A Welfare Analysis of the Central Bank Balance Sheet*

William Pagel[†‡]

This draft is a work in progress. Please do not circulate onwards.

Abstract

During the period when interest rates were constrained at the effective lower bound, central banks relied on balance sheet expansions to set the stance of monetary policy. But in the long-run, when rates are away from their lower-bound, these balance sheet policies are not needed to provide stimulus. This paper asks: what is the socially optimal size for the central bank balance sheet in the long-run? I introduce a central bank into a simple endowment economy model in which banks are liquidity mismatched and prone to sudden "bank runs". In this framework, the long-run provision of central bank reserves can reduce the likelihood of bank runs, at the cost of crowding out more productive investments. I calibrate the model to conditions before the financial crisis, and demonstrate that the model can reproduce important non-linearities in the data on the demand curve for reserves, as reserve supply moves from scarce to abundant. I compute the socially optimal size of the central bank and find that it is optimal for the central bank to expand supply of reserves almost to the point where their marginal liquidity benefit is zero, and reserve demand is satiated, but no further.

JEL Classification: E44, E58, G21, G33.

*Any views expressed are solely those of the author and should not be taken to represent those of the Bank of England or of its policy committees. I would like to thank Francesco Zanetti, Andrea Ferrero, Ricardo Reis, Ozge Akinci, and participants at the Bank of England Agenda for Research Conference (February 2025) for their helpful comments on this work.

†University of Oxford

‡Bank of England

1 Introduction

After the great financial crisis, central banks around the world underwent a massive, and unprecedented expansion. In the UK, the Bank of England balance sheet grew from roughly 5% of GDP to over 40% at its peak. The Federal Reserve, the ECB, and other central banks followed similar paths. This shift was a consequence of policies like quantitative easing (QE) and other liquidity interventions in which central banks purchased large volumes of assets from the private sector to provide emergency stimulus while interest rates were constrained at their effective lower bound (ELB). These purchases were largely financed by the creation of new central bank liabilities in the form of reserves. As short-term rates exit the ELB, the volume and composition of assets held on the central bank balance sheet are no longer dictated by the need for these 'unconventional' stimulus policies. In light of this, central banks have now begun to shrink their balance sheets, either through active asset sales or by allowing asset portfolios to mature. Since these programmes began, the sizes of central bank balance sheets have already fallen significantly (see Figure 1). The Fed, the ECB, and the Bank of England have indicated they do not intend to return to the 'lean' balance sheet regime of before the crisis, in which reserves were scarcely supplied in excess of regulatory minima. In 2024, these central banks indicated that they intend to normalise their balance sheets in a way that continues to ensure demand for reserves is satiated.[1] But despite an intense policy debate, the existing economic literature on central bank asset purchases and other balance sheet policies has said relatively little about how central banks *should* choose the size of their balance sheets, or whether the course being taken is optimal.

Against this backdrop this paper asks: what size of central bank balance sheet is optimal in the long run, or 'steady state'? When interest rates are above their lower bound, and the central bank balance sheet is not needed as an active policy tool to stabilise the economy, what size of central bank balance sheet would a social planner choose in order to maximise welfare? A larger central bank balance sheet implies a greater supply of central bank reserves. These central bank reserves are the most liquid asset available in the economy, providing banks with a near-frictionless means

[1] Federal Reserve - May 1, 2024; ECB - March 14, 2024; Bank of England - May 21, 2024

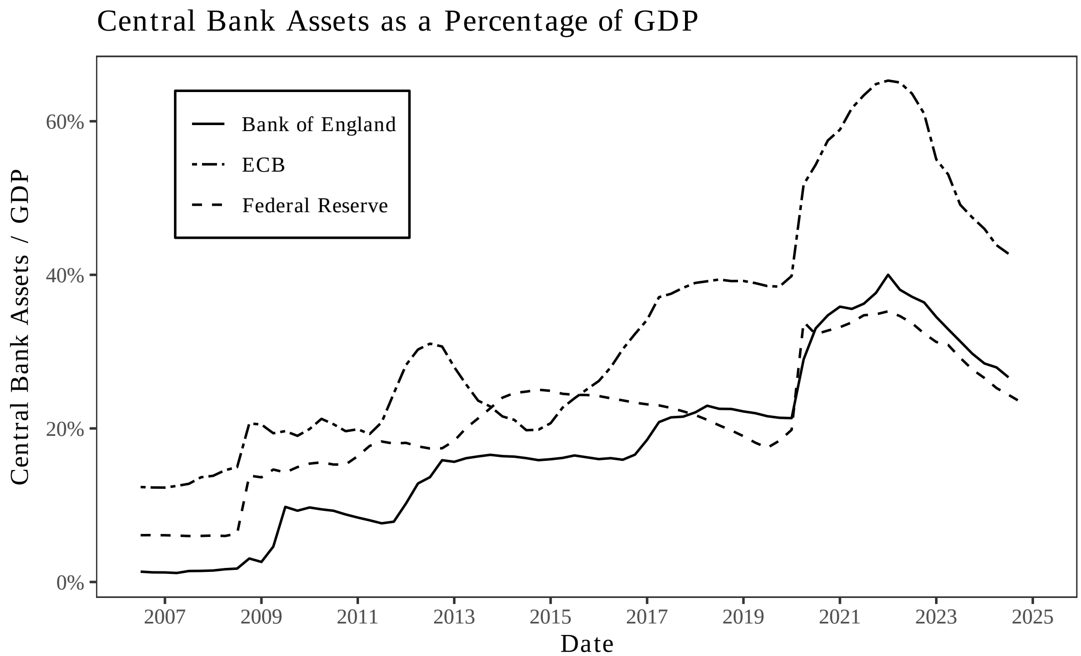

Figure 1: Central Bank Assets as a Percentage of GDP. Source: FRED, ECB, Bank of England

of payment, and a perfect store of nominal value. At the heart of my analysis are the welfare benefits *and costs* of supplying the financial sector with these uniquely liquid assets. My research question can therefore be restated: what is the socially optimal long-run provision of central bank reserves?

To answer this question, I embed a central bank which supplies reserves in a model of a simple endowment economy in which banks are exposed to bank runs as a result of a liquidity mismatch between their assets and liabilities, similar to Gertler and Kiyotaki (2015) (GK15 hereafter). In these models, stylised bank runs exist as an additional rational expectations equilibrium of this model, which emerge as a ‘sunspot’, whenever the banking sector is sufficiently fragile. In this framework, I demonstrate that central bank reserve supply in steady state can affect social welfare, without assuming that any of the agents inherently value liquidity services in their utility functions. Instead, reserves affect welfare through their impact on the likelihood of costly bank runs arising. At the margin, more central bank reserves mean private banks hold a greater quantity of liquid assets, making it easier for bankers to pay depositors in the event of a run and so, less likely that runs occur.

In addition to this benefit, I explicitly model the costs of operating a larger central bank. I assume some capital mis-allocation occurs when the central bank (rather than the private financial sector) intermediates assets, as more profitable investments are crowded out. So, operating a larger central bank implies a positive marginal cost to society.

The key contribution of this work is to find a unique, interior solution[2] to the problem of finding the optimal size of the central bank balance sheet in steady state. My results satisfy the logic of a Friedman rule: the optimal balance sheet sets the marginal social benefit of supplying reserves equal to its marginal cost. However, unlike in Friedman's original formulation, the marginal social cost of expanding supply in my framework is not zero. As a result, a well-defined optimum level of reserves exists, beyond which supplying additional liquidity is inefficient. Existing studies which analyse the welfare implications of the size of the (long-run) central bank balance sheet find corner solutions (Eren, Jackson, and Lombardo 2024; Arce et al. 2020).

In section 2, I derive the model.

In section 3, I develop a method to compute the social welfare function of this model by finding expressions for expected steady state utility in this model, accounting for the indeterminacy of the equilibrium due to the possibility of bank runs as sunspot equilibria, and introducing aggregate uncertainty to the steady state. I show that the central bank balance sheet policy acts on welfare by reducing the likelihood and severity of bank runs.

In section 4, I find the full non-linear solution to my model and show that the model can reproduce important non-linearities in the data on reserve demand as reserve supply moves from scarce to abundant. I calculate the optimal size of the central bank balance sheet and find that the optimal balance sheet falls close to the point of reserve satiation, where the marginal benefit is close to zero. Households prefer to expand the central bank and forgo steady state consumption in order to reduce the risk of financial panics, up to the point where the risk of bank runs is

[2]That is, the optimal size for central bank in steady state is neither zero, nor is it to hold the entire available supply of eligible assets.

almost entirely eliminated. These results therefore support the case for maintaining 'ample' reserve supply regimes, which keep reserve supply at or close to the point of satiating demand but not above. In fact, in my framework, even when the marginal cost of operating the central bank is very small, it is not optimal to set steady state supply of reserves in excess of this satiation point. Section 5 discusses the policy implications of these results in more detail.

Related Literature. My paper relates most closely to a new literature on the optimal size of the central bank balance sheet in steady state. Rather than studying central bank balance sheet policies as tools for stabilisation at the ELB (which has been explored in a vast body of literature), the objective of this field of research is to determine what kind of *regime* the central bank should operate in the long-run. Afonso et al. (2023) show that the supply of central bank reserves which maximises the central bank's control over interest rates in the interbank market is greater when reserve demand is uncertain. Kumhof and Salgado-Moreno (2024) study the effects of long-run reserve supply on a range of macro-variables and show that central bank balance sheet policy can affect equilibrium interest rates, lending, and GDP. But in contrast to this paper, these works do not study the *welfare* implications of long-run balance sheet policy.

Vissing-Jørgensen (2023) empirically estimates the supply of reserves which maximises the total surplus created by the availability of very liquid assets ("convenience") for the ECB and the Fed. The paper emphasises that optimal quantity of reserves depends on both the marginal convenience of both reserves, and the assets purchased by the central bank, implying that the optimal size of the central bank depends on its asset portfolio. Whereas that work assumes reserves have some convenience benefits that agents value, in this paper, I develop a micro-foundation for these benefits by considering how reserves can mitigate the risk of liquidity panics. My theoretical results are able to match empirical findings of strong non-linearities in demand for reserves, which is satiated beyond some endogenous level (Afonso, Duffie, et al. 2022; Lopez-Salido and Vissing-Jorgensen 2023; Borio 2023). This non-linearity emerges endogenously in my framework because the marginal benefit of reserves is decreasing, and falls to zero when banks already hold an ample supply.

Other theoretical welfare studies of long-run central bank balance sheet policy have found corner solutions for the optimal size of the central bank. Eren, Jackson, and Lombardo (2024) study this question in a canonical New Keynesian model with financial frictions and a rich set of financial assets. In addition, they study how the size of the central bank balance sheet affects the efficacy of monetary policy. Arce et al. (2020) compare the welfare and stabilisation properties of a large balance sheet regimes against the pre-crisis lean balance sheet in New Keynesian framework with an interbank market. In both of these frameworks, reserves alleviate financial frictions. As a result, steady-state welfare is monotonically increasing in the size of the central bank, implying that the problem of characterising the 'optimal' size for the steady state balance sheet is a corner solution: at empirically relevant calibrations, it is optimal for the central bank to hold the entire stock of eligible assets, and provide an abundant supply of reserves (in which reserves are supplied beyond the point of satiating demand for reserves). In contrast to these papers, the welfare effects of larger balance sheets need not be monotonic in my model. As a result, the optimal size of central bank balance sheet is an interior solution: at the optimum, the central bank holds a non-zero fraction of assets, but does not hold the maximum eligible amount.

In-line with existing literature, I assume some capital mis-allocation occurs when the central bank expands the supply of risk-free assets and intermediates risky financial assets. As the footprint of the central bank expands, a crowding out effect occurs. Although this cost takes a reduced form in my framework, existing literature on the optimal quantity of public debt can provide micro-foundations for this kind of 'crowding out' effect (Aiyagari and McGrattan 1998; Angeletos, Collard, and Dellas 2023).

The modelling strategy of this paper is based on recent research combining the incentive compatibility-based approach to financial frictions (Kiyotaki and Moore 1997; Gertler and Kiyotaki 2010; Gertler and Karadi 2011) with the possibility of financial panics ('bank runs'), led by Gertler and Kiyotaki (2015) and extended in Gertler, Kiyotaki and Prestipino (2016, 2020). In the spirit of Diamond and Dybvig (1983), these models capture the inherent susceptibility to bank runs of leveraged financial institutions, such as banks and shadow banks, which issue short-term liabilities and invest in more illiquid long-term assets. In these models, bank runs emerge as possible 'sunspot' equilibria. As a result, this class of model allows for endogenous, micro-founded financial crises to arise, without recourse to productivity shocks.[3]

I contribute to this field of research by developing a method to derive a steady-state welfare function for this class of model, in the case where a bank run equilibrium exists as a sunspot. I show that expected steady state utility depends on the endogenous probability of runs arising, and the severity of the outcomes which occur in those runs. Expected utility must therefore account for multiple possible equilibrium paths, including the possibility of recurrent runs, even in the absence of exogenous shocks. Previous papers have largely focussed on steady states where the endogenous probability of a run is negligible in steady state.

Finally, I note two areas that have received significant attention already, and which are beyond the scope of my paper. First, I do not consider the effect of supplying reserves on the implementation of monetary policy, and overnight interest rate volatility.[4] Equally, my framework abstracts from the implications of central bank balance sheet risk, or its implications for central bank independence from the fiscal authority.[5]

2 Model

Why Bank Runs? Before deriving the model, it is helpful to consider why considering systemic bank runs is important for assessing the welfare implications of the central bank balance sheet.

As noted above, central bank reserves are special insofar as they provide the private financial sector with instant liquidity. In order to model how this liquidity service has consequences for social welfare as a whole, the model must capture how an undersupply of private liquidity can create a negative externality. Bank runs arise due

[3]For a complementary account of endogenous financial crises, emphasising the fragility of financial markets (rather than intermediaries) see Boissay et al. (2021).

[4]For a recent summary of this topic and it's interaction with different balance sheet regimes, see Borio (2023) and Schnabel (2023).

[5]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.

to an inherent fragility in the business models of banks and which lead to contractions in lending, output and consumption. As such, they are a clear example of such an externality. Central bank reserve policy is crucial to determining system-wide liquidity, and vulnerability to runs, because banks need reserves to meet short-term liquidity outflows. This makes a model of bank runs a natural setting to study the welfare implications of reserve supply.

The model in this paper develops on the endowment economy model in GK15 with a banking sector that is exposed to bank runs. I augment their framework with a central bank that supplies reserves in exchange for capital. Commercial banks are identical to those in GK15, except for two differences. First, in my model, banks hold central bank reserves in addition to capital assets. Second, I assume bankers are members of the representative household (rather than independent agents). This assumption implies that aggregate welfare is simply the expected utility of the representative household, which simplifies normative analysis.

As in GK15, regular household members are less efficient than bankers at managing capital. This difference in efficiency will drive a liquidity mismatch: households will not in general wish to purchase bank assets at the same price as the bank, because those assets are less productive in their hands. As a result, if the bank experiences a run and needs to sell assets, it cannot liquidate the assets on its balance sheet at their normal market value. If this liquidation price is sufficiently low, even a solvent bank can go bankrupt in the event of a run and fail to repay depositors. When banks are sufficiently liquidity mismatched, this model will have two rational expectations equilibria: a 'normal times' equilibrium and a 'bank run' equilibrium.

The central bank supplies reserves by purchasing assets from financial intermediaries. Like households, the central bank is also inefficient at managing capital. This means that greater reserve supply comes at the opportunity cost of a more efficient allocation of assets. I will show that the introduction of a central bank which sets long-run reserve supply affects with the key mechanism in the model determining the likelihood and severity of bank runs: the liquidity mismatch on commercial banks balance sheet.

There are two goods determined by the market allocation: a non-durable consumption good and capital. The level of reserves is set exogenously by the central bank.

2.1 Capital

The stock of capital is durable, and does not depreciate. Households, banks and the central bank can all hold capital, and their total capital holdings equals the supply, normalised to 1.

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

where K_t^b, K_t^{cb} and K_t^h are the capital holdings of banks, the central bank and households, respectively.

When banks intermediate capital at period t, they retain their original capital (with no depreciation) and earn an income of $Z_{t+1}K_t^b$ in the following period, $t+1$, where Z_t is an aggregate productivity shock:

$$K_t^b \rightarrow \begin{cases} Z_{t+1}K_t^b & \text{output} \\ K_t^b & \text{capital} \end{cases} \tag{2}$$

Households and central banks are less efficient than banks at intermediating capital. Formally, when these agents hold capital at time t, they pay an upfront 'management cost' $f_i(K^i)$ in that period, before the return is realised at $t+1$:

$$\left.\begin{matrix} K_t^i & \text{capital} \\ f_i(K_t^i) & \text{goods} \end{matrix}\right\} \rightarrow \begin{cases} Z_{t+1}K_t^i & \text{output} \\ K_t^i & \text{capital} \end{cases} \tag{3}$$

In the case of households, these management costs are intended to reflect the fact that commercial banks have special expertise, relative to households, in allocating and monitoring their investments. Consequently, they are able to earn higher returns when making investments. In this way, $f(.)$ can be thought of as a dead-weight loss which emerges as a result of capital mis-allocation. I assume that $f(.)$ is strictly increasing and convex in capital holdings, K^i:

$$f_i(K_t^i) = \frac{\alpha^i}{2}(K_t^i)^2 \quad \text{for} \quad i \in \{h, cb\} \tag{4}$$

In the case of central banks, the management costs may reflect similar factors to households, as well as other operational costs of running a central bank balance sheet and liquidity facilities. More important, however, is the fact that central banks are not profit-maximisers. Central banks operate their balance sheets in order to meet their policy mandates, not to maximise private returns. For this reason, when central banks (rather than private banks) hold financial assets, some degree of crowding out of more productive investment occurs. Versions of this assumption are a common feature of the literature on balance sheet interventions.[6] Recent work on the optimal supply of public debt has provided a number of possible micro-foundations for the idea that supply of risk-free assets by the public sector can crowd out investment and under some circumstances reduce welfare (Angeletos, Collard, and Dellas 2023).

This formulation of "management costs" is reduced form. In reality, the extent of these dead-weight losses would depend on characteristics of the assets such as complexity or riskiness. For the purposes of this analysis, the crucial assumption is simply that capital produces lower returns when held households and the central bank relative to banks. But, the precise size of these effects is not crucial to the qualitative properties of the model. In section 4.4, I demonstrate this point by examining the robustness of my results to this parameter.

2.2 Central Bank

The central bank issues reserves, m, by purchasing capital. Since the aim of this work is to analyse the effect of different steady state policies, I hold the central bank balance sheet fixed, as a policy parameter. The central bank values the items on its balance sheet at their steady state values (i.e. 'mark-to-book' accounting). This ensures that the quantity of reserves does not vary through the cycle, in the absence of any actions

[6] See for instance Gertler and Karadi (2013) and Gertler, Kiyotaki, and Prestipino (2016).

by the central bank. As a result, I drop time subscripts on m and K^{cb}. The balance sheet is therefore given by the identity:

$$Q_{ss}K^{cb} = m \tag{5}$$

where m is the quantity of reserves issued by the central bank, and Q_{ss} is the steady state price of capital. The central earns a return on capital and remunerates reserves at the risk-free rate, R_t^f.

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

$$\tau_t = (Q_{ss} + Z_t)K^{cb} - R_t^f m - f_{cb}(K^{cb}) \tag{6}$$

2.3 Households

Households choose consumption and a level of savings. Households save either by lending to the bank in the form of deposits, D_t, or holding capital directly, K_t^h. Deposits are one period bonds issued by the bank. If there is no bank run, they pay a non-contingent gross rate $\bar{R}_{t+1}$. If there is a bank run, banks are fully liquidated, and depositors all face an equal haircut on their claims on the bank: they receive the return $x_{t+1}\bar{R}_{t+1}$ on their deposits, where $x_{t+1} \in [0, 1]$ is the recovery rate in the event of bank run. The deposit rate households receive can be expressed:

$$R_{t+1}^d = \begin{cases} \bar{R}_{t+1} \text{ if no bank run} \\ x_{t+1}\bar{R}_{t+1} \text{ if bank run} \end{cases} \tag{7}$$

When $x_t < 1$, depositors are not paid the full value of their deposit contract if a bank run occurs.

In addition to the income they receive from their savings, households receive an endowment of the consumption good which is proportional to aggregate productivity, Z_tW^h, and the lump-sum transfer, τ_t, rebating profits or losses from the central bank.

Within each household some fraction, f, of the members are bankers, who operate banks (which I discuss in the following section). On top of the variable endowment, the household receives a fixed endowment w^b which it pays to the newly entering banks as start-up funds. The household earns Π_t dividends from the bankers. The household takes prices, its endowment, transfers and dividends as given.[7]

Households share consumption equally amongst members, and maximise expected utility:

$$U_t = \mathbb{E}_t \sum_{i=0}^{\infty} \ln C_{t+i} \tag{8}$$

subject to the budget constraint

$$C_t + D_t + K_t^h + f_h(K_t^h) = Z_t W^h + w^b + R_t^d D_{t-1} + (Q_t + Z_t)K_{t-1}^h + (\Pi_t - w^b) + \tau_t \tag{9}$$

where Q_t denotes the market price of capital. The household funds consumption and asset holdings from their endowment, income from their asset portfolio, dividends (net of transfers to new bankers), and the lump-sum transfer from the central bank.

Let p_t denote the probability households assign to the bank run occurring at $t+1$, and use * to denote variables in the bank run state. For the purposes of exposition, it will be helpful to write out the probabilities of a bank run explicitly (rather than including them within the expectation operator). The households' first order condition for deposits implies that:

$$\bar{R}_{t+1} = \Big[(1-p_t)\mathbb{E}_t\Lambda_{t,t+1} + p_t\mathbb{E}_t(\Lambda^*_{t,t+1}x_{t+1})\Big]^{-1} \tag{10}$$

where the stochastic discount factor takes different values in the normal and bank run states:

$$\Lambda_{t,t+1} = \beta\frac{C_t}{C_{t+1}} \quad \text{and} \quad \Lambda^*_{t,t+1} = \beta\frac{C_t}{C^*_{t+1}}$$

[7]My treatment of bankers *as members of the household* mirrors Gertler and Karadi (2013). See the following sub-section for more details.

It is informative to compare $\bar{R}_{t+1}$ with R^f_{t+1}, given by:

$$R^f_{t+1} = \Big[(1-p_t)\mathbb{E}_t\Lambda_{t,t+1} + p_t\mathbb{E}_t\Lambda^*_{t,t+1}\Big]^{-1} \tag{11}$$

When the recovery rate equals 1, the deposit rate is identical to the risk-free rate. If the recovery rate falls below 1 ($x_t < 1$) and the probability of a bank run occurring is positive, $p_t > 0$, households demand a risk premium in the return on their deposits to compensate them for the risk of a bank run such that $\bar{R}_{t+1} > R^f_{t+1}$.

Conditional on the state of the economy, the return households can earn on direct capital holdings is given by:

$$R^{k,h}_{t+1} = \frac{Q_{t+1} + Z_{t+1}}{Q_t + \alpha^h K^h_t} \tag{12}$$

where $\alpha^h K^h_t$ captures the marginal cost households must pay for holding capital directly. The return households earn on capital depends on whether or not the economy enters the bank run equilibrium in the subsequent period. In the bank run equilibrium, the price of capital falls to its 'liquidation price', Q^*_{t+1}, as banks have to fire-sell the assets on their balance sheet to meet deposit withdrawals. I discuss how to determine the liquidation price of capital in section 2.6.

Households' first-order condition for capital is:

$$1 = \mathbb{E}_t \left((1-p_t)\Lambda_{t,t+1} \underbrace{\frac{Q_{t+1} + Z_{t+1}}{Q_t + \alpha^h K^h_t}}_{R^{k,h}_{t+1}} + p_t\Lambda^*_{t,t+1} \underbrace{\frac{Q^*_{t+1} + Z_{t+1}}{Q_t + \alpha^h K^h_t}}_{R^{k,h*}_{t+1}} \right) \tag{13}$$

Whenever households retain some direct holdings of capital, this first order condition will play a role in determining the market price of capital. This means that asset prices will be falling in the marginal cost of holding capital directly, $\alpha^h K^h_t$. This mechanism will be crucial to driving a fall in asset prices during bank runs, in which households absorb all capital previously held in the banking sector.

2.4 Banks

Banks issue deposits to fund their investment in capital. Banks will face an endogenous leverage constraint limiting the amount of deposits they can issue. As is common in the literature on financial frictions, I assume banks are not infinitely lived. This prevents them from saving sufficient funds such that the leverage constraint never binds. Bankers stochastically exit the market with probability $1-\sigma$, such that on average a banker expects to operate their bank for $\frac{1}{1-\sigma}$ periods. As a result, $(1-\sigma)f$ bankers exit and pay their retained earnings back to the household. Exiting bankers are replaced one-for-one with new entering bankers such that the total number remains constant. The household provides new bankers with a one-off start-up fund, $\frac{w^b}{(1-\sigma)f}$, per new banker, amounting an aggregate value of w^b.

2.4.1 Bankers' optimisation problem

Bankers maximise the expected dividend paid back to the household when they exit the market, equal to their net worth, n_t, upon exit. Unlike GK15, I assume bankers are members of the households. As a result, they discount their future dividends with the same marginal rate of substitution as the household, and are risk-averse. The bankers' value function can be written:

$$\begin{aligned} V_t =& (1-p_t)\mathbb{E}_t\left(\Lambda_{t,t+1}(1-\sigma)n_{t+1}+\sigma V_{t+1}\right)+ \\ & p_t\mathbb{E}_t\left(\Lambda^*_{t,t+1}(1-\sigma)n^*_{t+1}+\sigma V^*_{t+1}\right) \end{aligned} \tag{14}$$

In the event of a bank run, the bank's net worth falls to zero and the bank is forced to exit with a pay-off of zero. As a result, $n^*_{t+1}=V^*_{t+1}=0$ and (14) simplifies to:

$$V_t = (1-p_t)\mathbb{E}_t\left(\Lambda_{t,t+1}(1-\sigma)n_{t+1}+\sigma V_{t+1}\right)$$

Each period, banks issue deposit liabilities to households. They use these deposits and their net-worth to finance asset holdings in capital and central bank reserves. The bank's balance sheet is described by:

$$Q_t k_t^b + m = n_t + d_t \tag{15}$$

For surviving bankers, net-worth at t is pre-determined by their net interest income on their balance sheet from the previous period, given by the gross return on assets net of the interest paid on deposits. We can express n_t as a flow of funds constraint:

$$n_t = (Q_t + Z_t)k_{t-1}^b + R_t^f m - \bar{R}_t d_{t-1} \tag{16}$$

After issuing deposits and choosing their asset holdings, bankers can choose whether to operate honestly. If the banker operates honestly, they hold assets to their maturity at time $t+1$ and repay deposits with the promised interest rate. If the banker is dishonest, they can seize a fraction, θ, of their private sector assets and divert them to their own accounts. As is standard in the literature, I assume that the process of diverting funds is time-consuming, meaning that bankers must decide whether to operate honestly before any of the uncertainty about the following period is realised - the banker does not know the value of Z_{t+1}, or whether the households will run to the bank at $t+1$. In addition, when the bank operates dishonestly at t, the depositors can force it into bankruptcy (and therefore exit) at the beginning of period $t+1$.

Bankers operate honestly at period t if and only if the present discounted value of continuing to do so, V_t, exceeds the value of the funds it is possible to divert, $\theta Q_t K_t^b$. Depositors, who form rational expectations, recognise this and only lend to the bank when they know the banker will prefer to operate honestly. As a result, the bank faces the incentive constraint:

$$\theta Q_t k_t^b \leq V_t \tag{17}$$

The banker chooses k_t^b and d_t to maximise the franchise value function in (14), subject to the incentive constraint, (17), the balance sheet condition (15), and the flow of funds constraint (16). I assume the bank takes the level of reserves as given, rather than choosing an optimal portfolio of both reserves and capital.

The bank's optimisation problem boils down to a choice of leverage. Define the variable ϕ_t as the ratio of total bank assets to net worth:

$$\phi_t = \frac{Q_t k_t^b + m}{n_t} \tag{18}$$

Combining the bank's balance sheet and flow of funds constraint, (15) and (16), the evolution of the bank's net-worth can be re-written as:

$$n_{t+1} = \left[(R_{t+1}^{k,b} - \bar{R}_t)\phi_t - (R_{t+1}^k - R_{t+1}^f)\frac{m}{n_t} + \bar{R}_{t+1} \right] n_t \tag{19}$$

where

$$R_{t+1}^{k,b} = \frac{Q_{t+1} + Z_{t+1}}{Q_t} \tag{20}$$

It is useful to define ψ_t as the ratio of the bank's franchise value to net worth, $\frac{V_t}{n_t}$. Intuitively, ψ_t can be thought of as the shadow value of funds held within the bank per unit net worth, or equivalently, the 'Tobin's Q ratio' of the bank. Dividing equation (14) and (19) by n_t, the bank's optimisation problem can be expressed as a choice of asset to net worth ratio, ϕ_t:

$$\begin{aligned} \max_{\phi_t} \psi_t &= \mathbb{E}_t\{(1-p_t)\Lambda_{t,t+1}\left(1-\sigma+\sigma\psi_{t+1}\right)\frac{n_{t+1}}{n_t}\} \\ &= \mathbb{E}_t\{\Omega_{t,t+1}[(R_{t+1}^{k,b} - \bar{R}_t)\phi_t - (R_{t+1}^k - R_{t+1}^f)\frac{m}{n_t} + \bar{R}_{t+1}]\} \end{aligned} \tag{21}$$

subject to the incentive constraint which can be re-written:

$$\phi_t \leq \frac{\psi_t}{\theta} + \frac{m}{n_t} \tag{22}$$

and the deposit rate given in (10).

The banks' stochastic discount factor is

$$\Omega_{t,t+1} = (1-p_t)\Lambda_{t,t+1}\left(1-\sigma+\sigma\psi_{t+1}\right) \tag{23}$$

$\Omega_{t,t+1}$ captures the present discounted shadow value of net-worth, weighted by the probability of exit, $(1-\sigma)$. In the event of a bank run, the bank defaults, with $V_t = n_t = 0$, and pays no dividend.[8] Consequently, the bank's value function only considers the non-bank run state, weighted by the probability of no run occurring at $t+1$, namely $(1-p_t)$.

I consider a symmetric equilibrium in which all banks choose the same asset to net worth ratio. Banks will choose the maximum possible leverage subject to the incentive constraint, as long as the *marginal return* on leverage at the optimum is strictly positive, $\frac{\partial \psi_t}{\partial \phi_t} > 0$. Noting that the banks funding cost, $\bar{R}_t$, depends on ϕ_t through the households first order conditions, the marginal return on ϕ_t is given by:

$$\frac{\partial \psi_t}{\partial \phi_t} = \mathbb{E}_t\left(\Omega_{t,t+1}(R^{k,b}_{t+1} - \bar{R}_t)\right) - (\phi_t - 1)\mathbb{E}_t\left(\Omega_{t,t+1}\right)\frac{\partial \bar{R}_t}{\partial \phi_t} \tag{24}$$

In my numerical analysis, I restrict attention to a parametrisation where the marginal return on leverage is strictly positive, such that the bank chooses the maximum possible leverage given the incentive constraint, and equation (22) binds with equality at the optimum. This assumption simplifies computation of the non-linear solution to the model.[9]

2.5 Aggregation

In a symmetric equilibrium, summing over all banks we have:

$$Q_t K^b_t = \phi_t N_t - m \tag{25}$$

Summing over both surviving banks' net interest income and the endowment of newly entering banks in the absence of bank runs, the aggregate quantity of net worth, N_t, evolves according to:

$$N_{t+1} = \sigma\left[R^{k,b}_{t+1} Q_t K^b_t + R^f_{t+1} m - \bar{R}_{t+1} D_t\right] + w^b \tag{26}$$

[8] As I show in section 2.6, the bank run equilibrium only exists when this is the case.

[9] A discussion of the case in which the incentive constraint does not bind in equilibrium of a similar model can be found in Gertler, Kiyotaki, and Prestipino (2020). In this case, banks choose lower leverage either due to a precautionary savings motive, or to lower their deposit rates.

Exiting bankers pay the remaining fraction of the net earnings on assets as dividends:

$$\Pi_{t+1} = (1-\sigma)\left[R^{k,b}_{t+1} Q_t K^b_t + R^f_{t+1} m + \bar{R}_{t+1} D_t\right] \tag{27}$$

Total output is the sum of output produced from capital and the endowments for households and bankers:

$$Y_t = Z_t K_t + Z_t W^h + w^b \tag{28}$$

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

$$Y_t = C_t + f_h(K^h_t) + f_{cb}(K^{cb}_t) \tag{29}$$

2.6 Bank Runs

I now turn to the case in which bank runs exist as a rational expectations equilibrium of this model. I consider only runs on the entire banking system, rather than individual banks.[10]

2.6.1 Existence Conditions for a Bank Run Equilibrium

Assume that at t, the bank is solvent and no bank run has occurred. A household who deposited at the bank at $t-1$ must decide at t whether to "roll over" their deposit and leave their funds in the bank, or "run" to the bank and withdraw it either for consumption or direct investment in capital.

A bank run equilibrium exists, as a 'sunspot', if and only if the household correctly believes that the bank would not be able to repay the full value of its obligation, $\bar{R}_t D_{t-1}$, if all other depositors were to run to the bank. Under these conditions, the

[10] In the absence of any friction, individual bank runs would simply lead depositors to move their funds from one institution to another. This would have no aggregate effects in the representative agent framework above.

depositor recognises that a run by the other depositors would leave the bank with a net worth of zero. The recovery rate for depositors in a run is given by the ratio of the *liquidation* value of the bank's assets to the value of its deposits:

$$x_t = \min\left[1, \frac{(Q_t^* + Z_t)K_{t-1}^b + R_t^f m}{\bar{R}_t D_{t-1}}\right] \tag{30}$$

where Q_t^* is the *liquidation* price of capital, defined in the following section. The recovery rate is bounded at 1: banks pay at most the promised value of the deposit contract. When the recovery rate is below 1, the net worth of a bank would be wiped out in the event of a run. Under this condition, it is rational for a depositor to run to the bank if and only if they expect all others to do the same. So, the bank run equilibrium exists as a sunspot whenever the following condition is satisfied:

$$x_t < 1 \tag{31}$$

The recovery rate is determined by endogenous variables, entailing that the possibility of bank runs depends on macro-economic conditions, and the reserve policy of the central bank, m.

2.6.2 The Liquidation Price of Capital

New banks cannot enter in the period of a run. This entails that if a bank run occurs, the banks pay their liabilities by selling off all their assets to households (rather than new banks).[11] Consequently, households must purchase the entire stock of banks' capital in the event of a run, entailing that:

$$K_t^{h*} = 1 - K^{cb} \quad \& \quad K_t^{b*} = 0 \tag{32}$$

Using this fact, the liquidation price of capital is determined by the household Euler equation for capital, (13). Given that the no new banks enter, the probability

[11]This assumption is necessary for the run to have an aggregate impact. The qualitative results of the model are robust to longer delays. But, bank runs become more costly as the time before new banks may enter increases, since households must continue to pay large management costs for a greater period of time.

of a run occurring in the subsequent period is zero. So, from equation (13), the liquidation price of capital is given by:

$$Q_t^* = \mathbb{E}_t\left(\Lambda^*_{t,t+1}(Q_{t+1} + Z_{t+1})\right) - \alpha^h(1 - K^{cb}) \tag{33}$$

The liquidation price is strictly lower than the market value of capital, whenever banks hold capital. This drives the liquidity mismatch in bank's balance sheets: bank assets cannot be liquidated at their full market price. The liquidation price (and its gap relative to Q_t) is therefore an index of liquidity conditions in the market.

It is instructive to consider the special case in which households believe the ex-ante probability of a run is 0 such that $p_t = 0 \forall t$. Call this the 'naive household' case. In this case, I abstract from the effect of p_t in the aftermath of a run. Equation (33) can then be iterated forward, and $Q_t^{*,\text{naive}}$ can be expressed as:

$$Q_t^{*,\text{naive}} = \mathbb{E}_t\left(\sum_{i=0}^{\infty} \Lambda_{t,t+i}(Z_{t+1} - \alpha^h K^h_{t+i})\right) - \alpha^h(1 - K^{cb}) \tag{34}$$

The liquidation price is decreasing in the size of the marginal household management cost, $\alpha^h K^{h*}$. This implies that liquidity mismatch is worse when households are more inefficient, or when they have to absorb a large quantity of capital in a run (i.e. when K^{h*} is high). From (32), the latter is a direct function of central bank balance sheet policy. When the central bank balance sheet is larger, a greater proportion of capital assets are intermediated by the central bank and do not need to be absorbed by households during a bank run. So, the liquidation price which clears the market for capital during a bank run is higher (that is, banks are more liquid) when the central bank is larger.

In contrast to capital, central bank reserves, m, are perfectly liquid. Reserves can be liquidated by the bank at full value and transferred directly to the household as 1 period risk-free government bonds in the event of a run. These bonds pay the equilibrium risk-free rate in equation (11). 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:

$$
\begin{aligned}
N_t^* &= 0 \\
N_{t+1} &= w^b + \sigma w^b
\end{aligned}
\tag{35}
$$

After this period, aggregate banks net-worth continue to evolve according to the law of motion defined in equation (26).

2.6.3 Probability of a Bank Run

The processes described above offer no mechanism to determine which equilibria is realised. An additional assumption is required to provide some function, g, which endogenises this probability and closes the model. Following GK15, I choose a simple functional form for p_t which links the probability of a bank run sunspot to the fundamental determinant of whether the bank run equilibrium exists, $\mathbb{E}_t(x_{t+1})$, and assigns a probability of zero whenever the bank run equilibrium does not exist:

$$
\begin{aligned}
p_t &= g(\mathbb{E}_t(x_{t+1})) \\
&= 1 - \mathbb{E}_t(x_{t+1})
\end{aligned}
\tag{36}
$$

Linking probability of a bank run to the recovery rate will amplify aggregate shocks to the economy, beyond any implication from the traditional financial friction in the model. Because asset prices and leverage move pro-cyclically, the recovery rate is pro-cyclical. From the equation above, the probability of runs is therefore endogenously counter-cyclical, with runs more likely in bad times. I describe this process in more detail in the numerical results in section 4.

Note that the probability of runs will also depend on the central bank balance sheet. When banks hold more reserves, a greater proportion of their assets can be liquidated at full value raising the recovery rate in the event of a run, as described above.

2.6.4 Determining the equilibrium

Even with a functional assumption for the probability of the sunspot bank run equilibrium, there is no criteria for setting out which equilibrium is realised. In this sense, the model remains indeterminate.

To assess the optimal steady state central bank balance sheet, I study how central bank policy affects the long-run expected utility in this model. An endogenous probability for the bank run state arising is sufficient to capture the effects of both the severity and the expected frequency of bank runs on expected utility.[12] As a result, I will not need to take a position on determining which equilibrium is realised at any point. I provide a full definition of the 'no bank run' equilibrium in appendix A.

However, to carry out a long-run simulation of the economy for the calculation of business cycle moments, or to model a specific scenario, a rule for determining the equilibrium would be needed. The simplest way to do so would be to determine the equilibrium as an exogenous process (that must be consistent with the endogenous existence conditions for the bank run equilibrium). For instance, Gertler, Kiyotaki, and Prestipino (2020) assume that a sunspot arises with a calibrated exogenous probability, and define the equilibrium probability of a bank run as the product of this value and an endogenous variable very similar to p_t in this framework. For the purposes of this work, however, such a rule is not necessary.

3 Steady State Expected Utility

In this section, I characterise the steady state of the model in which there is aggregate uncertainty with respect to whether or not the economy will enter a 'bank run' state. I demonstrate that even in the absence of exogenous shocks, agents cannot rule out the possibility of endogenous dynamics arising as a result of bank run sunspots, and I derive a method to calculate expected steady state utility which accounts for this possibility. I apply this method in the numerical analysis in section 4.

3.1 Defining a steady state with aggregate uncertainty

I use the term 'steady state' to mean the case where the only exogenous variable (Z_t) is held constant at its long run expected value, Z_{ss}, and the vector of all endogenous

[12] In my dynamic simulations in Appendix D, I focus on the normal times equilibrium, and how the *possibility* of bank runs serves to amplify exogenous productivity shocks.

variables, denoted X_t, is at the fixed point of the system of equations defining the 'normal times' equilibrium, set out in Appendix A. Call this vector X_{ss}. If $p_{ss} > 0$, the possibility of a bank-run, as a sunspot equilibrium, means that even in steady state, households do not rule out the possibility of endogenous dynamics arising as a result of runs, when forming their expectations. The steady state remains ex-ante uncertain with the respect to which of the two possible equilibria will be realised.

To illustrate this point, consider the system in 'steady state' at the end of some period $t-1$: $X_{t-1} = X_{ss}$ with $p_{ss} > 0$. Assume a deterministic, and constant process for exogenous shocks given by $Z_t = Z_{ss}\ \forall t$.

With probability $1 - p_{ss}$, the economy remains in the no bank run steady state. In this case, the values of all endogenous variables remain constant, and period t is identical to the previous period, with the same aggregate uncertainty *vis-a-vis* $t+1$.

With the complementary probability p_{ss}, a bank run occurs. In this case, all endogenous state variables are "reset" to their run values. By the assumptions laid out in section 2.6, these values do not depend on the values of any other variables at the end of the previous period. Therefore any period in which a bank run occurs is identical (conditional on the contemporaneous value of Z_t), and independent of the entire past. In this sense the bank run state can be said to be 'memory-less'.

In the context of the steady state, where Z_t is held constant, any period in which a bank run occurs from steady state can be denoted by the time subscript, τ, without any loss of generality. The value of the endogenous state variables at τ are given by:

Proposition 3.1. *The endogenous state variables in a period with a bank run run are independent of all previous periods, and take the following values:*

The households absorb all capital not held by the central bank: $K^h_\tau = 1 - K^{cb}$

The banks are bankrupt, such that: $D_\tau = 0 \quad \& \quad N_\tau = 0$

The price of capital is the liquidation price: $Q_\tau = Q^*_\tau(Z_{ss})$

These values determine the full vector of endogenous variables X_τ. Note, however, that X_τ is not a fixed point of the system of equilibrium conditions. So, in the period after a bank run, $\tau+1$, all variables jump onto the saddle path back to the system's steady state. This path, $\{X_t\}_{t=\tau+1}^{\infty}$, is the solution to the set of equilibrium conditions of the model (in Appendix A), given a path for the exogenous shock which is held constant in this case. If no further bank runs occur, the system remains on this saddle path back to the steady state, X_{ss}.

However, along this saddle path, the probability of further bank runs may be positive. The probability of bank runs continues to be determined by the method in equation (36). If another run were to occur, the variables would return to the bank run state, X_τ. By the "memoryless" property of bank runs, each time a run reoccurs, the process repeats with identical dynamics.

3.2 Computing Steady State Expected Utility

The expected utility of the representative household in the steady state averages over the infinite set of possible paths the economy can take from X_{ss} holding $Z_t = Z_{ss} \forall t$, weighted by their ex-ante probability.

To derive an expression for this value which can be calculated numerically, I rely on two pieces of information to simplify the state space. First, all bank runs from steady state, and the subsequent saddle paths to X_{ss} are identical (Proposition 3.1). Second, I guess and verify numerically that:

Proposition 3.2. *Conditional on no subsequent bank runs occurring and the deterministic path of the exogenous shocks, the saddle path back from the bank run state re-converges to the steady state, X_{ss}, in a finite number of periods, T.*

Let $\mathcal{L}$ denote the ex-ante[13] expected utility of the representative household at some period t where the system begins in steady state. This value can be expressed in the form of a recursive value function:

$$\mathcal{L} = (1 - p_{ss})V_{ss} + p_{ss}V_{run,\tau} \tag{37}$$

where V_{ss} expected utility for the household (or value) if no run occurs at t, and $V_{run,\tau}$ is the value if a run does occur at t, and the economy jumps to the run state.

[13] Ex-ante here denotes households do not yet know whether a run will occur at t

The value of the no run state is given by:

$$V_{ss} = \ln(C_{ss}) + \beta(1 - p_{ss})V_{ss} + \beta p_{ss} V_{run,\tau} \tag{38}$$

If the economy stays in steady state, the value function is given by the flow utility from consumption at t plus the expected value at $t+1$, which takes expected utility of the two possible states at $t+1$ weighted by their probabilities.

The value of the run state is given by:

$$V_{run,\tau} = \ln(C_{\tau}) + \beta(1 - p_{\tau})V_{run,\tau+1} + \beta p_{\tau} V_{run,\tau} \tag{39}$$

where $V_{run,\tau+i}$ is the value i periods after the run occurs, conditional on no further runs occurring in that time. The value in the run state is the flow utility of household consumption, conditional on a run occurring, plus the expected value at $t+1$. $V_{run,\tau+1}$ is the value in the case where no run occurs in the following period. $V_{run,\tau}$ is the expected value in the case where a subsequent run occurs. In effect, here the economy "returns" to state τ. Each is weighted by its respective probability.

Equations (38) and (39) define a system of 2 equations in 3 unknowns. To solve these, eliminate one unknown by noting that Proposition 3.2 implies that the value to the household T periods after a run, conditional on no subsequent runs occurring during that time, is the same as the value in steady state:

$$V_{run,\tau+T} = V_{ss} \tag{40}$$

Substituting this into (39), I then solve (39) backwards from period $\tau + T$. For instance at $\tau + T - 1$:

$$V_{run,\tau+T-1} = \ln(C_{\tau+T-1}) + \beta(1 - p_{\tau+T-1})V_{ss} + \beta p_{\tau} V_{run,\tau} \tag{41}$$

Iterating backwards, and rearranging, $V_{run,\tau}$ can be expressed as a linear function of V_{ss}:

$$V_{run,\tau} = \Gamma_0 + \Gamma_1 V_{ss} \tag{42}$$

where Γ_0 and Γ_1 are functions of the discount factor and the saddle paths for consumption and the probability of bank runs in the periods after a run occurs, $\{C_t\}_{t=\tau}^{\tau+T}$ and $\{p_t\}_{t=\tau}^{\tau+T}$:

$$\Gamma_0 = \left(1 - \sum_{i=1}^{T} \beta^i p_{\tau+T+i-(T+1)} \prod_{j=1}^{i-1}(1 - p_{\tau+T+j-(T+1)})\right)^{-1} \sum_{i=1}^{T} \beta^{i-1} \ln C_{\tau+T+i-(T+1)} \prod_{j=1}^{i-1}(1 - p_{\tau+T+j-(T+1)})$$

$$\Gamma_1 = \left(1 - \sum_{i=1}^{T} \beta^i p_{\tau+T+i-(T+1)} \prod_{j=1}^{i-1}(1 - p_{\tau+T+j-(T+1)})\right)^{-1} \left[\beta^T \prod_{i=1}^{T}(1 - p_{\tau+T-i})\right]$$

Plugging (42) into (38) provides an expression V_{ss} in terms of steady state consumption, steady state probability of a run, Γ_0, and Γ_1:

$$V_{ss} = \frac{\ln(C_{ss}) + \beta p_{ss}\Gamma_0}{1 - \beta + \beta p_{ss}(1 - \Gamma_1)} \tag{43}$$

Substituting this result into equation (37), I obtain an expression for long run expected utility, $\mathcal{L}$:

$$\mathcal{L} = p_{ss}\Gamma_0 + (1 - p_{ss} + p_{ss}\Gamma_1)\frac{\ln(C_{ss}) + \beta p_{ss}\Gamma_0}{1 - \beta + \beta p_{ss}(1 - \Gamma_1)}$$

Partial derivatives confirm that $\mathcal{L}$ is increasing steady state consumption, Γ_0, and Γ_1, and decreasing with p_{ss}.[14]

These results are intuitive. Γ_0 is proportional to present discounted value of flow utility of consumption along the saddle path from the run equilibrium. So, the welfare is decreasing as the severity of the impact of runs on consumption increases. Γ_1 is proportional to the probability that the economy successfully returns to steady state with no further runs (given by the product on the numerator) implying that the

[14]Parameter signs are inferred from numerical results, and therefore these statements apply in the neighbourhood of the parametrisation used in the numerical section.

welfare falls as the risk of recurrent runs increases. Finally, the welfare is decreasing in the probability of bank runs.

Welfare is an increasing function of steady state consumption, given by:

$$C_{ss} = Y_{ss} - \frac{\alpha^h}{2}(K^h_{ss})^2 - \frac{\alpha^{cb}}{2}(K^{cb})^2$$

This is where the cost of a larger central bank acts directly on welfare. The efficiency cost incurred by holding assets in the central bank, rather than in the financial sector, lowers the level of steady state consumption at the margin.

As $p_{ss} \to 0$ and all aggregate uncertainty is removed, V_{ss} and $\mathcal{L}$ both converge to the present discounted value of utility from an infinite consumption stream of C_{ss}. This implies that when $p_{ss} = 0$, central bank balance sheet policy only effects welfare via its effect on steady state consumption.

4 Numerical Examples

In this section, I discuss the results of some numerical examples to demonstrate the workings of the model. Given the simplicity of the model, these numerical results should be taken as illustrative, rather than as quantitative estimates.

I solve the model in a setting where agents have perfect foresight with respect to the exogenous state variable (Z). In line with the results from the previous section, I do not assume agents have perfect foresight over whether or not runs will occur in the steady state. I implement a numerical procedure, similar to that in GK15, which finds a full non-linear solution. This is important because the model inherently displays quantitatively important non-linearities. Furthermore, the equilibrium studied is not unique (because the model allows for bank runs), ruling out the possibility of using perturbation. I outline the key steps in this procedure in Appendix B.

4.1 Calibration

I calibrate the model to match the financial conditions in the US on the eve of the financial crisis, prior to the expansion of the central bank balance sheet. I set the quarterly discount factor to 0.99, and set the size of the central bank liabilities to roughly 5% of the size of the private financial sector. I normalise the steady state values of Z such that the price of capital equals 1 in the calibrated steady state.

I jointly determine the entering endowments of households and bankers, W^h and w^b, the exogenous survival probability of the bankers, σ, and the the seizure rate, θ to target the following empirical moments: I target estimates from Philippon (2015) that before the financial crisis, commercial banks operated with a leverage ratio of 8, and a return on bank assets of 200bps. I target a long-run probability of default equal to 4% in steady state, to match the average frequency of financial crises in Reinhart and Rogoff (2009). Finally, I target a spread between deposits and reserves of 30bps in line with estimates from Krishnamurthy and Vissing-Jørgensen (2012) of the safety component of the convenience yield on treasuries.

To determine the combination of values for these parameters, I solve numerically for the values that minimise a loss function which penalises deviations of the relevant steady state moments from their target values (all with equal weights).[15] The simulated steady state moments were the leverage ratio, $\frac{QK^b}{N}$, the returns on bank assets and deposits, $R^{k,b}$ and $\bar{R}$, and the probability of bank runs, p. The results yield parameter values close to those used in existing literature. They imply a positive probability of bank runs in the steady state, and consequently, banks pay a small risk premium on the deposits. The model hits the calibration targets, with the exception of deposit spreads, which it fails to accurately match. Deposit spreads in the calibrated steady state are 8bps. I discuss the causes of this in section 4.3.

Following GK15, I set the parameter that governs the inefficiency of households relative to the financial sector (α^h) such that households are willing to buy capital in the bank run equilibrium at a price that generates an increase in credit spreads of 200bps, consistent with the evidence from the financial crisis (Gilchrist and Zakrajšek

[15] This method could be thought of as an application of simulated method of moments, outlined in Ruge-Murcia (2007).

2012).

The efficiency of the central bank. The only novel parameter in my framework is α^{cb} which governs the size of the management fee incurred by the central bank when it holds capital. This parameter is highly uncertain, given the lack of empirical evidence capital misallocation as a result of central bank balance sheet policy.

For my baseline analysis, I assume that the central bank is no more efficient than households at intermediating capital, such that $\alpha^{cb} = \alpha^h$. Existing literature on central bank liquidity interventions typically assumes the central bank is less efficient than financial intermediaries at making loans, but better than households (e.g. Gertler and Karadi (2013), Gertler, Kiyotaki, and Prestipino (2016)). My assumption is therefore relatively pessimistic on the efficiency of the central bank at financial intermediation. As such, it produces an upper-bound estimate of dead-weight loss which arises as a result of the management of financial assets by the central bank, and a lower bound estimate of the optimal size of the central bank.

However, given the uncertainty of this parameter, Section 4.4 shows robustness analysis for $\alpha^{cb} \in (0, \alpha^h]$ and demonstrates that the quantitative results of the model are not sensitive to this parameter.

Table 1 displays the baseline calibration.

Table 1: Parameter Values

Parameter	Value	Description
β	0.990	Discount rate
σ	0.929	Bankers survival probability
θ	0.205	Seizure rate
α^h	0.014	Household managerial cost
α^{cb}	0.014	Central Bank managerial cost
Z	0.016	Steady state productivity
ZW^h	0.059	Household endowment
w^b	0.002	Bankers start-up funds

4.2 Optimal Balance Sheet

Having parametrised and solved the model, I carry out comparative statics with the respect to the size of the central bank balance sheet. Since expected utility, and the values of all endogenous state variables depend on the entire saddle path for endogenous variables in the event of a run, statics must be computed numerically. I discuss the results and then consider implications for policy. Figure 2 shows ex-ante expected utility, $\mathcal{L}$, as derived in section 3.2 evaluated at different sizes of steady state central bank balance sheet, K^{cb}. In addition to the results from the baseline model (shown in the solid black line), I consider two additional cases to illustrate the roles of the novel elements of my theory, namely the risk of bank runs as motivation for reserve supply, and the explicit treatment of the cost of the central bank. The y axis of Figure 2 is normalised to show welfare in consumption equivalent terms (pp), relative to the optimal size of the central bank in the baseline model.

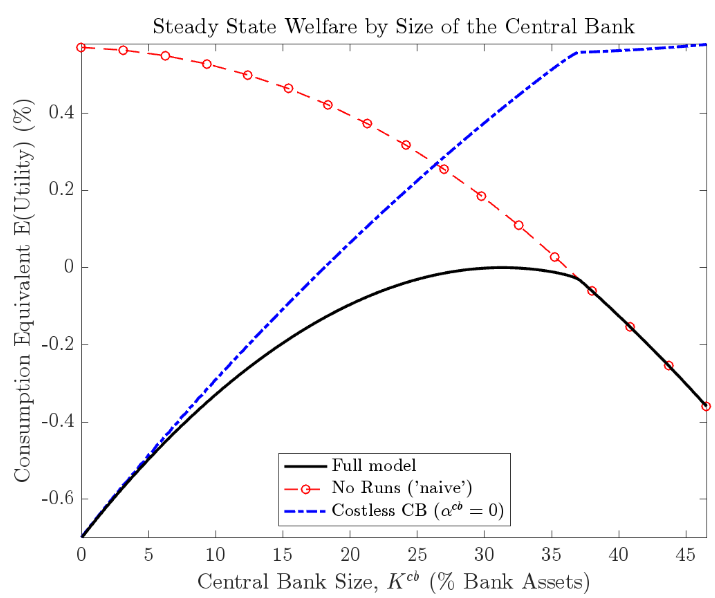

Figure 2: Comparative Statics: Ex-Ante Expected Utility. Notes: The figure plots the expected utility in 3 cases: the solid black line shows the full model. The hollow-dotted line shows results for the case where households assign zero probability to the bank run state, regardless of the banks' financial conditions (are 'naive'). The grey-dashed line shows the case where $\alpha^{cb} = 0$

In the case with 'naive households' (shown as the hollow dotted in line in Figure

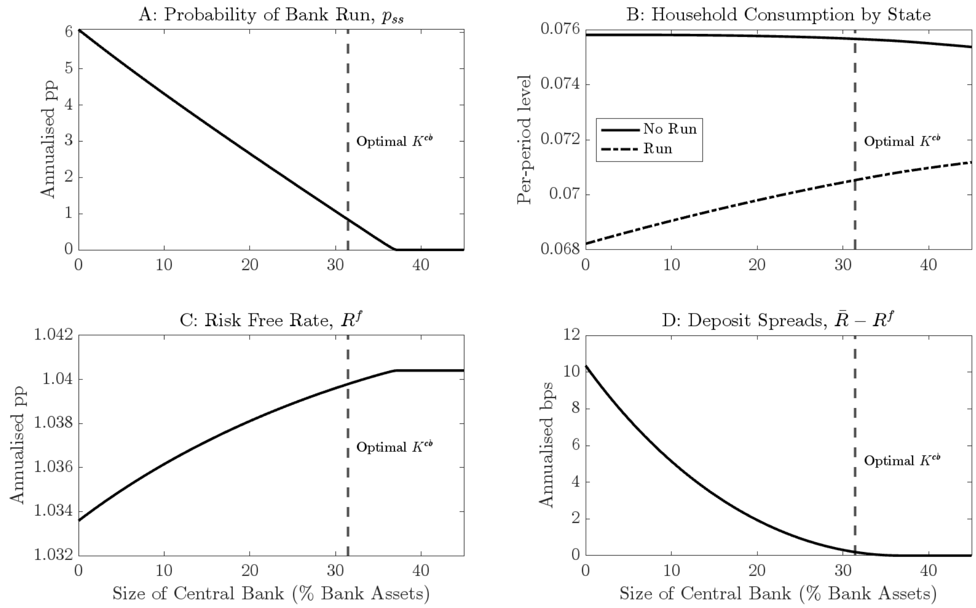

Figure 3: Comparative Statics: Endogenous Variables. Notes: The optimal size of the central bank balance sheet is marked by a dotted line for reference.

2), depositors always perceive the ex-ante probability of bank runs to be zero. This effectively switches off the bank run component of this model in expectation, and the model becomes isomorphic to a standard model of financial frictions with an endogenous leverage constraint and 'financial accelerator' dynamics. The optimal size of the central bank balance sheet is zero - the size which minimises the management cost incurred by the central bank. Intuitively, households do not recognise the liquidity mismatch on banks balance sheets or the fragility this entails, so the additional liquidity created by the provision of central bank reserves has no welfare benefit in expectation.[16] But, when the central bank supplies reserves by purchasing capital, it increases the total dead-weight loss by moving capital from (efficient) financial intermediaries to the (relatively less efficient) central bank. At this parametrisation, welfare is therefore monotonically decreasing in the size of the central bank.

[16] Note that if a run were to occur as an exogenous shock (assuming the conditions for a run were satisfied ex-post such that $x_t < 1$), then a larger central bank balance sheet would reduce the severity of the fall in consumption, because banks would have a greater stock of liquid assets with which to pay households. However, because households assign a probability of zero to this event, it has no effect on expected utility.

In the full model, households account for the probability of a bank run when forming expectations. The welfare results for this case are shown by the solid black line in Figure 2. Figure 3 shows comparative statics for key endogenous variables. Increasing the provision of reserves reduces liquidity mismatch on bank balance sheets, raising the recovery rate and lowering the probability of a bank run (Fig. 3, Panel A). This reduces the risk premium banks must pay their depositors to compensate them for the risk of runs, pushing down deposit spreads (Fig. 3, Panel D). Lower bank funding costs relax the banks' leverage constraint and push down on lending spreads, allowing the banks to increase deposits and asset holdings. Consequently, financial intermediation by the banking sector increases, making the asset allocation more efficient. More importantly, the household expects fewer runs to occur as the size of the balance sheet increases. Consumption in the bank run state, C_{τ}, is also increasing in the size of the central bank balance sheet implying that a larger central bank makes runs less severe, if they do occur (Fig. 3, Panel B). The reduction in likelihood and severity of bank runs generates substantial improvements in expected utility for the household. As a result of the reduction in aggregate uncertainty, households have a lower demand for precautionary savings and the risk-free rate rises in the size of the central bank balance sheet (Fig. 3, Panel C). This matches findings in Arce et al. (2020) that equilibrium short-term rates are increasing in the size of the central bank, albeit via a different mechanism.

When reserve supply becomes large enough to rule out the bank run sunspot equilibrium (i.e. $x_{ss} = 1$ and $p_{ss} = 0$), the marginal liquidity benefit of reserves is zero: there is no risk of bank runs in steady state and as a result further increases in steady state reserve supply have no further effect on the risk of bank runs - as a probability, p_t cannot fall below zero (Fig. 3, Panel A). This point can be thought of as the threshold above which reserve demand is satiated. Beyond this point, banks' marginal cost of funding becomes insensitive to reserve supply. Banks already have a sufficient volume of liquid assets to remain solvent in a run, so additional reserves simply replace high yielding capital assets, with low yielding reserves. When reserve demand is satiated, the full model becomes isomorphic to the 'naive household' case.

The optimal-sized balance sheet lies just below the point where reserve demand is satiated. When the risk of bank runs is sufficiently low, the marginal cost of increasing the size of the central bank outweighs the expected utility benefits of further reducing the likelihood of runs. Beyond this point, further increases in the size of the central bank reduce welfare as a result of the inefficiency of the central bank in managing capital.

Comparing the full model to the naive household case illustrates a key contribution of this paper: explicitly modelling financial panics is key to understanding the liquidity value of reserve supply. This model provides a simple micro-foundation for the welfare derived from very liquid assets (even in a representative agent framework) that has been studied in the extensive literature on convenience yields (Krishnamurthy and Vissing-Jørgensen (2012), Lopez-Salido and Vissing-Jorgensen (2023), Vissing-Jørgensen (2023)). This model matches key features of the convenience benefits in the theoretical literature. For instance, it is commonly assumed that the utility derived from such assets ('convenience') is both concave and satiated (that is, beyond some exogenously determined point, the marginal benefit of additional liquid assets is zero).[17] Both of these features emerge endogenously in the expected utility function of my framework. In consumption terms, the effect of reserve supply on the level of welfare is small, but not negligible. The representative household would sacrifice approximately 0.7 percentage points of their per-period consumption to move from the steady state with $K^{cb} = 0$ to the steady state with the optimal balance sheet, and avoid the risk of bank runs.[18] In the US in 2024, this corresponds to roughly $35 per month on average, comparable to the monthly cost of a mobile phone contract.[19]

Finally, to illustrate the importance of the parameter α^{cb} in normative analysis, the grey-dashed line in Figure 2 considers the edge case where the central bank is no less efficient at financial intermediation than banks. This assumption implies that no mis-allocation of assets occurs as a result of public intervention in private credit markets, and that there are no negative externalities from operating large central bank balance sheets. In this case, larger central bank balance sheets monotonically increase

[17]This kind of preference is also commonly assumed in macro-finance models or models with a role for money. See for instance Abadi, Brunnermeier, and Koby (2023), Benigno and Benigno (2022) and Afonso et al. (2023) for recent examples.

[18]To calculate this value, I compare the *certainty equivalent* levels of per period consumption for each level of expected life-time utility.

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

welfare, reproducing results from previous theoretical work on the optimal size of the central bank (Arce et al. (2020); Eren, Jackson, and Lombardo (2024)). More reserves initially increase welfare through the mechanisms described above. But now, when the ex-ante probability of bank runs falls to zero, additional reserves continue to increase welfare by relaxing the leverage constraint on the financial intermediary, and allowing them to increase deposits and asset holdings. While this mechanism is present in the full model and the naive household case, its effect on welfare is dominated by the negative effect on consumption from the greater dead-weight loss. The size of the α^{cb} and the functional form of the dead-weight losses which arise from large central bank asset holdings are quantitatively important for assessments of the optimal balance sheet, and should be a priority for future research and policy analysis.

A Figure showing statics for other endogenous variables is included in Appendix C. And, Appendix D shows the response of the economy to productivity shocks under 3 different sizes of balance sheet, corresponding to pre-crisis levels, optimal size, and a ‘too large’ calibration.

4.3 Empirical validity

Next, I evaluate the model’s key mechanisms by comparing the key theoretical predictions of the model to data from the USA.

It is now well-documented empirically that the demand for reserves becomes satiated above some threshold level of reserves. Beyond this point, the demand curve for reserves is flat and the spread at which reserves are lent and borrowed in the market becomes insensitive to changes in supply.

The left panel of Figure 4 presents econometric estimates of the reserve demand elasticity (the slope of the demand curve for reserves) from Afonso, Giannone, et al. (2022), published at daily frequency on FRED. These estimates show how the spread (in basis points) between the Effective Federal Funds Rate (EFFR) and the interest paid on reserves responds to a shock in reserve supply equal to 1pp of bank assets. I plot this elasticity against the size of the central bank balance sheet (normalised by the size of the banking sector). The right panel compares these data to their theoretical counterpart in the model: the elasticity of the spread between the banks’ marginal cost of funds (i.e the deposit rate) and the rate paid on reserves, to a 1pp change in reserve supply.[20]

The model qualitatively matches the key non-linearity in the data: when reserves are scarce, the elasticity is large (and negative). Banks' marginal cost of funds are more sensitive to changes in the supply of reserves. But as reserve supply increases, the demand curve flattens. When it reaches zero, the satiation point is reached, and the spread at which banks can borrow becomes invariant to reserve supply.

Second, I compare the model's key prediction on deposit spreads to data on quarterly deposit spreads (Figure 5). The key mechanism in the model relies on the idea that when reserves are scarce, banks are more illiquid and therefore more risky. As a result depositors demand higher compensation and deposit spreads are higher. The model predicts that these spreads are decreasing in the size of the balance sheet, up to the point where reserve demand is satiated, beyond which point they are zero. Data on deposit spreads shows a high degree of variance, since deposit spreads respond to a range of factor through the business cycle. For this reason, the data are grouped into periods which help contextualise this variance in terms of different reserve supply regimes. The data loosely demonstrate a similar non-linearity to that shown in the model.[21] The data points coloured blue and purple, in which spreads are elevated and reserve supply is high appear, at first glance, to contradict my results. These data correspond to periods in which the Federal Reserve was largely shrinking its balance sheet. In this sense, these elevated spreads may represent a *dynamic* response to contracting reserve supply, which is qualitatively in line with the mechanisms set out in the model, but not directly comparable to the *static* results presented above.

In both cases, the model matches the qualitative, non-linear properties of these relationships. However, it fails to capture their magnitudes. This is due to the fact that the model dramatically underestimates deposit spreads. This failure to capture risk premia is not surprising. It is well-known, since Mehra and Prescott (1985), that

[20] I do not explicitly model an interbank market for reserves, and hence have no direct counterpart to EFFR-IORB spread. However, given that this spread measures the short term cost at which banks can borrow, the deposit spread in my model is a reasonable proxy for it.

[21] Furthermore, deposits spreads rarely fall as low as zero in practice, since the reference rates (3 Month Treasury Bills) contain a convenience yield of their own, as demonstrated in Krishnamurthy and Vissing-Jørgensen (2012) and Lopez-Salido and Vissing-Jorgensen (2023), which are not present in the model.

simple consumption-based asset pricing fails to quantitatively capture risk premia and risk-free rates. While this inaccuracy can be reduced with a number of 'fixes' such as more sophisticated preferences (see, for instance, Cochrane (2017) for a review), I have used log utility in order to maintain tractability and facilitate welfare analysis.[22]

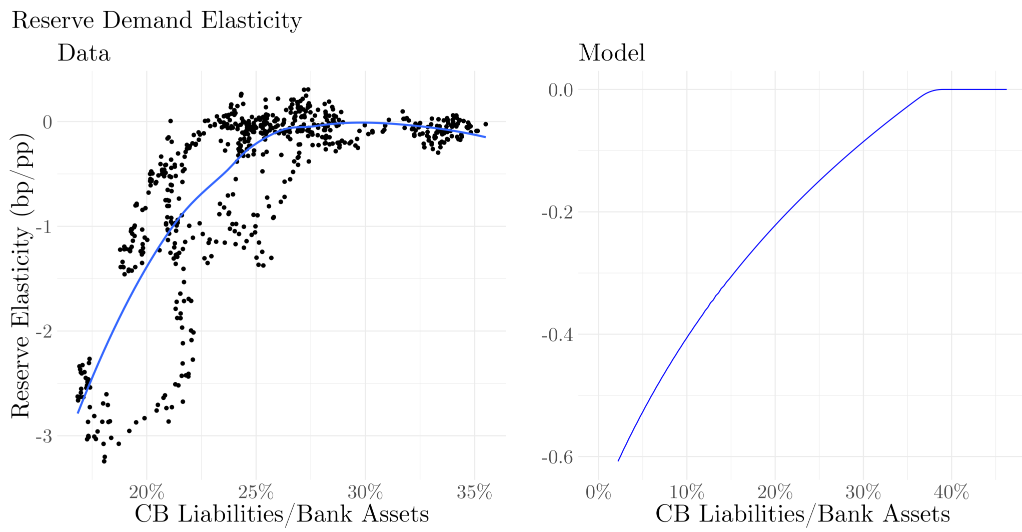

Figure 4: Reserve Demand Elasticity (Model vs Data). Source: Afonoso et al. (2022; Revised 2024), FRED, Author calculations.

4.4 Robustness to the Cost of the Central Bank

In this model, the inefficiency of the central bank, relative to private financial intermediaries, governs the marginal social cost of increasing the size of the central bank. As noted in Section 4.1, the baseline assumption that $\alpha^{cb} = \alpha^h$, implying that households and central banks face the same management costs when holding capital, is relatively pessimistic relative to the existing literature. As such, it implies a lower bound estimate of the optimal size of the central bank.

In this section, I consider how the optimal size of the balance sheet varies as a function of α^{cb}. Figure 6 shows social welfare as a function of the size (K^{cb}) and

[22]For instance, various types of 'recursive' utility functions can help to solve equity risk premium - risk free rate puzzles. However, these would introduce history dependence to the bank run state, altering the key property of the model that I exploit in order to calculate social welfare (Proposition 3.1).

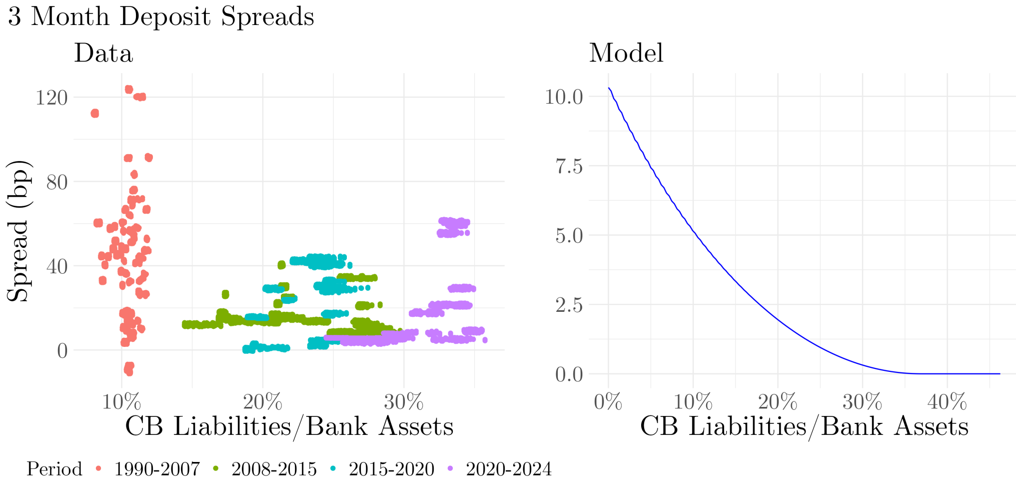

Figure 5: Deposit Spreads (Model vs Data). Source: FRED, Author calculations.

inefficiency (α^{cb}) of the central bank. The red line shows the welfare-maximising size of the central bank balance sheet as a function of α^{cb}.

The model-implied optimal size of the central bank is not highly sensitive to the *size* of the cost parameter, α^{cb}, provided it is not zero. As α^{cb} tends towards zero, the optimal size of the central bank balance sheet tends towards the point where the reserve supply is large enough to make the probability of bank runs equal zero. For any value of α^{cb}, this point lies at the 'kink' in the plotted surface in figure 6. Beyond this point, the marginal liquidity benefit provided by additional reserves is zero - runs cannot be made less likely! So, for any non-zero cost, expansions in the central bank are strictly welfare reducing. Since the location of this point is not sensitive to α^{cb}, nor is the optimal size of the balance sheet.

In the baseline analysis, with $\alpha^{cb} = \alpha^{h} = 0.014$, the optimal size of the central bank is $K^{cb} = 0.21$, or 31% of bank assets. This increases to 0.25, or 37% of bank assets, when $\alpha^{cb} = 0.001$. Given these estimates represent two edge cases, this move is relatively small. This analysis shows that the estimates of the optimal central bank balance sheet are robust to a wide range of assumptions about the marginal social cost of central bank asset holdings. Even if the central bank's cost of managing capital, relative to private banks, is extremely close to zero, the optimal size of the central bank will not exceed the point where reserve supply is large enough to rule out bank runs in steady state.

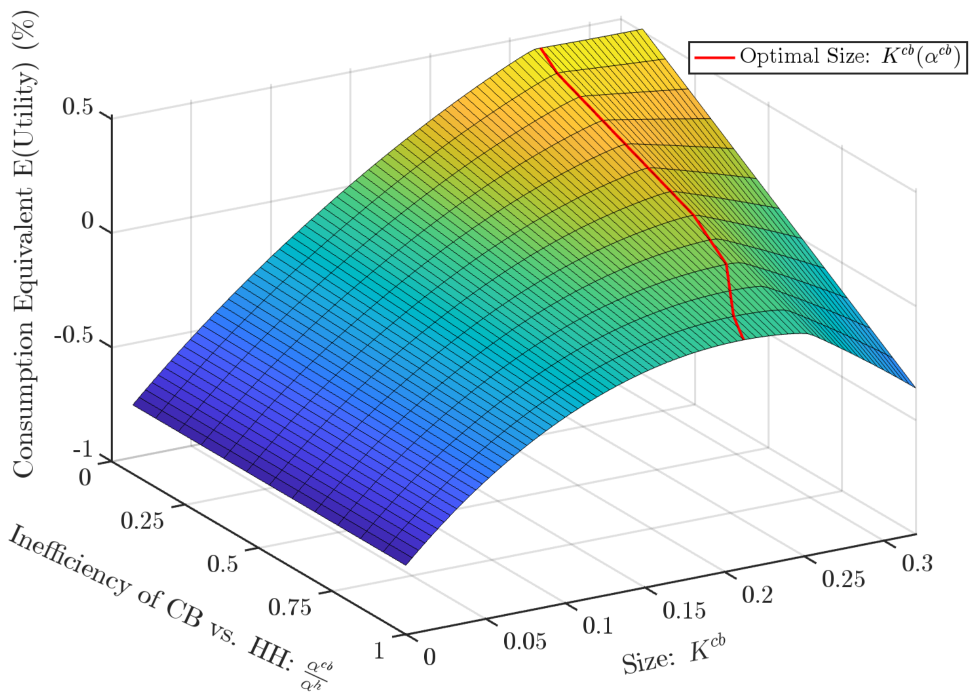

Figure 6: Social Welfare by Central Bank Size and Efficiency. Source: Author calculations.

5 Implications for Central Bank Balance Sheet Policy

In the numerical example, the welfare maximising size of the central bank occurs where $K^{cb} = 0.214$. To put this number in context, this implies the ratio of central bank liabilities[23] to banking sector assets of 31%, comparable to the Fed in 2023.[24]

Given the simplicity of the model, quantitative results are purely illustrative. Nevertheless, these results imply a clear qualitative insight for central bank balance sheet policy. The welfare-maximising size of the central bank balance sheet lies where the gross marginal benefit of additional reserves is close to zero, and reserve demand is almost fully satiated. Even under an upper-bound estimate of the efficiency cost of the central bank, households prefer to expand the central bank, forgoing some steady state consumption, to reduce the risk of bank runs to the point where the risk is almost eliminated. This result implies that the cost of insuring against liquidity risk, in the form of operating a large central bank, is small relative to the expected welfare gains of greater liquidity.

This model therefore provides a parsimonious, welfare-based rationalisation for ‘ample’ reserve regimes. Such regimes supply enough reserves to satisfy aggregate demand for reserves, in recognition of the fact that prior to the crisis, the supply of reserves was inefficiently low. For example, in 2024, the Bank of England has stated that it intends to continue to supply the “minimum level of reserves that satisfies commercial banks’ aggregate demand, both to settle their everyday transactions and to hold cash as a precaution against potential outflows in times of stress.”[25] This point corresponds to the region of the reserve demand curve where the the elasticity of demand is close to zero. Importantly, this type of reserve regime implies smaller balance sheets than the ‘abundant’ reserve regime seen on the far right of the charts in 4.3, where central banks supply reserve far in excess of the amount required to satiate demand.

The optimal central bank plays an insurance-like role, enhancing financial stability

[23] Net of government deposits

[24] Source: FRED Board of Governors of the Federal Reserve System (US) (2024)

[25] Bailey (2024)

by increasing the supply of liquid assets. The optimal sized central bank balance sheet is almost twice the size of that which maximises steady state consumption in the normal times equilibrium (Figure 3, Panel B). This is because agents are risk-averse and want to smooth consumption across the two possible equilibria. So, they prefer to give up steady state consumption to insure against the risk of large consumption falls in systemic bank runs. Even in the baseline analysis, the marginal cost of additional insurance (given by $\alpha^{cb}K^{cb}$) remains low enough that households prefer to insure against the vast majority of bank run risk.

Finally, section 4.4 shows that even if the management cost faced by the central bank is arbitrarily small, households never prefer to expand the long-run size of the central bank beyond the point where there is no risk of bank runs and the marginal benefits of additional reserves is zero. This suggests that 'ample' reserve supply regimes, where reserves are supplied significantly in excess of the volume required to satiate demand are unlikely to be optimal long run balance sheet policy.

6 Conclusion

This paper presents a simple endowment economy model where the central bank supply of reserves has a non-monotonic effect on aggregate welfare, allowing me to investigate the socially optimal long-run size of the central bank balance sheet. Welfare benefits arise in this model as a result of the liquidity of central bank reserves, by reducing the probability and severity of financial panics. Even with a pessimistic assumption about the efficiency of central bank credit intermediation, my numerical results suggest that these benefits come cheaply, relative to the marginal cost of increasing the size of the central bank. So, it is optimal for central banks to provide reserves up to the point where their marginal benefit to liquidity conditions is almost zero. This is consistent with an ample supply of reserves in steady state. Beyond this point, additional reserves have no effect on liquidity conditions and further expansion of the central bank crowds out the private financial sector and reduces aggregate welfare.

A number of extensions to this work would improve policy analysis. The most pressing areas concern the negative externalities, or moral hazards, which may arise from supplying abundant reserves. My results show that under some circumstances, reserves relax constraints on banks and allow them to take on more leverage. This increases their sensitivity to shocks and undoes some of the static liquidity benefits studied in this paper. Allowing for endogenous capital accumulation and explicitly modelling the interaction (or substitutability) of reserve supply with other macroprudential tools would facilitate this kind of analysis. This may help to account more fully for the risks of over-supplying reserves.

7 References

Abadi, Joseph, Markus Brunnermeier, and Yann Koby. 2023. "The Reversal Interest Rate." *American Economic Review* 113 (8): 2084–2120. https://doi.org/10.1257/aer.20190150.

Acharya, Viral V., Rahul Singh Chauhan, Raghuram G. Rajan, and Sascha Steffen. 2022. "Liquidity Dependence: Why Shrinking Central Bank Balance Sheets Is an Uphill Task." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4216001.

Afonso, Gara, Marco Cipriani, Adam Copeland, Anna Kovner, Gabriele La Spada, and Antoine Martin. 2020. "The Market Events of Mid-September 2019." *Federal Reserve Bank of New York Staff Reports* 918 (March).

Afonso, Gara, Darrell Duffie, Lorenzo Rigon, and Hyun Song Shin. 2022. "How Abundant Are Reserves? Evidence from the Wholesale Payment System." *NBER Working Paper*, no. 30736.

Afonso, Gara, Domenico Giannone, Gabriele La Spada, and John C. Williams. 2022. "Scarce, Abundant, or Ample? A Time-Varying Model of the Reserve Demand Curve." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4121823.

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}). 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.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2023. "Public Debt as Private Liquidity: Optimal Policy." *Journal of Political Economy* 131 (11): 3233–64. https://doi.org/10.1086/725170.

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.

Bailey, Andrew. 2024. "The Importance of Central Bank Reserves." https://www.bankofengland.co.uk/speech/2024/may/andrew-bailey-lecture-london-school-of-economics-charles-goodhart.

Benigno, Gianluca, and Pierpaolo Benigno. 2022. "Managing Monetary Policy Normalization." *Federal Reserve Bank of New York Staff Reports*, no. 1015. https://doi.org/10.2139/ssrn.4104789.

Board of Governors of the Federal Reserve System (US). 2024. "Total Assets, All Commercial Banks." *FRED, Federal Reserve Bank of St. Louis*. https://fred.stlouisfed.org/series/TLAACBW027SBOG.

Boissay, Frederic, Fabrice Collard, Jordi Galí, and Cristina Manea. 2021. "Monetary Policy and Endogenous Financial Crises." *NBER Working Papers*, December. https://ideas.repec.org//p/nbr/nberwo/29602.html.

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

Cochrane, John H. 2017. "Macro-Finance*." *Review of Finance* 21 (3): 945–85. https://doi.org/10.1093/rof/rfx010.

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 W., 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.

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

Gertler, Mark, and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics*, Carnegie-Rochester Conference Series on Public Policy: The Future of Central Banking April 16-17, 2010, 58 (1): 17–34. https://doi.org/10.1016/j.jmoneco.2010.10.004.

———. 2013. "QE 1 Vs. 2 Vs. 3. . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *International Journal of Central Banking* 9 (1): 5–53. https://ideas.repec.org//a/ijc/ijcjou/y2013q0a1.html.

Gertler, Mark, and Nobuhiro Kiyotaki. 2010. "Financial Intermediation and Credit Policy in Business Cycle Analysis." In *Handbook of Monetary Economics*, edited by Benjamin M. Friedman and Michael Woodford, 3:547–99. Elsevier. https://doi.org/10.1016/B978-0-444-53238-1.00011-9.

———. 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.

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.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105 (2): 211–48. https://doi.org/10.1086/262072.

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.

Kumhof, Michael, and Mauricio Salgado-Moreno. 2024. "Quantitative Easing and Quantitative Tightening: The Money Channel." *Bank of England Staff Working Paper*, August. https://www.bankofengland.co.uk/working-paper/2024/quantitative-easing-and-quantitative-tightening-the-money-channel.

Lopez-Salido, David, and Annette Vissing-Jorgensen. 2023. "Reserve Demand, Interest Rate Control, and Quantitative Tightening." Rochester, NY. https://doi.org/10.2139/ssrn.4371999.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15 (2): 145–61. https://doi.org/10.1016/0304-3932(85)90061-3.

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.

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.

Ruge-Murcia, Francisco J. 2007. "Methods to Estimate Dynamic Stochastic General Equilibrium Models." *Journal of Economic Dynamics and Control* 31 (8): 2599–2636. https://doi.org/10.1016/j.jedc.2006.09.005.

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.

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.

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

Appendix

A Equilibrium Conditions

Here I define the recursive competitive equilibrium with no bank run. It consists of aggregate quantities for:

$$(K_t^b, K_t^h, C_t, D_t, N_t, Y_t, \tau_t, \Pi_t)$$

prices:

$$\left(Q_t, R_{t+1}^{k,b}, \bar{R}_{t+1}, R_{t+1}^f, f_h(K_t^h)\right)$$

bank-related variables:

$$\left(\Omega_t, \frac{V_t}{n_t}, \phi_t, p_t\right)$$

as a function of state variables:

$$\left(K_{t-1}^b, K_{t-1}^h, R_t^d D_{t-1}, R_t^f m, Z_t\right)$$

which satisfy eqs. (1), (4), (6), (10), (11), (13), (18), (20) to (23), (25) to (29) and (36) where the constraint in (22) binds with equality.

To evaluate household expectations, we also need counter-factual bank run values for consumption and prices which are used in households' expectations of the bank run equilibrium:

$$(C_{t+1}^*, Q_{t+1}^*)$$

These values must satisfy (29) evaluated in the bank run state, and (33). This yields a system of 19 equations and 19 unknowns. The household budget constraint is satisfied by Walras' Law.

B Computation

Steady State. To solve for the steady state of the model, I guess a liquidation price of capital, $Q^*(Z_{ss})$, and solve for the steady state values of all endogenous variables (which depend on saddle paths for capital, consumption and the probability of bank runs in the event of run) numerically. I then verify that the guess of $Q^*(Z_{ss})$ is consistent with the household first order condition in equation (33), given the computed steady state values. If not, I update the guess and repeat.

Impulse responses. Once the steady state is found, I solve for the impulse response functions to a productivity shock at t = 1, which follows an AR(1), in the following steps:

1. Guess a time T after which all endogenous state variables have returned to steady state.
2. At each period $t^* = 1, 2, ...T$ calculate the saddle path of the economy back to steady state *if a run were to occur at* t^*. Using propositions 3.1 and 3.2, I impose initial and boundary conditions on all endogenous state variables to facilitate this computation.
3. To calculate IRFs, solve numerically for the vector of endogenous state variables $\{X_t\}_{t=0}^T$ which solve the system of equations defined by the equilibrium conditions of the model, given the exogenous shock process and the counter-factual values of the endogenous variables if a run were to occur the following period. Impose terminal conditions $X_T = X_{ss}$, and solve backwards from $t = T$ to $t = 1$. Verify the guess for T by checking the paths are consistent with the system starting in steady steady and the shock in period 1. If not, update the guess and repeat.

C Additional Comparative Statics

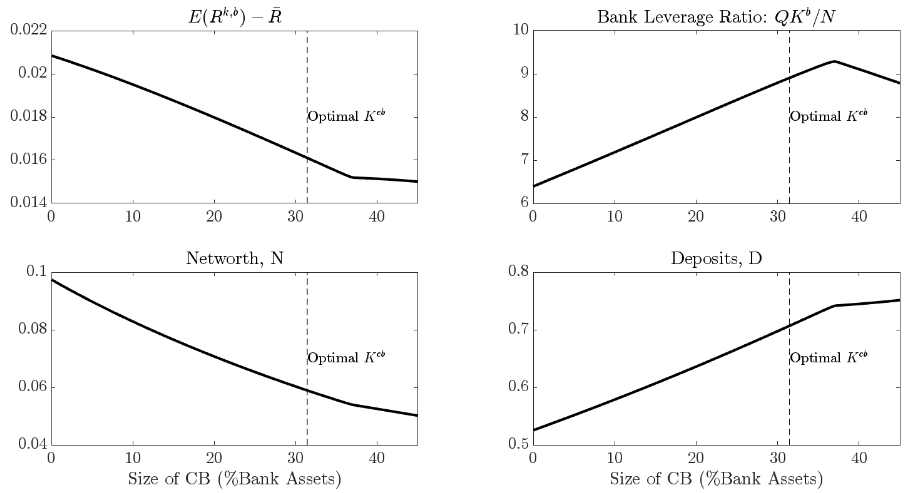

Figure 7: Additional Comparative Statics. Notes: The optimal size of the central bank is marked by a dotted line for reference

D Productivity Shocks

Figure 8 shows the responses of the model to a negative 3% shock to productivity, Z_t, which follows an AR(1) process with a serial autocorrelation of 0.95. I assume the economy stays in the no bank run equilibrium (ex post). In this model the shock is amplified by the interaction between the financial accelerator mechanism and the risk of bank runs. The fall in productivity pushes down asset prices. This pushes down the banks net worth, increases leverage, and forces them to contract deposits and sell assets (financial accelerator channel). Households absorb some capital, forcing them to incur larger management costs, and amplifying the size of the recession. The liquidation price of capital Q^* also falls, reducing banks' liquidation value and pushing down the recovery rate. This intensifies the risk of a bank run. As a result, deposit spreads rise (bank run risk channel), further weakening the bank's balance sheet and forcing them to further contract deposits and lending, in a feedback cycle with the financial accelerator.

The figure compares the effects of the shock for three sizes of central bank. A

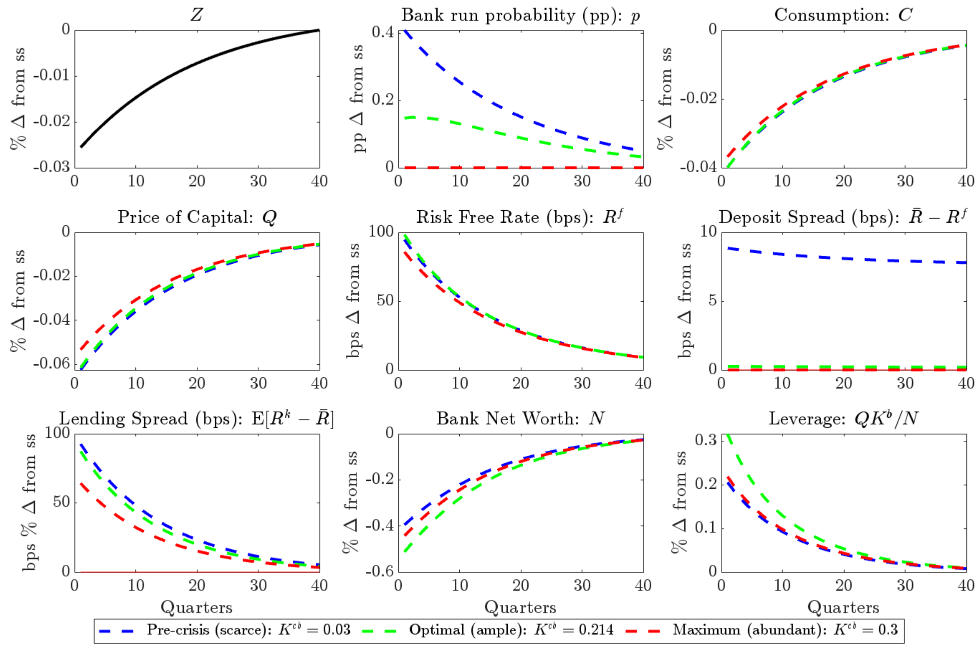

Figure 8: Recessionary Shocks by Size of Central Bank Balance Sheet

pre-crisis regime of $K^{cb} = 0.03$ (in blue), the 'Optimal' size identified in the previous section (in green), and the case where the central bank balance sheet is at its historical 'Maximum' (in red) and inefficiently large. In effect these three cases correspond to scarce, ample and abundant reserve supply regimes respectively.

The amplification of the shock to consumption and lending spreads is decreasing with the size of the central bank. This is because the recovery rate (equation (30)) is less sensitive to shocks when the level of reserves is higher. Consequently, the probability of bank runs increases by less when the central bank is larger, even beyond the optimal steady state balance sheet. In the largest balance sheet considered, the shock is not large enough to cause the recovery rate to fall below one. The probability of a bank run remains at zero throughout the recession, and there is minimal amplification beyond the direct financial accelerator dynamics. In general, when banks hold a greater stock of reserves, amplification via the bank run risk channel is weakened.

In my results, larger central bank balance sheets dampen the sensitivity of the economy to productivity shocks. But, important exceptions to this are bank leverage and net-worth. These variables are most sensitive to shocks in the case where the central bank is at the optimal size, and reserve supply is abundant. This is because banks are most leveraged in the optimal steady state. In my static results, leverage is increasing in the level of reserve supply, up to the point where the risk of bank runs is eliminated entirely at the optimum (see the figure in Appendix C). For higher levels of reserve supply, steady state bank leverage falls as the central bank expands. When the banks enter the shock with greater leverage, the financial accelerator mechanism is stronger, offsetting much of the benefits of greater liquidity. While this channel is not strong enough to cause spreads to rise more than in the 'pre-crisis' balance sheet case, it is notable that the impulse response functions for spreads, consumption, and the risk-free rate are almost identical in the two cases, despite the consequential difference in the level of reserve supply.

The finding that ample reserve supply may increase the sensitivity of financial intermediaries to shocks has empirical support. Recent work has found that ample central bank provision of liquidity has caused banks to disproportionately increase their risk appetite and leverage (Acharya et al. (2022)), *increasing* the sensitivity of the financial sector to liquidity shocks.[26] Moments of extreme illiquidity in money markets, such as the 2019 'Taper Tantrum' when the Fed began to shrink its balance sheet (Afonso et al. (2020)) and the 'Dash for Cash' at the onset of the Covid-19 pandemic in 2020 illustrate this effect. My results do not suggest that this increase in the strength of the financial accelerator channel is strong enough to fully undo the liquidity benefits of reserve supply in the context of dynamic shocks. However, they do suggest that a purely static view may overstate the benefits of long-run reserve supply, and that further research should investigate whether abundant reserve supply may contribute to financial volatility, or even create a moral hazard problem by making it possible for banks to take on excessive leverage.

[26]Schnabel (2023) notes that increases in regulatory requirements for banks to hold high quality liquid assets, and changes in risk management have also increased the structural demand for reserves since the financial crisis. This fact may partially explain the increase in sensitivity to liquidity shocks.