative to the asset pricing literature. The model therefore likely delivers a conservative estimate of the welfare gains from eliminating risk. Table 2 shows the model moments compared to the target moments.

is chosen to approximate the variation of aggregate book leverage of US commercial banks.

Table 1: Parameter Values

Parameter	Value	Description
Households		
β	0.99	Discount factor
α^h	0.0106	Household managerial cost
W^h	0.0406	Household endowment
Bankers		
σ	0.9104	Banker survival probability
θ	0.1781	Seizure rate
w^b	0.00185	Banker start-up funds
Production		
A	0.0449	TFP
α	0.33	Output elasticity of capital
λ	21.8	Fund manager yield sensitivity
χ	0.29	Measure of Eligible firms
Run Beliefs		
δ_0	5.93	Run belief coefficient
δ_1	1.45	Run belief exponent
ρ_ξ	0.940	AR1 coefficient of ξ
σ_ξ	0.003	Std. dev. of ξ shock

Table 2: Model vs. Target Moments

Moment	Model	Target	Source
Run probability (ann., pp)	4.00	4.0	JST (2017)
Convenience yield (bp)	-26	-25	EFFR $-$ T-bill, 2002–08
Run spread increase (bp)	269	250	GZ (2012)
Bank leverage, ϕ	10.42	10	GKP (2020)
Bank spread, $R^k - R^d$ (bp)	161	[146, 235]	Philippon (2015)
HH capital holdings, K^h	0.38	Untargeted	—
Price of capital, Q	1.00	1 (normalisation)	—

Point targets are weighted by their standard errors and the interval target is penalised only outside its range; K^h is not targeted. The reserve-demand elasticity curve is matched separately at three anchor points. JST (2017): Jordà, Schularick and Taylor (2017).

5.2 Empirical validity

The parametrised model matches its target data on reserve demand, especially in the post-2014 period. Figure 3 compares the timeseries of reserve demand elasticity estimates from Afonso et al. (2024), and the corresponding measure from the model at equivalent levels of reserve supply. Qualitatively, the model distinguishes between periods of abundant (when the RDE was zero) and more scarce reserve supply with high accuracy. Quantitatively, the ergodic distribution of the model fits the post-2014 sample closely, accurately capturing the increasing scarcity of reserves which crystallised in the 2019 repo market turmoil. But in the early post-crisis period, the model under-estimates the extent of scarcity. The difficulty fitting both periods is unsurprising. Afonso et al. (2024) document low-frequency shifts in the demand curve for reserves over time, with a notable shift occurring in 2014-15 with adjustments to the ONRRP, which a single model parametrisation cannot reflect.

Figure 4 shows the observed quarterly volatility of the convenience yield on reserves as a function of the level of reserve supply, used to calibrate the shock process. The model reproduces the negative relationship between volatility in the convenience yield on reserves and the level of reserve supply and matches the level of volatility.

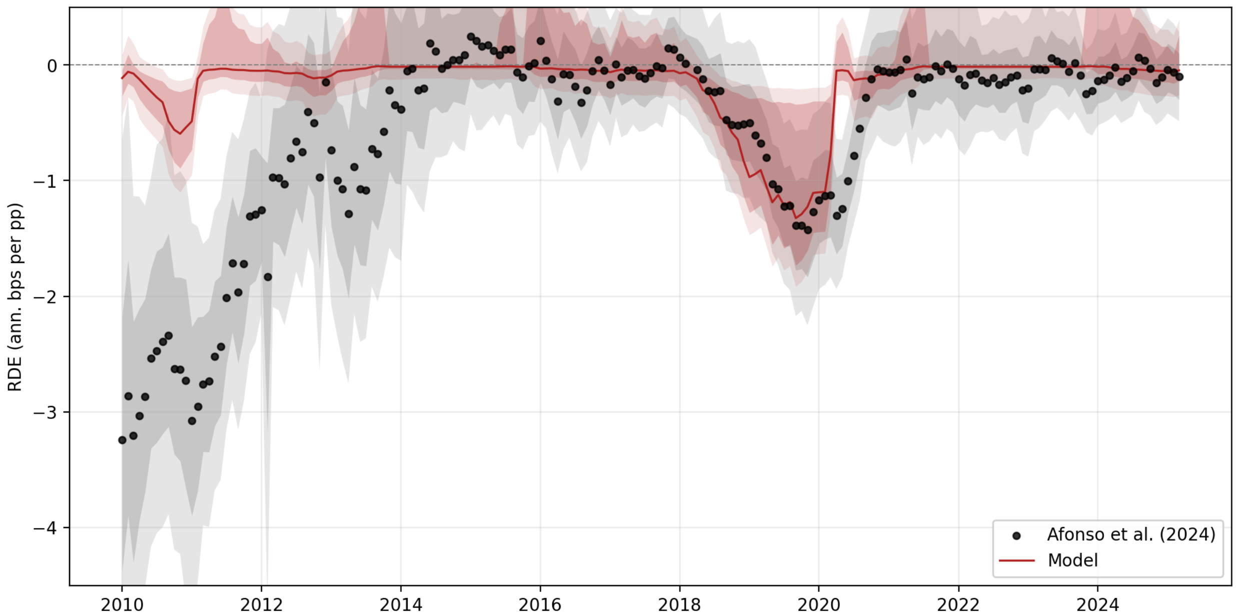

Figure 3: *Slope of the reserve demand curve.* Notes: The model-implied reserve demand elasticity (RDE) since 2010 (red) against the rolling econometric estimates of Afonso et al. (2024) (grey, points). For both the model and the data, the darker and lighter shades mark the 68th and 95th percentiles respectively.

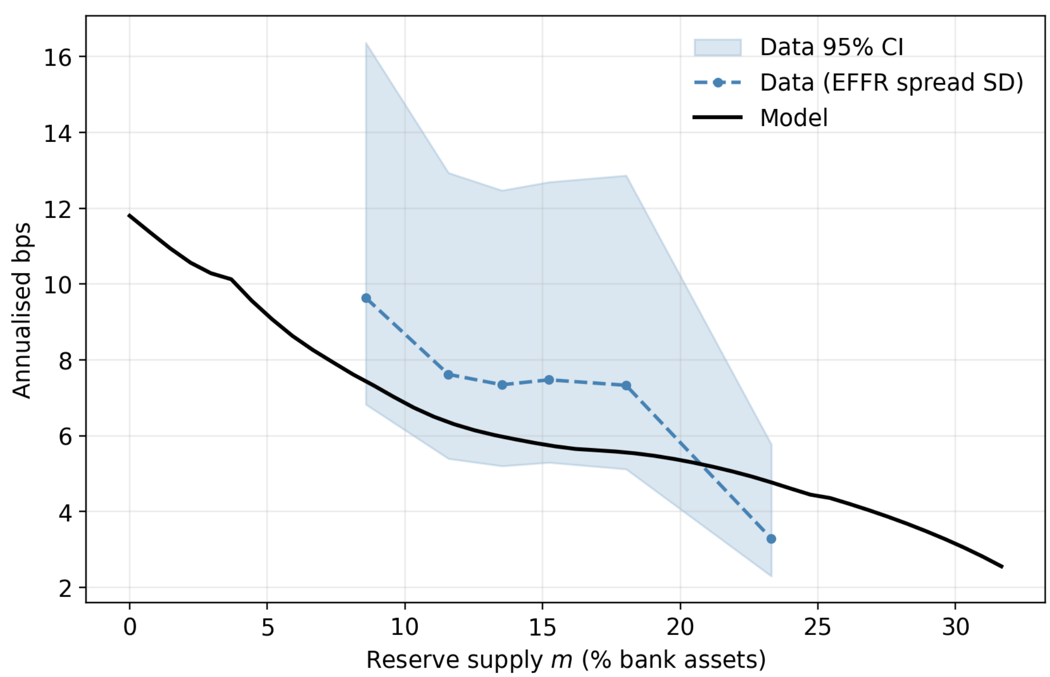

Figure 4: *Volatility of the convenience yield.* Notes: The standard deviation of the reserve convenience yield at different levels of reserve supply in model and data (source: FRED); quarterly volatility is estimated by aggregating observations into six bins of roughly twelve each.

5.3 Optimal reserve supply

Having parametrised and solved the model, I carry out comparative statics with respect to the size of the central bank balance sheet. Throughout, "steady state" refers to the risky steady state, as in Coeurdacier, Rey, and Winant (2011). This is the fixed point of the economy when shocks realise at their mean, but agents' decision rules (and the value functions used for welfare) continue to price the full distribution of future shocks, including runs. Unlike a deterministic steady state, in which anticipated risk would be absent, this measure preserves the precautionary behaviour that gives reserves their value.

Figure 5 shows steady state expected utility, $\mathcal{W}$, evaluated at different levels of reserve supply, converted into consumption-equivalent terms (pp) relative to the calibrated steady state in the baseline model. Figure 6 shows comparative statics for key endogenous variables, tracing the channels of the social planner's FOC.

The optimal balance sheet in the calibrated model sits in the ample reserve supply regime: the optimum is at the kink, with p approaching zero from above, and a near-zero convenience yield. In the parametrised model, the optimal-sized reserve supply is roughly 15% of bank assets. Given its location very close to the kink (rather than to the right of it, in *abundant* reserve supply), some volatility in the convenience yield on reserves will persist (standard deviation of around 6bps). Intuitively, consumers would have the central bank almost entirely eliminate run risk, at the cost of lower consumption in good times. But when the risk of runs becomes sufficiently remote, further expansions to eliminate extreme tail risks have a higher welfare cost than the benefit of further liquidity insurance.

To illustrate what drives this result, consider two limiting cases and the comparative statics of other key endogenous variables in the full model.

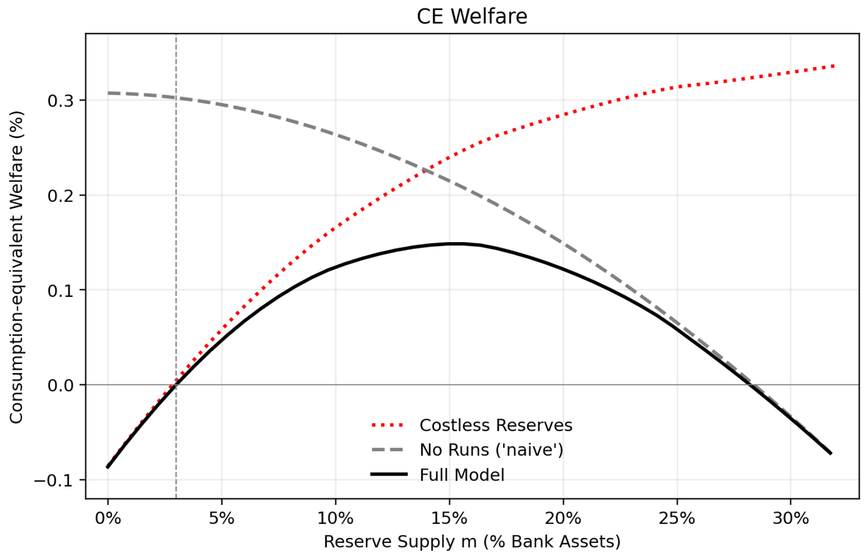

Figure 5: *Comparative Statics: Ex-Ante Expected Utility.* Notes: The figure plots the expected utility in 3 cases: the solid black line shows the full model. The grey, dashed line shows results for the case where households assign zero probability to the bank run state, regardless of the banks' financial conditions (are 'naive'). The red, dotted line shows the case where $\lambda = \infty$.

In the special case with 'no runs' ($\delta_0 = 0$), shown by the grey-dashed line in Figure 5, depositors fail to anticipate runs, and always perceive the probability of bank runs to be zero. Liquidity risk has no effect on agents' expected utility and the provision of central bank reserves confers no liquidity benefit. When there is no liquidity risk, the optimal size of the central bank balance sheet is zero – the size which minimises capital misallocation. Scarce reserves are the optimal supply regime.

The red-dotted line considers the special case where the central bank creates no capital misallocation because asset markets are perfectly elastic ($\lambda = \infty$). Because additional reserves do not come with any cost in the form of capital misallocation, they monotonically increase welfare by reducing the risk of runs and alleviating financial frictions. Abundant reserves are optimal.

In the full parametrised model, households account for the probability of a bank run when forming expectations. Increasing the provision of reserves reduces liquidity mismatch on bank balance sheets, raising the recovery rate and lowering the probability of a bank run (Figure 6, Panel A), reducing the risk premium banks must pay their depositors and lowering lending spreads.[18] In the bank run state, consumption is increasing with reserve supply because

[18] A figure showing statics for other endogenous variables is included in Appendix E.

banks are more liquid, meaning depositors recover a greater share of their funds (Figure 6, Panel B). The initial expansions in reserve supply generate substantial improvements in expected utility for the household, largely driven by the *liquidity benefits* described in Section 4.

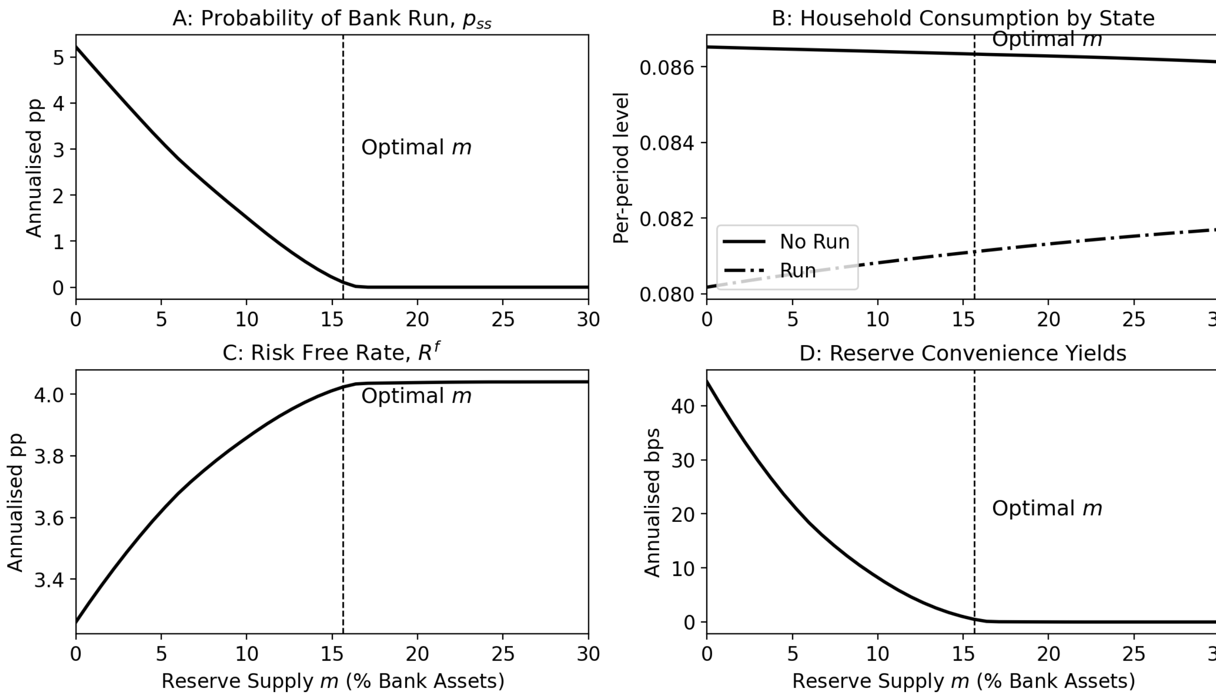

Figure 6: *Comparative Statics: Endogenous Variables.* Notes: The optimal size of the central bank balance sheet is marked by a dotted line for reference.

The reduction in aggregate uncertainty also pushes down demand for precautionary savings and the unconditional risk-free rate in (4) rises as reserve supply expands (Figure 6, Panel C). This is driven by a familiar mechanism: public liquidity crowds out precautionary saving and raises equilibrium rates (Aiyagari and McGrattan 1998; Nuño 2025).[19]

As reserve supply increases and the likelihood of runs falls, the marginal benefit of additional reserves is decreasing. This is, in effect, because they provide insurance against a form of liquidity risk which is becoming increasingly remote. As the marginal impact of reserves on aggregate risk falls, so too does their convenience yield (Figure 6, Panel D): banks demand higher compensation to continue to increase their reserve holdings as the marginal benefits of doing so dwindle, consistent with the empirical evidence on convenience yields.

[19]This echoes findings that short-term rates increase in the size of the central bank (Arce et al. 2020) and suggestions by policymakers that balance sheet expansions may raise r* (Schnabel 2025).

When liquidity risk is sufficiently low, the marginal benefit of additional reserves has almost no effect on the risk of runs.

When reserve supply is large enough to rule out the run equilibrium (i.e. $x = 1$ and $p = 0$), the marginal liquidity benefit of reserves is (numerically) zero. The convenience yield falls to zero (Figure 6, Panel D) and reserve demand is perfectly flat. The liquidity insurance value of reserves collapses to zero: beyond this point, banks have a sufficient volume of liquid assets to remain solvent in a run, and will only hold additional reserves at marginal cost.

The full model becomes isomorphic to the 'naive household' case when reserve supply is large enough to rule out runs in the region of the steady state, and create a liquidity buffer so large that shocks big enough to generate the risk of runs in future also have a probability close to zero. At that point, the black and grey lines in Figure 5 coincide, when reserve supply reaches almost 30% of bank assets.

In consumption terms, the effect of reserve supply on the level of welfare is small, but not negligible. Moving from the pre-crisis steady state to the optimal balance sheet is worth about 0.15pp of (certainty equivalent) per-period consumption. In the US in 2026Q2, this corresponds to roughly \$100 per year on average.[20]

6 Robust optimal supply of reserves

If the policymaker incorrectly estimates the model's parameters, they risk mismeasuring the forces which determine the optimal reserve supply, such as the location of the reserve demand curve, the cost of reserve supply, and the probability of liquidity crises. The analysis so far treats the calibrated parameters as known to the policymaker. In this section, I relax that assumption.

In practice, the data do not perfectly distinguish the point estimate from a set of neighbouring calibrations that match the targets almost as well. This section asks two questions. First, how robust is the optimal policy result from Section 5 to parameter uncertainty? Second, how costly would it be to operate a reserve supply policy which is more robust? To answer these questions, I construct a set of empirically admissible economies and then compare reserve policies by the "regret" that they generate.

I consider the question of robustness from two perspectives, which differ by the nature

[20] Personal consumption expenditures per capita (A794RC0Q052SBEA) in 2026. Source: U.S. Bureau of Economic Analysis (2024)

of policymaker uncertainty. First I consider a policymaker who acknowledges uncertainty in parameter values, but trusts that parameter vectors which fit the targets more closely are more likely. I then consider a policymaker who takes a more agnostic view, unwilling to attach probabilities to parametrisations that lie within a broadly defined 'admissible set'. In this exercise, the policymaker simply seeks a policy which is robust across the range of economies which cannot be ruled out.

6.1 The admissible set

To set the scope of the robustness exercise, begin by defining a wide set of empirically plausible parametrisations.

Let $\tilde{\Theta} = (\theta, \alpha^h, W^h, w^b, \sigma, \delta_0, \delta_1, \lambda, \chi)$ collect every parameter a policymaker should doubt: the estimated structural parameters of Section 5.1 together with the two parameters governing the cost of supplying reserves, λ and χ. I explore this space with a quasi-random sample of roughly 1500 draws of $\tilde{\Theta}$. These samples are spread evenly over the following space: the estimated structural parameters vary within $\pm 22\%$ around $\hat{\tilde{\Theta}}$,[21] λ spans the full range of investor-level elasticity estimates [10.85, 32.76] (Eren, Schrimpf, and Xia 2023), and χ its observed range, [0.25, 0.45] since 1990.

To construct the 'admissible set', I retain the 484 economies whose implied moments stay within the ranges the data have actually visited: bank leverage within the range of 8–20,[22] return on assets within the observed historical range of the cost of intermediation, 1.46–2.35% (Philippon 2015); run spreads within 150–350 basis points, and in which runs occur at a positive annual rate ($p \geq 0.2$ per cent) no greater than 13.3 per cent.[23]

6.2 Regret and robust decision criteria

For each admissible economy, let $M^*(\tilde{\Theta})$ denote the optimal level of reserve supply a central bank would choose under that economy's true parameter vector. The "regret" of running a common policy M in economy $\tilde{\Theta}$ is the gap between what that economy could have had and what it gets: the certainty-equivalent consumption cost to households of a central bank that

[21] Except the banker survival rate σ ($\pm 5\%$): it is a probability, so wide variation would exceed one, and in expected banker lifetime $1/(1-\sigma)$ even $\pm 5\%$ spans 7–22 quarters around the baseline's 11.

[22] This wider range reflects the significant heterogeneity in leverage across the banking sector, with wholesalers typically operating at higher leverage ratios (Gertler, Kiyotaki, and Prestipino 2016).

[23] This value represents the most crisis-prone fifteen-year period in US history—the 1893–1907 banking panics (Jordà, Schularick, and Taylor 2017).

acted on the wrong model of their economy,[24]

$$r(M;\tilde{\Theta}) = \mathrm{CE}\Big[\mathcal{W}(M^*(\tilde{\Theta});\tilde{\Theta})\Big] - \mathrm{CE}\Big[\mathcal{W}(M;\tilde{\Theta})\Big].$$

Using this measure, I consider two approaches to robustness, which differ by the nature of policymaker uncertainty. First consider an approach driven by the calibration data: each economy is weighted by its conformity with the calibration targets, $\propto e^{-J_i/2}$.[25] Economies that match the targets more closely are treated as more likely. Second, consider a more agnostic policymaker who wants the chosen supply to be robust not only to error in the estimated parameters but to the possibility that the model or its targets are misspecified or mismeasured. In this case, a closer fit is no guarantee of being closer to the truth, so the policymaker attaches equal weight to every economy that cannot be ruled out, however implausible the fit may judge it (Hansen and Sargent 2008). The two policymakers therefore differ only in the baseline weights they place on the admissible economies: proportional to fit for the first, equal for the second.

Given its baseline weights, each type of policymaker evaluates policy robustness against a family of decision rules, indexed by a single parameter, τ, which measures the depth of the adverse tail against which the policymaker insures. Decision rule τ will select the level of reserve supply that minimises the expected regret across the worst economies that jointly account for a $(1-\tau)$ share of the probability weight. The higher the value of τ, the higher the level of robustness sought by the policymaker. Formally, this value is the conditional value-at-risk (CVaR) of regret at level τ (Rockafellar and Uryasev 2000). At $\tau = 0$ the rule reduces to expected regret under the policymaker's own weights: the fit-weighted expectation for the first type, simple averaging for the second. As τ rises the rule selects a policy robust to an increasingly severe set of economies.

[24]Welfare levels are not comparable across candidate economies. For instance, a low-endowment parametrisation is poorer at every reserve supply, for reasons no policy can affect. The regret measure differences out these effects, comparing only the loss the policy could have avoided (Savage 1951; Manski 2004).

[25]Economies are weighted by $e^{-J_i/2}$, which falls with the calibration criterion. As a result, better-fitting parametrisations count for more. The exact weight corresponds to the Gaussian weight obtained by treating the targets as independent normal measurements with standard errors σ_k. It is a simple device for rewarding fit rather than a formal posterior, and the conclusion that fit-weighting concentrates the recommendation near the point estimate does not depend on its precise form.

6.3 Results

Under the fit-weighted expectation, which weights parametrisations by how closely they fit the data, the estimated policy is robust. The expected-regret-minimising supply is the full information optimum itself. Insuring against tail outcomes implies a negligible buffer of no more than one percentage point against the worst 5% of outcomes (Appendix Table 4).[26]

Now consider the problem of the more agnostic policymaker. Table 3 reports the results from decision rules at different τ-levels, as well as the "Premium" the policymaker pays to achieve each level of robustness. That premium measures the cost of deviating from the full information optimal policy (expressed as a percentage of the original welfare gain from reserve supply). Figure 7 shows the CVaR regret curves at different levels of reserve supply. This depicts how poorly any given level of reserve supply performs against the $1 - \tau$ worst cases. The policies selected by the robust decision rules are marked with vertical lines.

Rules that insure against deeper tail risks raise optimal supply. The averaging rule raises the supply by 1.4 percentage points of bank assets relative to the central-case optimum; insuring against the worst tenth of economies raises it by 1.8, and against the worst one per cent by 2.4. These precautionary buffers are positive because the costs of deviating from optimal policy are not symmetric: too few reserves raise the economy's exposure to liquidity crises, while too many misallocate capital. Averaged across the admissible set, the welfare loss from supplying two percentage points too few reserves is about 1.6 times larger than that from supplying two too many. A "robust" reserve policy therefore errs on the side of 'too many' reserves, rather than too few.

[26]Under fit-weighting I report τ levels up to 0.95. The fit weights imply an effective number of economies well below n; at $\tau = 0.99$ the adverse tail retains too few effective draws to estimate the CVaR.

Table 3: Robust reserve supply: decision rules and the price of caution

Policy	Supply (% assets)	ΔM^* (pp)	95% CI (pp)	Premium (% of policy value)
Central (point est.)	15.6	0.0	—	0
Averaging	17.0	+1.4	$[+0.8, +1.8]$	3
CVaR $\tau = 0.80$	17.2	+1.6	$[+1.0, +2.0]$	4
CVaR $\tau = 0.90$	17.4	+1.8	$[+1.2, +2.2]$	5
CVaR $\tau = 0.95$	17.6	+2.0	$[+1.2, +2.8]$	5
CVaR $\tau = 0.99$	18.0	+2.4	$[+1.4, +2.8]$	7

Equal weight over the $n = 484$ admissible economies; bootstrap 95% confidence intervals in brackets. Premium: the central economy's regret at the rule's supply, as a share of the reserve policy's own value (0.15% of lifetime consumption).

These buffers provide effective insurance against severe outcomes. At the full information optimum, the least favourable 1% of admissible economies implies an average regret of 0.11 per cent of consumption, equivalent to three-quarters of the welfare gain from adopting the optimal reserve supply. For those economies, calibration error would forfeit much of the value of the policy. The robust supply reduces that exposure by roughly a third.

Guarding against the most adverse tail risks forgoes between 3 (averaging) and 7 (CVaR $\tau = 0.99$) per cent of the consumption-equivalent gain from supplying reserves, driven by the capital-misallocation cost of a larger central bank footprint (See 'Premium' in Table 3).

There are two implications for policymakers. Weighting the likelihood of parameters by their fit to the calibration targets, the full information optimal policy is already robust. However, a policymaker who distrusts the model-implied probability ranking has a well-defined response: a modest buffer of 2.4 percentage points ensures robustness to the worst 1% of plausible outcomes. Even the most risk-averse rules therefore prescribe a high degree of restraint relative to open-ended, abundant reserve supply. Robustness to wide-ranging uncertainty in the costs and benefits of reserve supply requires only a small expansion beyond the point of ample reserve supply.

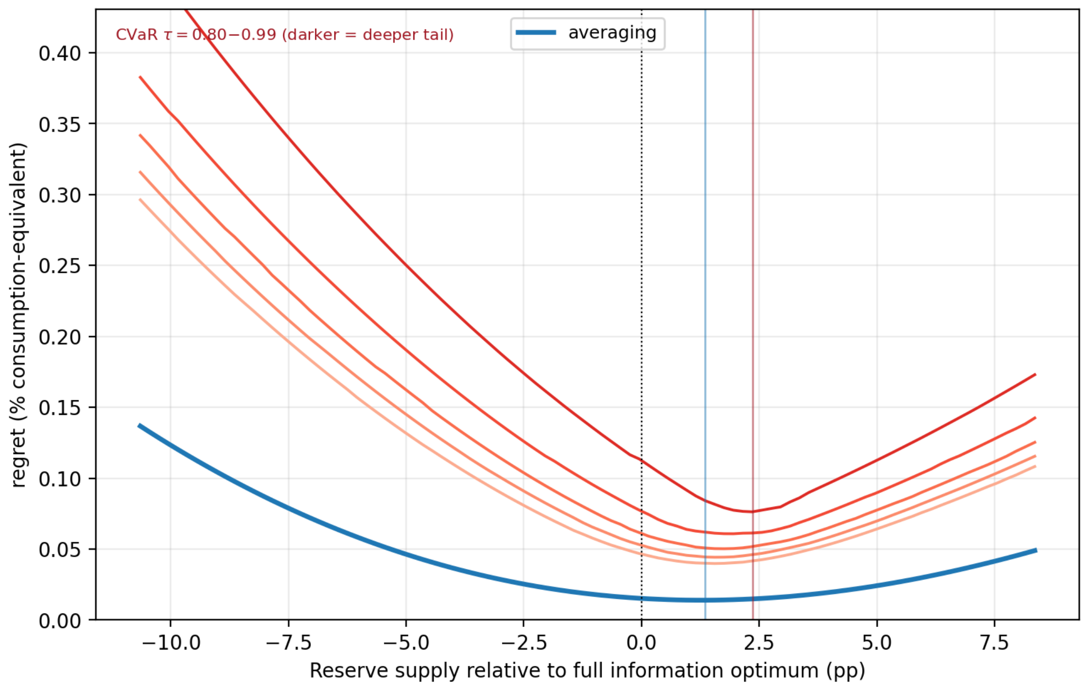

Figure 7: Robust decision criteria. Mean regret across the admissible set (averaging), CVaR_τ for τ between 0.80 and 0.99 (swathe; darker shades insure deeper tails), as functions of the reserve supply measured relative to the central-case optimum. Vertical lines mark the averaging and $\mathrm{CVaR}_{0.99}$ optima; the dotted line is the central case, at zero.

Appendix D exploits the wide-ranging parameter robustness exercise used in this section to assess which parameters drive both the level of optimal reserves and the asymmetry in costs that generate the precautionary motives for robust policy results.

7 Conclusion

This paper has asked what the socially optimal long-run supply of central bank reserves is, in a framework where both sides of the trade-off arise endogenously. Reserves reduce the likelihood and severity of liquidity crises, while the assets the central bank holds to supply them crowd out private intermediation and misallocate capital. Estimated on US data, the model reproduces the non-linear reserve demand curve at the centre of current policy debates.

There are three policy relevant lessons. First, in the estimated model, it is optimal to operate an ample reserve supply regime. The welfare-maximising balance sheet supplies reserves up to the point where demand approaches satiation: enough that reserves no longer carry a measurable liquidity premium (‘convenience yield’), but short of the open-ended provision of abundant reserves. Second, robustness to uncertainty requires higher reserve supply. However, high degrees of robustness can be achieved with only a small increase in reserve supply. In practice a highly risk averse policymaker would only choose to run a small buffer above the ample optimum, not a return to abundant reserve supply far in excess of what is required to satiate demand. Finally, the composition of the central bank's assets has only a muted effect on the optimal reserve supply at the empirically relevant calibration, but bears on the cost of supplying reserves nonetheless. Policies which reduce the distortions from operating a large balance sheet by widening the set of assets the central bank uses to back reserves can deliver meaningful efficiency improvements. Such policies are welfare improving in all states of the world, and reduce the cost of expanding reserve supply to ensure robustness against more adverse cases. This result offers a welfare-based rationale for recent decisions by the Bank of England and the ECB to broaden the set of collateral against which they supply reserves (Saporta 2025; Schnabel 2024).

The simple framework presented abstracts from several margins that qualify these conclusions. The model treats the level of reserve supply as fixed, implying a vertical supply curve. The framework is therefore silent on interactions with standing lending facilities and with other policy tools (fiscal, macro-prudential). The model also focusses on a particular channel through which reserve supply is costly in steady state. Other effects such as liquidity dependence, or moral hazard arising from the public supply of liquidity in crisis are also relevant factors to both the level of reserve supply and the optimal transition towards the eventual long-run balance sheet (Acharya et al. 2023; Acharya and Rajan 2024; Acharya, Rajan, and Shu 2025). I leave these extensions to future work.

References

Abadi, Joseph. 2025. "Monetary Policy with Inelastic Asset Markets." Working paper 25-15. Federal Reserve Bank of Philadelphia. https://doi.org/10.21799/frbp.wp.2025.15.

Acharya, Viral V., Ryan Banerjee, Matteo Crosignani, Tim Eisert, and Renée Spigt. 2025. "Exorbitant Privilege? Quantitative Easing and the Bond Market Subsidy of Prospective Fallen Angels." *Journal of Financial Economics* 170 (August): 104084. https://doi.org/10.1016/j.jfineco.2025.104084.

Acharya, Viral V., Rahul S. Chauhan, Raghuram Rajan, and Sascha Steffen. 2023. "Liquidity Dependence and the Waxing and Waning of Central Bank Balance Sheets." Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w31050.

Acharya, Viral V., and Raghuram Rajan. 2024. "Liquidity, Liquidity Everywhere, Not a Drop to Use: Why Flooding Banks with Central Bank Reserves May Not Expand Liquidity." *The Journal of Finance* 79 (5): 2943–91. https://doi.org/10.1111/jofi.13370.

Acharya, Viral V., Raghuram Rajan, and Zhi Quan (Bill) Shu. 2025. "When Is Less More? Bank Arrangements for Liquidity Vs Central Bank Support." Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w34099.

Afonso, Gara, Domenico Giannone, Gabriele La Spada, and John C. Williams. 2024. "Scarce, Abundant, or Ample? A Time-Varying Model of the Reserve Demand Curve." Staff {Reports}. Federal Reserve Bank of New York. https://doi.org/10.2139/ssrn.4121823.

Afonso, Gara, Kyungmin Kim, Antoine Martin, Ed Nosal, Simon Potter, and Sam Schulhofer-Wohl. 2025. "Monetary Policy Implementation with Ample Reserves." Staff {Reports}. Federal Reserve Bank of New York.

Afonso, Gara, Gabriele La Spada, Thomas M. Mertens, and John C. Williams. 2023. "The Optimal Supply of Central Bank Reserves Under Uncertainty." Staff {Reports}. Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr1077.pdf.

Aiyagari, S. Rao, and Ellen R. McGrattan. 1998. "The Optimum Quantity of Debt." *Journal of Monetary Economics* 42 (3): 447–69. https://doi.org/10.1016/S0304-3932(98)00031-2.

Alexandrov, Andrey, and Markus K Brunnermeier. 2026. "Optimal (Un)Conventional Monetary Policy." *Working Paper*, February.

Arce, Óscar, Galo Nuño, Dominik Thaler, and Carlos Thomas. 2020. "A Large Central Bank Balance Sheet? Floor Vs Corridor Systems in a New Keynesian Environment." *Journal of Monetary Economics* 114 (October): 350–67. https://doi.org/10.1016/j.jmoneco.2019.05.001.

Bianchi, Javier, and Saki Bigio. 2022. "Banks, Liquidity Management, and Monetary Policy." *Econometrica* 90 (1): 391–454. https://doi.org/10.3982/ECTA16599.

Bindseil, Ulrich. 2016. "Evaluating Monetary Policy Operational Frameworks." *Proceedings of the Jackson Hole Economic Symposium of the Federal Reserve Bank of Kansas City on "Designing Resilient Monetary Policy Frameworks for the Future"*, August.

Borio, Claudio. 2023. "Getting up from the Floor." *BIS Working Paper*, no. 1100 (May).

Chen, Han, Vasco Cúrdia, and Andrea Ferrero. 2012. "The Macroeconomic Effects of Large-scale Asset Purchase Programmes." *The Economic Journal* 122 (564): F289–315. https://doi.org/10.1111/j.1468-0297.2012.02549.x.

Coeurdacier, Nicolas, Hélène Rey, and Pablo Winant. 2011. "The Risky Steady State."

American Economic Review P&P 101 (3): 398–401. https://doi.org/10.1257/aer.101.3.398.

Copeland, Adam, Darrell Duffie, and Yilin (David) Yang. 2025. "Reserves Were Not So Ample After All." *The Quarterly Journal of Economics* 140 (1): 239–81. https://doi.org/10.1093/qje/qjae034.

Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki. 2017. "The Great Escape? A Quantitative Evaluation of the Fed's Liquidity Facilities." *American Economic Review* 107 (3): 824–57. https://doi.org/10.1257/aer.20121660.

Del Negro, Marco, and Christopher A. Sims. 2015. "When Does a Central Bank's Balance Sheet Require Fiscal Support?" *Journal of Monetary Economics*, Carnegie-Rochester-NYU Conference Series on Public Policy "Monetary Policy: An Unprecedented Predicament" held at the Tepper School of Business, Carnegie Mellon University, November 14-15, 2014, 73 (July): 1–19. https://doi.org/10.1016/j.jmoneco.2015.05.001.

Diamond, Douglas, and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401–19. https://www.jstor.org/stable/1837095.

Diamond, William, Zhengyang Jiang, and Yiming Ma. 2024. "The Reserve Supply Channel of Unconventional Monetary Policy." *Journal of Financial Economics* 159 (September): 103887. https://doi.org/10.1016/j.jfineco.2024.103887.

Eren, Egemen, Timothy Jackson, and Giovanni Lombardo. 2024. "The Macroprudential Role of Central Bank Balance Sheets." *BIS Working Paper*, March.

Eren, Egemen, Andreas Schrimpf, and Fan Dora Xia. 2023. "The Demand for Government Debt." Rochester, NY: Social Science Research Network. https://doi.org/10.2139/ssrn.4466154.

Friedman, Milton. 1969. *The Optimum Quantity of Money*. 2. print. New Brunswick, NJ: Aldine Transaction.

Gabaix, Xavier, and Ralph S. J. Koijen. 2023. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." {SSRN} {Scholarly} {Paper}. Rochester, NY. https://doi.org/10.2139/ssrn.3686935.

Gertler, Mark, and Nobuhiro Kiyotaki. 2015. "Banking, Liquidity, and Bank Runs in an Infinite Horizon Economy." *The American Economic Review* 105 (7): 2011–43. https://www.jstor.org/stable/43495449.

Gertler, Mark, Nobuhiro Kiyotaki, and Andrea Prestipino. 2016. "Wholesale Banking and Bank Runs in Macroeconomic Modeling of Financial Crises." In *Handbook of Macroeconomics*, edited by John B. Taylor and Harald Uhlig, 2:1345–425. Elsevier.

https://www.sciencedirect.com/science/article/pii/S1574004816000100.

———. 2020. "A Macroeconomic Model with Financial Panics." *The Review of Economic Studies* 87 (1): 240–88. https://doi.org/10.1093/restud/rdz032.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–1720. https://doi.org/10.1257/aer.102.4.1692.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein. 2018. "The Federal Reserve's Balance Sheet as a Financial-Stability Tool." In *World Scientific-Now Publishers Series in Business*, 15:63–124. World Scientific/NOW Publishers. https://www.worldscientific.com/doi/abs/10.1142/9789813236592_0003.

Greenwood, Robin, and Dimitri Vayanos. 2010. "Price Pressure in the Government Bond Market." *American Economic Review* 100 (2): 585–90. https://doi.org/10.1257/aer.100.2.585.

Hall, Robert E., and Ricardo Reis. 2015. "Maintaining Central-Bank Financial Stability Under New-Style Central Banking." Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w21173.

Hansen, Lars Peter, and Thomas J. Sargent. 2008. *Robustness*. Course Book. Princeton, N.J: Princeton University Press. https://doi.org/10.1515/9781400829385.

Harenberg, Daniel, Stefano Marelli, Bruno Sudret, and Viktor Winschel. 2019. "Uncertainty Quantification and Global Sensitivity Analysis for Economic Models." *Quantitative Economics* 10 (1): 1–41. https://doi.org/10.3982/QE866.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India*." *The Quarterly Journal of Economics* 124 (4): 1403–48. https://doi.org/10.1162/qjec.2009.124.4.1403.

Ihrig, Jane, Zeynep Senyuz, and Gretchen C. Weinbach. 2020. "Implementing Monetary Policy in an 'Ample-Reserves' Regime: Maintaining an Ample Quantity of Reserves (Note 2 of 3)," August. https://www.federalreserve.gov/econres/notes/feds-notes/implementing-monetary-policy-in-an-ample-reserves-regime-maintaining-an-ample-quantity-of-reserves-note-2-of-3-20200828.html.

Jordà, Òscar, Moritz Schularick, and Alan Taylor. 2017. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016, Edited by Martin Eichenbaum and Jonathan A. Parker*. Vol. 31. Chicago: University of Chicago Press.

Karadi, Peter, Dominik Thaler, and Oreste Tristani. 2024. "The Central Bank's Balance Sheet in the Long Run: A Macro Perspective." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4975240.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy* 127 (4): 1475–515. https://doi.org/10.1086/701683.

Krishnamurthy, Arvind, and Annette Vissing-Jørgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2): 233–67. https://doi.org/10.1086/666526.

Lagos, Ricardo, and Gaston Navarro. 2026. "Reserve Demand Estimation with Minimal Theory." Working {Paper}. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w34972.

Lopez-Salido, David, and Annette Vissing-Jørgensen. 2025. "Reserve Demand, Interest Rate Control, and Quantitative Tightening." Rochester, NY. https://doi.org/10.2139/ssrn.4371999.

Manski, Charles F. 2004. "Statistical Treatment Rules for Heterogeneous Populations." *Econometrica* 72 (4): 1221–46. https://doi.org/10.1111/j.1468-0262.2004.00530.x.

Nuño, Galo. 2025. "Three Theories of Natural Rate Dynamics." *Working Paper*, May.

Philippon, Thomas. 2015. "Has the US Finance Industry Become Less Efficient? On the Theory and Measurement of Financial Intermediation." *American Economic Review* 105 (4): 1408–38. https://doi.org/10.1257/aer.20120578.

Poole, William. 1968. "Commercial Bank Reserve Management in a Stochastic Model: Implications for Monetary Policy." *The Journal of Finance* 23 (5): 769–91. https://doi.org/10.2307/2325906.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. "This Time Is Different: Eight Centuries of Financial Folly." In *This Time Is Different.* Princeton University Press. https://doi.org/10.1515/9781400831722.

Rockafellar, R. Tyrrell, and Stanislav Uryasev. 2000. "Optimization of Conditional Value-at-Risk." *The Journal of Risk* 2 (3): 21–41. https://doi.org/10.21314/JOR.2000.038.

Sack, Brian, and Joseph E. Gagnon. 2019. "Recent Market Turmoil Shows That the Fed Needs a More Resilient Monetary Policy Framework." https://www.piie.com/blogs/realtime-economic-issues-watch/recent-market-turmoil-shows-fed-needs-more-resilient-monetary.

Saporta, Victoria. 2025. "Let's Get Ready to Repo!" https://www.bankofengland.co.uk/speech/2024/july/victoria-saporta-speech-at-afme-seminar.

Savage, L. J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association* 46 (253): 55–67. https://doi.org/10.2307/2280094.

Schnabel, Isabel. 2023. "Back to Normal? Balance Sheet Size and Interest Rate Control." *ECB Speech*, March. https://www.ecb.europa.eu/press/key/date/2023/html/ecb.

sp230327_1~fe4adb3e9b.en.html.

———. 2024. "The ECB's Balance Sheet Reduction: An Interim Assessment," November.

———. 2025. "No Longer Convenient? Safe Asset Abundance and r*." *Speech at Bank of England Agenda for Research Conference*, February. https://www.ecb.europa.eu/press/key/date/2025/html/ecb.sp250225~49a8dd911a.en.html.

U.S. Bureau of Economic Analysis. 2024. "Personal Consumption Expenditures Per Capita." *FRED, Federal Reserve Bank of St. Louis*. https://fred.stlouisfed.org/series/A794RC0Q052SBEA.

Vayanos, Dimitri, and Jean-Luc Vila. 2021. "A Preferred-Habitat Model of the Term Structure of Interest Rates." *Econometrica* 89 (1): 77–112. https://doi.org/10.3982/ECTA17440.

Vissing-Jørgensen, Annette. 2023. "Balance Sheet Policy Above the ELB." *Presented at ECB Forum on Central Banking*.

Appendix

A Proofs

A.1 Proof of Proposition 3.1

($\Leftarrow$) Suppose $K^{cb} = 0$. Guess the symmetric allocation $k_E = k_N = 1$ and $z_E = z_N = z^\star$. The firm FOC is satisfied by construction. The fund managers' portfolio ratio condition (20) gives $k_E^p / k_N^p = e^{\lambda(z_E - z_N)} = e^0 = 1$, so $k_E^p = k_N^p$. The budget constraint then yields $k_E^p = k_N^p = 1$. Market clearing is satisfied and MPKs are equalised across sectors, confirming the first best.

($\Rightarrow$) Suppose $K^{cb} > 0$ and, towards a contradiction, that the equilibrium achieves the first best, so that $k_E = k_N = 1$ and $z_E = z_N$. Market clearing for eligible firms requires $k_E = k_E^p + K^{cb}/\chi = 1$ so that $k_E^p = 1 - K^{cb}/\chi < 1$. Market clearing for ineligible firms requires $k_N = k_N^p = 1$. But equal yields imply the portfolio ratio $k_E^p / k_N^p = e^{\lambda(z_E - z_N)} = 1$, which requires $k_E^p = k_N^p = 1$. This contradicts $k_E^p < 1$. Hence the first-best allocation cannot be sustained when $K^{cb} > 0$. $\square$

A.2 Proof of Proposition 3.2

For a fixed total portfolio $\bar{K}^{cb}$ and finite $\lambda > 0$, expanding the central bank's eligibility set strictly reduces the aggregate output loss from capital misallocation.

Suppose instead of allocating its entire portfolio to E, that the central bank holds a fixed total portfolio $\bar{K}^{cb}$, allocated with intensity q_E per firm in E and q_N per firm in N, such that $\chi q_E + (1-\chi) q_N = \bar{K}^{cb}$. For $\lambda > 0$ and finite, the aggregate output loss from capital misallocation is strictly decreasing in q_N for $q_N < q_E$ and equals zero if and only if $q_E = q_N$.

Proof. Capital in each sector is $k_E = k_E^p + q_E$ and $k_N = k_N^p + q_N$. The firm FOC gives $z_i = A\alpha k_i^{\alpha-1}$, so $z_E < z_N$ whenever $k_E > k_N$. The portfolio FOC requires $k_E^p / k_N^p = e^{\lambda(z_E - z_N)}$.

When $q_E > q_N$, eligible firms are inefficiently over-capitalised, so $z_E < z_N$ and $k_E^p / k_N^p < 1$: the private manager tilts away from eligible assets, partially offsetting the central bank's concentration. However, for finite λ, the offset is less than one-for-one with the central bank. Suppose, toward a contradiction, that private rebalancing fully restores $k_E = k_N$. Then $z_E = z_N$, implying $k_E^p = k_N^p$. But then $k_E - k_N = q_E - q_N > 0$, a contradiction. Hence $k_E > k_N$ in any equilibrium with $q_E > q_N$.

Now consider a reallocation $dq_N > 0$, $dq_E = -\frac{(1-\chi)}{\chi} dq_N$ holding $\bar{K}^{cb}$ fixed. This directly reduces $q_E - q_N$. The private manager responds: the narrower yield gap raises k^p_E/k^p_N toward 1, shifting private capital back toward eligible assets. But this rebalancing only partially undoes the direct effect. The net effect is a reduction in $|k_E - k_N|$. By strict concavity of k^α, the output loss $|\Delta Y| = |A[\chi k_E^\alpha + (1-\chi)k_N^\alpha - 1]|$ is strictly increasing in $|k_E - k_N|$, so the distortion falls. At $q_E = q_N$, the unique equilibrium is symmetric ($k_E = k_N = 1$, $z_E = z_N$) and the loss is zero. $\square$

B The Ramsey problem: derivation and characterisation

This Appendix derives the planner's first-order condition (26) and characterises the three regimes in which the optimal balance sheet can fall.

The recursion below evaluates the economy with the exogenous shock ξ held at its mean; aggregate risk enters through the endogenous probability of runs. The quantitative analysis of Section 5 maximises welfare in the risky steady state, where agents price in the full distribution of ξ, in which p, C_n and $\mathcal{W}_r$ all vary with ξ. The FOC derived here serves to decompose the trade-off.

B.1 The planner's problem

The planner chooses K^{cb} to maximise $\mathcal{W}(K^{cb}) = \mathcal{W}_n(K^{cb})$, where the no-run value satisfies

$$\mathcal{W}_n = u(C_n) + \beta\Big[(1-p)\mathcal{W}_n + p\mathcal{W}_r\Big], \tag{27}$$

with $\mathcal{W}_r$ the value of entering a run from the no-run steady state. Differentiating gives

$$\Big[1 - \beta(1-p)\Big]\frac{d\mathcal{W}_n}{dK^{cb}} = u'(C_n)\frac{dC_n}{dK^{cb}} + \beta(\mathcal{W}_n - \mathcal{W}_r)(-p_K) + \beta\, p\,\frac{d\mathcal{W}_r}{dK^{cb}}. \tag{28}$$

The last two terms are the run-risk benefits, terms (i) and (ii) in (26). The remainder of this Appendix characterises dC_n/dK^{cb}.

Steady-state consumption is, from (23) and (21),

$$C_n = \chi y_E + (1-\chi)y_N + W^h + w^b - F(K^h),$$

with $y_i = Ak_i^\alpha$ and total capital fixed at one, $K^h + K^b + K^{cb} = 1$. Differentiating and using the firm FOCs $z_i = \alpha A k_i^{\alpha-1}$,

$$\frac{dC_n}{dK^{cb}} = \chi z_E \frac{dk_E}{dK^{cb}} + (1-\chi) z_N \frac{dk_N}{dK^{cb}} \; - \; \alpha^h K^h \frac{dK^h}{dK^{cb}}. \tag{29}$$

The first two terms are the production-side effect of changing the allocation of capital; the last is the household-side effect of changing how much capital households manage themselves.

B.2 Capital misallocation cost

Because total capital is fixed at one, differentiating $\chi k_E + (1-\chi)k_N = 1$ gives

$$\chi \frac{dk_E}{dK^{cb}} + (1-\chi) \frac{dk_N}{dK^{cb}} = 0. \tag{30}$$

Substituting into the first two terms of (29),

$$\chi z_E \frac{dk_E}{dK^{cb}} + (1-\chi) z_N \frac{dk_N}{dK^{cb}} = -(z_N - z_E)\, \Phi(K^{cb}), \qquad \Phi(K^{cb}) \equiv \chi \frac{dk_E}{dK^{cb}}. \tag{31}$$

The factor Φ records the equilibrium incidence of central-bank asset purchases on eligible firms, net of private fund-manager rebalancing. The following section derives a closed form and shows $\Phi \in (0,1]$ with $\Phi \to 1-\chi$ as $\lambda \to 0$ and $\Phi \to 0$ as $\lambda \to \infty$.

The capital misallocation cost is therefore $u'(C_n)\,(z_N - z_E)\,\Phi$: the marginal-utility-weighted product of the MPK gap across sectors and the share of the central bank's purchase that lands on eligible firms. It is zero when $K^{cb} = 0$ or $\lambda \to \infty$, continuous, and strictly positive otherwise. In the calibrated model it is increasing in K^{cb} throughout the relevant range.

B.2.1 Incidence of central-bank purchases

The production block (Section 3) is a system of four equations in (k_E, k_N, z_E, z_N):

$$\begin{aligned} z_i &= A\alpha\, k_i^{\alpha-1}, \quad i \in \{E, N\}, \\ \frac{k_E^p}{k_N^p} &= e^{\lambda(z_E - z_N)}, \\ \chi k_E + (1-\chi)k_N &= 1, \\ k_E^p &= k_E - K^{cb}/\chi, \qquad k_N^p = k_N. \end{aligned} \tag{32}$$

The financial block enters (32) only through total capital, which is normalised to one, so Φ is a deterministic function of K^{cb} and the production primitives $(\alpha, A, \chi, \lambda)$.

Differentiate (32) totally with respect to K^{cb} and express every derivative in terms of $\Phi \equiv \chi\, dk_E/dK^{cb}$. First, adding-up (30) gives the sectoral responses,

$$\frac{dk_E}{dK^{cb}} = \frac{\Phi}{\chi}, \qquad \frac{dk_N}{dK^{cb}} = -\frac{\Phi}{1-\chi}.$$

Second, the definitions of private holdings, $k_E^p = k_E - K^{cb}/\chi$ and $k_N^p = k_N$, give the private responses: the eligible one carries the direct displacement $-1/\chi$,

$$\frac{dk_E^p}{dK^{cb}} = \frac{\Phi - 1}{\chi}, \qquad \frac{dk_N^p}{dK^{cb}} = -\frac{\Phi}{1-\chi}.$$

Third, the firm FOCs give the yield responses, $dz_i/dK^{cb} = -(1-\alpha)(z_i/k_i)\, dk_i/dK^{cb}$.

Now differentiate the portfolio condition in log form, $\log k_E^p - \log k_N^p = \lambda(z_E - z_N)$:

$$\frac{dk_E^p/dK^{cb}}{k_E^p} - \frac{dk_N^p/dK^{cb}}{k_N^p} = \lambda\left(\frac{dz_E}{dK^{cb}} - \frac{dz_N}{dK^{cb}}\right), \tag{33}$$

and substitute the three sets of responses:

$$\frac{\Phi - 1}{\chi k_E^p} + \frac{\Phi}{(1-\chi)k_N^p} = -\lambda(1-\alpha)\,\Phi\left[\frac{z_E}{\chi k_E} + \frac{z_N}{(1-\chi)k_N}\right].$$

This is linear in Φ; solving,

$$\Phi(K^{cb}) = \frac{1/(\chi k_E^p)}{\dfrac{1}{\chi k_E^p} + \dfrac{1}{(1-\chi)k_N^p} + \lambda(1-\alpha)\left[\dfrac{z_E}{\chi k_E} + \dfrac{z_N}{(1-\chi)k_N}\right]}. \tag{34}$$

All denominator terms are strictly positive for $\lambda \geq 0$, $\alpha \in (0,1)$, $K^{cb} \in [0, \chi)$, and the first equals the numerator, so $\Phi \in (0,1]$ unconditionally. As $\lambda \to 0$ the third denominator term vanishes and, at $K^{cb} = 0$, $\Phi = 1 - \chi$: private capital does not rebalance and the purchase concentrates fully on eligible firms. As $\lambda \to \infty$, $\Phi \to 0$: private rebalancing fully offsets the central bank's direct effect. In this limit the yield gap also closes, $z_N - z_E = \log(k_N^p/k_E^p)/\lambda \to 0$, so the misallocation cost $(z_N - z_E)\Phi$ vanishes at rate $1/\lambda^2$. Φ is continuous and decreasing in λ at any $K^{cb} > 0$: more elastic private demand implies lower capital misallocation.

B.3 The capital-reallocation term

The third term of (29), $-\alpha^h K^h \, dK^h/dK^{cb}$, depends on how much of the central bank's purchase the household ends up absorbing. Total capital is fixed, $K^h + K^b + K^{cb} = 1$, so

$$\frac{dK^h}{dK^{cb}} = -\Big(1 + \frac{dK^b}{dK^{cb}}\Big), \tag{35}$$

and, writing $s_b \equiv -\,dK^b/dK^{cb}$ for the intensity with which banks shed capital as reserves rise,

$$-\,\alpha^h K^h \frac{dK^h}{dK^{cb}} = \alpha^h K^h \Big(1 + \frac{dK^b}{dK^{cb}}\Big) = \alpha^h K^h \, (1 - s_b). \tag{36}$$

The 1 is the direct displacement of household capital that would occur if banks held their portfolios fixed; the $-s_b$ nets off the bank's endogenous reallocation. Together they define the capital-reallocation term of the planner's FOC,

$$\mathrm{MB}_K \;\equiv\; u'(C_n)\, \alpha^h K^h \, (1 - s_b), \tag{37}$$

term (iii). It is a management-cost *saving* when $s_b < 1$ (banks under-absorb, so households are the residual sellers) and a *cost* when $s_b > 1$ (banks substitute into reserves faster than the central bank removes capital, so households must take up the excess).

Sign of s_b. Two opposing forces determine the sign of s_b. Reserves relax the incentive constraint $QK^b = (\psi/\theta)n$ by raising the franchise value ψ, which raises intermediated capital *in value* (pushing $s_b < 0$); but reserves are also the bank's preferred safe asset, so banks rebalance toward them, and as Q rises they hold fewer units of capital for any given value (pushing $s_b > 1$). Because the management cost is convex in units, the unit response governs (36). In the calibrated model the substitution force dominates: $s_b \approx 1.1$–$1.6 > 1$ at every reserve level, so $\mathrm{MB}_K < 0$ – the reallocation term is a small net cost, not a saving. The sign is consistent with the empirical evidence in W. Diamond, Jiang, and Ma (2024), who find that reserve injections reduce, rather than expand, bank credit.

B.4 The planner's FOC

Combining (31), (36), and (29),

$$\frac{dC_n}{dK^{cb}} = \underbrace{-(z_N - z_E)\, \Phi(K^{cb})}_{\text{misallocation cost}} + \underbrace{\alpha^h K^h \, (1 - s_b)}_{\text{capital reallocation } (\mathrm{MB}_K/u')} . \tag{38}$$

Substituting into (28) recovers (26), with MB_K of (37) as term (iii).

Define the marginal cost and marginal benefit schedules,

$$\mathrm{MC}(K^{cb}) \equiv u'(C_n)\,(z_N - z_E)\,\Phi(K^{cb}), \tag{39}$$

$$\mathrm{MB}(K^{cb}) \equiv \beta(\mathcal{W}_n - \mathcal{W}_r)(-p_K) + \beta p\,\frac{d\mathcal{W}_r}{dK^{cb}} + \mathrm{MB}_K(K^{cb}). \tag{40}$$

MC is continuous and strictly positive for $K^{cb} > 0$; in the calibrated model MC is increasing and MB decreasing over the relevant range, so the crossing is unique. MB is continuous at $\bar{K}$ but falls steeply approaching it. With $p = \delta_0\big[1 - \mathbb{E}(x)\big]^{\delta_1}$ and $\delta_1 > 1$, terms (i) and (ii) vanish as $\varepsilon^{\delta_1 - 1}$ and ε^{δ_1} respectively in the distance $\varepsilon = \bar{K} - K^{cb}$: for $\delta_1 \in (1, 2)$ the level of MB is continuous at $\bar{K}$ while its left-derivative diverges. Past $\bar{K}$ only MB_K survives; in the calibration $\mathrm{MB}_K < 0$, so MB $<$ MC and the planner does not expand beyond satiation.

The steep approach also explains why the location of the optimum is insensitive to the level of marginal cost. Setting the leading term of MB equal to MC, the distance of the crossing from satiation scales as $\varepsilon^* \propto \Big[\mathrm{MC}(\bar{K}) - \mathrm{MB}_K(\bar{K})\Big]^{1/(\delta_1 - 1)}$, an exponent of roughly 2.2 at the calibrated $\delta_1 = 1.45$. Any marginal cost below the scale of the liquidity benefit therefore delivers an optimum concentrated just below $\bar{K}$, and variation in the cost within that range moves the optimum little.

C Computation

The computational strategy relies on a time-iteration algorithm based closely on that described in Gertler, Kiyotaki, and Prestipino (2016), generalised to find the global solution under uncertainty over the shock ξ_{t+1}.

For computation, write the bank's franchise value per unit of net worth, the aggregate counterpart of (8), as

$$\psi_t = \frac{1 - p_t}{N_t}\,\mathbb{E}_t\left[\Omega_{t,t+1}\,\frac{(N_{t+1} - w^b)_+}{\sigma}\right],$$

where $\Omega_{t,t+1}$ is the exit-adjusted discount factor of Section 2.3 and, by (11), $(N_{t+1} - w^b)/\sigma$ is next period's net worth of the banks already operating at t, stripped of the transfer to entrants.

Let the state of the economy if a run has **not** happened be denoted by

$$S = (N, \xi, m),$$

and the state in case a run **has** happened be

$$S^* = (0, \xi, m).$$

The belief shock ξ of (15) follows an AR(1) process,

$$\xi_t = \rho_\xi \, \xi_{t-1} + \sigma_\xi \, \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0,1),$$

which is discretised by the Rouwenhorst method into states $\Xi = \{\xi_1, \ldots, \xi_{n_Z}\}$ and transition matrix P.

Policy functions are approximated as piecewise linear functions over the state space:

$$\{\, Q(S),\ C^h(S),\ \psi(S),\ N'(S,\xi'),\ p(S),\ \bar{R}(S),\ R^f(S),\ V(S),\ \Gamma(S)\,\}, \quad S \in [w^b, \bar{N}] \times \Xi \times [0, \bar{m}],$$

and, in run states,

$$\{\, Q^*(S^*),\ C^{h*}(S^*),\ V^*(S^*),\ \Gamma^*(S^*)\,\}, \quad S^* \in \{0\} \times \Xi \times [0, \bar{m}].$$

Here $\bar{R}(\cdot)$ is the deposit rate (2), $R^f(\cdot)$ the risk-free rate (4), and $\Gamma(\cdot)$ and $\Gamma^*(\cdot)$ denote the transition laws for next-period states implied by equilibrium conditions.

Computational algorithm (time iteration)

1. **Approximation space.** Choose a functional space for equilibrium objects, and define state grids:
$$G \subset [w^b, \bar{N}] \times \Xi \times [0, \bar{m}], \qquad G^* \subset \{0\} \times \Xi \times [0, \bar{m}].$$

2. **Initial guesses.** Set $i = 0$ and guess

$$\mathrm{NRPol}_i = \{\, Q_i(S),\ C_i^h(S),\ \psi_i(S),\ N_i'(S,\xi'),\ p_i(S),\ \bar{R}_i(S),\ R_i^f(S),\ V_i(S),\ \Gamma_i(S)\,\}_{S \in G},$$

$$\text{RPol}_i = \{\, Q_i^*(S^*),\ C_i^{h*}(S^*),\ V_i^*(S^*),\ \Gamma_i^*(S^*)\,\}_{S^* \in G^*}.$$

For each current shock node ξ_ℓ, define per-node successor states using transition probabilities $P_{\ell k}$:

$$S_{i,k}^{NR}(S) = \Big(N_i'(S, \xi_k),\, \xi_k,\, m\Big), \qquad S_{i,k}^{R}(S) = \Big(0,\, \xi_k,\, m\Big), \qquad k = 1, \ldots, n_Z.$$

For zero-net-worth states, the next-period law implied by the balance-sheet block is

$$N' = (1+\sigma) w^b.$$

3. **Interpolate.** Given NRPol_i and RPol_i, construct approximating functions that match the stored policy values at every grid point and interpolate linearly between them. Expectations over the shock are formed by

$$\mathbb{E}_i[f(S') \mid S] = \sum_{k=1}^{n_Z} P_{\ell k}\, f(S_k'), \qquad S = (N, \xi_\ell, m).$$

4. **Update policy functions.** Compute updated objects $\widehat{\text{NRPol}}_i$ and $\widehat{\text{RPol}}_i$ by imposing equilibrium conditions with next-period objects replaced by their interpolated values at $\{S_{i,k}^{NR}, S_{i,k}^{R}\}_{k=1}^{n_Z}$.

   For any next-period object F, the run-weighted expectation is

$$\mathbf{F}_i(S) = \sum_{k=1}^{n_Z} P_{\ell k} \left[(1 - p_i(S)) F_i\Big(S_{i,k}^{NR}(S)\Big) + p_i(S) F_i^*\Big(S_{i,k}^{R}(S)\Big)\right], \qquad S = (N, \xi_\ell, m).$$

   The run probability is updated from the noisy recovery belief (15), expressed per unit of net worth:

$$\tilde{x}_i(S) = \sum_{k=1}^{n_Z} P_{\ell k} \left(\frac{R_{i,k}^{run}(S)}{\bar{R}_i(S)\,[\phi_i(S) + \gamma_i(S) - 1]} + \xi_k \right),$$

   where $R_{i,k}^{run}(S)$ is the liquidation value of the bank's assets at node k per unit of net worth and $\phi_i(S) + \gamma_i(S) - 1$ is deposits per unit of net worth, so that the ratio is the numerator of (15) over its denominator. The run probability (16) is then

$$p_i(S) = \delta_0 \Big[\max\{0,\, 1 - \tilde{x}_i(S)\}\Big]^{\delta_1},$$

the maximum imposing the cap $\mathbb{E}(x) \leq 1$ on the recovery rate.

The deposit rate solves the fixed point implied by (2),

$$\frac{1}{\bar{R}_i(S)} = (1 - p_i(S))\, \mathbb{E}_i[\Lambda(S')] + p_i(S)\, \mathbb{E}_i[\Lambda^*(S^{*\prime})\, x'] \,,$$

where x' is the realised recovery rate of Section 2.4.

5. **Convergence and damping.** Compute the distance between the updated objects and the current iterates,

$$d_{NR} = \max_{S \in G} \left\| \widehat{\mathrm{NRPol}}_i - \mathrm{NRPol}_i \right\|, \qquad d_R = \max_{S^* \in G^*} \left\| \widehat{\mathrm{RPol}}_i - \mathrm{RPol}_i \right\|.$$

If both are below tolerance (e.g. 10^{-6}), accept; otherwise update with damping weight $\omega \in (0, 1)$,

$$\mathrm{NRPol}_{i+1} = \omega\, \mathrm{NRPol}_i + (1 - \omega)\, \widehat{\mathrm{NRPol}}_i,$$

$$\mathrm{RPol}_{i+1} = \omega\, \mathrm{RPol}_i + (1 - \omega)\, \widehat{\mathrm{RPol}}_i,$$

and return to step 3.

D Further robustness results

This Appendix reports two results supporting Section 6. The first, Table 4, gives the fit-weighted counterpart to the decision rules of the main text: rather than weighting every admissible economy equally, each is weighted by its conformity with the calibration targets, $\propto e^{-J_i/2}$, so that economies matching the data more closely count for more. The buffers then collapse toward the central optimum — the averaging rule returns it almost exactly, and even the deepest reliable tail ($\mathrm{CVaR}_{0.95}$) raises supply by less than one percentage point — because the fit-weighted mass concentrates near the point estimate. The equal-weight rules of the main text are, in this sense, the conservative end of the analysis; disciplining the set by fit shrinks the precautionary buffer to almost nothing. The second result, to which the rest of the Appendix turns, asks which parameters drive the optimal supply and the cost asymmetry behind the precautionary buffer.

Table 4: Fit-weighted decision rules

policy	supply (% assets)	ΔM^* (pp)	premium (% of policy value)
central (point est.)	15.6	0.0	0
averaging	15.8	+0.2	0
CVaR $\tau = 0.80$	16.2	+0.6	1
CVaR $\tau = 0.90$	16.4	+0.8	1
CVaR $\tau = 0.95$	16.6	+1.0	2

Each admissible economy is weighted by its conformity with the calibration targets, $\propto e^{-J_i/2}$ (moment criterion), in place of equal weights; reported to $\tau = 0.95$ (the deepest tail with reliable support under fit-weighting, whose mass concentrates near the point estimate). Premium expressed as a percentage of the value of the reserve policy (0.15% of lifetime consumption).

This exercise is a global sensitivity analysis in the sense of Harenberg et al. (2019): parameter uncertainty is described by bounds, propagated through the model, and summarised by the resulting distribution of the optimal supply. Unlike the local measures, it does not evaluate derivatives at the point estimate. This distinction matters here since the model displays strong non-linearities.

For each admissible economy I compute the own welfare optimum M_i^* and the own satiation point (the supply beyond which the reserve demand elasticity stays within 0.25 basis points per percentage point of bank assets). Parameter influence (Table 5) is measured by least squares on the nine standardised parameters, alongside the partial R^2 from drop-one regressions. Regression-based measures are used because the admissibility screens make the retained draws mutually dependent, ruling out variance decompositions that require independent inputs. The exercise asks how the optimum varies as parameters move across their full plausible ranges.

The location of each economy's optimum is governed chiefly by the banking-sector parameters — the household's cost of holding capital α^h, the seizure rate θ, and banker survival σ — with the run-belief exponent δ_1, the household endowment W^h, and the eligibility share χ as secondary drivers. The demand elasticity λ, the run-belief coefficient δ_0, and banker start-up funds w^b barely move it. The same parameters — chiefly α^h, θ, χ and δ_1 — drive the asymmetry between the costs of under- and over-supplying reserves, and with it the strength of the precautionary motive for robust policy.

Table 5: What moves the optimum, and what moves the cost of missing it

	effect per 1 s.d.	partial R^2
A. Location of the own optimum M_i^* (pp of bank assets; $R^2 = 0.95$)		
θ	-4.28	0.89
α^h	+3.72	0.92
σ	+2.65	0.75
δ_1	-2.26	0.71
W^h	-1.37	0.66
χ	+1.28	0.65
δ_0	+0.83	0.26
λ	+0.40	0.15
w^b	-0.19	0.04
B. Regret asymmetry (under-/over-supply cost ratio; $R^2 = 0.83$)		
θ	-0.79	0.60
α^h	+0.73	0.72
σ	+0.60	0.47
δ_1	-0.50	0.41
χ	+0.42	0.53
W^h	-0.23	0.24
δ_0	+0.22	0.12
λ	+0.14	0.11
w^b	-0.04	0.01

OLS of each outcome on the nine standardised parameters across the $n = 484$ admissible economies. "Effect per 1 s.d." is the ceteris-paribus change in the outcome when the parameter rises by one standard deviation *as observed within the admissible set* (post-selection spread, not the sampling box); partial R^2 is from drop-one regressions. Because the moments identify some parameters more tightly than others (e.g. σ's within-set s.d. is ~2% of its mean), the partial R^2 column is the more directly comparable measure across parameters. Sub-grid effects (|effect| $\lesssim$ 0.2pp in Panel A) are bounded by solution-grid resolution.

E Additional comparative statics

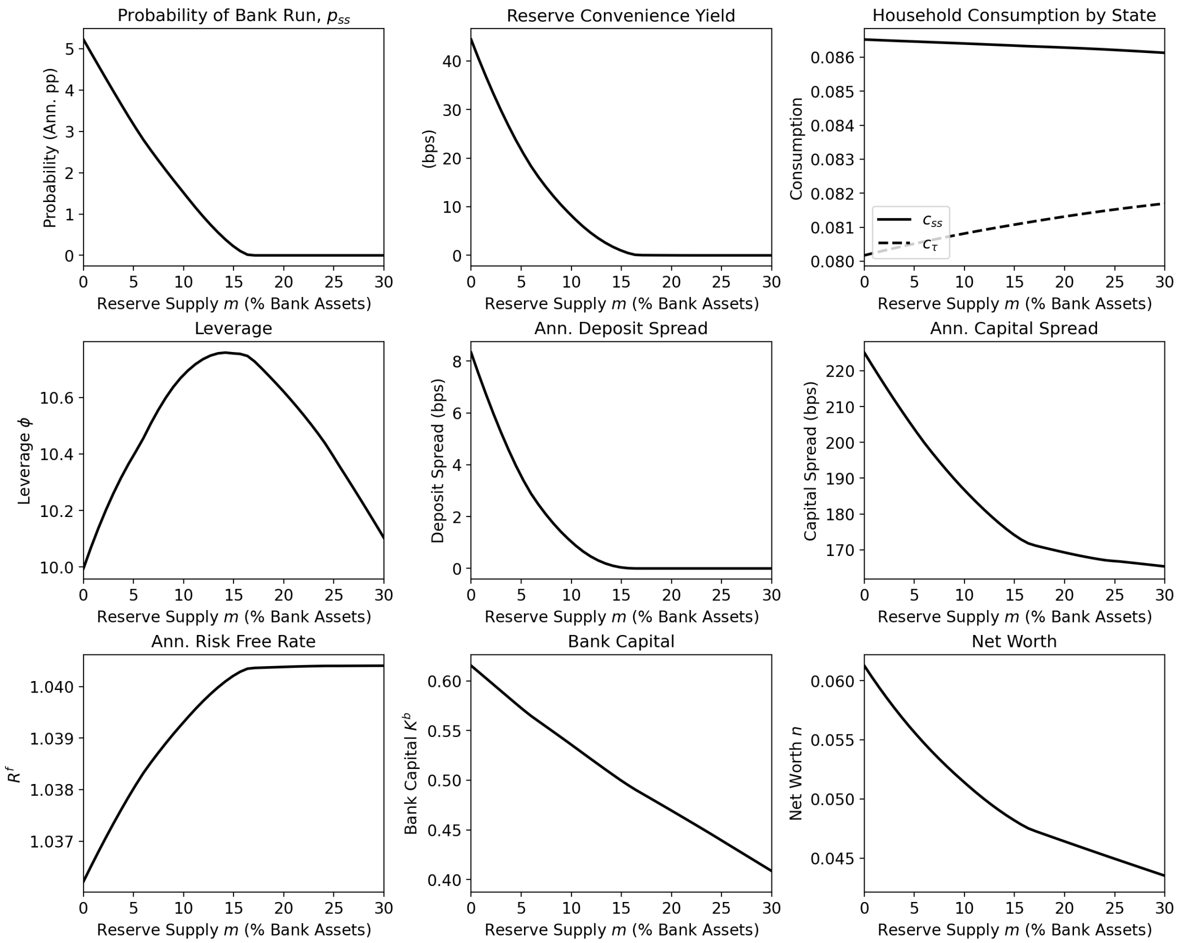

Figure 8: Additional Comparative Statics.

F Data

US

- Banking Sector Assets - Federal Reserve Economic Data, TLAACBW027SBOG.
- Demand Deposits - Federal Reserve Economic Data, WDDNS.
- Effective Federal Funds Rate - Federal Reserve Economic Data, EFFR.
- Federal Reserve Total Assets - Federal Reserve Economic Data, WALCL.
- GDP - Federal Reserve Economic Data, GDP.
- Interest on Reserves (2008 - 2021) - Federal Reserve Economic Data, IORR.

- Interest on Reserve Balances (2021 - 2025) - Federal Reserve Economic Data, IORB.
- ONRRP rate - Federal Reserve Economic Data, RRPONTSYAWARD.
- ONRRP - Federal Reserve Economic Data, RRPONTSYD.
- Other Checkable Deposits - Federal Reserve Economic Data, OCDNS.
- Other Liquid Deposits - Federal Reserve Economic Data, MDLM.
- Reserves - Federal Reserve Economic Data, WRESBAL.
- Savings Deposits - Federal Reserve Economic Data, SAVINGS.
- Reserve Demand Elasticity - Afonso et al. (2024).

UK

- Central Bank Assets - Sum over all line items listed under Assets in the Bank of England Weekly Report Table B1.1.3
- GDP - Federal Reserve Economic Data, UKNGDP.

Euro-system

- Central Bank Assets - Federal Reserve Economic Data, ECBASSETSW.
- GDP - Federal Reserve Economic Data, EUNNGDP.