BANK OF ENGLAND

Staff Working Paper No. 605

The macroeconomics of central bank issued digital currencies

John Barrdear and Michael Kumhof

July 2016

BANK OF ENGLAND

Staff Working Paper No. 605

The macroeconomics of central bank issued digital currencies

John Barrdear[(1)] and Michael Kumhof[(2)]

Abstract

We study the macroeconomic consequences of issuing central bank digital currency (CBDC) — a universally accessible and interest-bearing central bank liability, implemented via distributed ledgers, that competes with bank deposits as medium of exchange. In a DSGE model calibrated to match the pre-crisis United States, we find that CBDC issuance of 30% of GDP, against government bonds, could permanently raise GDP by as much as 3%, due to reductions in real interest rates, distortionary taxes, and monetary transaction costs. Countercyclical CBDC price or quantity rules, as a second monetary policy instrument, could substantially improve the central bank's ability to stabilise the business cycle.

Key words: Distributed ledgers, blockchain, banks, financial intermediation, bank lending, money creation, money demand, endogenous money, countercyclical policy.

JEL classification: E41, E42, E44, E51, E52, E58, G21.

(1) Bank of England. Email: john.barrdear@bankofengland.co.uk
(2) Bank of England. Email: michael.kumhof@bankofengland.co.uk

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

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Tables

Figures

I. Introduction

This paper studies the macroeconomic consequences of a central bank granting universal, electronic, 24x7, national-currency-denominated and interest-bearing access to its balance sheet via the issuance, according to well-specified policy rules, of a *central bank digital currency* (CBDC). To study this issue we use a monetary-financial DSGE model, calibrated to match the United States in the pre-crisis period, that models CBDC as an imperfect substitute for bank deposits in the provision of monetary transaction services, and that models bank deposits as being created through loans or asset purchases as in Jakab and Kumhof (2015).

A monetary regime with CBDC has never existed anywhere, a major reason being that the technology to make it feasible and resilient has until now not been available. There is therefore very little historical or empirical material that could help us to understand the costs and benefits of transitioning to such a regime, or to evaluate the different ways in which monetary policy could be conducted under it. We therefore choose the alternative approach of using a theoretical model as a laboratory where we can systematically study these issues. The model we present is detailed rather than stylised, both in order to make the exercise credible for policymakers and in order to avoid prejudging what are the most important economic mechanisms to determine the effectiveness of CBDC. Nevertheless, the modifications needed to introduce CBDC into this model are kept to a minimum in order to understand the key transmission channels more clearly. We have made considerable progress using this strategy, but much more work needs to be done.

As a baseline, we consider a setting in which an initial stock of CBDC equal to 30% of GDP is issued against an equal amount of government debt, and is then, subject to countercyclical variations over the business cycle, maintained at that level. We choose 30% because this is an amount loosely similar to the magnitudes of QE conducted by various central banks over the last decade. We do not examine the question of the optimal steady-state stock of CBDC, but we note that it ought to be large enough to avoid problems with a "quantity zero lower bound" in the conduct of countercyclical policies.

Our simulations suggest that this policy has a number of beneficial effects. First, it leads to an increase in the steady-state level of GDP of almost 3%, due to reductions in real interest rates, in distortionary tax rates, and in monetary transaction costs that are analogous to distortionary tax rates. Second, a CBDC regime can contribute to the stabilisation of the business cycle, by giving policymakers access to a second policy instrument that controls either the quantity or the price of CBDC in a countercyclical fashion. This second policy instrument becomes especially effective in response to shocks to private money demand and private money creation, and if the substitutability between CBDC balances and bank deposits in the production of monetary transaction services is low. Financial stability considerations also generally favour the issuance of CBDC, provided that the issuance mechanism ensures that CBDC is only issued or withdrawn against government debt. On the negative side, there remains a clear concern with the proper management of the risks involved in transitioning to a different monetary and financial regime.[1]

Our work is motivated by the recent emergence of *private* digital currencies that offer participants access to both an alternative unit of account that is governed by predetermined money supply rules, and to a new payment system that is claimed to be superior to that offered by existing

[1] We should also mention that our model does not admit the possibility of bank runs in the style of Diamond and Dybvig (1983). It is however not clear whether such runs would be more or less likely under a CBDC regime.

banking systems. Bitcoin was the first such system, being launched in January 2009 following the earlier release of a white paper describing its operation (Nakamoto (2008)).[2] A substantial number of alternatives have since been developed,[3] but at present Bitcoin remains the largest such system in operation.

The monetary aspects of private digital currencies – a competing currency with an exogenous, predetermined money supply – are commonly held to be undesirable from the perspective of policymakers, but the innovation embodied in the payment systems of such schemes is held of some interest. It is important to stress that these two aspects of private digital currencies are logically distinct: it would be technically possible to implement a distributed payment system in the style of Bitcoin that nevertheless remained denominated in a traditional currency. The question that naturally emerges is whether it might be socially beneficial to do so,[4] and it is towards this question that our paper is addressed.

To the best of our knowledge, the only other theoretical work on the subject of private digital currencies is that of Fernández-Villaverde and Sanches (2016), who explore the conditions under which multiple units of account could exist in stable currency competition, by extending the model of Lagos and Wright (2003). They demonstrate the possibility of the value of privately-issued currencies being constant over time, but show that, absent productive capital, an efficient allocation as the unique global equilibrium requires driving private money out of the economy. However, the presence of productive capital reverses this result, so that equilibria that feature the value of private money converging to zero are ruled out. In our model, we abstract from all considerations of currency competition, instead supposing an economy with a single, government-defined unit of account. Privately-created money does exist in our model, but only in the form of bank deposits that maintain a 1-to-1 exchange rate with government money.

The rest of this paper is organized as follows. Section II gives a brief overview of digital currencies and the distributed ledgers that underlie them. Section III provides a general and non-technical overview of the pros and cons of CBDC. Section IV contains a detailed presentation of the theoretical model. Section V discusses the calibration of this model. Section VI presents simulation results, and discusses policy lessons. Section VII concludes.

II. Electronic Money, Digital Currencies, Distributed Ledgers

The phrase "digital currency" is, perhaps, a regrettable one, as it may invite a number of misunderstandings among casual readers. Most importantly, there is no innovation in the provision of an electronic *form* of money, as the vast majority of money in a modern economy is already electronic and has been for some time. In the United Kingdom, for example, physical currency (notes and coin) in public circulation represented only 4% of broad money balances in

[2] Earlier systems of electronic money, or e-money, existed as precursors of Bitcoin, including ecash in the USA and Mondex in the UK, but such offerings (i) still represented tradable claims on an issuing entity, while bitcoins are pure tokens and not the liability of any other party; and (ii) were generally implemented as physical cards with no public record of transactions, while Bitcoin implements a distributed ledger.

[3] The website http://coinmarketcap.com lists several hundred privately-developed digital currencies.

[4] The public research agenda of the Bank of England, which was published in February 2015, invites research into why a central bank might choose to issue a digital currency (Bank of England (2015)).

February 2016.[5] By broad money, we refer to the Bank of England's M4x measure, which equals notes and coin held by the non-bank public plus sight and time deposits held by households, private non-financial corporations and non-intermediary other financial corporations. Records of such financial instruments have been held electronically, if perhaps inefficiently, since the advent of the mainframe computer.

If the definition of money is allowed to expand further, then the share of the total held in physical form will naturally be still less. Indeed, in our formal model, we will refer to all non-equity[6] items on the liability side of the aggregate balance sheet of the entire financial system as deposits, because all of them represent "safe, information-insensitive financial assets" in the sense of Gorton, Lewellen and Metrick (2012). Specifically, they are used for various real and financial transaction purposes precisely because of their perceived safety, with low-interest checking accounts (current accounts in the UK) being particularly useful as a transaction medium, but longer-term and higher-interest liabilities still serving important transaction services. Because of our treatment of financial institutions as a single sector, for ease of exposition we use the term "financial institution" interchangeably with "bank" unless the distinction is necessary.

Nor is there particular innovation in the provision of electronic *access* to money, as debit and credit cards, internet banking and their union in online shopping have all been available for some time. Instead, the innovations proposed by existing private digital currencies, beyond the advocacy of new units of account and hard money supply rules, are particular to the manner in which electronic records of money and its exchange are implemented. In particular, digital currencies propose a distributed ledger and a payment system, in other words a process to update the ledger, that is decentralised, with copies of the ledger distributed across many agents and with no individual entity being indispensable in order for any given payment to be processed. Consequently, in this paper we define "digital currency" as any electronic form of money, or medium of exchange, that features a distributed ledger and a decentralised payment system.

By comparison, existing electronic payment systems are tiered and therefore *centralised*, with central banks typically at their centre. Private non-financial agents gain access to the system by holding claims on specific financial institutions. Payments between agents at the same bank are settled across that institution's ledger, payments between agents at separate non-clearing banks require settlement across those institutions' own accounts at a common clearing bank, and payments between agents at separate clearing banks require settlements across those banks' own settlement accounts at the central bank. Such systems are inherently centralised, as participants must, in general, rely on the fair, timely and accurate operation of the institutions that grant them access. In order to ensure agents' trust in the system, banks are regulated and subject to capital, leverage and liquidity requirements. Although necessary to ensure financial stability, these regulations represent barriers to entry and thereby grant banks pricing power, including power over the pricing of their liabilities, which serve as the economy's primary transaction medium.

Digital currency systems can offer a number of *benefits*, such as improved competition, accessibility and resiliency. We explore these in more detail in Section III, along with additional benefits specific to CBDC. But such systems also incur additional *costs*. This includes costs related to the storage and synchronisation of the multiple ledgers, even when all copies can be trusted as accurate and their operators as honest. However, much higher costs can arise when

[5]See tables A1.1.1 and A2.2.3 of the Bank of England's Bankstats database, available at http://www.bankofengland.co.uk/statistics/Pages/bankstats/default.aspx

[6]Following the literature, in our model section we will refer to bank equity as the "net worth" of the bank.

participants cannot trust each other to operate correctly. For some systems, these costs can be prohibitively high from a societal perspective.

The reason for this last concern is that the designer of any distributed system must confront the issue of how to repeatedly arrive at a consensus between participants within a reasonable timeframe when the honest and correct operation of any one agent cannot, in general, be guaranteed.[7] For example, if anybody may freely enrol in the system, a simple vote between participants is not sufficient as it is, in general, trivially easy to create an arbitrary number of nodes on a computer network (a 'Sybil Attack').

We note that communication between participants in this setting represents a cheap talk problem in the sense of Farrell (1987): the construction and delivery of a message between computers is effectively costless on the margin; messages are non-binding as agents can drop out of the network and freely re-enrol; and the validity of a transaction is not fully verifiable because without access to all copies of the ledger, no agent can be sure that the transaction doesn't represent "double spending", in which the same unit of money is spent multiple times.[8] To address this problem, it is necessary to alter one of these three features, and it is in the choice of this alteration that different digital currencies may be distinguished.

A. Cryptocurrencies

The best-known digital currencies yet established, including Bitcoin, require that, in addition to the fundamental cost of operation, any agent proposing an addition to the ledger must also demonstrate that it was costly for them to put forward that proposal.[9] *Cryptocurrencies* implement this by forcing would-be transaction verifiers to compete against each other in searching for a cryptographic *proof of work* – a verifiable demonstration that they have paid a cost in computation time – to accompany their candidate block of transactions. The first agent to successfully demonstrate their proof has their proposed transaction block accepted, and all agents then move on to build new candidate blocks and repeat the process.

A proof of work system has the benefit of allowing a participating agent to prove that they have paid the necessary cost without a need to disclose their identity or otherwise interact with other agents. This then allows cryptocurrencies to maintain complete decentralisation of their network, with free entry and exit of agents to the verification process.

In a proof of work system, the probability that any one verifier will be successful is proportional to the amount of computing capacity they deploy. To ensure that the average time between transaction blocks remains stable – a feature necessary to allow news of each success to be transmitted to the entire network – cryptocurrencies adjust the difficulty of their proof of work problems over time, scaling them to represent the total computational capacity of their network.

[7] This is commonly referred to as the 'Byzantine Generals Problem' in the field of Computer Science. See Lamport, Shostak and Pease (1982).

[8] Digital signatures, based on public-key cryptography, do permit the verification that a given single transaction is being initiated by the person who legitimately controls the source account or 'wallet'.

[9] We present only a brief overview of the design of cryptocurrencies here. For a more substantive introduction, see Ali, Barrdear, Clews and Southgate (2014a,b) or Bohme, Christin, Edelman and Moore (2015).

To compensate transaction verifiers for their underlying costs of operation and of demonstrating proofs of work, cryptocurrencies award the successful verifier any transaction fees associated with the payments they verify, and an allocation of new currency for each block of accepted transactions.[10] This second form of payment acts as a subsidy to payments, and has allowed transaction fees in most cryptocurrencies to remain exceedingly low in most cases (often well below 0.05% of the value of transactions).

However, since the award per block of transactions is a winner-takes-all game, going to only one of the competing verifiers, and since the probability of "winning" a block is increasing in one's own computational capacity but decreasing in the network's total capacity, a negative externality emerges. Individual verifiers have an incentive to overinvest in their own computing capacity and, since coordination is not possible, in equilibrium a socially inefficient excess of computing capacity will be deployed.

The magnitude of this externality can be substantial, to say the least. For example, in its original implementation, and still at the time of writing, Bitcoin is limited to between 7 and 10 transactions per second, or roughly 3,500 transactions per hour, perhaps sufficient to provide electronic payment services to a medium-sized town.[11] Despite this, the real resource cost of maintaining the Bitcoin network is on the scale of entire national economies. O'Dwyer and Malone (2014) estimate that the total electricity consumption of the Bitcoin network in early 2014 was comparable to that of Ireland (roughly 5GW). Deetman (2016) further estimates that, at current growth rates in computing efficiency and popular uptake, the Bitcoin network could potentially consume as much as 15GW by 2020, similar to the consumption rate of Denmark in 2014.

On this basis alone, we believe that in order for a digital currency system to be socially beneficial over the long run, an alternative method of addressing the cheap talk problem in transaction verifications needs to be developed and adopted.

B. Central-Bank-Issued Digital Currencies

By CBDC, we refer to a central bank granting universal, electronic, 24x7, national-currency-denominated and interest-bearing access to its balance sheet. We conceive of a world in which the majority of transaction balances would continue to be held as deposits with commercial banks and subject, where relevant, to existing deposit protection arrangements. Credit provision would remain the purview of existing intermediaries, and commercial banks would continue to be the creators of the marginal unit of money in the economy. In short, we imagine a world that implements Tobin's (1987) proposal for "deposited currency accounts".

From a macroeconomic perspective, the use of distributed ledgers is not strictly required for the operation of a CBDC system, but we contend that it would be necessary as a practical matter, in order to ensure the resiliency of a system that would clearly be of critical importance to the financial stability of the economy. There are several ways in which a such decentralised system could be implemented. A central bank could maintain all of the copies of the ledger itself, several

[10] This allocation is how cryptocurrencies introduce new currency into circulation. The magnitude of the allocation is typically pre-determined at the time of the system's inception in order to deliver the preferred total money supply into the system.

[11] By comparison, major card payment processors report peak capacity in the order of several tens of thousands of transactions per second (Visa, 2015).

public institutions could maintain copies for each other, or private sector agents could be involved in collaboration with the central bank. For this paper, we remain entirely agnostic about the technical implementation of any such system, beyond an assumption that it can be implemented without the explosive costs of existing cryptocurrencies.[12] Instead we focus on the macroeconomic implications of its adoption.

As a matter of practical implementation of CBDC issuance, the central bank could set the interest rate paid on CBDC and allow the private sector to determine its quantity by offering to buy and sell CBDC in exchange for well-defined asset classes, or it could set the quantity of CBDC and allow the private sector to bid the CBDC interest rate up or down until the market clears. Other, more complicated arrangements can also be conceived,[13] but we do not consider these here. We also emphasise that in our model, all financial contracts are for one period only and we therefore do not consider potential implications of CBDC for the dynamics of the yield curve.

Of course, central banks have always had the ability to grant universal access to their balance sheets through the issuance of banknotes.[14] However, banknotes require storage and physical exchange for payment, and pay an interest rate of zero. And the existence of an interest-free financial asset, after accounting for storage and transaction costs, represents the basis for a lower bound (ZLB) on central banks' policy rates. Various proposals for materially circumventing the ZLB have been put forward, including, among others, a tax on banknotes (Gesell (1916)), a managed exchange rate between cash and electronic forms of money (Agarwal and Kimball (2015)) and the abolition of cash altogether (Rogoff (2015)). However, no central bank has attempted to implement one of these schemes to date.

This paper focuses on the macroeconomic effects of CBDC during normal economic times. It therefore abstracts away from the ZLB, and also from the existence of cash, which is a quantitatively negligible and systemically non-constitutive[15] component of the broad money supply. We emphasise, though, that any issuance of CBDC need not be predicated on the withdrawal of banknotes from circulation. It would be perfectly plausible for the two to operate in tandem alongside commercial bank deposits. Furthermore, as our simulations will show, there are reasons to believe that, when searching for expansionary monetary policy options at the ZLB, the injection of additional CBDC would represent a promising alternative to negative policy rates, thereby removing part of the rationale for the withdrawal of bank notes.

[12] One obvious candidate to achieve this would be to render statements binding by implementing a "permissioned" system, where only certified and audited agents could take part in the verification process. Since no verifier would need to act as a deposit taker for the CBDC sector of the financial system, the regulatory burden imposed upon verifiers could be materially lower than for banks, and could be limited to what is necessary to ensure the operators' veracity. For an example of a proposal for a permissioned system, see Danezis and Meiklejohn (2016).

[13] For example, the central bank could simultaneously set the interest rate on CBDC and impose per-user quantity caps that varied by type of user. This might ensure "equitable" access to central bank liabilities and might potentially also help in a period of stress, depending on the policy rule deployed. But it would naturally limit the set of transactions for which CBDC could be used, and so could potentially decrease its attractiveness to end users.

[14] Although the share of total transactions conducted in cash has been trending down for some time, demand for banknotes remains and is expected to remain for the foreseeable future (Fish and Whymark (2015)).

[15] This means that, away from the zero lower bound, the rest of the financial system would function exactly as it does today if cash were removed entirely.

III. The Pros and Cons of CBDC

The analytical and quantitative part of our paper, in Sections IV, V and VI, studies a large set of macroeconomic issues associated with CBDC. However, because that analysis is based on a DSGE model, which inherently has a limited ability to address questions outside of macroeconomics, it can necessarily only study a subset of the issues. This section therefore provides a broader, qualitative analysis of the pros and cons of a CBDC regime. In our discussion we distinguish between structural issues, price and output stability issues and financial stability issues.

A. Structural Issues

Risk and the level of interest rates. Empirical studies for the pre-crisis period have found that real interest rates on US government debt are increasing in the size of the US government's (defaultable) debt burden. We incorporate this feature into our model, by allowing the CBDC-financed acquisition of government bonds to lower the equilibrium real interest rate.[16] When the central bank issues money and holds government debt against it, this has two effects of relevance in this matter. First, it lowers the government's interest burden, as profits made from the central bank's net interest margin are remitted back to the government, thereby making any given stock of debt more sustainable. That interest margin must be positive on average, because of the greater usefulness for transaction purposes, and therefore the greater non-pecuniary benefits, of holding CBDC. Second, government bonds held by the central bank are, in effect, non-defaultable as any losses imposed by a government default must trigger a replacement of central bank capital by the government.[17] And from the perspective of the private sector, claims against the central bank, such as CBDC, are not repayable by anything other than other claims against the central bank. Consequently, the CBDC-financed acquisition of government bonds serves to lower the stock of defaultable debt in the hands of the private sector, and hence lowers perceptions of the associated credit risk, thereby leading to lower interest rates on government debt.[18] Because the interest rate on government debt, by arbitrage, anchors the economy's entire interest rate structure, this translates, ceteris paribus, into lower borrowing costs for private borrowers, with obvious benefits for capital accumulation and economic growth.

However, there is an important countervailing force, in that upon the introduction of CBDC a substantial portion of retail transaction balances might be expected to switch from bank deposits to CBDC, thereby leaving a larger portion of bank financing dependent on the wholesale market, at higher interest rates. This would reduce the spread between the interest rates on government bonds and on bank deposits. Because banks finance new loans by creating new deposits, it is the rate on bank deposits, rather than the policy rate itself, that constitutes the marginal cost of funding for the banking system, and therefore determines the likely evolution of borrowing costs.

These two forces behind the evolution of interest rates under CBDC will be carefully studied in the formal and quantitative part of this paper. Our analysis suggests that, overall, real deposit interest rates will decline upon the introduction of CBDC, implying positive output effects.

[16] The relevant empirical literature for the size of this effect is cited in Section V below.

[17] We are implicitly ruling out the possibility of hyperinflationary episodes here.

[18] We do not explicitly model risky government debt in this paper, as one might do in an ideal (and much more complex) model. Instead we incorporate a small friction in the household Euler equation as a reduced-form device that moves as a function of the stock of government debt held by the private sector.

CBDC stock issuance arrangements, balance sheets, and real economic outcomes. The formal analysis of this paper will assume that both the initial stock issuance of CBDC, and subsequent countercyclical stock injections and withdrawals, will exclusively take the form of exchanges against government bonds, in outright open market operations or repo transactions. However, while this is the most obvious choice, and for countercyclical operations probably the superior choice, this is by no means the only arrangement through which additional stocks of CBDC could be injected into or withdrawn from circulation, and the precise arrangements could have real economic consequences. Alternatives include the acquisition by the central bank of financial assets other than government bonds, the payment of a citizen's dividend (see Benes and Kumhof (2012)), or the lending of additional CBDC to private banks for on-lending to the private sector. These possibilities are not currently studied in the paper, but they are within the scope of what can be studied using the formal model.

CBDC flow issuance arrangements, fiscal policy, and real economic outcomes. In addition to the *stock* injections and withdrawals of CBDC during the initial issuance and subsequent countercyclical operations, the establishment of a CBDC regime would also be associated with a permanent increase, ceteris paribus, in ongoing consolidated (government plus central bank) fiscal income flows, due to reductions in net interest expenses. This would permit an increase in spending or a lowering of tax rates by the fiscal authority, at unchanged deficit and debt targets. The precise real effects depend on which category of spending or taxation is affected. Our paper addresses some of these issues when studying the transition to an economy with a CBDC regime, but it can only choose and evaluate one out of many possible fiscal arrangements, namely a deployment of the savings towards a reduction in distortionary taxation. We find that this would be associated with output gains that are almost as large as those attributable to lower real interest rates. The issuance of CBDC inherently has a strong fiscal dimension, and these monetary-fiscal interactions will therefore have to be explored in much more detail in future work.

Reduced cost of providing transaction services. Advocates of the technologies underlying digital currencies commonly suggest that there are substantial opportunities for reducing the total cost and, potentially, the marginal cost, of operating the payment system in developed economies.[19] We find it difficult to distinguish how much of any potential gain might reasonably be attributable to the adoption of a distributed payment system as opposed to, for example, a streamlining of multiple overlapping record-keeping systems. But we are on more solid ground when it comes to predicting which type of implementation of a CBDC system is most promising in terms of its potential for cost reductions.

Since the operation of computing servers exhibits increasing returns to scale, the cheapest alternative for running a CBDC system would clearly be a fully centralised architecture, but this would come with increased resiliency risks that are likely to be deemed unacceptable. Among distributed systems, a fully permissionless system with no barriers to entry or exit for agents verifying transactions, like in existing private digital currencies, would almost certainly involve the *addition* of prohibitive societal costs if the system is to attain a macroeconomically significant scale, at least as it is currently designed. But intermediate options that adopted a distributed, but "permissioned" architecture would, we contend, provide an improvement in the efficiency of settlement and serve to improve resiliency relative to the status quo, both of which would represent a reduction in costs to the real economy.

[19] See, for example, Autonomous Research (2016).

Quite apart from any technological efficiency gains offered by a distributed payment system, however, economy-wide transaction costs are likely to be lower in a world with CBDC because of the likely overall increase in the total quantity of monetary transaction balances (bank deposits plus CBDC) following its adoption. We estimate that the effect of this increase in liquidity could be almost of the same magnitude as the effects of lower distortionary taxation.

Competition in payment services. With a lower entry hurdle to becoming a transactions verifier in a distributed system than to becoming a member bank in a tiered system, we would expect more intense competition in the provision of payment services. To the extent that existing systems grant pricing power to member institutions, this should ensure that transaction fees more accurately reflect the marginal cost of verification (net of any public subsidy, if one were to be provided), and allow more rapid adoption of new technologies that sit on top of the payment infrastructure.

Competition in accounts services. From the perspective of households and firms, CBDC would be economically equivalent to the establishment of an online-only, reserve-backed, narrow bank alongside the existing commercial banking system and, as such, would represent an expansion of competition in the market for deposit accounts. This should again lead to a more rapid adoption of innovative technologies and account offerings.

Final settlement and collateral. Any tiered payment system necessarily features counterparty risk between its members – particularly *credit risk*, in that a bank may become insolvent with outstanding liabilities to other members of the system, and *liquidity risk*, in that an otherwise solvent member may not possess the liquid assets necessary to settle a payment at a given moment. To guard against these, tiered systems require that collateral be posted by more "junior" members of the system with the clearing member that sits above them. Since these typically require that assets with high liquidity and low credit risk (such as government bonds and bills, or money) be posted, this represents an opportunity cost to member banks if they would otherwise seek to hold other, higher-yielding assets. A CBDC system, by contrast, would allow final settlement directly between payee and payer, across the central bank's balance sheet. Counterparty risk is therefore avoided altogether, so that collateral need not be posted to guard against it.[20] This would free up significant amounts of collateral for non-settlement transactions. To the extent that there is a shortage of good collateral in financial markets today, as some commentators have suggested, this could have important macroeconomic and financial stability benefits. A quantitative evaluation of this issue is currently beyond the scope of our model.

Time of operation. Because of the criticality of their operation, central banks in tiered payment systems often have specific times of operation to allow for maintenance.[21] Although payment systems that rest on top of the central bank's settlement system may continue to operate 24 hours a day, this arrangement exposes banks to counterparty risk overnight. As a result, overnight transactions are limited to low-value transactions. In a decentralised setting, no individual transaction verifier is essential to the operation of the system as a whole. As such, any of them may be deactivated to allow for maintenance without the system ever needing to stop.

[20] We note, however, that one possible form of regulation by a central bank to ensure the correct operation of delegated private transactions verifiers may be to require the posting of collateral from which fines would be drawn in the event of an error or of fraud. Nevertheless, the need for collateral would likely be much smaller than in the current environment.

[21] For example, the Bank of England's Real Time Gross Settlement (RTGS) system currently operates from 05:45 to 18:20 on business days (Bank of England, 2014).

Summary. On the basis of this preliminary discussion it seems safe to assume that the implementation of a CBDC system would be a net positive for the steady-state economy, through the alleviation of a number of frictions. Our formal quantitative analysis will confirm this, even though it cannot consider all of the issues listed above.

B. Price and Output Stability Issues

A second policy instrument. CBDC would be an imperfect substitute for other types of financial assets. This is because of its role in providing specific types of transaction services, potentially at lower cost than bank deposits, and potentially also with attractive additional features. This imperfect substitutability implies that the central bank can control an additional policy instrument, in addition to the traditional policy interest rate, and that this policy instrument can be used in a countercyclical fashion that supports the policy rate. In this paper, we examine the two most obvious alternatives for this second policy instrument: a quantity rule whereby the central bank fixes the quantity of CBDC relative to GDP, and a price rule whereby the central bank fixes the spread between the policy interest rate and the (lower) CBDC interest rate. Countercyclical versions of these rules also respond to deviations of inflation from its target, either by reducing the ratio of CBDC to GDP during an economic expansion, or by increasing the spread between the policy rate and the CBDC rate during an economic expansion (and vice versa during a contraction). Whether this countercyclical use of CBDC is effective depends on a number of factors, most importantly the nature of shocks (managing the quantity or price of monetary transaction balances is primarily effective in response to shocks that affect the supply of and demand for such balances), the substitutability between CBDC and other types of monetary transaction balances (injecting or withdrawing CBDC is far more effective when bank deposits cannot easily substitute for CBDC), and interactions with fiscal policy (injecting or withdrawing CBDC is more effective when their budgetary effects do not trigger countercyclical changes in distortionary tax rates or in government spending). Studying these issues accounts for a large part of the formal quantitative analysis performed in this paper.

The zero lower bound on nominal policy interest rates. Monetary policy starts to have difficulties in providing additional stimulus to the economy once nominal interest rates drop to near or below zero, because at that point households and firms are presumed to want to substitute into (zero-interest) cash rather than continuing to hold other (negative-interest) financial instruments. As mentioned in Section II.B, some students of this issue have suggested that central banks remove the ZLB altogether, but to date no central bank has adopted their proposed measures. Instead, when policy rates have reached their effective lower bounds, central banks have sought to provide additional stimulus by expanding their balance sheets and, in particular, by engaging in large-scale asset purchases funded through the issuance of reserves.[22]

The main proposed transmission channel for such "quantitative easing" (QE) policies relies on relative price changes in asset markets, where the movement of government bonds to the central bank's balance sheet serves to lower the outstanding stock of defaultable debt in the hands of the general public, and so reduces the equilibrium real interest rate.[23] However, to the extent that the

[22] The U.S. Federal Reserve, the Bank of England, the European Central Bank, the Swiss National Bank and the Bank of Japan have all increased their balance sheets significantly following the financial crisis of 2007/08 and the Euro crisis of 2010/11. The largest shares of purchases have been of government debt, although significant purchases of private sector securities have been made.

[23] In this context, it bears re-emphasising that central bank money, while found on the liability side of the central

transmission of QE also relies on lowering monetary transaction costs via the injection of liquidity into the economy, its efficacy will depend on an interaction with the incentives of commercial banks. Since only commercial banks may hold reserves at the central bank, any reserves-funded acquisition of securities from the non-bank private sector requires intermediation through a commercial bank, whose balance sheet increases in the process. To the extent that banks have other incentives that lead them to (partially) offset any increase in their balance sheets caused by QE,[24] any benefit from the injection of real money balances to the economy will correspondingly be reduced.

CBDC-funded asset purchases on the other hand do not require any involvement from commercial banks, as they would involve a direct exchange of central bank money with the non-bank sellers of the assets being purchased. These asset purchases therefore need not face the above-mentioned risk of offsetting decisions by commercial banks. As a result, when faced with a contractionary shock, a countercyclical injection of CBDC may potentially have greater efficacy than traditional QE.

Increasing interdependence of monetary and fiscal policies. The issuance of CBDC has a major and unavoidable fiscal dimension. As already discussed, this is likely to be beneficial in steady state, because lower government interest expenses permit the government to reduce distortionary taxation. Over the cycle, the fiscal aspect could however become politically contentious. Also, the interaction of CBDC and fiscal policy rules could complicate the conduct of either policy. While this cannot be entirely avoided, it can at least be minimized, through an appropriate design of the fiscal policy rule, which we will discuss at length, and through the proper design of automatic revenue sharing rules between the central bank and the fiscal authority.

More, and more timely, data. A key feature of distributed-ledger payment systems is that the entire history of transactions is available to all transaction verifiers, and potentially to the public at large, in real time. CBDC would therefore provide vastly more data to policymakers, including the ability to observe the response of the economy to shocks or to policy changes almost immediately. This would be helpful for macroeconomic stability management.

Summary. Again, it seems safe to assume that the implementation of a CBDC system would be a net positive for price and output stability, through the availability of a second monetary policy instrument that complements the policy rate. Our quantitative analysis will confirm this. Welfare analysis will have to be left for future work.

C. Financial Stability Issues

Under the previous two subheadings we concluded that a CBDC regime would have almost exclusively positive consequences. However, in preparing this paper we consistently encountered two main concerns regarding CBDC. Both fall primarily under the heading of financial stability. The first is the set of risks inherent to the transition to a new and untested monetary and

bank's balance sheet, is neither defaultable nor redeemable, and is therefore different from the common conception of debt.

[24] For example, if the banking system's binding constraint is a non-risk based leverage ratio and bank capital is slow or costly to acquire, then an increase in banks' balance sheets through QE could potentially be offset through a contraction in other parts of the balance sheet, so that broad measures of money may fail to fully reflect the increase in reserves.

financial environment, and the second is the risk of an economy-wide run from bank deposits to CBDC that would leave the banking system "short of funds". We consider these two topics first, and then discuss a number of other financial stability issues.

Transition risks. A key question in any major monetary reform, including the introduction of a CBDC regime, is whether the risks of the transition are justified by the benefits once that transition is complete. Much of the remainder of this paper will discuss the macroeconomic benefits of transitioning to a CBDC regime, and how to maximize these benefits through an appropriate choice of CBDC issuance arrangements, CBDC policy rules, and fiscal policy choices. This transition would, of course, have to be designed with careful attention to detail, given the complexity of today's financial markets and financial products. The issues here are not exclusively macroeconomic, and include the careful design and market-testing of new CBDC account products, the need for a thoroughly-tested and reliable digital infrastructure, including cyber security and protection against hacking, appropriate prior training of the human operators of such a system, careful analysis of the legal implications (including, if necessary, introduction of additional legislation), appropriate changes to financial sector regulation, full coordination with foreign central banks and foreign financial institutions, and many others. But in the final analysis the question will be whether the risks of mismanaging these issues are so large as to outweigh the benefits found in this paper.

Risk of a run from bank deposits to CBDC. The other major potential concern with a CBDC regime is with the possibility of a run from bank deposits to CBDC. We do not claim to have evaluated this risk exhaustively, and this remains an important topic for future research. In our model, the calibration is such that minimum capital adequacy requirements, together with the voluntary buffers that banks optimally add to avoid the penalty for breaching them, are sufficient to make deposits effectively risk-free.[25] This removes a major reason for a run on deposits, and also for risk-driven increases in bank funding costs that could affect lending rates. Furthermore, we would expect that a CBDC regime would retain deposit protection for bank deposits, so that under all but the most extreme circumstances there would be no incentive for a systemic, as distinct from an institution-specific, run (at least for retail deposits).

Nevertheless, our model does feature the possibility of a shock to households' *relative* preference for holding CBDC versus bank deposits, and we present a simulation of this scenario in section VI.B.3. As might be expected, the dynamics of the economy in this simulation depend critically on the nature of the central bank's policy response.

Under a quantity rule for CBDC, the central bank would decline to issue additional CBDC in response to the increase in demand, instead allowing the private sector to bid the interest rate on CBDC down until the market cleared. The magnitude of this movement, including the possibility of a negative rate of interest on CBDC, would then simply reflect a combination of the elasticity of substitution between CBDC and bank deposits and the size of the relative preference shock.[26] However, it is important to emphasise the necessity of a flexible CBDC interest rate under a CBDC quantity rule. Specifically, to minimise the impact of fluctuations in CBDC demand on banks' funding costs via deposit rates, it is essential that the CBDC market exhibits a highly flexible price finding mechanism, so that fluctuations in CBDC demand can be accommodated

[25] Furthermore, as mentioned in footnote 1, even if deposits were not effectively risk-free, it is not clear whether bank runs would be more or less likely under a CBDC regime.

[26] In practice, a lower bound on the CBDC interest rate would endogenously emerge at the point where the negative interest rate exceeded the carrying cost of zero-interest banknotes.

through changes in the CBDC interest rate rather than the deposit rate, and/or that the marginal demand for deposits is highly interest sensitive, so that the deposit rate is mainly driven by the policy rate. Our model satisfies both conditions.

Under a price rule for CBDC, in which the central bank sets the interest rate on CBDC and allows the private sector to determine the quantity, the issuance arrangements for CBDC become of critical importance. In our model, we suppose that the central bank only issues CBDC against government bonds, either indefinitely via open market operations or temporarily via repurchase agreements. In this case, a private-sector agent that wishes to switch their bank deposit to CBDC must first acquire government bonds in order to offer these to the central bank. In making this acquisition, the deposits do not leave the banking system as a whole – they are simply transferred to the seller of the bonds – and so funding cannot, in aggregate, be "lost". The key to this result is that the central bank declines to fund commercial banks directly, thus forcing agents to first exchange their deposits for assets that are not bank liabilities. In our model, which features frictionless capital markets, deposit interest rates, which are the marginal cost of banks' funding, do not materially change in this setting, due to an arbitrage condition with the policy rate.

If the central bank were willing to fund commercial banks directly, then a deposit holder could, in effect, exchange their claim on the banking system directly for CBDC and so cause a reduction in aggregate deposit funding, but again without bank funding, in aggregate, being "lost". In this case, banks' average cost of funding might increase if central bank funding was provided at the policy rate, but the marginal cost of funding would continue to be given by the deposit rate, which would continue to be pinned down at a nearly constant discount relative to the policy rate.

For completeness, we should mention that it is possible for bank balance sheets to shrink significantly during a run to CBDC. But the mechanism is such that it would again not leave the aggregate banking system exposed to a shortage of funds. This case arises if banks themselves are significant holders of government bonds (they are not in our formal model, but this is only for simplicity), and if they decide to sell them to households who then wish to trade them against CBDC at the central bank. This sale amounts to a cancellation of deposits on the liability side of banks' balance sheet against bonds on the asset side. It has no direct negative implications for banks unless they need those bonds elsewhere as liquid collateral, or to satisfy regulations. But in that case they presumably would not sell them in the first place, so that households would again have to turn to financial investors to acquire the necessary government bonds.

Potential partial removal of Too Big To Fail concerns. The protection of retail deposits, and of the corresponding payment systems, in the event of a bank insolvency is commonly regarded as a critical requirement in the design of regulatory frameworks.[27] So long as banks remain necessary gatekeepers to the payment system, they will have the capacity to achieve systemically-important status if they manage to obtain a sufficient market share. But if a universally accessible and sufficiently large CBDC payment system were to be established as an alternative alongside the existing bank-based payment system, the hypothetical failure of any individual bank, however large, need not necessarily cause an amplification through the economy by impairing payments.

In the event of an actual bank failure, CBDC could also help to ensure an orderly resolution. The payout process of deposit protection systems typically takes several days, during which time depositors cannot access payment services through their accounts. Depending on the design of

[27] See, for example, Independent Commission on Banking (2011).

resolution arrangements and the central bank's willingness to adopt temporary credit risk, CBDC could offer the alternative of transferring the affected deposit balances to the central bank, thereby offering continuity of service within hours.

We emphasise, however, that other reasons for judging a financial institution to be systemically important would remain in a world with CBDC, such as the provision of credit to the real economy, the provision of other vital banking services including some classes of derivative contract, and interconnectedness between the institution and other financial entities. The establishment of CBDC could not, and would not, remove the need to address these concerns in protecting an economy's financial stability.

Resiliency of the payment system. CBDC as envisaged in this paper would likely increase payment system resiliency relative to existing payment systems, both because end users will have an additional alternative in case one system fails, and because the underlying technology would deploy a distributed architecture. Tiered payment systems are subject to *operational risk*, and associated reputation risk for the central bank, when a member bank ceases to function, either temporarily (e.g. because of an IT failure) or permanently. Since payments must be routed through member banks, this can have the effect of temporarily removing access to the payment system altogether from end users banking with that member bank (until they migrate to another, still operational, member bank). Distributed payment systems are robust to this risk by design, as any one transaction verifier may confirm the validity of any transaction. On the other hand, if CBDC were implemented with a centralised technology, cyber attacks would represent a much bigger risk.

More data on interconnectedness. In addition to data on the flow of economic activity, a CBDC payment system would improve policymakers' understanding of interconnectedness in the financial system. This would especially be true if a system of 'smart contracts' – protocols by which the execution or enforcement of a contract is partially automated – were to be implemented on top of the payment system.[28]

Summary. The net benefits of a system of CBDC for financial stability are not as clear cut as for steady output and for price and output stability. While it seems clear that some risks would be mitigated, other risks would emerge, and it is not certain which would be greater. In this context some limitations of our formal analysis also need to be kept in mind, including the absence of credit risk on bank deposits. However, there is one very clear financial stability risk, that of mismanaging the transition to a new and as yet untested monetary and financial environment. Mainly for this reason, policymakers would clearly have to carry out a very careful due diligence before deciding on the transition to a CBDC regime. But if that due diligence found the transition risks to be manageable, a CBDC regime could be considered a serious option, given the many other sizeable benefits of CBDC identified in this paper.

[28] For example, a swap between two parties could be programmed to automatically pay out when an agreed third party digitally signed a confirmation of the outcome referenced in the contract.

IV. The Model

The model economy consists of four sectors, households, financial investors, unions and banks, and a government that determines fiscal and monetary policies. It features a number of nominal and real rigidities, including sticky nominal prices and wages, habit persistence in consumption, and investment adjustment costs. The model of banking is based on a combination of Benes and Kumhof (2012) and Jakab and Kumhof (2015). As in Benes and Kumhof (2012), banks make private sector loans in four different lending markets, namely consumer loans, mortgage loans, working capital loans and investment loans. And as in Jakab and Kumhof (2015) and Benes and Kumhof (2012), banks create deposits to fund loans (or asset purchases), rather than borrowing funds from depositors and lending them to borrowers. Deposits themselves can serve as additional collateral in the four lending markets, thereby endogenously boosting the ability of banks to lend.

There are several advantages to using a model with this degree of real and financial sector detail. First, this provides an integrated framework where all of the differences between economies with and without CBDC emerge simultaneously, without prejudging which of them are or are not important. Second, it generates a realistic transition scenario to a CBDC regime, based on an accurate (as far as possible) estimate of the sizes of initial balance sheets and interest rate spreads. Third, it makes it possible for model calibration to be guided by the DSGE literature, which has identified a number of nominal and real rigidities that are critical for the ability of such models to generate reasonable impulse responses.

Two types of private agents interact directly with banks. Financial investors, whose share in the population is small at 5%, consume, supply labour and invest in bank deposits and government debt. Their main role is to generate an arbitrage condition between bank deposits and government debt.[29] Households, whose share in the population is large, consume, supply labour, produce goods and physical capital, hold physical capital and land, and borrow from banks, against several different types of asset stock and income flow collateral, in order to obtain deposits that facilitate transactions related to goods consumption, physical investment, input purchases in production, and real estate purchases.

This class of models can be thought of as a natural extension of traditional Sidrauski-Brock models of monetary economies (Sidrauski (1967), Brock (1975)). Such models have, from the point of view of our analysis, two essential features. First, money is valued because a representative agent is subject to a cash-in-advance constraint, a transaction cost technology, or money in the utility function. Second, the money that is valued by the representative agent is exogenously supplied by the government. These models have been and continue to be extremely useful in the development of monetary theory. What we argue is that, for the development of a realistic model of monetary and financial systems, the problematic feature of such models is not the first but the second. In fact the first feature, the representative agent assumption, turns out to be essential.

The assumption of exogenous government money is problematic for two reasons. First, as discussed in Jakab and Kumhof (2015), government-supplied money as it exists today, which includes cash and reserves, is fully endogenous. This means that during normal times (this

[29] It is not strictly necessary to include financial investors in the model. One alternative is to assume monopolistic competition in the issuance of bank deposits, as in Jakab and Kumhof (2015), with banks as the sole holders of government debt. However, this makes it harder to endogenize movements in the spread between the policy rate and the rate on bank deposits, which play a critical role in our simulation of a transition to a CBDC regime.

excludes economies operating at the zero lower bound for nominal interest rates) it is supplied by the government on demand, with demand coming either from households and firms (cash) or banks (reserves). Second, in all modern economies, and again during normal times, government-supplied money accounts for only a very small fraction of broad monetary aggregates, around 4% in the recent UK context. For both of these reasons, government-supplied money is omitted from our model altogether.

Our objective is to model supply and demand for the other 96% of broad monetary aggregates, the part that is created privately by financial institutions, with equilibrium quantities determined by the interaction of the profit and utility maximization objectives of financial institutions and of their customers, households and financial investors. In our exposition, we will use the terms money, monetary transaction balances and liquidity interchangeably.

The model economy experiences a constant positive technology growth rate $x = T_t/T_{t-1}$, where T_t is the level of labour augmenting technology. When the model's nominal and real variables, say V_t and v_t, are expressed in real normalized terms, we divide by the level of technology T_t and, for nominal variables, by the price level P_t. We use the notation $\check{v}_t = V_t/ (T_t P_t) = v_t/T_t$, with the steady state of $\check{v}_t$ denoted by $\bar{v}$. The population shares of financial investors and households are given by ω and $1 - \omega$. Because of their central role in the economy, we start our exposition with banks.

A. Banks

1. Bank Balance Sheets

Bank loans to households consist of four different classes of loans, consumer loans (superscript c) that are secured on a combination of labour income and monetary transaction balances, mortgage loans (superscript a) that are secured on a combination of land and monetary transaction balances, working capital loans (superscript y) that are secured on a combination of sales revenue and monetary transaction balances, and investment loans (superscript k) that are secured on a combination of physical capital and monetary transaction balances. In each case monetary transaction balances include bank deposits and, under a CBDC regime, CBDC balances.

Bank deposits are modelled as a single homogenous asset type with a one-period maturity. In our calibration this corresponds to all non-equity items on the liability side of the consolidated financial system's balance sheet, all of which are interpreted as near-monies with some measure of liquidity. Refinements of the model that break homogenous deposits into different deposit types, for example retail and wholesale deposits, are perfectly feasible and interesting, but they are not essential for the study of CBDC.

While all bank deposits are therefore identical from the point of view of banks, their customers use deposits for several distinct functions, namely financial investor deposits (superscript u, to denote that this group of agents is financially unconstrained) that reduce transaction costs related to financial investors' consumption, consumption deposits (superscript c) that reduce transaction costs related to households' consumption and that provide additional collateral for consumer loans, real estate deposits (superscript a) that reduce transaction costs in the real estate market and that provide additional collateral for mortgage loans, working capital deposits (superscript y) that reduce transaction costs related to producers' payments to their suppliers and that provide additional collateral for working capital loans, and investment deposits (superscript k) that reduce transaction costs related to investment and that provide additional collateral for investment loans.

Banks maintain positive net worth because the government imposes official minimum capital adequacy requirements (henceforth referred to as MCAR). These regulations are modelled on the current Basel regime, by requiring banks to pay penalties if they violate the MCAR.[30] Banks' total net worth exceeds the minimum requirements in equilibrium, in order to provide a buffer against adverse shocks that could cause net worth to drop below the MCAR and trigger penalties.

As we will explain in more detail below, banks are assumed to face heterogeneous realizations of non-credit risks, and are therefore indexed by i. We use the superscript x to represent the different categories of loan and deposit contracts. Households' and financial investors' nominal and real per capita loan stocks between periods t and $t+1$ are $L_t^x(i)$ and $\ell_t^x(i)$, $x \in \{c, a, y, k\}$, their deposit stocks are $D_t^x(i)$ and $d_t^x(i)$, $x \in \{c, a, y, k, u\}$, and bank net worth is $N_t^b(i)$ and $n_t^b(i)$. Banks' nominal and ex-post real gross deposit rates are given by $i_{d,t}$ and $r_{d,t}$, where $r_{d,t} = i_{d,t-1}/\pi_t$, $\pi_t = P_t/P_{t-1}$, and P_t is the consumer price index. Their gross wholesale lending rates, which are at a premium over deposit rates because they offset the cost of regulation, are given by $i_{\ell,t}^x$ and $r_{\ell,t}^x$, $x \in \{c, a, y, k\}$. Banks' gross retail lending rates, which add a credit risk spread to the wholesale rates, are given by $i_{r,t}^x$ and $r_{r,t}^x$. We denote the ex-post "return" on nominal cash flows between periods $t-1$ and t by $r_{n,t} = 1/\pi_t$. This return is needed to value flow collateral for loans. Finally, gross nominal and real interest rates on government debt are denoted by i_t and r_t. Bank i's balance sheet in real terms is given by

$$\ell_t^\ell(i) = d_t(i) + n_t^b(i) \quad . \tag{1}$$

Here total loans are $\ell_t^\ell(i) = \Sigma_{x\in\{c,a,y,k\}} \ell_t^{xx}(i)$, where $\ell_t^{xx}(i) = (1-\omega)\, \ell_t^x(i)$. Total deposits are $d_t(i) = \omega d_t^u(i) + (1-\omega)\, \Sigma_{x\in\{c,a,y,k\}} d_t^x(i)$. Banks can make losses on each of their four loan categories. Loan losses are given by $\mathfrak{L}_t^b(i) = (1-\omega)\left(\mathfrak{L}_t^c(i) + \mathfrak{L}_t^a(i) + \mathfrak{L}_t^y(i) + \mathfrak{L}_t^k(i)\right)$. Each individual bank is assumed to hold a fully diversified portfolio of loans, such that its share in loans to an individual borrower is equal to its share in aggregate loans. As a consequence, each bank's share in aggregate loan losses is proportional to its share in aggregate loans.

2. Risk and Regulation

Our model focuses on bank solvency risks and ignores liquidity risks. Banks are therefore modelled as having no incentive, either regulatory or precautionary, to maintain reserves or cash at the central bank. Furthermore, for households cash is dominated in return by bank deposits. In this economy there is therefore no demand for government-provided real money balances.

Banks face pecuniary costs of falling short of official minimum capital adequacy ratios. The regulatory framework we assume introduces a discontinuity in outcomes for banks. In any given period, a bank either remains sufficiently well capitalized, or it falls short of capital requirements and must pay a penalty. In the latter case, bank net worth suddenly drops further. The cost of such an event, weighted by the appropriate probability, is incorporated into the bank's optimal capital choice. Modelling this regulatory framework under the assumption of homogenous banks would lead to outcomes where all banks simultaneously either pay or do not pay the penalty. A

[30] Furfine (2001) and van den Heuvel (2005) contain a list of such penalties, according to the Basel rules or to national legislation, such as the U.S. Federal Deposit Insurance Corporation Improvement Act of 1991.

more realistic specification therefore requires a continuum of banks, each of which is exposed to idiosyncratic shocks, so that there is a continuum of ex-post capital adequacy ratios across banks, and a time-varying small fraction of banks that have to pay penalties in each period.

Specifically, banks are assumed to be heterogeneous in that the return on their loan book is subject to an idiosyncratic shock ω_{t+1}^b that is lognormally distributed, with $E(\omega_{t+1}^b) = 1$ and $Var(\ln(\omega_{t+1}^b)) = \left(\sigma^b\right)^2$, and with the density function and cumulative density function of ω_{t+1}^b denoted by $\mathfrak{f}_{t+1}^b = \mathfrak{f}_t^b(\omega_{t+1}^b)$ and $\mathfrak{F}_{t+1}^b = \mathfrak{F}_t^b(\omega_{t+1}^b)$. This risk can reflect a number of individual bank characteristics that are not directly related to credit extension, such as differing success at raising non-interest income and minimizing non-interest expenses, where the sum of these over all banks equals zero.

The regulatory framework stipulates that banks must pay a real penalty of $\chi \ell_t^\ell(i)$ at time $t+1$ if the sum of the gross returns on their loan book, net of gross interest expenses on deposits and of realized loan losses, is less than a fraction Υ of the gross risk-weighted returns on their loan book. Different risk-weights ζ^x will be one of the determinants of equilibrium interest rate spreads. The penalty cutoff condition for bank i is given by

$$\Sigma_{x\in\{c,a,y,k\}} r_{\ell,t+1}^x \ell_t^{xx}(i)\omega_{t+1}^b - r_{d,t+1}d_t(i) - \mathfrak{L}_{t+1}^b(i) < \Upsilon \Sigma_{x\in\{c,a,y,k\}} r_{\ell,t+1}^x \zeta^x \ell_t^{xx}(i)\omega_{t+1}^b \ . \tag{2}$$

Because the left-hand side equals pre-dividend (and pre-penalty) net worth in $t+1$, while the term multiplying Υ equals the value of risk-weighted assets in $t+1$, Υ represents the minimum capital adequacy ratio of the Basel regulatory framework. We denote the realized (ex-post) cutoff idiosyncratic shock to loan returns below which the MCAR is breached by $\bar{\omega}_t^b$. It is given by

$$\bar{\omega}_t^b = \frac{r_{d,t}d_{t-1}(i) + \mathfrak{L}_t^b(i)}{\Sigma_{x\in\{c,a,y,k\}} \left(1-\zeta^x\Upsilon\right) r_{\ell,t}^x \ell_{t-1}^{xx}(i)} \ . \tag{3}$$

3. Wholesale Lending Rates and Net Worth

Banks choose their loan volumes to maximize their expected pre-dividend net worth, which equals gross returns on the loan book minus the sum of gross interest expenses on deposits, loan losses, and penalties:

$$\underset{\ell_t^{xx},\ x\in\{c,a,y,k\}}{Max} E_t\left\{\Sigma_{x\in\{c,a,y,k\}} r_{\ell,t+1}^x \ell_t^{xx}(i)\omega_{t+1}^b - r_{d,t+1}d_t(i) - \mathfrak{L}_{t+1}^b(i) - \chi\ell_t^\ell(i)\mathfrak{F}_t^b(\bar{\omega}_{t+1}^b)\right\} \ . \tag{4}$$

This yields four optimality conditions. For consumer loans we have

$$\begin{aligned} 0 \ = \ & E_t\Bigg\{ r_{\ell,t+1}^c - r_{d,t+1} - \chi\mathfrak{F}_{t+1}^b \\ & -\chi\mathfrak{f}_{t+1}^b \left(\frac{\ell_t^\ell(i)}{n_t^b(i)}\right) \frac{r_{d,t+1}r_{\ell,t+1}^c\left(1-\zeta^c\Upsilon\right) + r_{d,t+1}\Sigma_{x\in\{a,y,k\}}\frac{\ell_t^{xx}(i)}{n_t^b(i)}\left[r_{\ell,t+1}^x\left(1-\zeta^x\Upsilon\right) - r_{\ell,t+1}^c\left(1-\zeta^c\Upsilon\right)\right]}{\left(\Sigma_{x\in\{a,y,k\}}\left(1-\zeta^x\Upsilon\right) r_{\ell,t+1}^x \frac{\ell_t^{xx}(i)}{n_t^b(i)}\right)^2} \Bigg\} \ , \end{aligned} \tag{5}$$

with similar conditions for the other three loan classes. Notice that each bank faces the same expectations for future returns $r_{\ell,t+1}^x$ and $r_{d,t+1}$, and the same risk environment characterized by the functions $\mathfrak{f}_{t+1}^b$ and $\mathfrak{F}_{t+1}^b$. Aggregation of the model over banks is therefore trivial because loans, deposits and loan losses are proportional to the bank's level of net worth. Indices i can therefore be dropped in both (3) and (5).

Condition (5) states that banks' wholesale lending rate $i^c_{\ell,t}$ is at a premium over the deposit rate $i_{d,t}$. We will refer to differences between wholesale lending rates $i^x_{\ell,t}$ and the deposit rate as regulatory spreads. The magnitude of these spreads depends on the size of the MCAR Υ, the penalty coefficient χ for breaching the MCAR, and expressions $\mathfrak{f}^b_{t+1}$ and $\mathfrak{F}^b_{t+1}$ that reflect the expected riskiness of banks and therefore the likelihood of a breach of the MCAR. Banks' retail lending rate $i^c_{r,t}$, on the other hand, whose determination is discussed below, is at another premium over $i^c_{\ell,t}$, to compensate banks for the bankruptcy risks of their borrowers. The correct interpretation of the wholesale rate in sector x is therefore as the rate that a bank would charge to a hypothetical sector x borrower (not present in the model) with zero default risk.

Note that the policy rate i_t does not enter these optimality conditions, because the marginal cost of banks' funds is given by the rate $i_{d,t}$ at which banks can create their own funds, which is lower than i_t because of the non-pecuniary benefits of holding monetary transaction balances, or liquidity.[31] As we will discuss below, financial investor arbitrage ensures that the deposit rate spread $i_t - i_{d,t}$ remains nearly constant over the business cycle, unless the government debt-to-GDP ratio changes substantially.

Another outcome of this optimization problem is banks' actual Basel capital adequacy ratio Υ^a_t. This is considerably above the minimum requirement Υ, because by maintaining an optimally chosen buffer, banks protect themselves against the risk of penalties while minimizing the cost of excess capital. There is no simple formula for Υ^a_t, which in general depends nonlinearly on a number of parameters. We will discuss the calibration of its steady state value $\bar{\Upsilon}^a$ below, and relate it to the Basel-III capital conservation buffer.

Banks' aggregate net worth n^b_t represents an additional state variable of the model. It equals the difference between the gross return on loans and the sum of gross interest expenses on deposits, loan losses, penalties and dividends. The ex-post cost of penalties, which is partly a lump-sum transfer to households and partly a real resource cost, is $\mathcal{M}^b_t = \chi \ell^\ell_{t-1} \mathfrak{F}^b_t$. Dividends, which will be discussed in more detail below, are given by $\delta^b n^b_t$. Once more using the fact that bank-specific indices can be dropped, we have

$$n^b_t = \Sigma_{x \in \{c,a,y,k\}} r^x_{\ell,t} \ell^{xx}_{t-1} - r_{d,t} d_{t-1} - \mathfrak{L}^b_t - \mathcal{M}^b_t - \delta^b n^b_t \ . \tag{6}$$

4. Retail Lending Rates and Loans, Deposit Rates and Deposits

Retail lending rates and equilibrium loan levels are determined by the interaction of banks' zero profit conditions and borrowers' optimality conditions for collateral assets and collateral income flows. Both will be discussed as part of the household optimization problem below.

Deposit rates and equilibrium deposit levels are determined by the interaction of banks' optimality conditions for loan pricing, which determine the spread between lending and deposit rates and thus the opportunity cost of holding deposits, and households' and financial investors' optimality conditions for bank deposits, which determine monetary transaction cost savings or

[31] The policy rate would represent the marginal cost of funds if banks were able to appropriate the funding cost advantage of deposits due to market power. See Jakab and Kumhof (2015).

utility gains and thus the benefits of holding deposits. Transaction cost technologies and, for financial investors, money in the utility function, will be discussed as part of the household and financial investor optimization problems below.

5. Bank Net Worth Ownership and Dividends

In the model the acquisition of fresh capital by banks is subject to market imperfections. This is a necessary condition for capital adequacy regulations to have non-trivial effects. We use the "extended family" approach of Gertler and Karadi (2011), whereby bankers transfer part of their accumulated net worth to households and financial investors at an exogenously fixed net rate.

Each household and financial investor represents an extended family that consists of two types of members, workers and bankers. Bankers enter their occupations for random lengths of time, after which they revert to being workers. There is perfect consumption insurance within each family. Workers supply labour, and their wages are returned to the family each period. Each banker manages a bank and transfers earnings back to the family at the time when his period as a banker ends. Before that time he retains accumulated earnings within the bank. This means that, while the family ultimately owns banks, net worth cannot be freely injected into or withdrawn from them. That in turn means that net worth and leverage matter for the decisions of banks.

Specifically, at a given point in time a fraction $(1-\mathfrak{f})$ of each family are workers and a fraction $\mathfrak{f}$ are bankers. Bankers stay in their occupations for one further period with unconditional probability p^b. This means that in each period $(1-p^b)\mathfrak{f}$ bankers exit to become workers, and the same number of workers is assumed to randomly become bankers. The shares of workers and bankers within households and financial investors therefore remain constant over time.

Finally, both households and financial investors supply start-up funds to their new bankers, and we assume that these represent a small fraction ι^b of the existing aggregate banking sector net worth. Each existing banker makes identical decisions that are proportional to his existing stock of accumulated net worth, so that aggregate decision rules are straightforward to derive. Therefore, the magnitude that matters for aggregate dynamics is the share of aggregate banking sector net worth n_t^b paid out to households each period, $(1-p^b)\mathfrak{f}n_t^b$, net of start-up funds, $\iota^b n_t^b$. As these terms are proportional to the aggregate stock of bank net worth, their net effects can be denoted by $\delta^b n_t^b$, and our calibration is therefore simply in terms of δ^b. These parameters can alternatively be thought of as fixed dividend policies of the banking sector, and for simplicity we will utilize this terminology in the remainder of the paper.

B. Households

1. Optimization Problem

The share of households in the economy equals $1-\omega$. Households have unit mass and are indexed by i. Household-specific variables are denoted by a superscript c, to denote that this group of agents is financially constrained. Households' utility at time t depends on an external consumption habit $c_t^c(i)-\nu c_{t-1}^c$, where $c_t^c(i)$ is the per capita consumption of household i and c_{t-1}^c is lagged average per capita consumption, and where consumption is a Dixit-Stiglitz CES aggregate over varieties, with elasticity of substitution θ_p. Utility also depends on hours worked $h_t^c(i)$ and land $a_t(i)$. Households' lifetime utility function is

$$Max \quad E_0 \sum_{t=0}^{\infty} \beta_c^t \left\{ S_t^c (1-\frac{v}{x}) \log(c_t^c(i) - v c_{t-1}^c) - \psi_h \frac{h_t^c(i)^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} + \psi_a \log(a_t(i)) \right\} , \tag{7}$$

$$c_t^c(i) = \left(\int_0^1 c_t^c(i,j)^{\frac{\theta_p - 1}{\theta_p}} dj \right)^{\frac{\theta_p}{\theta_p - 1}} , \tag{8}$$

where β_c is the discount factor, S_t^c is a shock to the marginal utility of consumption, v determines the degree of habit persistence, η is the labour supply elasticity, and ψ_h and ψ_a fix the utility weights of hours worked and land.

Household income, other than income from asset returns, consists of real labour income $w_t^{hh} h_t^c(i) (1-\tau_{L,t})$, where w_t^{hh} is the real wage paid to households by unions and $\tau_{L,t}$ is the labour income tax rate, and of real lump-sum income $\iota \Omega_t / (1-\omega)$, where Ω_t is aggregate real lump-sum income, ι is the share of this income received by households, with the remainder received by financial investors, and with each member of these two groups receiving an equal per capita share of this income. Real lump-sum income is in turn given by

$$\Omega_t = \delta^b n_t^b + \Pi_t^u + trf_t - \tau_t^{ls} \overline{\mathfrak{ls}_t} + (1-\mathfrak{r}) (\mathcal{M}_t + \mathcal{T}_t) \ . \tag{9}$$

The first two items are dividends and profits received from banks and unions. The next two items are lump-sum transfers from the government and lump-sum taxes paid to the government.[32] For lump-sum taxes, $\overline{\mathfrak{ls}_t}$ is an exogenous component that grows at the gross growth rate of exogenous technological progress x, and which is therefore fixed in real normalized terms, while τ_t^{ls} is a component with mean one that can be made time-varying according to a fiscal policy rule. The final item in (9) is the lump-sum transfer component of monitoring costs $\mathcal{M}_t$ and of monetary transaction costs $\mathcal{T}_t$. The other component of these costs, $\mathfrak{r}(\mathcal{M}_t + \mathcal{T}_t)$, is a real resource cost and therefore enters the goods market clearing condition rather than the expression for transfer income.

Households invest in a number of real and financial assets. Land $a_t(i)$ has a real price of p_t^a and a real return of

$$ret_{a,t} = p_t^a / p_{t-1}^a \ . \tag{10}$$

Capital $k_t(i)$ has a real market price of q_t (Tobin's q) and a real return of

$$ret_{k,t} = (r_{k,t} + (1-\Delta) q_t - \tau_{k,t} (r_{k,t} - \Delta q_t)) / q_{t-1} \ , \tag{11}$$

where $r_{k,t}$ is the user cost of capital, Δ is the capital depreciation rate, and $\tau_{k,t}$ is the capital income tax rate. Bank deposits $d_t^x(i) = D_t^x(i)/P_t$, $x \in \{c,a,y,k\}$, have a real return of $r_{d,t} = i_{d,t-1}/\pi_t$. Under a CBDC regime, CBDC balances $m_t^x(i)$, $x \in \{c,a,y,k\}$, have a real return of $r_{m,t} = i_{m,t-1}/\pi_t$. The desired deposits are created by banks through loans, with real loan levels

[32] We do not net these items against each other so as to be able to calibrate the fiscal accounts in a realistic fashion. This becomes important in cases where, in response to shocks, all types of taxes, including lump-sum taxes, are assumed to be varied in a countercyclical fashion, so that we need the lump-sum tax component of total tax revenue to correspond to its real world counterpart. This is only possible if expenditures are specified to include both government spending on goods and services and government transfers, as in the real world.

denoted by $\ell_t^x(i) = L_t^x(i)/P_t$, $x \in \{c, a, y, k\}$. Households face real quadratic adjustment costs that penalize rapid changes in the level of loans, and that are proportional to the real level of loans:[33]

$$\mathfrak{C}_{\ell,t}^x(i) = \ell_t^x(i) \frac{\varphi_x}{2} \left(\frac{L_t^x(i)}{P_t T_t} - \frac{L_{t-1}^x}{P_{t-1} T_{t-1}} \right)^2 . \tag{12}$$

The empirical literature has found that US equilibrium real interest rates exhibit a small but positive elasticity with respect to the level of government debt – see the discussion in Section V. The model replicates this feature by assuming that households face financial assets transaction costs that are proportional to their real holdings of financial assets[34],

$$\mathfrak{C}_{f,t}^x(i) = \left(d_t^x(i) + m_t^x(i)\right) \phi_b \left(b_t^{rat} - \bar{b}_{ss}^{rat}\right) , \tag{13}$$

where $b_t^{rat} = B_t / \left(4GDP_t\right) = b_t / \left(4gdp_t\right)$ is the government debt-to-GDP ratio, and $\bar{b}_{ss}^{rat}$ is the corresponding steady state value. This cost is treated as exogenous by households, and is rebated back to households as part of lump-sum transfers $\Psi_t^c(i)$. The budgetary effect is therefore neutral, while marginal conditions are affected. Interest rates on all financial assets are assumed to be affected in an identical fashion, so that a change in the government debt-to-GDP ratio, ceteris paribus, will affect the level of interest rates but not the structure of spreads.

Households face monetary transaction costs $s_t^x(i)$, $x \in \{c, a, y, k\}$, whose level is a function of their holdings of real monetary transaction balances. The functional form of these transaction costs follows Schmitt-Grohé and Uribe (2004), and will be discussed in more detail in subsection IV.B.3. After-tax real consumption expenditure including transaction costs equals $c_t^c(i)\left(1 + \tau_{c,t}\right)\left(1 + s_t^c(i)\right)$, where $\tau_{c,t}$ is the consumption tax rate, real investment expenditure including transaction costs equals $I_t(i)\left(1 + s_t^k(i)\right)$, the real cost of land including transaction costs equals $p_t^a a_t(i)\left(1 + s_t^a(i)\right)$, and finally real expenditure on production inputs including transaction costs equals $\left(w_t^{pr} h_t^h(i) + r_{k,t} K_{t-1}(i)\right)\left(1 + s_t^y(i)\right)$, where w_t^{pr} is the real wage charged to producers by unions, and $h_t^h(i) = h_t(i) / \left(1 - \omega\right)$ and $K_{t-1}(i)$ are per capita labour and capital inputs into production.

Households borrow against one sector-specific form of real collateral and one or two forms of financial collateral in each of the four sectors $x \in \{c, a, y, k\}$. In the pre-CBDC economy bank deposits are the only form of financial collateral, while under a CBDC regime CBDC balances are a second form of financial collateral. Banks are willing to base their lending decisions on a fraction κ_t^{xr} of real collateral and a fraction $\kappa_t^{xr}\kappa_t^{xf}$ of financial collateral. In the rest of the paper we will refer to these coefficients as willingness-to-lend coefficients. The lending contract, which will be discussed in more detail in subsection IV.B.2 below, specifies that banks in period $t + 1$ will receive a fraction $\Gamma_{x,t+1}$, $x \in \{c, a, y, k\}$, of the value of collateral, where $\Gamma_{x,t+1}$ covers both the expected interest on performing loans and the expected residual value, after monitoring costs, of defaulting loans. Collateral includes the gross expected returns on financial assets $r_{d,t+1} d_t^x(i)$, $x \in \{c, a, y, k\}$, and, under CBDC, $r_{m,t+1} m_t^x(i)$, $x \in \{c, a, y, k\}$. The two real collateral assets are gross expected returns on land $ret_{a,t+1} p_t^a a_t(i)$ and on capital $ret_{k,t+1} q_t k_t(i)$. The two real

[33]This cost is introduced for numerical reasons only, and will be calibrated to be very small. A larger loan adjustment cost could be motivated by appealing to the administrative, time and reputational costs of negotiating a rapid change in loan levels.

[34]The assumption of transaction costs that are quadratic in the debt-to-output ratio is commonly used in other literatures. The main example is the open economy literature with incomplete asset markets (Schmitt-Grohé and Uribe (2003), Neumeyer and Perri (2005)). In the closed economy literature, Heaton and Lucas (1996) have used the same device.

collateral income flows are the expected real values of annualised current sales revenue $r_{n,t+1}4y_t(i)$ and of annualised current labour income $r_{n,t+1}4w_t^{hh}h_t^c(i)\left(1-\tau_{L,t}\right)$. Annualised quarterly flows are used as collateral for easier comparison with asset stock values during calibration.

As in Schmitt-Grohé and Uribe (2004), each household is the monopolistic producer of one variety of intermediate goods, and as such maximizes the difference between sales revenue and costs, where the former equals $P_t(i)y_t(i) = P_t(i)y_t\left(P_t(i)/P_t\right)^{-\theta_p}$, and the latter consist of real wage payments $w_t^{pr}h_t^h(i)\left(1+s_t^y(i)\right)$, real payments to capital $r_{k,t}K_{t-1}(i)\left(1+s_t^y(i)\right)$, and quadratic inflation rate adjustment costs as in Ireland (2001):

$$\mathfrak{C}_{p,t}^x(i) = \frac{\phi_p}{2}y_t\left(\frac{\frac{P_t(i)}{P_{t-1}(i)}}{\pi_{t-1}}-1\right)^2 . \tag{14}$$

Each household also produces capital goods, with profits equal to the difference between the market value of new investment goods $Q_tI_t(i)$ and the sum of the effective purchase price of new investment goods $P_tI_t(i)\left(1+s_t^k(i)\right)$ and quadratic investment adjustment costs[35]

$$\mathfrak{C}_{I,t}^x(i) = \frac{\phi_I}{2}I_t\left(S_t^i\frac{\left(I_t(i)/x\right)}{I_{t-1}(i)}-1\right)^2 . \tag{15}$$

Households' budget constraint in period t, in real terms, and for the case of a CBDC regime, where $m_t^x(i)>0$, $x\in\{c,a,y,k\}$, is therefore

$$\begin{aligned}
& \Sigma_{x\in\{c,a,y,k\}}\left(d_t^x(i)+m_t^x(i)\right)\left(1+\phi_b\left(b_t^{rat}-\bar{b}^{rat}\right)\right)+p_t^a a_t(i)\left(1+s_t^a(i)\right)+q_tk_t(i) \qquad (16)\\
& -\Sigma_{x\in\{c,a,y,k\}}\ell_t^x(i)\left(1-\frac{\varphi_x}{2}\left(\frac{L_t^x(i)}{P_tT_t}-\frac{L_{t-1}^x}{P_{t-1}T_{t-1}}\right)^2\right)-\Psi_t^c(i)\\
= \; & \Sigma_{x\in\{c,a,y,k\}}\left(r_{d,t}d_{t-1}^x(i)+r_{m,t}m_{t-1}^x(i)\right)+ret_{a,t}p_{t-1}^a a_{t-1}(i)+ret_{k,t}q_{t-1}k_{t-1}(i)\\
& -\kappa_{t-1}^{xr}\kappa_{t-1}^{xf}\Sigma_{x\in\{c,a,y,k\}}\Gamma_{x,t}\left(r_{d,t}d_{t-1}^x(i)+r_{m,t}m_{t-1}^x(i)\right)\\
& -\kappa_{t-1}^{ar}\Gamma_{a,t}ret_{a,t}p_{t-1}^a a_{t-1}(i)-\kappa_{t-1}^{kr}\Gamma_{k,t}ret_{k,t}q_{t-1}k_{t-1}(i)\\
& -\kappa_{t-1}^{cr}\Gamma_{c,t}r_{n,t}4w_{t-1}^{hh}h_{t-1}^c(i)\left(1-\tau_{L,t-1}\right)-\kappa_{t-1}^{yr}\Gamma_{y,t}r_{n,t}4y_{t-1}\left(\frac{P_{t-1}(i)}{P_{t-1}}\right)^{1-\theta_p}\\
& -c_t^c(i)(1+\tau_{c,t})\left(1+s_t^c(i)\right)+w_t^{hh}h_t^c(i)\left(1-\tau_{L,t}\right)+\frac{\iota}{1-\omega}\Omega_t(i)\\
& +y_t\left(\frac{P_t(i)}{P_t}\right)^{1-\theta_p}-\left(w_t^{pr}h_t^h(i)+r_{k,t}K_{t-1}(i)\right)\left(1+s_t^y(i)\right)-\frac{\phi_p}{2}y_t\left(\frac{\frac{P_t(i)}{P_{t-1}(i)}}{\pi_{t-1}}-1\right)^2\\
& +q_tI_t(i)-I_t(i)\left(1+s_t^k(i)\right)-\frac{\phi_I}{2}I_t\left(S_t^i\frac{\left(I_t(i)/x\right)}{I_{t-1}(i)}-1\right)^2 .
\end{aligned}$$

The constraint that demand for output variety i must equal supply is

$$y_t\left(\frac{P_t(i)}{P_t}\right)^{-\theta_p} = \left(T_tS_t^a h_t^h(i)\right)^{1-\alpha}\left(K_{t-1}(i)\right)^{\alpha} . \tag{17}$$

[35] Production of capital goods also involves the purchase of old depreciated capital and its resale inside period t. Because both of these transactions take place at the time t price q_t, they drop out of the optimization problem.

Finally, capital accumulation is

$$k_t(i) = (1-\Delta)k_{t-1}(i) + I_t(i) \ . \tag{18}$$

The representative household maximizes (7) subject to (8), (16), (17), (18), and banks' zero profit conditions for the four different loan types. The latter will be derived in the next subsection, along with the optimality conditions for all assets and income flows that serve as loan collateral. The optimality condition for individual goods varieties is given by

$$c_t^c(j) = c_t^c \left(\frac{P_t(j)}{P_t}\right)^{-\theta_p} \ . \tag{19}$$

The optimality conditions for aggregate consumption and investment, after normalizing by trend growth and exploiting symmetry across households, are given by

$$\frac{S_t^c(1-\frac{v}{x})}{\check{c}_t^c - \frac{v}{x}\check{c}_{t-1}^c} = \check{\lambda}_t^c \left(1+\tau_{c,t}\right)\left(1+s_t^c+s_t^{c'}v_t^c\right) \ , \tag{20}$$

$$q_t = \left(1+s_t^k+s_t^{k'}v_t^k\right) + \phi_I S_t^i \left(\frac{\check{I}_t}{\check{I}_{t-1}}\right)\left(S_t^i\frac{\check{I}_t}{\check{I}_{t-1}}-1\right) - \beta_c E_t\left\{\frac{\check{\lambda}_{t+1}^c}{\check{\lambda}_t^c}\phi_I S_{t+1}^i\left(\frac{\check{I}_{t+1}}{\check{I}_t}\right)^2\left(S_{t+1}^i\frac{\check{I}_{t+1}}{\check{I}_t}-1\right)\right\} \ . \tag{21}$$

The normalized optimality conditions for labour and capital input choices are

$$\check{w}_t^{pr}\left(1+s_t^y+s_t^{y'}v_t^y\right) = \check{m}c_t \frac{(1-\alpha)\,\check{y}_t}{h_t^h} \ , \tag{22}$$

$$r_t^k\left(1+s_t^y+s_t^{y'}v_t^y\right) = \check{m}c_t \frac{\alpha\check{y}_t}{\check{K}_{t-1}/x} \ , \tag{23}$$

where marginal cost is

$$\check{m}c_t = A\left(\frac{\check{w}_t^{pr}}{S_t^a}\right)^{1-\alpha}\left(r_t^k\right)^{\alpha}\left(1+s_t^y+s_t^{y'}v_t^y\right) \ . \tag{24}$$

2. Lending Technologies

Households at time t choose an optimal combination of real bank loans $\ell_t^x(i)$ in order to obtain an optimal combination of real bank deposits $d_t^x(i)$, where in both cases $x \in \{c,a,y,k\}$. The time t pledged[36] collateral $\mathfrak{c}_t^x(i)$ that is used to secure these loans consists, in each case, of a combination of pledged real collateral $\kappa_t^{xr}\mathfrak{c}_t^{xr}(i)$ and pledged financial collateral $\kappa_t^{xr}\kappa_t^{xf}\mathfrak{c}_t^{xf}(i)$:

$$\begin{aligned}
\mathfrak{c}_t^c(i) &= \kappa_t^{cr} r_{n,t+1} 4 w_t^{hh} h_t^c(i)\left(1-\tau_{L,t}\right) + \kappa_t^{cr}\kappa_t^{cf}\left(r_{d,t+1}d_t^c(i) + r_{m,t+1}m_t^c(i)\right) \ , \text{(collaterals)} \\
\mathfrak{c}_t^a(i) &= \kappa_t^{ar} ret_{a,t+1} p_t^a a_t(i) + \kappa_t^{ar}\kappa_t^{af}\left(r_{d,t+1}d_t^a(i) + r_{m,t+1}m_t^a(i)\right) \ , \\
\mathfrak{c}_t^y(i) &= \kappa_t^{yr} r_{n,t+1} 4 y_t \left(P_t(i)/P_t\right)^{1-\theta_p} + \kappa_t^{yr}\kappa_t^{yf}\left(r_{d,t+1}d_t^y(i) + r_{m,t+1}m_t^y(i)\right) \ , \\
\mathfrak{c}_t^k(i) &= \kappa_t^{kr} ret_{k,t+1} q_t k_t(i) + \kappa_t^{kr}\kappa_t^{kf}\left(r_{d,t+1}d_t^k(i) + r_{m,t+1}m_t^k(i)\right) \ .
\end{aligned}$$

[36]The term pledged is used here because collateral in general is not simply equal to the total value of the underlying collateral asset or income flow, but instead to the portion of collateral value that is pledged to the bank in support of the loan. This equals the total value of the respective sector-specific collateral multiplied by the corresponding willingness-to-lend coefficient.

At the beginning of period $t+1$, each borrower i draws an idiosyncratic shock which changes $\mathfrak{c}_t^x(i)$ to $\omega_{t+1}^x \mathfrak{c}_t^x(i)$, where ω_{t+1}^x is a unit mean lognormal random variable distributed independently over time and across borrowers. The standard deviation of $\ln(\omega_{t+1}^x)$, $S_{t+1}^z \sigma^x$ is the risk shock of Christiano et al. (2014), a stochastic process that will play a key role in our analysis. It has an aggregate component S_{t+1}^z that is common across all loan categories, and a type-specific and (for the simulations in this paper) constant component σ^x, $x \in \{c, a, y, k\}$. The density function and cumulative density function of ω_{t+1}^x are given by $\mathfrak{f}_{t+1}^x = \mathfrak{f}_t^x(\bar{\omega}_{t+1}^x)$ and $\mathfrak{F}_{t+1}^x = \mathfrak{F}_t^x(\bar{\omega}_{t+1}^x)$.

We assume that each borrower receives a standard debt contract from the bank. This specifies a nominal loan amount $L_t^x(i)$, the percentages of collateral value against which the bank is willing to lend κ_t^{xr} and $\kappa_t^{xr}\kappa_t^{xf}$, and a gross nominal retail rate of interest $i_{r,t}^x$ to be paid if ω_{t+1}^x is sufficiently high to avoid default. We will refer to the differences between the retail and wholesale lending rates $i_{r,t}^x - i_{\ell,t}^x$ as risk spreads, to be distinguished from the regulatory spreads $i_{\ell,t}^x - i_{d,t}$. There are three important differences between our model and those of Bernanke et al. (1999) and Christiano et al. (2014). First, bank deposits themselves (and also CBDC) can serve as financial collateral, $\kappa_t^{xf} > 0$. This is not only realistic, it also multiplies the ability of banks to create deposits. Second, the collateral coefficients on real collateral in general differ from one, $\kappa_t^{xr} \neq 1$. Third, the interest rate $i_{r,t}^x$ is assumed to be pre-committed in period t, rather than being determined in period $t+1$ after the realization of time $t+1$ aggregate shocks.[37] The latter, conventional assumption ensures zero ex-post profits for banks at all times, while under our debt contract banks make zero expected profits, but realized ex-post profits generally differ from zero. It is this ability of banks to make losses that justifies the presence of MCAR in the model (and in the real world), as a way to protect depositors from losses, thereby ensuring trust in the economy's primary medium of exchange.

Borrowers who draw ω_{t+1}^x below a cutoff level $\bar{\omega}_{t+1}^x$ cannot pay the contractual interest rate $i_{r,t}^x$ and enter bankruptcy. They must hand over all of their assets (or income flows) to the bank, but the bank can only recover a fraction $(1-\xi^x)$ of the collateral value of such borrowers. The remaining fraction represents monitoring costs. The latter are assumed to partially (fraction $1-\mathfrak{r}$) represent lump-sum income paid out to households, with the remaining portion (fraction $\mathfrak{r}$) representing real resource costs. Banks' ex-ante zero profit condition for borrower group x, in real terms, and using generic expressions for collateral, is given by

$$E_t \left\{ \left[\left(1 - \mathfrak{F}_t(\bar{\omega}_{t+1}^x)\right) r_{r,t+1}^x \ell_t^x(i) + (1-\xi^x) \int_0^{\bar{\omega}_{t+1}^x} \mathfrak{c}_t^x(i) \omega^x \mathfrak{f}_t^x(\omega^x) d\omega^x \right] - r_{\ell,t+1} \ell_t^x(i) \right\} = 0 \, . \tag{25}$$

This states that the expected payoff to lending must equal expected wholesale interest charges $r_{\ell,t+1}\ell_t^x(i)$. The first term in square brackets is the gross real interest income on loans to borrowers whose idiosyncratic shock exceeds the cutoff level, $\omega_{t+1}^x \geq \bar{\omega}_{t+1}^x$, and who can therefore pay the contractual nominal interest rate $i_{r,t}^x$. The second term is the amount collected by the bank in case of the borrower's bankruptcy, where $\omega_{t+1}^x < \bar{\omega}_{t+1}^x$. This cash flow is based on the idiosyncratic (given the presence of ω) value of collateral, but multiplied by the factor $(1-\xi^x)$ to reflect a proportional bankruptcy cost ξ^x.

The ex-post cutoff productivity level is determined by equating, at $\omega_t^x = \bar{\omega}_t^x$, the gross interest charges due in the event of continuing operations $r_{r,t}^x \ell_{t-1}^x(i)$ to the gross return on the part of the

[37] See Bernanke et al. (1999): "... conditional on the ex-post realization of R_{t+1}^k, the borrower offers a (*state-contingent*) *non-default* payment that guarantees the lender a return equal in expected value to the riskless rate."

borrower's assets that is pledged as collateral, $\mathfrak{c}^x_{t-1}(i)\bar{\omega}^x_t$. This yields the condition

$$\bar{\omega}^x_t = \frac{r^x_{r,t}\ell^x_{t-1}(i)}{\mathfrak{c}^x_{t-1}(i)} \; . \tag{26}$$

Using this equation, we can replace (25) by

$$E_t\left\{\mathfrak{c}^x_t(i)\left(\Gamma_{x,t+1} - \xi^x G_{x,t+1}\right) - r_{\ell,t+1}\ell^x_t(i)\right\} = 0 \;\; , \tag{27}$$

where $\Gamma_{x,t+1}$ is the bank's gross share in the value of collateral,

$$\Gamma_{x,t+1} = \Gamma^x_t(\bar{\omega}^x_{t+1}) \equiv \int_0^{\bar{\omega}^x_{t+1}} \omega^x_{t+1}\mathfrak{f}^x_t(\omega^x_{t+1})d\omega^x_{t+1} + \bar{\omega}^x_{t+1}\int_{\bar{\omega}^x_{t+1}}^{\infty} \mathfrak{f}^x_t(\omega^x_{t+1})d\omega^x_{t+1} \;\; ,$$

and $\xi^x G_{x,t+1}$ is the proportion of collateral value that the lender has to spend on monitoring costs,

$$\xi^x G_{x,t+1} = \xi^x G^x_t(\bar{\omega}^x_{t+1}) = \xi^x \int_0^{\bar{\omega}^x_{t+1}} \omega^x_{t+1}\mathfrak{f}^x_t(\omega^x_{t+1})d\omega^x_{t+1} \;\; .$$

In other words, the bank will set the unconditional nominal lending rate $i^x_{r,t}$ such that its expected gross share in collateral values (gross earnings of collateral assets plus real collateral income flows) is equal to the expected sum of monitoring costs and the opportunity cost of the loan.

The borrower selects optimal loan and collateral levels by solving the optimization problem described in the previous subsection, that is the borrower maximizes (7) subject to (8), (16), (17), (18) and (27). We observe that each borrower in sector x faces the same expectations of future returns, and the same risk environment represented by the functions $\Gamma_{x,t+1}$ and $G_{x,t+1}$. Aggregation of the model over borrowers is therefore trivial because both borrowing and collateral levels are proportional to borrower net worth. Borrower-specific indices i can therefore henceforth be dropped.

The optimality condition for *loans*, in real normalized form and exploiting symmetry across borrowers, is identical for all four loan categories, and given by

$$\check{\lambda}^c_t\left(1 - \frac{\varphi_x}{2}\left(\check{\ell}^x_t - \check{\ell}^x_{t-1}\right)^2 - \varphi_x\check{\ell}^x_t\left(\check{\ell}^x_t - \check{\ell}^x_{t-1}\right)\right) = \frac{\beta_c}{x}E_t\left\{\check{\lambda}^c_{t+1}r^x_{\ell,t+1}\tilde{\lambda}^x_{t+1}\right\} \;\; , \tag{28}$$

where $x \in \{c, a, y, k\}$, $\tilde{\lambda}^x_{t+1} = \Gamma^\omega_{x,t+1}/\left(\Gamma^\omega_{x,t+1} - \xi^x G^\omega_{x,t+1}\right)$, and $\Gamma^\omega_{x,t+1}$ and $G^\omega_{x,t+1}$ are the derivatives of $\Gamma_{x,t+1}$ and $G_{x,t+1}$ with respect to $\bar{\omega}^x_{t+1}$.

The optimality condition for *capital* is

$$\check{\lambda}^c_t = \frac{\beta_c}{x}E_t\left\{\check{\lambda}^c_{t+1}ret_{k,t+1}\left[\left(1 - \kappa^{kr}_t\Gamma_{k,t+1}\right) + \tilde{\lambda}^k_{t+1}\kappa^{kr}_t\left(\Gamma_{k,t+1} - \xi^k G_{k,t+1}\right)\right]\right\} \;\; , \tag{29}$$

which can be combined with (28) to yield a condition for the optimal loan contract that, for $\kappa^{kr}_t = 1$ and $\varphi_x = 0$, is identical to Bernanke et al. (1999). For the remaining types of loan collateral there are differences to the basic form of (29) that will be explained below.

For land there are two differences. First, land enters the utility function of households, so that the optimality condition contains an additional utility-related term. Second, land enters one of the four transaction cost technologies whereby a higher value of land requires higher monetary transactions balances, so that there are additional terms related to the derivatives of that technology.

The optimality condition for *land* is therefore

$$\check{\lambda}_t^c\left(1+s_t^a+s_t^{a'}v_t^a\right)-\frac{\psi_a}{\check{p}_t^a a_t}=\frac{\beta_c}{x}E_t\left\{\check{\lambda}_{t+1}^c ret_{a,t+1}\left[\left(1-\kappa_t^{ar}\Gamma_{a,t+1}\right)+\tilde{\lambda}_{t+1}^a\kappa_t^{ar}\left(\Gamma_{a,t+1}-\xi^a G_{a,t+1}\right)\right]\right\} \ . \tag{30}$$

The optimality condition for *goods prices* is modified by a term that relates to the use of sales revenue as flow collateral for working capital loans:

$$\begin{aligned}\mu_p mc_t-1 &= \phi_p\left(\mu_p-1\right)\left(\frac{\pi_t^p}{\pi_{t-1}^p}-1\right)\frac{\pi_t^p}{\pi_{t-1}^p}-\beta_c E_t\left\{\frac{\check{\lambda}_{t+1}^c}{\check{\lambda}_t^c}\frac{\check{y}_{t+1}}{\check{y}_t}\phi_p\left(\mu_p-1\right)\left(\frac{\pi_{t+1}^p}{\pi_t^p}-1\right)\frac{\pi_{t+1}^p}{\pi_t^p}\right\}\\ &\quad+\frac{\beta_c}{x}E_t\left\{\frac{\check{\lambda}_{t+1}^c}{\check{\lambda}_t^c}r_{n,t+1}4\left(-\kappa_t^{yr}\Gamma_{y,t+1}+\tilde{\lambda}_{t+1}^y\kappa_t^{yr}\left(\Gamma_{y,t+1}-\xi^y G_{y,t+1}\right)\right)\right\} \ . \end{aligned} \tag{31}$$

The optimality condition for *hours worked* is modified by a term that relates to the use of after-tax labour income as flow collateral for consumer loans:

$$\begin{aligned}\psi_h\left(h_t^c\right)^{\frac{1}{\eta}} &= \check{\lambda}_t^c\check{w}_t^{hh}\left(1-\tau_{L,t}\right)\\ &\quad+\frac{\beta_c}{x}E_t\left\{\check{\lambda}_{t+1}^c r_{n,t+1}4\check{w}_t^{hh}\left(1-\tau_{L,t}\right)\left(-\kappa_t^{cr}\Gamma_{c,t+1}+\tilde{\lambda}_{t+1}^c\kappa_t^{cr}\left(\Gamma_{c,t+1}-\xi^c G_{c,t+1}\right)\right)\right\}\end{aligned} \tag{32}$$

The optimality conditions for *deposits* and, if applicable, *CBDC*, are modified by terms that relate to the financial assets transaction cost technologies $\mathfrak{C}_{f,t}^x(i)$, $x\in\{c,a,y,k\}$ and the monetary transaction cost technologies $s_t^x(i)$, $x\in\{c,a,y,k\}$:

$$\begin{aligned}&\check{\lambda}_t^c\left(1+\phi_b\left(b_t^{rat}-\bar{b}_{ss}^{rat}\right)-s_t^{x'}\left(v_t^x\right)^2 f_t^{x'dep}\right)\\ =\ &\frac{\beta_c}{x}E_t\left\{\check{\lambda}_{t+1}^c r_{d,t+1}\left[\left(1-\kappa_t^{xr}\kappa_t^{xf}\Gamma_{x,t+1}\right)+\tilde{\lambda}_{t+1}^x\kappa_t^{xr}\kappa_t^{xf}\left(\Gamma_{x,t+1}-\xi^x G_{x,t+1}\right)\right]\right\} \ ,\end{aligned} \tag{33}$$

$$\begin{aligned}&\check{\lambda}_t^c\left(1+\phi_b\left(b_t^{rat}-\bar{b}_{ss}^{rat}\right)-s_t^{x'}\left(v_t^x\right)^2 f_t^{x'mon}\right)\\ =\ &\frac{\beta_c}{x}E_t\left\{\check{\lambda}_{t+1}^c r_{m,t+1}\left[\left(1-\kappa_t^{xr}\kappa_t^{xf}\Gamma_{x,t+1}\right)+\tilde{\lambda}_{t+1}^x\kappa_t^{xr}\kappa_t^{xf}\left(\Gamma_{x,t+1}-\xi^x G_{x,t+1}\right)\right]\right\} \ .\end{aligned} \tag{34}$$

The terms $s_t^{x'}\left(v_t^x\right)^2$ are the derivatives of monetary transaction costs with respect to total liquidity f_t^x, while $f_t^{x'dep}$ and $f_t^{x'mon}$ are the derivatives of f_t^x with respect to bank deposits and CBDC. In the pre-CBDC economy, (34) is absent.

Banks' real aggregate ex-post loan losses in sector x, which correspond to gains for their borrowers, are given by

$$\mathfrak{L}_t^x=r_{\ell,t}^x\ell_{t-1}^x-\left(\Gamma_{x,t}-\xi^x G_{x,t}\right)\mathfrak{c}_{t-1}^x \ . \tag{35}$$

These are positive if wholesale interest expenses, which are the opportunity cost of banks' retail lending, exceed banks' net (of monitoring costs) share in borrowers' collateral value. This will be the case if a larger than anticipated number of borrowers defaults, so that, ex-post, banks find that they have set their pre-committed retail lending rate at an insufficient level to compensate for lending losses.

Banks' real aggregate ex-post monitoring costs in sector x are given by

$$\mathcal{M}_t^x = \xi^x G_{x,t} \mathfrak{c}_{t-1}^x \ . \tag{36}$$

Total ex-post monitoring costs $\mathcal{M}_t$, which also include default penalties on banks $\mathcal{M}_t^b$, therefore equal

$$\mathcal{M}_t = (1-\omega)\Sigma_{x\in\{c,a,y,k\}} \mathcal{M}_t^x + \mathcal{M}_t^b \ . \tag{moco\_total}$$

3. Transaction Cost Technologies

The Transaction Cost Function and Velocity The specification of households' monetary transaction costs, $s_t^x(i) = s_t^x(v_t^x(i))$, $x \in \{c, a, y, k\}$, follows Schmitt-Grohé and Uribe (2004). Here $v_t^x(i) = e_t^x(i)/f_t^x(i)$ is velocity, $e_t^x(i)$ is a generic expression for different types of spending (or asset holding), $f_t^x(i)$ are the monetary transaction balances that are used to carry out that type of spending, and $s_t^{x'} > 0$. The general expression for transaction costs is

$$s_t^x(i) = s_t^x\left(v_t^x(i)\right) = S_t^{md} A_x v_t^x(i) + \frac{B_x}{v_t^x(i)} - 2\sqrt{A_x B_x} \ . \tag{37}$$

The shock S_t^{md} changes the demand for total liquidity. An increase in S_t^{md}, one of the key shocks in our business cycle analysis, can be thought of as a flight to safety, meaning a higher demand for liquid assets for a given volume of real economic transactions. In equilibrium a higher S_t^{md} leads to a combination of a lower velocity of circulation and a lower level of transactions.

There are four different expressions e_t^x for the numerators of the velocity expressions. They relate to consumption, investment, input purchases by producers, and real estate transactions:

$$e_t^c(i) = c_t^c(i)\left(1+\tau_{c,t}\right) \ , \tag{38}$$

$$e_t^k(i) = I_t(i) \ , \tag{39}$$

$$e_t^y(i) = w_t^{pr} h_t^h(i) + r_{k,t} K_{t-1}(i) \ , \tag{40}$$

$$e_t^a(i) = p_t^a a_t(i) \ . \tag{41}$$

An important implication of this specification is that the effective purchase prices of consumption goods, investment goods, inputs into production and real estate, which can be found in optimality conditions (20), (21), (22), (23) and (30), are higher than their market prices by mark-ups that are due to monetary frictions. Furthermore, there is an equivalence between these mark-ups and changes in distortionary tax rates. We will therefore refer to these mark-ups as liquidity taxes, and we will use a notation that emphasizes the similarity of their economic effects to those of distortionary tax rates:

$$\tau_{x,t}^{\ell iq} = 1 + s_t^x + s_t^{x'} v_t^x \ . \tag{42}$$

The equivalence can be seen clearly in the optimality conditions. In (20), the equivalence is between changes in $\tau_{c,t}^{\ell iq}$ and changes in the consumption tax rate $\tau_{c,t}$ and, if (20) is combined with (32) to obtain an expression for the marginal rate of substitution between consumption and hours worked, between changes in $\tau_{c,t}^{\ell iq}$ and changes in the labour income tax rate $\tau_{L,t}$. The expression for the equilibrium rate of return to capital (11) combined with (21) shows the equivalence between changes in $\tau_{k,t}^{\ell iq}$ and changes in the capital income tax rate $\tau_{k,t}$. For input purchases in optimality conditions (22) and (23), $\tau_{y,t}^{\ell iq}$ is equivalent to a tax on inputs into production, which is not a part of our formal model. And in (30), $\tau_{a,t}^{\ell iq}$ is equivalent to a tax on land, which is also not a part of our formal model. In our calibration, $\tau_{c,t}^{\ell iq}$ and $\tau_{k,t}^{\ell iq}$ will be quantitatively far more important than the other two liquidity taxes. The effects of changes in the size of bank balance sheets or in the quantity of CBDC, and therefore in the available quantity of monetary transaction balances, will therefore be transmitted to the real economy in a formally equivalent way to changes in labour income, capital income and consumption tax rates, all of which are distortionary. The distortion in the case of monetary transaction balances is a shortage, relative to the Friedman rule, of such balances, a shortage that can never be completely eliminated because the cost of the creation of bank deposits can never go to zero. Because monetary transaction balances play a critical role in the ability of businesses and households to engage in economic trades, we will also refer to this friction as the "cost of doing business".

Because $\tau_{c,t}^{\ell iq}$ and $\tau_{k,t}^{\ell iq}$ are by far the most important liquidity taxes in our model, our figures will report the following average as "Average Liquidity Tax":

$$\tau_{avg,t}^{\ell iq} = 1 + \frac{c_t^c\left(\tau_{c,t}^{\ell iq} - 1\right) + I_t\left(\tau_{k,t}^{\ell iq} - 1\right)}{c_t^c + I_t} \ . \tag{43}$$

Finally, aggregate monetary transaction costs are given by

$$\mathcal{T}_t = (1-\omega)\left(c_t^c\left(1+\tau_{c,t}\right)s_t^c + I_t s_t^k + p_t^a a_t s_t^a + w_t^{pr} h_t^h s_t^y + r_t^k K_{t-1} s_t^y\right) \ . \tag{44}$$

The Liquidity Generating Function (LGF) Households require monetary transaction balances, or liquidity, $f_t^x(i)$, $x \in \{c,a,y,k\}$, to lower their transaction costs.[38] The optimal choice of liquidity can be broken down into two key trade-offs, the first involving the costs and benefits of borrowing from banks to create bank deposits, and the second involving the costs and benefits of holding CBDC versus bank deposits.

When households borrow from banks in order to obtain deposits, they trade off the benefits, reduced transaction costs due to the additional deposits, against the cost, increased net interest expenses (spread between borrowing and deposit interest rates) due to the additional loans. Borrowing and deposit interest rates adjust to equate these costs and benefits. This trade-off is present in any "financing through money creation" model of banking, see Jakab and Kumhof (2015).

When households hold CBDC instead of bank deposits, they trade off the cost of holding CBDC, which equals the difference between the deposit interest rate and the (lower) interest rate on CBDC, against three benefits. First, the creation of CBDC does not require debt, so that households do not need to pay a spread between borrowing and deposit interest rates. Second, when the government supplies a strictly limited quantity of CBDC, and when CBDC and bank deposits are imperfect substitutes, the scarcity of CBDC gives it an additional non-pecuniary benefit. Third, CBDC may have technological advantages over bank deposits, whereby each unit of CBDC generates more liquidity services than a unit of bank deposits. The interest rate on CBDC (for CBDC quantity rules) or the quantity of CBDC (for CBDC price rules) adjusts to equate these costs and benefits.

[38] The liquidity demand of financial investors will be discussed in the next subsection.

In this paper, in order to study different aspects of the introduction of CBDC, we make use of three different model variants that differ in their specification of the LGF. The first and second model variants are needed to study the transition from the pre-CBDC economy to a CBDC regime, while the third variant studies the business cycle properties of an economy that has fully transitioned to a CBDC regime, and therefore uses CBDC as one of its two monetary transaction media.

In the first model variant only bank deposits enter the LGF. This model variant is used to determine the pre-CBDC steady state of the economy, which is required as the initial condition for a simulation of the transition to the CBDC regime.

In the second model variant bank deposits and CBDC jointly generate liquidity through an additively separable and non-linear (but close to linear) LGF. Some nonlinearity is required in order to numerically pin down the steady state portfolio shares of CBDC and bank deposits, but the functional form and calibration that we will assume nevertheless implies a very high elasticity of substitution between CBDC and bank deposits. Additive separability is critical for the simulations of the transition from the pre-CBDC economy to the CBDC regime, which needs to account for the fact that the transition starts at a zero stock of CBDC. Using a conventional non-separable LGF for such a simulation would imply that the initial liquidity aggregate, despite a large volume of bank deposits, would be equal to zero, which implies infinite velocity and therefore infinite monetary transaction costs. With an additively separable and near-linear LGF, CBDC simply adds to the preexisting stock of liquidity that is provided through bank deposits.

Clearly, additive separability is no longer critical once business cycle simulations are performed around the steady state of an economy that has fully transitioned to the partial use of CBDC as a monetary transaction medium. For these simulations we will therefore use a more flexible third model variant with a non-separable LGF, specifically with a CES LGF. This functional form allows us to explore the implications of different degrees of substitutability between bank deposits and CBDC, and thereby of different interest semi-elasticities of the demand for CBDC relative to bank deposits. Seen in this light, the second model variant can be seen as a version of the third model variant that features a very high (but not infinite) elasticity of substitution. The assumption of a high elasticity of substitution in the second model variant is a necessary compromise that allows us to perform a meaningful transition simulation that starts from zero CBDC balances.[39]

Liquidity generation in the **first model variant**, without CBDC, is given by

$$f_t^x = (T_t)^{1-\theta} \left(d_t^x\right)^{\theta} \ , \tag{45}$$

where the term $(T_t)^{1-\theta}$ ensures consistency with balanced growth, θ is equal across all four sectors, and $\theta < 1$ but close to 1.

Liquidity generation in the **second model variant**, with CBDC and additive separability, is given by

$$f_t^x = (T_t)^{1-\theta} \left(\left(d_t^x\right)^{\theta} + \left(T^{fintec} m_t^x\right)^{\theta} \right) \ , \tag{46}$$

where θ is equal to its value in the first model variant, and $T^{fintec} > 1$ is a relative productivity coefficient for CBDC. The size of T^{fintec} quantifies the extent to which a unit of CBDC is more

[39] It can be shown that simulations of the second and third model variants are very similar once the elasticity of substitution of the third model variant is set equal to that of the second.

productive at generating monetary transactions services than a unit of bank deposits. Ceteris paribus, for a higher T^{fintec} households are willing to hold CBDC at a lower interest rate relative to the interest rate on bank deposits. The elasticity of substitution between CBDC and bank deposits for this functional form of the LGF equals $1/(1-\theta)$.

Liquidity generation in the **third model variant**, with CBDC and CES LGF, is given by

$$f_t^x = \left((1 - S_t^{mm}\gamma_x)^{\frac{1}{\epsilon}} (d_t^x)^{\frac{\epsilon-1}{\epsilon}} + (S_t^{mm}\gamma_x)^{\frac{1}{\epsilon}} \left(T^{fintec} m_t^x \right)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} . \tag{47}$$

The role of the technology coefficient T^{fintec} is the same as in (46). The elasticity of substitution between bank deposits and CBDC is in this case simply equal to ϵ, which will be translated, in Section V, into a value for the interest semi-elasticity of the demand for CBDC relative to bank deposits.

The shock S_t^{mm} is to the demand for CBDC relative to bank deposits. An increase in S_t^{mm} can be thought of as a flight to the super-safety of CBDC at the expense of bank deposits. The equilibrium effects of a higher S_t^{mm} will depend on whether the monetary authority elastically satisfies the additional demand for CBDC.

C. Financial Investors

The share of financial investors in the economy equals ω. They have unit mass and are indexed by i. Their utility at time t depends on an external consumption habit $c_t^u(i) - vc_{t-1}^u$ (where $c_t^u(i)$ is the per capita consumption of financial investor i and c_{t-1}^u is lagged average per capita consumption), labour hours $h_t^u(i)$ and liquidity $f_t^u(i)$. Their lifetime utility function is

$$Max \quad E_0 \sum_{t=0}^{\infty} \beta_u^t \left\{ S_t^c (1 - \frac{v}{x}) \log(c_t^u(i) - \nu c_{t-1}^u) - \psi_h \frac{h_t^u(i)^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} + \psi_f \frac{\left(\frac{f_t^u(i)}{T_t}\right)^{1-\frac{1}{\vartheta}}}{1-\frac{1}{\vartheta}} \right\} , \tag{48}$$

where, apart from the discount factor β_u, all other common utility function parameters (v, η, ψ_h) are identical to those of households. All financial investors have identical initial endowments. The index i is therefore only required for the distinction between $c_t^u(i)$ and c_{t-1}^u, and can henceforth be dropped.

This utility function exhibits two differences to those of households that require discussion. The first is the absence of land, and the second is the presence of monetary transactions balances in the utility function, rather than as an argument of a monetary transaction cost technology.

We begin our discussion with the absence of land from the utility function. We have simulated model variants where financial investors hold and value (financially unencumbered) land, and where they can trade this land with households in a frictionless market. This however gives rise to two problems. First, it has the strong counterfactual implication that variations in households' demand for land would be almost entirely reflected in portfolio shifts (trades of land against deposits) between financial investors and households, with minimal effects on the price of land. On the other hand, under the assumption that only households own and value land, variations in their demand for land would be entirely reflected in changes in the price of land rather than in reallocations of land across sectors. When forced to make a choice between these polar extremes, we found the second assumption to be more plausible.[40] Second, upon the initial injection of CBDC in exchange for government bonds in our baseline transition scenario, financial investors increase their holdings of deposits by roughly 50%. If financial investors also held and traded land, this would imply that households would want to purchase a very large additional quantity of land from them, due to CBDC relaxing their financial constraints. As a result, financial investors would end up with their total deposits increasing by almost 100%. For financial investors to be willing to hold these additional deposits, this would require a much higher deposit rate relative to the policy rate than in our baseline transition scenario. We find that qualitatively most of the results reported in this paper would continue to hold if financial investors held and traded land. But quantitatively the results for the alternative assumption are more plausible, and this is therefore our baseline.

We now turn to the presence of monetary transactions balances in financial investors' utility function. The key reason for including financial investors as a second group of asset-holding agents in our model is to generate an arbitrage condition between bank deposits and government bonds. We recall that our concept of bank deposits corresponds to the entirety of items on the liability side of the banking system's balance sheet, with the exception only of net worth. We are thinking of financial investors as the primary holders of large and comparatively less liquid liabilities of financial institutions such as bonds, wholesale funds and large term deposits, which in general would also pay a higher interest rate than retail deposits, and which, crucially, would exhibit a much higher interest semi-elasticity of deposit demand than retail deposits, where the sectorial interest semi-elasticities of deposit demand ε_x^d, $x \in \{c, a, y, k, u\}$, are defined as the percentage changes in demand for bank deposits resulting from a one percentage point increase in the opportunity cost of bank deposits. Our model does not distinguish between wholesale and retail deposit interest rates in order to avoid additional (and unnecessary) complications, but it does require a quantification of the change in the deposit rate relative to the policy rate in response to shocks that lead to large changes in desired deposit holdings. This spread is a critical input into our simulations, because it is one of the two main determinants[41] of the change in banks' average funding costs when CBDC is introduced into the economy via an open market exchange against government bonds. The group of agents which determines this spread at the margin is the one whose demand function for bank deposits exhibits the highest interest semi-elasticity of deposit demand, which as we have just argued is financial investors. This demand function enters into financial investors' arbitrage condition between bank deposits and government bonds. We therefore need a functional form of that demand function which permits the calibration and simulation of a sufficiently high steady state ε_u^d, in order to match empirically plausible estimates of changes in marginal bank funding costs relative to the policy rate. While this would appear to be possible in principle using the same Schmitt-Grohé and Uribe (2004) formulation for investors as that used for households, in practice this functional form runs into numerical limits, at sufficiently low ε_u^d, where simulations break down. This is the sole reason why we adopted the alternative money-in-the-utility-function specification (48) for financial investors, which does not run into similar numerical limits.

[40] This is equivalent to assuming that financial investors do hold and value land, but that sales of land between the two groups are subject to prohibitively large real frictions.

[41] The other determinant is the response of the real policy rate itself to changes in government debt, which depends on the parameter ϕ_b in the financial assets transaction cost technology (13).

Specifying a significantly higher ε_u^d for financial investors also has an important implication for the steady state sectorial demands for CBDC. High-elasticity financial investors are highly sensitive to relative returns, including relative returns on CBDC versus bank deposits. Because the equilibrium return on CBDC is significantly lower than that on bank deposits, the simulations drive the share of CBDC that financial investors wish to hold towards zero. This again leads to numerical convergence problems, so that we have opted to simplify the model further by excluding financial-investor-held CBDC from the model altogether.[42] For the above-mentioned first and second model variants that are used for simulating the transition to a CBDC regime, we therefore have the specification of the financial investor LGF

$$f_t^u = (T_t)^{1-\theta} (d_t^u)^\theta \ , \tag{49}$$

where θ is calibrated at the same value as for households. For model variant three, on the other hand, we simply have[43]

$$f_t^u = d_t^u \ . \tag{50}$$

The real assets held by financial investors are domestic government bonds $b_t^u(i)$ and bank deposits $d_t^u(i)$. Their real consumption spending equals $c_t^u(i)(1+\tau_{c,t})$, and their after-tax real labour income is $w_t^{hh}h_t^u(i)(1-\tau_{L,t})$. They also receive the fraction of other income Ω_t that is not received by households. Their overall budget constraint is

$$\begin{aligned} & \left(b_t^u(i) + d_t^u(i)\right)\left(1 + \phi_b\left(b_t^{rat} - \bar{b}_{ss}^{rat}\right)\right) \\ = \ & r_t b_{t-1}^u(i) + r_{d,t} d_{t-1}^u(i) + \Psi_t^u(i) \\ & - c_t^u(i)(1+\tau_{c,t})(1+s_t^u(i)) + w_t^{hh} h_t^u(i)(1-\tau_{L,t}) + \frac{1-\iota}{\omega}\Omega_t \ , \end{aligned} \tag{51}$$

Financial investors maximize (48) subject to (51). The normalized optimality conditions for consumption, hours worked, government bonds and bank deposits are

$$\frac{S_t^c(1-\frac{v}{x})}{\check{c}_t^u - \frac{v}{x}\check{c}_{t-1}^u} = \check{\lambda}_t^u \left(1+\tau_{c,t}\right) \ , \tag{52}$$

$$\psi_h \left(h_t^u\right)^{\frac{1}{\eta}} = \check{\lambda}_t^u \check{w}_t^{hh} \left(1-\tau_{L,t}\right) \tag{53}$$

$$\check{\lambda}_t^u \left(1 + \phi_b\left(b_t^{rat} - \bar{b}_{ss}^{rat}\right)\right) = \frac{\beta_u}{x} E_t\left(\check{\lambda}_{t+1}^u r_{t+1}\right) \tag{54}$$

$$\check{\lambda}_t^u \left(1 + \phi_b\left(b_t^{rat} - \bar{b}_{ss}^{rat}\right)\right) - \psi_f \left(\check{f}_t^u\right)^{-\frac{1}{\vartheta}} f_t^{u'dep} = \frac{\beta_u}{x} E_t\left(\check{\lambda}_{t+1}^u r_{d,t+1}\right) \tag{55}$$

D. Unions

Unions have unit mass and are indexed by i. They are managed by households and financial investors, and their intertemporal marginal rate of substitution is an average, weighted by current labour supplies, of the intertemporal marginal rates of substitution of these two groups. Each union buys homogenous labour from households at the nominal household wage W_t^{hh}, and sells

[42] Of course this does not mean that financial investors could not trade their government bonds against CBDC, but it does imply that they would then instantaneously swap their CBDC against bank deposits.

[43] Recall that for this model variant θ is not a model parameter.

labour variety i to producers at the nominal producer wage $W_t^{pr}(i)$. Each producer demands a CES composite of labour varieties, with elasticity of substitution θ_w, so that unions' steady state gross mark-up of $W_t^{pr}(i)$ over W_t^{hh} equals $\mu_w = \theta_w/(\theta_w - 1)$. The aggregate nominal producer wage is given by W_t^{pr}. Unions face wage adjustment costs $\mathfrak{C}_{w,t}(i)$ that, similar to price adjustment costs above, make it costly to change the rate of wage inflation. These costs are assumed to be received as lump-sum income by households and financial investors. Lump-sum reimbursement of these costs to unions is denoted by $\Psi_t^w(i)$. Unions' optimization problem is therefore given by

$$\underset{\{W_t^{pr}(i)\}_{t=0}^{\infty}}{Max} E_0 \left\{ \Sigma_{t=0}^{\infty} \mathfrak{D}_{0,t} \left[\left(W_t^{pr}(i) - W_t^{hh} \right) h_t(i) - W_t^{pr} \mathfrak{C}_{W,t}(i) + \Psi_t^w(i) \right] \right\} , \tag{56}$$

$$\mathfrak{D}_{0,t} = \frac{\omega h_0^u}{h_0} \beta_u^t \frac{\Lambda_t^u}{\Lambda_0^u} + \frac{(1-\omega)\, h_0^c}{h_0} \beta_c^t \frac{\Lambda_t^c}{\Lambda_0^c} , \tag{57}$$

$$\mathfrak{C}_{w,t}(i) = \frac{\phi_w}{2} h_t T_t \left(\frac{\frac{W_t^{pr}(i)}{W_{t-1}^{pr}(i)}}{\frac{W_{t-1}^{pr}}{W_{t-2}^{pr}}} - 1 \right)^2 , \tag{58}$$

where $h_t(i) = h_t \left(W_t^{pr}(i)/W_t^{pr} \right)^{-\theta_w}$ and Λ_t^u and Λ_t^c are the multipliers of the nominal budget constraints of financial investors and households. The optimization problem yields a familiar New Keynesian Phillips curve for wages. In real normalized form, this is given by

$$\mu_w \frac{\check{w}_t^{hh}}{\check{w}_t^{pr}} - 1 = \phi_w \left(\mu_w - 1 \right) \left(\frac{\pi_t^w}{\pi_{t-1}^w} - 1 \right) \frac{\pi_t^w}{\pi_{t-1}^w} - E_t \left\{ \mathfrak{D}_{t,t+1} \frac{\check{w}_{t+1}^{pr}}{\check{w}_t^{pr}} \frac{h_{t+1}}{h_t} \phi_w \left(\mu_w - 1 \right) \left(\frac{\pi_{t+1}^w}{\pi_t^w} - 1 \right) \frac{\pi_{t+1}^w}{\pi_t^w} \right\} . \tag{59}$$

The real aggregate profits of unions, which enter into the expression for lump-sum income Ω_t, are given by $\Pi_t^u = (w_t^{pr} - w_t^{hh}) h_t$.

E. Fiscal Policy

1. Government Budget Constraint

The government budget constraint is

$$b_t^g + m_t^g = r_t b_{t-1}^g + r_{m,t} m_{t-1}^g + g_t + trf_t - \tau_t , \tag{60}$$

where government issuance of CBDC implies $m_t^g > 0$, government spending is g_t, and tax revenue is

$$\tau_t = \tau_t^{ls} \overline{\mathfrak{ls}_t} + \tau_{c,t} c_t + \tau_{L,t} w_t^{hh} h_t + (1-\omega)\, \tau_{k,t} \left(r_{k,t} - \Delta q_t \right) k_{t-1} . \tag{61}$$

For future reference we note that the steady state relationship between the debt-to-GDP ratio $\overline{b^{rat}}$ and the deficit-to-GDP ratio $\overline{gd^{rat}}$ is given by the accounting relationship $\overline{b^{rat}} = \left(\overline{gd^{rat}}/4 \right) \frac{x\bar{\pi}^p}{x\bar{\pi}^p - 1}$, where $\bar{\pi}^p$ is the inflation target of the central bank, and where the factor of proportionality 1/4 is due to the fact that our model is quarterly.

For all but one of our simulations we assume that government spending g_t is equal to a fixed fraction s_g of GDP:

$$g_t = s_g gdp_t . \tag{62}$$

2. Fiscal Policy Rule

Fiscal policy follows a structural deficit rule that differs for the pre-CBDC economy and the CBDC regime.

Pre-CBDC The fiscal deficit-to-GDP ratio for the pre-CBDC economy is

$$gd_t^{rat} = 100\frac{B_t^g - B_{t-1}^g}{GDP_t} = 100\frac{b_t^g - b_{t-1}^g/\pi_t^p}{gdp_t} = 100\frac{\check{b}_t^g - \check{b}_{t-1}^g/\left(x\pi_t^p\right)}{g\check{d}p_t} = 100\frac{g\check{d}_t}{g\check{d}p_t} \ , \tag{63}$$

where GDP_t is nominal GDP. The government's long-run target for the deficit-to-GDP ratio is fixed at a structural deficit target of $\overline{gd^{rat}}$. Automatic stabilizers allow the deficit to fluctuate with the output gap $\ln(g\check{d}p_t/gdp_{ss})$, with a response coefficient d^{gdp}. We have the rule

$$gd_t^{rat} = \overline{gd^{rat}} - 100d^{gdp}\ln\left(\frac{g\check{d}p_t}{gdp_{ss}}\right) \quad . \tag{64}$$

This rule, by itself, does not state how specific fiscal tax or spending categories adjust. Our paper studies different possibilities for this. The details can be found in subsection IV.E.3 below.

CBDC Under a CBDC regime, it is critical that the government should insulate its budget, and thereby tax rates and/or government spending, from the budgetary effects of potentially highly volatile seigniorage revenue from CBDC creation.[44] We therefore define the adjusted budget deficit ratio as $gdx_t^{rat} = gd_t^{rat} + gdm_t^{rat}$, where

$$gdm_t^{rat} = 100\frac{M_t^g - M_{t-1}^g}{GDP_t} = 100\frac{m_t^g - m_{t-1}^g/\pi_t^p}{gdp_t} = 100\frac{\check{m}_t^g - \check{m}_{t-1}^g/\left(x\pi_t^p\right)}{g\check{d}p_t} = 100\frac{g\check{d}m_t}{g\check{d}p_t} \ . \tag{65}$$

In this case we have the rule

$$gdx_t^{rat} = \overline{gdx^{rat}} - 100d^{gdp}\ln\left(\frac{g\check{d}p_t}{gdp_{ss}}\right) \quad . \tag{66}$$

In other words, for the purpose of the fiscal rule, changes in the stock of CBDC are added to changes in the stock of government debt to arrive at the relevant deficit ratio. Exchanges of CBDC against government debt therefore have no direct effect on the fiscal deficit. They do of course have many indirect effects, due to changes in economic activity brought about by CBDC issuance or withdrawal.

[44] In all but one of the simulations of business cycle shocks shown in this paper, the sole countercyclical instrument is the lump-sum tax rate, in which case these fiscal effects are less important. They become much more important when the fiscal rule varies distortionary tax rates and/or government spending countercyclically.

3. Determination of Individual Fiscal Instruments

The above fiscal rule is sufficient to endogenise one primary countercyclical fiscal instrument. Additional fiscal instruments can be endogenised by adding auxiliary fiscal policy rules. Our paper considers three different assumptions about the choice of the primary countercyclical instrument and about the corresponding auxiliary rules.

First, when we simulate the transition to a CBDC regime, we assume that the revenue gains associated with CBDC issuance, including direct gains from a lower average cost of government financing, and indirect gains from the resulting economic stimulus, are applied towards a reduction of three distortionary tax rates. Specifically, the fiscal rule (66) is assumed to endogenise the labour income tax rate $\tau_{L,t}$, while two auxiliary fiscal rules endogenise the other two distortionary tax rates $\tau_{c,t}$ and $\tau_{k,t}$ such that they follow labour income tax rates in proportional fashion,

$$\left(\tau_{c,t} - \bar{\tau}_c\right)/\bar{\tau}_c = \left(\tau_{L,t} - \bar{\tau}_L\right)/\bar{\tau}_L \ , \tag{67}$$

$$\left(\tau_{k,t} - \bar{\tau}_k\right)/\bar{\tau}_k = \left(\tau_{L,t} - \bar{\tau}_L\right)/\bar{\tau}_L \ . \tag{68}$$

Here $\bar{\tau}_L$, $\bar{\tau}_c$ and $\bar{\tau}_k$ are the pre-CBDC steady state tax rates on labour, consumption and capital. Lump-sum taxes and lump-sum transfers are held constant for this simulation, and government spending follows the rule (62).

Second, for simulations of business cycle shocks we prefer to separate the monetary and financial effects of those shocks from any fiscal effects due to variations in distortionary tax rates. The reason is that it can otherwise become difficult to interpret the results clearly. For these simulations, we therefore assume that the fiscal rule (66) endogenises the lump-sum tax rate τ_t^{ls}, while all other taxes and transfers are held constant at their steady state values, and government spending again follows (62).

Third, in one instance we discuss the effects of replacing lump-sum taxes by a full countercyclical fiscal response that uses all fiscal instruments in almost exact proportion to their steady state values. In this case, the labour income tax rate is endogenised by (66), the auxiliary rules (67) and (68) are used, and in addition we have

$$\left(\tau_t^{ls} - \bar{\tau}^{ls}\right)/\bar{\tau}^{ls} = \left(\tau_{L,t} - \bar{\tau}_L\right)/\bar{\tau}_L \ , \tag{69}$$

$$\left(g_t - s_g gdp_t\right)/s_g gdp_t = -\left(\tau_{L,t} - \bar{\tau}_L\right)/\bar{\tau}_L \ . \tag{70}$$

By distributing the fiscal response across lump-sum taxes and government spending in addition to distortionary taxes, we ensure that the large real effects of exclusive countercyclical variations of distortionary tax rates are moderated.

F. Monetary Policy

Monetary policy under CBDC is partly very recognizable, because the policy interest rate would be used as it is today, and partly unfamiliar, because monetary policy would control a second instrument that determines the manner in which CBDC is issued. We discuss both of these in turn.

1. The Policy Interest Rate

Monetary policy in our model follows a conventional inflation forecast-based interest rate rule, with interest rate smoothing and a countercyclical response to deviations of three-quarters-ahead annual inflation from the inflation target:

$$i_t = (i_{t-1})^{i_i} \left(\frac{x\bar{\pi}^p \left(1 + \phi_b \left(b_t^{rat} - \bar{b}^{rat}\right)\right)}{\beta_u} \right)^{(1-i_i)} \left(\frac{\pi^p_{4,t+3}}{(\bar{\pi}^p)^4} \right)^{\frac{(1-i_i)i_{\pi p}}{4}} . \tag{71}$$

The second term on the right-hand side is the steady state nominal interest rate, which takes into account that the steady state real interest rate is increasing in the government debt-to-GDP ratio. The third term on the right-hand side is the response to inflation, where $\pi^p_{4,t} = \pi^p_t \pi^p_{t-1} \pi^p_{t-2} \pi^p_{t-3}$.

2. The Second Monetary Policy Instrument under CBDC

For monetary policy as currently practised, the debate about the appropriate policy tool was settled decades ago, in favour of controlling a short-term interest rate. In the past there were debates, most recently in the 1980s, about whether a monetary quantity aggregate should be controlled instead, but this was eventually rejected, based on a hierarchy of three arguments. First, because there are many possible definitions of "money", this would pose problems in defining the economically relevant monetary aggregate. Second, even if the relevant monetary aggregate could be defined, the argument is that there would be problems in controlling it effectively, given that all but the narrowest monetary aggregates are under the control of private banks rather than the central bank.[45] However, none of these two arguments apply to CBDC as envisaged in this paper, which would be an economically relevant monetary aggregate as long as the quantity outstanding is sufficiently large and its substitutability with other monetary transactions media is sufficiently low, and which would be very straightforward to define and, if desired, to control. This leaves us with the third argument, which concerns the relative economic benefits of controlling the quantity versus the price of CBDC, in the spirit of Poole (1970). In this case, the argument against controlling a CBDC quantity aggregate is that this will increase aggregate volatility if shocks to money demand are sufficiently important. This argument is not as easily dismissed. In fact, our simulations will offer arguments against a CBDC quantity rule precisely on this basis, even though the results are far less pronounced than in Poole (1970), because under a CBDC regime banks remain the creators of the marginal unit of money, while in the main argument of Poole (1970) all money is created by government. In order to generate simulations that are capable of studying this question, we need to first specify CBDC quantity and price rules.

Quantity Rule for CBDC We assume that under a CBDC quantity rule, the central bank fixes the ratio of CBDC to GDP at an average value of $\overline{m^{rat}}$ over the cycle, and then permits it to vary countercyclically according to

$$m_t^{rat} = \overline{m^{rat}} S_t^{ms} - 100 m_{\pi^p} E_t \ln \left(\frac{\pi^p_{4,t+3}}{(\bar{\pi}^p)^4} \right) , \tag{72}$$

[45] This includes the possibility, known as Goodhart's Law (Goodhart (1975)), that private sector behavior may change in response to changes in the targeted monetary aggregate. This is closely related to the Lucas critique (Lucas (1976)).

where $m_t^{rat} = 100\left(m_t^g / (4gdp_t)\right)$, $\overline{m^{rat}}$ is the target and therefore steady state value of m_t^{rat}, $m_{\pi^p} \geq 0$, and S_t^{ms} is a money supply shock that represents discretionary monetary stimulus (or tightening). The baseline version of this rule, with $m_{\pi^p} = 0$ and $S_t^{ms} \equiv 1$, implies a completely fixed quantity of CBDC relative to GDP, so that any changes in demand for CBDC will be reflected in $i_{m,t}$, the interest rate on CBDC, alone (except to the extent that they affect GDP).

For $m_{\pi^p} > 0$, when inflation is expected to be above target, this rule removes CBDC from circulation, through a central bank sale to the private sector of government debt against CBDC. This has a countercyclical effect that goes beyond the effects of the policy rate i_t. It withdraws liquidity that agents can only replace through additional bank deposits, which requires additional bank loans. Bank deposits are not only more expensive than CBDC, they are also imperfect substitutes for CBDC in the LGF. As a result, there is an increase in transaction costs and in liquidity taxes $\tau_{x,t}^{\ell iq}$, which, similar to an increase in distortionary tax rates, causes a decrease in real activity.

The relationship between the equilibrium quantity of and the equilibrium interest rate on CBDC is such that a reduced quantity of CBDC leads to a *lower* interest rate on CBDC. There are two intuitive explanations for this result. First, we offer a general explanation that does not refer to the specific monetary nature of CBDC. As with any asset that is imperfectly substitutable for other assets, when the supply of that asset declines at a given demand, the remaining stocks of the asset become more valuable, so that potential holders are willing to pay a higher price for them. In other words, they are willing to accept a lower return, or a lower interest rate. Second, we offer a specific explanation that does refer to the monetary nature of CBDC as the reason for its imperfect substitutability. Abstracting from risk, the return on an asset that serves as a monetary transaction medium always consists of two parts, the non-pecuniary benefit of holding liquidity and the financial return on the asset. By arbitrage, the sum of these two has to equal the opportunity cost of holding the asset, in this case the policy rate, which is paid on a non-monetary asset, in other words on an asset that does not yield non-pecuniary benefits. When CBDC is withdrawn from circulation, liquidity becomes more scarce, which increases the non-pecuniary benefits of holding the remaining CBDC. For a given policy rate, this means that the financial return on CBDC, $i_{m,t}$, can fall.

Price Rule for CBDC We assume that under a CBDC price rule, the central bank varies the nominal interest rate paid on digital currency according to

$$i_{m,t} = \frac{i_t}{\mathfrak{sp}} \left(\frac{\pi_{4,t+3}^p}{(\bar{\pi}^p)^4} \right)^{-i_{\pi^p}^m} . \tag{73}$$

The baseline version of this rule, with $i_{\pi^p}^m = 0$, implies a fixed spread $\mathfrak{sp} > 1$ of the policy rate relative to the interest rate on CBDC, so that any changes in demand for CBDC will be reflected in m_t^g, the quantity of CBDC, alone (except to the extent that they affect the policy rate).

For $i_{\pi^p}^m > 0$, when inflation is expected to be above target, the interest rate on digital currency is *lowered* relative to the policy rate. This, ceteris paribus, makes CBDC less attractive to hold, so that agents will exchange it for government bonds. This endogenous reduction in liquidity has the same effects as the direct withdrawal of liquidity under the countercyclical quantity rule discussed above.

For any CBDC price rule, and also for a countercyclical CBDC quantity rule, to be effective, the steady state quantity of CBDC has to be large enough to make sizeable withdrawals of CBDC feasible without hitting a "quantity zero lower bound". This is one of the reasons why we will assume that the steady state stock of CBDC equals 30% of GDP. Another reason is that the technological and competitive advantages of CBDC are unlikely to be realized unless private agents can use it with sufficient frequency. And finally, the steady state efficiency gains of CBDC issuance are large and, at least over the range that is likely to be initially considered, monotonically increasing in the amount issued. It is reasonable to expect that policymakers would want to take at least some advantage of this.

G. Equilibrium and Market Clearing

In equilibrium, each group of agents maximizes its respective objective function subject to constraints, the government follows a set of fiscal and monetary policy rules, and markets clear.

1. Individual Optimality

Households maximize (7) subject to (8), (16), (17), (18) and (27), financial investors maximize (48) subject to (51), banks maximize (4) subject to (2), and unions maximize (56) subject to (57) and (58).

2. Policy Rules

Fiscal policy follows the rule (64) in the pre-CBDC economy, and the rule (66) under the CBDC regime. The individual fiscal instruments are determined according to one of the three schemes detailed in subsection IV.E.3. Monetary policy follows the inflation-forecast-based rule (71) at all times, and one of the rules (72) or (73) under the CBDC regime.

3. Market Clearing

We define aggregate consumption as $c_t = \omega c_t^u + (1-\omega)\, c_t^c$. Then the goods market clearing condition is given by

$$(1-\omega)\, y_t = c_t + (1-\omega)\, I_t + g_t + \mathfrak{r}\left(\mathcal{M}_t + \mathcal{T}_t\right) \ . \tag{74}$$

The labour market clearing condition is

$$h_t = \omega h_t^u + (1-\omega)\, h_t^c \ . \tag{75}$$

The fixed supply of land is denoted by a. Then the land market clearing condition is

$$a_t = a \ . \tag{76}$$

The capital market clearing condition is

$$k_t = K_t \ . \tag{77}$$

The bonds market clearing condition is

$$b_t^g = \omega b_t^u \ . \tag{78}$$

The market clearing condition for digital balances is

$$m_t^g = (1-\omega)\left(m_t^c + m_t^a + m_t^y + m_t^k\right) \ . \tag{79}$$

Finally, GDP is defined as

$$gdp_t = c_t + (1-\omega) I_t + g_t \ . \tag{80}$$

H. Shocks

The shocks S_t^x, $x \in \{a, c, i, md, ms, mm\}$ are autoregressive, and follow the processes

$$\ln S_t^x = \rho_x \ln S_{t-1}^x + \varepsilon_t^x \ . \tag{81}$$

The shock S_t^z multiplies each standard deviation of borrower riskiness, σ_t^x, $x \in \{c, a, y, k\}$, and therefore captures a generalized change in borrower riskiness. We assume that it consists of two components, $S_t^z = S_t^{z1} S_t^{z2}$. The first component consists of the sum of news shocks ε_t^{news} received over the current and the preceding 12 quarters, $\ln S_t^{z1} = \Sigma_{j=0}^{12} \varepsilon_{t-j}^{news}$, while the second component is autoregressive and given by $\ln S_t^{z2} = \rho_z \ln S_{t-1}^{z2} + \varepsilon_t^{z2}$.

V. Calibration

This paper studies the effects of introducing CBDC into the US economy, under the assumption that it operates away from the zero lower bound on nominal policy interest rates. We therefore calibrate the steady state of our model economy based on US data for the period 1990-2006. One period corresponds to one quarter. We begin our discussion with the calibration of the pre-CBDC economy, almost all of which carries over to the calibration of the CBDC economy. We then turn to a discussion of the calibration features that are specific to CBDC.

Tables 1-4 show a complete listing of the calibrated parameter values.[46] The calibration targets of the second column of Tables 1-4 are held constant across model variants except where indicated otherwise. The columns "Pre-CBDC" show parameter values for the pre-transition baseline economy, with LGF (45). The columns "Post-CBDC" show parameter values for the post-transition CBDC regime, with CES LGF (47), which is used to simulate business cycle shocks. Parameter values for the post-transition CBDC regime with LGF (46) are not shown. However, the non-CBDC parameters for all three of these economies are very similar. This is evident for the two regimes shown in Tables 2 and 3.

[46] The tables use the following acronyms: HH = households, FI = financial investors, WCAP = working capital, TA Cost = transaction cost, MCAR = minimum capital adequacy ratio.

A. Pre-CBDC Economy

The trend real growth rate is calibrated at 2% per annum and the inflation target at 3% per annum, through the appropriate choices of x and $\bar{\pi}^p$. The relevant historical average for the real interest rate on short-term US government debt equals around 2.5%-3.0% per annum. We calibrate it at 3% per annum, through the choice of the discount factor of financial investors β_u. The population share of financial investors is assumed to equal 5%, or $\omega = 0.05$. The steady state ratio of per capita household and financial investor consumption is for simplicity fixed at 1:1, through the choice of the parameter ι, which determines these groups' relative income levels. The parameter $\mathfrak{r}$, which determines the share of monitoring and transaction costs that represent real resource costs rather than lump-sum payments, is fixed at $\mathfrak{r} = 0.25$.

The labour income share is calibrated at 61% by fixing α. In BLS data for the US business sector this share has exhibited a declining trend over recent decades, and we therefore base our calibration on the more recent values. The private investment to GDP ratio is set to 19% of GDP, roughly its average in US data. The implied depreciation rate, at close to 10% per annum, is in line with much of the literature. The investment adjustment cost parameter, at $\phi_I = 2.5$, follows Christiano et al. (2005). The price and wage mark-ups of monopolistically competitive manufacturers and unions are fixed, in line with much of the New Keynesian literature, at 10%, or $\mu_p = \mu_w = 1.1$. Together with the assumptions for the price and wage inflation stickiness parameters of $\phi_p = 200$ and $\phi_w = 200$, this implies an average duration of price and wage contracts of 5 quarters in an equivalent Calvo (1983) setup with full indexation to past inflation. This is similar to the results of Christiano et al. (2005).

The government accounts are calibrated in considerable detail, because many of the effects of transitioning to a CBDC regime are of a fiscal nature. This includes a lower average cost of government financing, due to lower market interest rates and a partial switch to CBDC financing, which can be applied towards lower taxation (or higher spending). We study the quantitative implications of these changes for interest rates and for distortionary taxes, the latter starting from levels that are consistent with the data. The calibrated value for the initial steady state government debt-to-GDP ratio is 80%, roughly equal to its value prior to the onset of the Great Recession. The government spending to GDP ratio is set to its approximate historical average of 18% of GDP. Tax rates on labour, capital and consumption are calibrated to reproduce the historical ratios of the respective tax revenues to GDP, which are 17.6% for labour income taxes, 3.2% for capital income taxes, and 4.6% for consumption taxes. The implied initial steady state tax rates are $\tau_L = 0.317$, $\tau_k = 0.248$ and $\tau_c = 0.073$. The share of taxes that cannot be classified as either labour income, capital income or consumption taxes is set at one third. This implies a ratio of government transfer payments to GDP equal to just under 20%, approximately equal to the value of this ratio in the data. Fiscal policy can be characterized by the magnitude of automatic stabilizer effects, in other words by the size of d^{gdp}. This has been quantified by the OECD (Girouard and André (2005)), whose estimate for the United States we adopt, at $d^{gdp} = 0.34$.

The calibration of the historical U.S. monetary policy reaction function is close to the coefficient estimates reported for the Federal Reserve Board's SIGMA model (Erceg et al. (2006)) and the IMF's Global Projection Model (Carabenciov et al. (2013)). The coefficients are $i_i = 0.7$ and $i_{\pi^p} = 2.0$.

Laubach (2009), Engen and Hubbard (2004) and Gale and Orszag (2004) report empirical estimates, for the United States, of the elasticity Ξ (in percent per annum) of the real interest rate r_t with respect to changes in the government debt-to-GDP ratio b_t^{rat} (in percentage points). They report a range of $\Xi \in [0.01, 0.06]$, in other words each percentage point increase in the debt ratio increases the real interest rate by between 1 and 6 basis points. Following the transition to a CBDC regime, the stock of total government financing, which equals government debt plus CBDC, remains at 80% of GDP. As discussed in Section III, there are nevertheless two reasons why the elasticity with respect to government debt alone should be considerably larger than zero. First, the interest burden of consolidated government financing is substantially reduced when the government (partly) replaces non-monetary debt instruments with monetary CBDC, whose real interest burden vis-a-vis the private sector is less than half of that on government debt. Second, with CBDC, unlike with government debt, there is no default risk, which makes the entirety of government financing inherently less risky. On balance, we therefore calibrate the elasticity conservatively at 2 basis points, or $\Xi = 0.02$, which requires $\phi_b = 0.00005$.

In household preferences, all common behavioural parameters except for the discount factors are identical across households and financial investors. We set the labour supply elasticity to $\eta = 1$, a common choice in the business cycle literature. Habit persistence is parameterised by $v = 0.7$. The utility weight on leisure ψ_h is set to normalize the steady state labour supply of financial investors to 1.

For banks, the parameter Υ that determines the Basel minimum capital adequacy ratio (MCAR) is set to 8% of risk-weighted assets, as under both the Basel-II and Basel-III regulations. Based on details of the Basel-III regulations, we set the risk-weight parameters to $\zeta^a = 0.5$ for mortgage loans, $\zeta^c = 0.75$ for consumer loans, $\zeta^y = 1.0$ for working capital loans, and $\zeta^k = 0.9$ for investment loans. The parameter δ^b is calibrated such that banks maintain an average actual capital adequacy ratio of 10.5%, which means that they maintain a capital conservation buffer of 2.5%, as envisaged under Basel-III. The steady state percentage of banks violating the MCAR is set to 2.5% of all banks per quarter, through calibration of the bank riskiness parameter σ^b.

Our calibration of the interest rate margin between the policy rate i_t and banks' deposit rate $i_{d,t}$ is based on Ashcraft and Steindel (2008), who find a margin of 1.34%, in 2006, between the average rate of commercial banks' portfolio of treasury and agency securities on the one hand, and their overall portfolio of liabilities on the other hand. We have repeated their computations for a longer time span, and found similar spreads in the 1990s and between 2004 and 2006, but lower spreads from 2001-2004. During that period the policy rate was lowered to unprecedented (at that time) levels while interest rates on bank liabilities, of which a significant portion is of longer duration, did not drop to the same extent. This in fact suggests that during times of steady policy rates the margin of Ashcraft and Steindel (2008) may be biased downwards because banks' portfolio of liabilities has a longer average duration, and thus a larger term premium, than their portfolio of government securities. On the other hand, our paper uses a broader concept of deposit liabilities than Ashcraft and Steindel (2008), who focus only on commercial banks, while we include all financial institutions that offer liquid liabilities, including the shadow banking system. The liabilities of non-bank financial institutions are less liquid on average than those of commercial banks, and therefore need to offer a higher interest rate. Our compromise calibration is to fix a steady state interest rate margin between i_t and $i_{d,t}$ of 1%, through the choice of the discount factor of households β_c. Finally, the steady state interest rate margin between the deposit rate $i_{d,t}$ and the rate that banks would charge on riskless[47] private mortgage loans $i^{a}_{\ell,t}$ is fixed at 1.5% per annum in steady state.[48] This means that the margin of the riskless private lending rate over the policy rate is 0.5%, which is roughly equal to the historical spread of the 3-month US$ LIBOR over the 3-month treasury bill rate. The implied value for penalty costs χ equals around 1.2% of the value of assets of those banks who are in violation of MCAR.

The overall size of the financial system's balance sheet is calibrated at 180% of GDP, with the net worth component determined through the Basel regulatory parameters. This figure is a compromise, in that it exceeds the 100% of GDP reported for 2006 by the Federal Financial Institutions Examination Council (2007), while it is well below the almost 350% of GDP reported by Gorton et al. (2012). Federal Financial Institutions Examination Council (2007) includes US commercial banks, US branches and agencies of foreign banks, thrift institutions, and credit unions, but it excludes the shadow banking system, while Gorton et al. (2012) include shadow banks. Similar numbers to Gorton et al. (2012) are provided by Pozsar et al. (2010), who use the Flow of Funds database to show that, just prior to the onset of the 2007 crisis, total liabilities of the US commercial banking sector equalled around 100% of GDP, while the size of the shadow banking sector was around 150% of GDP, and despite a large subsequent contraction it still exceeded the size of the commercial banking sector in 2010. Gorton (2010) emphasizes that a key component of the shadow banking system is the repo market, but that its size is very hard to estimate reliably due to a lack of comprehensive data coverage. Nevertheless, available estimates for the relevant period range between US$ 10 trillion (gross) according to Hördahl and King (2008) and US$ 12 trillion according to Gorton (2010). Our reason for including at least a conservative estimate of the liabilities of the shadow banking system is that these liabilities perform money-like functions that must not be omitted from a model of the modern US financial system. This is emphasized by Gorton et al. (2012), who describe the functions of such liabilities as follows: "To the extent that debt is information-insensitive, it can be used efficiently as collateral in financial transactions, a role in finance that is analogous to the role of money in commerce." A large share of financial system debt can therefore command an interest rate discount below the policy rate by yielding these financial, rather than purely goods market, transaction services. Because, in terms of our model, the full menu of safe assets considered by Gorton et al. (2012)[49], and also by Pozsar et al. (2010), includes items not always intermediated by the financial system and items that would represent double-counting in a model with a single aggregated banking system, we adopt a compromise calibration of a total balance sheet size of 200% of GDP[50]. We then reduce this figure to 180% by excluding the 20% of GDP of treasury debt that is held by the banking system, assuming instead that all of treasury debt is held by the non-bank private sector. This has modelling advantages, while the effects on our results are small.

[47] While in the model no borrower is charged the riskless rates $i^{x}_{\ell,t}$, $x \in \{c, a, y, k\}$, these rates are important as the base lending rates to which risk-based spreads are added.

[48] We choose the wholesale mortgage rate because mortgages have the lowest risk-weighting, which makes them the closest equivalent in the model to loans to blue chip, or near-zero-risk, borrowers.

[49] The asset categories included in safe assets by Gorton et al. (2012), who use the Federal Reserve's Flow of Funds database, include bank deposits, money market mutual fund shares, commercial paper, federal funds, repurchase agreements, short-term interbank loans, treasuries, agency debt, municipal bonds, securitized debt and high-grade financial sector corporate debt.

[50] Primarily due to the inclusion of the shadow banking system, this figure is much larger than traditional measures of the money supply such as M2, MZM or M3 (discontinued in 2006), with even M3 only reaching around 80 percent of GDP in 2006.

A second reason for choosing the 180% figure is that it is approximately consistent with Flow of Funds information on the size of borrowing exposures of the US corporate and household sectors. Just before the onset of the crisis, the total amount of credit market debt outstanding for non-financial businesses reached 80% of GDP. Ueda and Brooks (2011) show that around 20%-25% of this was short-term debt with a maturity of up to one year. We therefore set the steady state values of short-term or working capital loans ℓ_t^y equal to 20% of GDP.[51] We allocate the remaining 60% of GDP to long-term or investment loans ℓ_t^k. Similarly, just prior to the crisis the ratio of residential mortgages to GDP reached around 80% of GDP. Our model does not feature housing investment, but rather a fixed factor referred to as land. A significant portion of housing investment does represent the acquisition of the underlying land, and a significant portion of the remainder represents purchases of preexisting houses where additional investment is only a minor consideration. On the other hand, housing construction does play an important role, and in fact in US statistics it is included as part of private investment. Out of the stock of residential mortgages, we therefore allocate one quarter, or 20% of GDP, to ℓ_t^k, to represent the housing construction component of investment loans, taking the total to 80% of GDP. We allocate the remaining 60% of GDP to mortgage loans ℓ_t^a. The Flow of Funds data show that short-term consumer loans reached just under 20% of GDP prior to the crisis, and we therefore set ℓ_t^c to 20% of GDP in steady state. All loans-to-GDP ratios are calibrated by fixing the steady state willingness-to-lend coefficients for real collateral κ^{xr}, $x \in \{c, a, y, k\}$, under the maintained assumption that the coefficients for financial collateral are equal those for real collateral, $\kappa^{xf} = 1$, $x \in \{c, a, y, k\}$. The resulting values are within a plausible range, at $\kappa^{cr} = 0.88$, $\kappa^{ar} = 1.17$, $\kappa^{yr} = 0.35$ and $\kappa^{kr} = 0.52$.[52]

We comment next on the calibration of balance sheet leverage. Ueda and Brooks (2011) contain information on the leverage or debt-to-net-worth ratio of all listed US companies. For the overall non-financials group this has fluctuated around 140% since the early 1990s, while for the core manufacturing and services sectors it has fluctuated around 110%. Leverage for unlisted companies is likely to have been lower on average due to more constrained access to external financing. We therefore choose a steady state leverage ratio of 100% for investment loans, and we interpret this as the ratio of investment loans to the difference between the value of physical capital and investment loans. The parameter that allows us to fix this ratio is the money-in-the-utility-function parameter of financial investors ψ_f, which determines the overall size of banks' balance sheets after all other categories of loans and deposits have been calibrated through other parameter choices. For mortgage leverage ratios we use data from the Flow of Funds and Fannie Mae to decompose the total value of the US housing stock into its net worth and mortgage-debt components. The resulting leverage ratio equalled around 80% in the two decades prior to the crisis. We adjust this figure by removing the portion of the housing stock that is owned outright, which is equal to around a third of the total, and which can be thought of as the housing stock owned by financial investors, under the maintained assumption (see our discussion in Section IV.C) that financial investors do not trade land with households. For the remaining housing stock, which can be thought of as the housing stock of households, this results in a leverage ratio of 200%, which we interpret as the ratio of mortgage loans to the difference between the value of land and mortgage loans. We calibrate this ratio by fixing the physical stock of land a. We also normalize the steady state price of land to 1, through the choice of the utility

[51] Bates, Kahle and Stulz (2008) show that, similar to our model, nonfinancial firms simultaneously borrow and hold large amounts of cash, reaching a cash-to-assets ratio of 23.2% in 2006.

[52] Note that $\kappa_t^{ar} > 1$ does not imply that loans exceed the value of land collateral. Steady state loans equal 60% of the value of total collateral, and 67% of the value of land. See below for a discussion of mortgage leverage ratios.

weight ψ_a. For consumer loans and working capital loans, balance sheet leverage is not an easily defined concept, because a large share of their collateral takes the form of flow income rather than valuable assets. We therefore do not calibrate leverage ratios for these two sectors.

Sectorial deposits-to-GDP ratios are harder to calibrate from the data than loans-to-GDP ratios, because deposits are fungible and can be used for multiple rounds of spending, and for multiple types of spending, during any given period. Our baseline assumes that steady state consumption and investment deposits equal 50% and 30% of GDP, and that working capital and real estate deposits each equal 10% of GDP. We calibrate these ratios using the transaction cost technology coefficients A_x, $x \in \{c, a, y, k\}$. The LGF parameter on deposits θ in (45) (and also in (46)) is fixed at $\theta = 0.95$.

Given the overall balance sheet size of the financial sector of 180% of GDP, and the fact that Basel regulations fix steady state bank capital at 14.385% of GDP, this implies that the deposits of financial investors account for 65.615% of GDP. As a result of these choices, the two most important liquidity taxes, for consumption and investment, equal $\tau_{c,t}^{\ell iq} = 1.049$ and $\tau_{k,t}^{\ell iq} = 1.071$. In other words, monetary transaction costs account for steady state liquidity taxes of around 5%-7%. Given their smaller deposits-to-GDP ratios, liquidity taxes for working capital and real estate are considerably smaller.

Steady state interest rate lending spreads are computed from average margins between different corporate and household borrowing rates and 3-month US treasury bill rates. Ashcraft and Steindel (2008) compute, for 2006, a 2% spread for real estate loans, a 3% spread for commercial and industrial loans, and a 5% spread for credit card and other consumption loans. We therefore fix the following steady state spreads: 2% for mortgage loans, 5% for consumer loans, 3% for working capital loans, and 1.5% for investment loans. Only the latter deviates from Ashcraft and Steindel (2008), principally because these authors only consider the commercial banking sector, while long-term corporate funding, to the extent that it does not come directly from capital markets, comes to a significant extent from the shadow banking system where spreads tend to be lower. For example, in the commercial paper market average spreads prior to the crisis averaged less than 0.5%. We calibrate these risk spreads by fixing the sector-specific borrower riskiness parameters, σ^x, $x \in \{c, a, y, k\}$.

We fix steady state loan default rates at levels that are, as much as possible, consistent with the literature. Ueda and Brooks (2011) show that the default rate for non-financial listed U.S. companies has averaged around 1.5% since the early 1990s. Default rates for smaller, non-listed companies are known to be higher. We therefore set the steady state default rate for investment loans to 1.5% of all firms per period, and the default rate for working capital loans to 3%. As for household loans, the average personal bankruptcy rate was just under 1% over the two decades preceding the Great Recession. But, as discussed in White (1998), only a fraction of households who default file for bankruptcy. For banks it is often more cost-efficient to simply write off the debt, especially for smaller personal loans. And even for mortgage loans, lenders may be willing to incur significant costs in restructuring the loan before forcing the borrower to resort to bankruptcy. In our model, the costs of such write-offs and restructurings, and not only outright bankruptcy, represent default events. We therefore calibrate the steady default rate on mortgages at 2.5%, and on short-term household loans at 4%. These high rates can also be justified by appealing to the U.S. credit score distribution and associated delinquency rates, where delinquency is a 90+ days late payment on any type of debt. In the United States, as of 2012, 15% of households fell into a score range that exhibited a delinquency rate of 50% or more, and another 12% into a range that exhibited a delinquency rate of around one third. We calibrate default rates by setting the sector-specific default cost parameters, ξ^x, $x \in \{c, a, y, k\}$.[53]

The interest semi-elasticities of deposit demand ε^d_x, $x \in \{c, a, y, k, u\}$, are the percent changes, evaluated at the steady state, in deposit demands in response to a one percentage point increase in the opportunity cost of holding those deposits. Traditional empirical studies have estimated money demand equations separately rather than as part of an overall general equilibrium model, and have found interest semi-elasticities of money demand of 5 (Ball (2001)) or even lower (Ireland (2007), O'Brien (2000)). We adopt Ball's estimate of 5 for all four deposit categories held by households, through an appropriate choice of the transaction cost technology parameters B_x, $x \in \{c, a, y, k\}$. However, we find that if that value was adopted across all five sectors, in a general equilibrium model where deposits are a very broad aggregate representing all transactions-related (goods *and* financial markets transactions) liabilities of the financial system, the implications become highly counterfactual. Specifically, a bank lending boom triggered by greater optimism about the creditworthiness of borrowers raises the volume of lending and therefore of deposits. With a generalized low interest semi-elasticity, despite the fact that lending spreads over the deposit rate fall, the overall cost of lending could rise dramatically, because the increase in deposits would require a very large increase of deposit interest rates relative to the policy rate. This has not been a prominent feature of interest rate data during lending booms. To the contrary, the data suggest that the *marginal* wholesale depositor is willing to increase his holdings of deposits substantially for a relatively modest increase in deposit rates relative to policy rates.[54] In other words, the marginal depositor's money demand exhibits a high interest semi-elasticity of deposit demand. The marginal depositor of our model is the financial investor, and for these agents we assume that $\varepsilon^d_u = 250$, which can be calibrated by fixing ϑ accordingly. The quantitative implications of this choice will be explained in the next subsection.

B. CBDC Parameters - Transition Simulations

For the transition simulations, realistic scenarios require that the non-CBDC-specific structural features and parameters of the post-transition CBDC economy should remain identical to those of the pre-transition economy. We therefore assume that under CBDC the structural model, other than for the additively separable presence of CBDC in the LGF, remains identical, and that all the main structural parameters also remain identical.

The introduction of CBDC is assumed to occur instantaneously, in an amount equal to 30% of GDP, through an exchange of government bonds against CBDC. The pre-transition LGF is given by (45), and the post-transition LGF by (46), in each case with $\theta = 0.95$. The latter implies an elasticity of substitution between CBDC and bank deposits of 20. For the purpose of the transition simulation we assume that the government keeps CBDC at its initial level relative to GDP throughout the entire transition, with $\overline{m^{rat}} = 30$ and $m_{\pi^p} = 0$. The introduction of CBDC reduces the ratio of government debt to GDP to 50% on impact. It is assumed that the government subsequently keeps the target for government debt at 50% of GDP, and thus the target for total government financing at an unchanged 80% of GDP, through the appropriate

[53] The assignment of parameters to calibrated data moments is of course somewhat arbitrary. An obvious example is σ^x and ξ^x, which are interchangeable.

[54] Recall that when we refer to deposits, what we have in mind is all non-equity items on the balance sheet of the aggregate financial system.

choice of $\overline{gdx^{rat}}$, but allowing for short-run deviations due to countercyclical fiscal policy. We discussed above that the baseline calibration for automatic stabilizers, for business cycle simulations, is $d^{gdp} = 0.34$. However, for the transition the policymaker knows that the economy will eventually experience a significant and permanent expansion in output. This can be sped up by allowing for a stronger countercyclical response of fiscal policy to the gap between current and long-run potential output. For the transition simulation we therefore set $d^{gdp} = 0.50$.

After the transition to a CBDC regime the structure of bank balance sheets adjusts endogenously. To ensure that the steady state capital conservation buffer remains at 2.5%, in other words that the steady state Basel ratio remains at 10.5%, we assume that banks' dividend payout parameter δ_b is adjusted accordingly. However, the implied change in δ_b is small.

The interest rate spread between the deposit rate and the CBDC rate is determined by the costs and benefits of holding CBDC instead of bank deposits that were discussed in Section IV.B.3. We calibrate the financial technology coefficient T^{fintec} such that in the post-transition steady state this spread equals 80 basis points. The implied value of T^{fintec} is 1.153.

C. CBDC Parameters - Business Cycle Simulations

For the business cycle simulations, we switch to the third model variant, with CES LGF (47). Wherever possible we calibrate the steady state to reproduce the calibration targets of the first or pre-transition model variant in Tables 1-3. However, there are new and unique aspects to the CES LGF specification, most importantly the allocation of CBDC to its multiple possible uses. To calibrate this, we proceed in steps.

First, we calibrate the parameters of both the second and third model variants using the calibration targets of the first model variant from Tables 1-3, while in each case setting the steady state government debt-to-GDP and CBDC-to-GDP ratios equal to 50% and 30%, as shown in Table 4. One implication is that for the business cycle simulations the steady state real policy rate is assumed to remain at 3% per annum.

Second, the second model variant is used to compute the allocation of CBDC to its multiple uses, under three additional assumptions. First, we again assume $\theta = 0.95$. Second, T^{fintec} is again calibrated to generate a post-transition steady state spread between the deposit and CBDC rates of 80 basis points. Third, the interest semi-elasticities of deposit demand ε_x^d, $x \in \{c, a, y, k, u\}$, remain unchanged at their original values of 5 for households and 250 for financial investors, while the formulas for ε_x^d are of course different once LGF (46) is used in place of LGF (45). On the basis of these assumptions, the calibration endogenously generates an allocation of CBDC to m^c, m^a, m^y and m^k.

Third, for the third model variant this allocation, together with the assumption that the elasticity of substitution between CBDC and bank deposits equals $\epsilon = 2$, and the assumption that the spread between CBDC rates and deposit rates remains at 80 basis points, is then used as an input into the calibration. This is done through an appropriate choice of the CES quasi-share parameters γ_x, $x \in \{c, a, y, k\}$ and of $\overline{m^{rat}}$. In this calibration, the values of the interest semi-elasticities of deposit demand ε_x^d are again assumed to remain unchanged at 5 and 250, and the formulas for these elasticities are again different once LGF (47) is used in place of LGF (46).

Our assumption that the elasticity of substitution between CBDC and bank deposits equals $\epsilon = 2$ can be shown to imply an interest semi-elasticity of the demand for CBDC relative to deposits of around 34 ($\varepsilon^m = 34$), which means that in response to a one percentage point increase in the CBDC interest rate relative to the deposit rate, and holding deposits constant, demand for CBDC would increase by around one third, or 10 percent of GDP. We have no guidance from the literature as to whether this is a reasonable assumption. We do note, however, that the elasticity of substitution across retail deposit accounts at different banks is unlikely to be very high. For example, Competition and Markets Authority (2015) reports that in a 2014 survey of UK households, almost 60% had been with the present provider of their main current account (roughly equivalent to a US checking account) for over 10 years, and only 10% had been with their present provider for under two years. This is despite the fact that in this period the interest rates offered on instant-access accounts, even excluding bonus 'teaser' rates, varied across banks by as much as 2.5 percentage points.[55]

This provides some loose evidence against very high substitutability, and may also be suggestive of potentially much lower substitutability. To illustrate the effects that this would have, we will repeatedly consider the alternative of $\epsilon = 0.075$, which corresponds to $\varepsilon^m = 1.6$. This is a world where any increase in liquidity provided through bank loans d^x requires a similar increase in CBDC liquidity m^x in order to increase effective liquidity f^x, $x \in \{c, a, y, k\}$.

VI. Results

Figures 1-17 show impulse response functions for the main variables of the model, first for the transition simulation over a period of 60 quarters, and then for a variety of business cycle shocks over a period of 32 quarters. All real variables (GDP, consumption, investment) are shown in percent deviations from trend, which means that any increase in GDP shown in the figures has to be added to an underlying trend real growth rate of 2% per annum. All interest rates are shown in percent per annum, and in levels, because their initial steady state values convey important information. The inflation rate and all fiscal and liquidity tax rates are shown in percentage point deviations from their initial steady state values, and the same is true for all balance sheet ratios relative to GDP. The acronym "EoS" in the title of several figures relates to the elasticity of substitution ϵ between CBDC and bank deposits, and "pp" denotes percentage points.

A. Steady State Effects of the Transition to CBDC

Figure 1 studies the effects of introducing, in period 0, CBDC equal to 30% of GDP, through a purchase of government bonds of the same value. This policy increases steady state GDP by 2.94%. The completion of the transition to this new steady state takes well over two decades. Figure 1 studies the first 15 years, or 60 quarters. To ensure that the final steady state values are visible in the figure, the red dotted line in each case displays the change, in period 0, of the long-run steady state of the respective variable. The solid line shows the actual transition path. The maintained assumptions for monetary policy are that the policy rate continues to follow the

[55] Our figure of 2.5 percentage points comes from the individual bank interest rates that the Bank of England uses to construct its published Quoted Rates series. For more detail, see http://www.bankofengland.co.uk/statistics/Pages/iadb/notesiadb/household_int.aspx.

inflation-forecast-based interest rate rule (71), and that the supply of CBDC is kept at 30% of GDP through a quantity rule (72) with $m_{\pi^p} = 0$. The maintained assumptions for fiscal policy are that the fiscal rule (66) endogenises the labour income tax rate $\tau_{L,t}$, and that the other two distortionary tax rates $\tau_{c,t}$ and $\tau_{k,t}$ follow labour income tax rates in proportional fashion according to rules (67) and (68). Government and transfer spending, and lump-sum taxes, are fixed, the first relative to GDP according to (62).

The long-run net beneficial effects of this CBDC issuance are driven by four main factors, reductions in real policy rates, increases in deposit rates relative to policy rates, reductions in distortionary tax rates, and reductions in liquidity tax rates due to increases in monetary transaction balances.[56] We will discuss each of these in turn, followed by a discussion of the transition dynamics.

The first two factors are reductions in real policy rates and increases in deposit rates relative to policy rates. Our calibration of $\Xi = 0.02$ implies that a 30 percentage points drop in the government debt-to-GDP ratio is associated with a 60 basis points drop in the real policy rate. We see this in Figure 1, where this rate drops from 3% initially to 2.4% in the long run. However, banks' funding costs are determined by the deposit interest rate, which is the rate that they have to pay on the funds that they themselves create to finance loans. The spread between the policy and deposit rates is determined by financial investors, who are the deposit holders with the highest interest semi-elasticity of deposit demand, $\varepsilon^d_u = 250$. Financial investors hold deposits equal to 65.615% of GDP immediately before the introduction of CBDC, 95.615% of GDP immediately thereafter as they first trade government debt against CBDC with the government and then CBDC against bank deposits with households, and 105.06% of GDP in the very long as they accumulate additional financial assets. Financial investors therefore eventually experience an increase in their deposit holdings of 65%. With $\varepsilon^d_u = 250$, this suggests that the deposit interest rate must rise by 26 basis points relative to the policy rate in order to make financial investors willing to hold the additional deposits.[57] We see this narrowing of the spread between policy and deposit rates in Figure 1, but it amounts to 30 basis points, not 26 basis points, with the policy and deposit rates dropping by 0.60% and 0.30%, to 2.4% and 1.7%, respectively.[58] An interpretation of the narrowing of the spread between the policy rate and the deposit rate is that banks' funding becomes more expensive because they have to rely on wholesale funding to a greater extent, with a sizeable share of retail monetary transaction services now being performed by CBDC instead. However, we caution that this interpretation does not have an exact counterpart in our model, which does not feature separate retail and wholesale deposits. There is no initial value for the long-run interest rate on CBDC, as there was no CBDC prior to the transition. The 80 basis points post-transition steady state real interest rate discount relative to the deposit rate was instead calibrated, and implies a long-run level of the real CBDC interest rate of 0.9%. As for lending volumes and lending interest rates, we observe that there is a long-run increase in bank lending and bank deposits equal to 5% of GDP[59], as banks satisfy a higher demand for deposit balances. The latter is triggered by a combination of increased

[56] We re-emphasize that these calculations cannot fully account for the transition risks, nor for some of the minor costs and benefits, listed in Section III.

[57] While one can argue about the exact size of this effect, a calibration of $\varepsilon^d_u < 100$ would clearly have unrealistic implications. It can be shown that in that case real deposit rates would actually increase relative to the pre-CBDC steady state, despite the substantial drop in policy rates.

[58] The calibration of ε^d_u is based on an approximation at the original steady state. The approximation error is due to the fact that the economy subsequently moves very far away from that steady state.

[59] Loans and deposits drop slightly relative to GDP on impact, but their absolute drop is close to zero.

economic activity, and of the increase in CBDC balances, which requires an increase in deposit balances due to imperfect substitutability between the two forms of monetary transaction balances. Despite this increase in lending, the average real wholesale lending rate declines from 3.94% to 3.65%, and thus follows the deposit rate almost exactly, with no significant increase in the regulatory spread. This is because the increase in bank lending is not accompanied by a significant increase in bank riskiness. The average retail lending rate however declines by only 15 basis points, from 5.22% to 5.07%, with the 15 basis points increase in the retail lending spread reflecting higher loan-to-value ratios among bank borrowers, principally for mortgage loans but also for investment loans. To summarize, the reduction in real deposit and wholesale lending rates of around 30 basis points, which is the net effect of a 60 basis points reduction in the real policy rate and a 30 basis points increase in deposit rates relative to policy rates, is the first major reason behind the observed increase in GDP.

The third factor behind the long-run output effects is a reduction in distortionary tax rates, which are shown in the sixth row of Figure 1. Because they start from different initial levels and change proportionally, the labour income tax rate drops by 132 basis points in the long run, and the capital and consumption tax rates by 103 and 30 basis points. There are several reasons for this drop. Most importantly, while the target for the sum of government debt and CBDC remains at 80% of GDP, the cost of financing government is considerably lower than in the pre-CBDC economy, first because the reduction in government debt to 50% of GDP lowers the steady state real and nominal interest rates on government debt by 60 basis points, to 2.4%, and second because the interest rate on the CBDC component of government financing ends up being 150 basis points below the interest rate on government debt, at 0.9%. The combined effect of these two savings in funding costs adds up to around 1.0% of GDP in the long-run. Furthermore, as transfers are held constant in real normalized terms, the long-run increase in GDP reduces the ratio of transfers to GDP by around 0.5%. The constancy of real government transfer payments, such as benefit payments, after a boost to GDP is an automatic stabilizer effect. With the sum of government debt and CBDC remaining constant relative to GDP, this means that the long-run ratio of tax revenue to GDP can fall by 1.5%, and this is accomplished by the above-mentioned reductions in distortionary tax rates. To summarize, the across-the-board reduction in distortionary tax rates is the second major reason behind the observed increase in GDP.

The fourth factor behind the long-run output effects is an increase of monetary transaction balances. Not only does CBDC increase from 0% to 30% of GDP, but also, due to positive synergies and economic stimulus effects, bank deposits increase by 5% of GDP in the long run. In addition to these very significant increases in the simple amounts of transaction balances, CBDC yields higher liquidity services dollar for dollar, due to the higher technical efficiency of the technology used by CBDC, with $T^{fintec} = 1.153$. As a result of this increase in liquidity, the liquidity taxes on consumption goods, investment goods, production inputs and land drop by 39, 77, 5 and 1 basis points in the long run. Because $\mathfrak{r} > 0$, this saves resources. But more importantly, the effects of reductions in these monetary frictions are equivalent to those of reductions in distortionary tax rates, and the size of their change is smaller but not dramatically smaller than that of tax rates. To summarize, the across-the-board reduction in liquidity taxes is the third major reason behind the observed increase in GDP.

The fact that bank deposits increase despite a substantial injection of CBDC should be emphasized. A common fear concerning the introduction of CBDC is that it might take business away from banks. However, as we have just seen, the beneficial output and efficiency effects of CBDC in fact lead to synergies whereby demand for banks' services increases substantially.

It is useful to comment on the connection between the foregoing and the Friedman rule. The Friedman rule (Friedman (1969)) states that, if possible, the money supply should be expanded to reach the satiation point of the money demand function, where the marginal benefit of money reaches the marginal cost of producing it, which in a world of exogenously created high-powered money equals zero. However, this is not possible in a world where money is created endogenously by the private banking system, because the marginal benefit of this money in equilibrium again has to equal the cost of creating it, but in this case that cost equals the spread between loan and deposit rates. Because of financial frictions and financial regulation, this spread will always be positive. This means that money holders will always be unable to reach the satiation point of their money demand function. The introduction of CBDC, which is created debt-free and therefore does not come with the cost of a banking spread, takes agents closer to the satiation point of their money demand function, and this explains part of the beneficial effects of CBDC.

The overall long-run output effect in Figure 1 is a very substantial GDP gain of 2.94%, with the consumption gain at 2.23% and the investment gain at 5.28%. Table 5 decomposes this into the contributions of the above-mentioned four effects, including all possible permutations of $\Xi = 0.02$ versus $\Xi = 0$ for real policy rate effects, and distortionary taxation versus pure lump-sum taxation for tax rate effects. Liquidity effects are more difficult to quantify directly, and will therefore be inferred indirectly from the other results in Table 5.

Table 5 shows that the output effects of CBDC for the case of zero real policy rate effects, $\Xi = 0$ instead of $\Xi = 0.02$, and lump-sum taxation instead of distortionary taxes, are a slightly negative -0.04%. We will comment on this number below, but for now take it as our point of comparison. Starting from this point, lower real interest rates account for the majority of the output gains, with $\Xi = 0.02$ instead of $\Xi = 0$ implying a steady state real policy rate of 2.4% per annum rather than 3.0%, and leading to an output gain of 1.47%. But the gains from lower distortionary tax rates are not very far behind, with distortionary taxes instead of lump-sum taxes implying lower rather than constant distortionary tax rates, and leading to an increase in the output gain of 0.76%. The sum of these two partial effects equals 2.23%, which is lower than the actual combined effects by 0.75 percentage points. The difference is due to synergies between the two, as part of the benefits of lower real interest rates comes through the government budget, with lower costs of government financing allowing the government to reduce distortionary taxes. As a rough approximation, the contributions of lower real policy rates and lower distortionary taxes equal 1.8% and 1.1%.

Finally, the GDP effect of -0.04%, at $\Xi = 0$ and a regime of lump-sum taxation, allows us to quantify, again approximately, the liquidity effects of CBDC. Recall that at this point the steady state real policy rate remains at 3.0% per annum. But the deposit interest rate still rises relative to the policy rate when CBDC is issued, which means that in this case it rises in absolute terms. We have simulated the model at $\Xi = 0$ and a regime of lump-sum taxation, and have found that the magnitudes of steady state real interest rate increases for deposit rates, average wholesale and average retail lending rates are 17, 17 and 8 basis points, respectively. The effect of CBDC issuance on interest rates therefore goes in the opposite direction of the simulation in Figure 1, and is approximately half as large. Again, as a rough approximation, this means that, absent liquidity effects, the GDP effects of CBDC issuance would now be negative, and roughly half as large as the positive effects of lower real interest rates in Table 5. The fact that the GDP effects are instead close to zero reflects the effects of additional monetary transaction balances, both CBDC and bank deposits, whose real effects are mainly due to reductions in liquidity taxes. While these are not as large as the effects of distortionary tax rates, especially when taking account of the synergistic effects of taxes with lower real interest rates, they are nevertheless very large, with a GDP gain of somewhere between 0.5% and 1.0%.

As in any exercise of this kind, the estimated output gain of approximately 3% is dependent on the details of the model calibration. We see this as a strength rather than a weakness, because it makes it possible to explore the sensitivity of our results to many different aspects of that calibration. As an illustration, previous versions of this paper reported an output gain of approximately 1% rather than 3%. The reasons for this change can be identified, and debated, precisely. There are two of them. First, the original calibration of the interest semi-elasticity of deposit demand of financial investors was much lower, at 50 rather than 250.[60] This meant that steady state deposit interest rates changed from 2.0% to 2.4% rather than to 1.7% following the introduction of CBDC, in other words the spread between policy and deposit rates narrowed by 100 rather than 30 basis points. We have discussed in section V.B why we consider this too large, and therefore why financial investors' interest semi-elasticity of deposit demand should be calibrated at a much higher value than 50. Making that change, which implies lower bank funding costs and therefore lending rates in the new steady state, took output effects from 1% to just over 2%. Second, in the original calibration financial investor consumption accounted for 19.4% of steady state GDP, while in the current version it accounts for 3.2%. Under the original calibration a large share of aggregate demand was therefore not subject to financial frictions, and did not benefit from the generalised reduction in financial frictions following the introduction of CBDC. To the contrary, this portion of aggregate demand was negatively affected by a reduction in returns on financial investments. We did not want to overemphasize this feature of the model, and therefore reduced the share of financial investor consumption in GDP. Making this change took output effects from just over 2% to just under 3%.

It remains to comment on some aspects of the shorter-run profile of the transition. One important feature is an increase in the rate of inflation immediately after the introduction of CBDC, which is sizeable at 0.8 percentage points, and quite persistent. One reason is that the 2.94% gains in aggregate output are only realized after a prolonged transition period, while aggregate demand picks up much more quickly due to the immediate realization of the associated wealth effects. With demand running ahead of supply, inflation rises, and with the policy rate reacting to inflation, this means that all real interest rates are elevated for some time. Because higher real interest rates also increase the interest cost component of the government budget, distortionary taxes have to temporarily remain above their lower long-run level in order to satisfy the fiscal rule. This of course further dampens activity in the short run, relative to the final steady state. However, tax rates always remain well below their initial levels. This is one reason why investment almost immediately grows substantially, and nearly reaches its long-run level after about one year, and why GDP immediately expands by around 1.5%, and thereafter temporarily stalls in years 2 and 3, but without declining significantly. The other reason is an immediate, strong and persistent drop in liquidity taxes, which is due both to the direct injection of liquidity through CBDC and the associated creation of additional bank deposits. This is shown in the final row of Figure 1.

[60] The main reason was numerical difficulties with calibrating high interest semi-elasticities using the Schmitt-Grohe and Uribe (2004) specification of money demands. See the discussion in Section IV.C.

B. Quantity Rules or Price Rules for CBDC?

1. Credit Cycle Shocks

In this subsection we compare the properties of strict ($m_{\pi^p} = 0$) quantity rules (72) versus strict ($i^m_{\pi^p} = 0$) price rules (73) for CBDC issuance. The exploration of countercyclical rules, with $m_{\pi^p} > 0$ or $i^m_{\pi^p} > 0$, is considered in subsection VI.C below.

Figure 2 assumes a standardized shock scenario that we will use repeatedly throughout this paper. Specifically, over a period of three years banks receive positive news shocks to the riskiness of their corporate borrowers, as in Christiano et al. (2014), with magnitudes that by the end of the third year reduce the standard deviation of borrower riskiness by 40%. At that time the news shocks are reversed, and banks receive a large negative shock to the riskiness of their borrowers, which thereafter unwinds, as a first-order autoregressive process with coefficient 0.8.

We begin with a discussion of the features of this boom-bust credit cycle that are common under both rules. Banks respond to lower credit risk among their borrowers through both a reduction in retail lending spreads and an increase in the stock of loans, where the additional loans are funded through the creation of additional deposits. Both lower spreads and increased lending increase economic activity and therefore inflation. As a result the policy rate rises, and so does the deposit rate, which is closely arbitraged with the policy rate by financial investors. The increase in the deposit rate, which is the marginal cost of funds to the banking system, is however not large enough to reduce lending, because retail spreads drop to initially more than offset the increase in the deposit rate. The increase in lending and thus in deposit creation, which reaches around 17% of GDP by the end of the third year, reduces liquidity taxes $\tau^{\ell iq}_{c,t}$ and $\tau^{\ell iq}_{k,t}$ by well over 2 percentage points. This is equivalent to consumption tax and capital income tax cuts of the same order of magnitude, and therefore has highly stimulative effects. This is the main reason for the substantial increase in GDP, by over 2% by the end of the third year. Investment is the main beneficiary, but consumption also increases. When the negative shock hits after three years, banks respond through a combination of higher spreads and a dramatic reduction in lending, and thus in the provision of liquidity to the economy. The resulting large increase in liquidity taxes is the main factor behind a drop in GDP of well over 4%.

The initial reduction in borrower riskiness raises the relative efficiency of creating liquidity through bank loans instead of CBDC. As a result, issuance of an unchanged quantity of CBDC relative to GDP, under a CBDC quantity rule, requires that CBDC pays a higher interest rate than before. In other words, the spread of the policy rate over the CBDC rate declines, by over 30 basis points by the end of the third year. Government debt declines, as the government's fiscal rule requires a countercyclical surplus during this period. When the financial cycle turns at the end of the third year, all of these developments are reversed.

On the other hand, maintenance of a fixed spread of the policy rate over the CBDC rate, under a CBDC price rule, implies that the quantity of CBDC demanded, and issued on demand by the central bank, declines during the boom phase of the cycle, by more than 4 percent of GDP by the end of the third year. During this period households therefore return CBDC to the government in exchange for government bonds, with the result that privately held government debt declines by much less during the boom phase, despite the fiscal surpluses. This endogenous reduction in the quantity of CBDC is countercyclical, in that it counteracts the increase in liquidity, and therefore in economic activity, generated by additional bank lending. There is a slight offset in that bank lending grows a little more strongly under a CBDC price rule, but the overall effect on liquidity is contractionary compared to a CBDC quantity rule. This can be seen in the behaviour of liquidity taxes, which decline less strongly under a price rule. This therefore has countercyclical effects on GDP, which grows less strongly during the boom phase under a CBDC price rule. But the difference between price and quantity rules is quantitatively small. A strict CBDC price rule therefore offers only a modest degree of countercyclicality relative to a strict CBDC quantity rule.

2. Demand Shocks and Technology Shocks

Impulse responses for a range of standard real shocks, including shocks to investment and consumption demand, and technology shocks, are not shown. The reason is that the differences between quantity and price rules are very small, because under all these real shocks the implied changes in the demand for monetary transaction balances are quite small in size. Therefore, because the marginal dollar of liquidity can always be endogenously created or cancelled by banks, the fact that a quantity rule for CBDC does not endogenously accommodate relatively modest-sized changes in demand for liquidity is of little consequence. This changes when the substitutability between CBDC and bank deposits is extremely low, a topic to which we return in subsection VI.B.4 below.

We now show that for shocks other than real shocks, specifically for shocks to the demand for total liquidity and especially to the demand for CBDC liquidity, the performance of CBDC quantity and price rules differs more substantially.

3. Shocks to the Demand for Total Liquidity and for CBDC Liquidity

Shocks to the Demand for Total Liquidity Figure 3 studies a shock to S_t^{md} that increases households' demands for total monetary transaction balances f_t^x, $x \in \{c, a, y, k\}$. This shock can be interpreted as a flight to safety, whereby existing transaction balances are used to a greater extent as safe saving vehicles rather than being spent on real economic transactions. This reduces the velocity of circulation of existing liquidity by between 5% and 10%, in other words it makes existing liquidity less effective at facilitating real economic transactions. As a result, the average liquidity tax increases by over 2.5 percent on impact, and GDP contracts by over 1.5% on impact. In order for the economy to adjust to this shock, there is therefore an urgent need to generate additional liquidity, by whichever means available. The means are additional bank loans that create additional bank deposits, and central bank injections of additional CBDC, through purchases of government bonds.

One key result of this simulation is that banks can respond to this shock extremely quickly, by expanding their balance sheets by well over 10% of GDP on impact. The main reason is that banks are modelled according to the financing through money creation (FMC) view of Jakab and Kumhof (2015) rather than according to the traditional intermediation of loanable funds (ILF) view, which in its simplest version would have required some agents to deposit real savings in banks in order for the banks' balance sheet to be able to grow.[61] As in the real world, FMC banks

[61] Some ILF models also allow bank balance sheets to grow through bank purchases of stocks and bonds. However, as shown in Jakab and Kumhof (2015), this channel only applies to a very small fraction of banks' overall activity,

have the ability to expand their balance sheets completely independently of real savings, and this is critical for their ability to react to money demand shocks. This does however come with a cost, because banks are only able to expand liquidity generation against collateral, and by charging higher lending spreads. This cost is another reason for the steep increase in liquidity taxes.

The other key result is that for this shock the CBDC policy rule makes a more significant difference. Demand for both bank deposits and CBDC balances increases following the shock. Under a quantity rule the demand for additional CBDC balances is not satisfied, and instead the real interest rate on CBDC is allowed to drop very substantially, from 1.2% to around -0.1%.[62] Under a price rule the CBDC interest rate instead follows the policy rate, which means that it remains significantly higher than under a quantity rule. At this rate households wish to hold an additional quantity of CBDC equal to more than 4% of GDP. This is supplied on demand by the central bank, and helps to satisfy the additional demand for total liquidity. As a result liquidity taxes rise by less than under a quantity rule, and this has noticeable, albeit not extremely large, effects on the size of the contraction in GDP. As we will show in subsection VI.B.4, these effects become stronger as the substitutability between CBDC and bank deposits decreases.

This result is therefore clearly in the spirit of Poole (1970), who found that when money demand shocks dominate, a price rule, meaning an interest rate rule, is more effective than a quantity rule at limiting macroeconomic volatility. However, there are important differences. We have already explained that, unlike Poole (1970), we do not find significant advantages for either rule when real shocks dominate, because the marginal unit of liquidity is created by banks, and banks can respond very flexibly to changes in liquidity demand unless those changes are extremely large. The latter continues to be true even under shocks to demand for total liquidity, although now the shock is large enough for a price rule to exhibit a noticeably better performance. We would have to be in a world closer to that of Poole (1970), where liquidity can *only* be provided by the central bank, for the flexible provision of central bank liquidity under a price rule to make a very large difference. Clearly, the more the shock to the demand for liquidity is specific to CBDC, and the lower the substitutability between CBDC and bank deposits, the closer we are to such a world. We will turn to these two issues, in that sequence, next.

Shocks to the Demand for CBDC Liquidity Figure 4 studies a shock to S_t^{mm} that increases households' demands for CBDC balances m_t^x, $\{x \in c, a, y, k\}$. This type of liquidity can only be provided by the central bank, and it is therefore for this shock that we observe the largest differences between CBDC quantity and price rules. But the caveat is that differences in the GDP effects are still not extremely large, mainly because in our model CBDC constitutes only a fraction of the economy's total liquidity.

The shock in Figure 4 increases the liquidity benefits of CBDC, and therefore increases the demand for CBDC at any given CBDC interest rate, or reduces the interest rate that the CBDC issuer needs to offer at any given CBDC quantity. Under a CBDC price rule the CBDC interest rate is directly tied to the policy rate, which barely moves in response to this shock. The increased liquidity benefits of CBDC therefore translate into an increase in CBDC issuance equal to around 8 percent of GDP on impact. This is a switch to a socially more efficient way of creating liquidity, because it circumvents the financial frictions and increases in spreads that accompany an increase

corporate lending, and even for this small fraction it is quantitatively not significant as an explanation for the observed large jumps of the consolidated US financial system's balance sheet.

[62] The nominal CBDC interest rate remains substantially above zero.

in liquidity creation through banks. Because it reduces liquidity taxes by around 0.2%, there is an immediate and persistent expansion of GDP, albeit of a modest magnitude of around 0.15%. Under a CBDC quantity rule the increase in demand for CBDC is not satisfied by the central bank, so that the increase in the liquidity benefits of CBDC instead translates into a significant drop in the interest rate on CBDC, from 1.2% to around 0.25% on impact. In this case liquidity taxes increase by over 0.3% on impact, because liquidity production through banks does not offer an alternative when the shock specifically affects demand for CBDC. As a result GDP drops, but again only by a modest magnitude approaching 0.2%. The quantity of bank deposits only changes little under either rule. Under a CBDC price rule, because of the expansion in GDP, bank deposits increase slightly in the long run, while the opposite is true under quantity rules.

This takes us to a very important matter. One of the main concerns with a CBDC regime is a perceived vulnerability of banks to sudden runs from bank deposits to CBDC, at least when CBDC is supplied by the central bank on demand as under a CBDC price rule. As shown clearly in Figure 4, this reflects a misconception. The key issue is the issuance arrangements for CBDC. What is assumed in the model is that the central bank only issues CBDC against government bonds. As seen very clearly in Figure 4, the counterpart of the increased CBDC issuance under a price rule is therefore a drop in privately held government bonds, with only a minor and transient drop in bank deposits. The "run" is therefore not from bank deposits to CBDC but from government debt to CBDC – because that is the only way in which households can obtain CBDC.

4. The Role of Substitutability Between CBDC and Bank Deposits

The magnitude of the elasticity of substitution ϵ between CBDC and bank deposits in the transaction cost technologies of households has two important implications for the behaviour of our model economy. First, under any given rule for CBDC issuance, the economy behaves very differently at low ϵ. We will illustrate this in Figure 5 below, which studies the behaviour of the economy under credit cycle shocks, under a CBDC quantity rule, and under our baseline assumption of $\epsilon = 2$ versus an alternative assumption of $\epsilon = 0.075$. Second, any differences between CBDC quantity and price rules are accentuated at very low ϵ. We will illustrate this in Figure 6 below, which studies the behaviour of the economy under shocks to the demand for total liquidity, under CBDC quantity versus price rules, and under $\epsilon = 0.075$. The relevant comparison here is across figures, namely with Figure 3, which shows the same shock at our baseline calibration of $\epsilon = 2$.

Credit Cycle Shocks Figure 5 studies the same sequence of shocks as in Figure 2, but limiting the analysis to a CBDC quantity rule. The solid line therefore corresponds to the quantity rule simulation of Figure 2, with $\epsilon = 2$. The dotted line simulates the same shock for $\epsilon = 0.075$, with all other calibration targets remaining at the same values as in Tables 1-4.

We recall that under this shock banks are much more willing to create additional liquidity through lending, because the riskiness of their borrowers has declined. Bank loans and bank deposits therefore grow strongly. However, under a CBDC quantity rule the central bank does not provide any additional CBDC beyond what is justified by additional GDP growth. With low substitutability between CBDC and bank deposits, this severely limits the liquidity benefits of additional liquidity creation by banks. As a result, average liquidity taxes, or the costs of doing business, decline by far less, around 0.6% by the end of the boom phase compared to 2.3% under our baseline. This means that the initial boost to GDP is about one quarter of the baseline. Lower GDP growth in turn has a negative feedback effect on bank lending, which grows significantly less under the low substitutability scenario, for the same sequence of shocks.

The largest differences between the two scenarios concern the behaviour of interest rates. The deposit rate, because of the high interest semi-elasticity of deposit demand ε_u^d of financial investors, always remains closely arbitraged with the policy rate. However, the CBDC rate behaves very differently from the deposit rate, and this becomes extreme at low ϵ. For our baseline substitutability of $\epsilon = 2$, the CBDC rate rises by slightly more than other rates during the boom phase, because under this shock the liquidity benefit of CBDC declines relative to that of bank deposits when substitutability is high enough. However, under low substitutability the much stronger complementarity between CBDC and bank deposits reverses this result, as it implies that any increase in bank deposit creation strongly increases the liquidity benefits of CBDC. When, under a quantity rule, this CBDC is not supplied on demand, the result is a much lower real interest rate on CBDC. In our simulation, the CBDC real interest rate therefore drops below zero during an extended part of the boom phase, and shoots up very strongly during the crisis, with spreads relative to the policy rate reaching almost 5% during the boom phase, and almost -3% during the crisis.

Shocks to Demand for Total Liquidity Figure 6 returns to a comparison of CBDC quantity and price rules. Both simulations assume that $\epsilon = 0.075$, with all other calibration targets kept at exactly the same values as in Tables 1-4. Other than the calibration of ϵ, this simulation is therefore identical to Figure 3.

When discussing Figure 3, we noted that the main benefit of a price rule is that it flexibly supplies additional CBDC liquidity when demand for total liquidity increases, thereby adding to the effects of the additional liquidity creation through banks. Figure 6 shows that these benefits of a price rule are significantly stronger when ϵ is lower, in other words when the complementarities between CBDC and bank deposits are stronger. This is because, with strong complementarities between CBDC and bank deposits, the increase in bank lending is much more effective at helping the economy cope with the shock when it is accompanied by an increase in CBDC issuance. Furthermore, because of these complementarities, and because of the beneficial output effects of a price rule, the increase in CBDC issuance actually permits a much larger expansion of bank lending. Taken together this implies a much more significantly reduced increase in liquidity taxes, and therefore a less severe reduction in GDP, than in Figure 3.

Finally, note that the larger quantity effects under a CBDC price rule are replaced by larger price effects under a CBDC quantity rule. Specifically, when CBDC is held constant relative to GDP following a shock to the demand for total liquidity, the liquidity benefits of CBDC increase very strongly when CBDC is highly complementary with bank deposits, and as a result the CBDC real interest rate under this scenario drops below -2% on impact, and remains negative for an extended period thereafter. The nominal CBDC interest rate remains positive, however.

5. The Role of Uncertainty about CBDC Demand

Before turning to countercyclical CBDC policy rules, we should mention that another important factor in choosing between quantity and price rules arises at the time of the initial introduction of CBDC. This is that the policymaker will initially find it very difficult to estimate the steady state interest rate spread sp between the policy rate and the CBDC rate that corresponds to the desired steady state ratio $\overline{m^{rat}}$ of CBDC to GDP. In the absence of a reliable estimate of sp, issuing CBDC under a price rule would therefore have unpredictable consequences for the quantity of CBDC demanded, and thus for the demand for government bonds and bills. In that case it might be preferable to initially issue CBDC under a quantity rule in order to let the market establish a reasonable range for CBDC interest rates. After an appropriate period of time, policy could then switch, if desired, to a price rule that could take these lessons into account.

C. Countercyclical CBDC Policy Rules

In this subsection we will demonstrate that countercyclical CBDC policy rules can make a significant contribution to stabilising the business cycle, over and above the stabilisation provided by traditional inflation-forecast-based interest rate rules. As discussed in Section IV.F.2, we study the properties of a specific class of CBDC quantity or price rules that respond to forward-looking inflation, in a similar fashion to the policy rate. One key result is that quantity and price rules have very similar potential for countercyclical policy responses to standard shocks, in the sense that for every m_{π^p} in the quantity rule (72) there is an $i^m_{\pi^p}$ in the price rule (73) that delivers similar countercyclical properties. The choice between quantity and price rules must therefore mostly be based on the fact that price rules, as mentioned above, perform better under money demand shocks. Another important result is that the relative performance of countercyclical quantity and price rules depends critically on the elasticity of substitution between CBDC and bank deposits, with a lower elasticity implying that smaller quantity responses and larger interest rate responses to inflation are required to achieve the same degree of countercyclicality.

1. Credit Cycle Shocks

Figures 7, 8 and 9 again study credit cycle shocks, first under quantity rules and then under price rules. In each case, the black solid line is identical to the corresponding line in Figure 2, in other words to quantity rules with $m_{\pi^p} = 0$ or to price rules with $i^m_{\pi^p} = 0$.

In Figure 7, the two alternatives represent $m_{\pi^p} = 4$ and $m_{\pi^p} = 8$. In other words, for a one percentage point deviation of expected inflation from target, the quantity of CBDC in circulation is reduced by 4 or 8 percent of GDP. Because inflation rises by just over 2 percentage points just before the collapse of the boom, this means that CBDC equal to around 8 and 16 percent of GDP is withdrawn under the two alternative scenarios, and we can see this clearly in the subplot "CBDC/GDP". The counterpart of this decrease in CBDC is an increase in privately held government debt, but the magnitude of this increase is smaller because the government runs a countercyclical surplus during the boom phase. After the boom collapses, inflation quickly drops, and additional CBDC is injected in countercyclical fashion.

Withdrawing CBDC during booms and injecting it during crashes dampens the cycle of aggregate liquidity creation. There is an offset from increased creation of bank deposits, but this is only partial. As a result, liquidity taxes decrease by far less during the boom and increase by less during the crash. The effect is similar to a sizeable (by close to 1 percentage point under $m_{\pi^p} = 8$) increase in capital income, labour income and consumption taxes during the boom, and a cut of the same taxes during the crash. This has a strong impact on GDP during the boom phase, with the size of the increase in GDP almost cut in half for $m_{\pi^p} = 8$. There is also a positive effect on the speed of recovery from the crash.

The changes in interest rates clearly illustrate the importance of the interest semi-elasticities of deposit demand ε^d and of the demand for CBDC relative to bank deposits ε^m. First, the interest semi-elasticity of deposit demand ε^d for the overall economy is determined at the margin by financial investors, as explained in Section V.B, and their high semi-elasticity of $\varepsilon_u^d = 250$ implies that the policy rate and the deposit rate are almost arbitraged one-for-one. Second, the spread between the policy rate and the CBDC rate, and therefore the spread between the deposit rate and the CBDC rate, evolves precisely as predicted by ε^m. Our baseline elasticity of substitution $\epsilon = 2$ was shown in Section V.C to imply that a 10 percentage point decrease in CBDC relative to GDP is accompanied by a one percentage point decrease in the CBDC rate relative to the deposit rate. We see this exactly for $m_{\pi^p} = 8$, where a decrease in CBDC of 16 percent of GDP is accompanied by an increase in the spread between the policy rate and the CBDC rate of 1.6 percentage points. The economic intuition is that the withdrawal of CBDC during the boom raises the non-pecuniary liquidity benefits of the remaining CBDC, so that the government can pay a lower interest rate on the remaining CBDC balances. While the switch from government financing via low-interest CBDC to high-interest government debt tends to drive up government interest payments relative to the baseline, the sizeable drop in the interest rate paid on the remaining CBDC, together with the smaller increase in real interest rates due to the countercyclical effect of this policy, means that the government saves significantly on interest payments, and is therefore able to lower taxes relative to the case of $m_{\pi^p} = 0$. This has more significant real effects when the taxes that are varied countercyclically are distortionary, as we will discuss in subsection VI.E.

In Figure 8, the two alternatives represent $i^m_{\pi^p} = 0.4$ and $i^m_{\pi^p} = 0.8$. In other words, for a one percentage point deviation of expected inflation from target, the interest rate on CBDC, which at $i^m_{\pi^p} = 0$ follows the policy rate one-for-one, is reduced by 40 or 80 basis points relative to the policy rate. The main observation is that the impulse responses for this case are very similar to Figure 7. This illustrates that for any choice of countercyclicality m_{π^p} of a quantity rule, there is a corresponding choice of countercyclicality $i^m_{\pi^p}$ of a price rule that delivers very similar countercyclical properties. But there are some small differences. The price rule in Figure 8 was calibrated so that the price and quantity of CBDC would behave similarly to Figure 7. However, the GDP effects relative to the baseline are weaker in Figure 8. This is because under the price rule the baseline itself, due to the endogenous contraction of CBDC, already has some countercyclical effect.

Figure 9 reproduces the most countercyclical scenario, with $i^m_{\pi^p} = 0.8$, of Figure 8. This figure is designed to illustrate the behaviour over the cycle of the two main policy-determined interest rates under a CBDC price rule, the nominal policy rate i_t and the nominal CBDC rate $i_{m,t}$. The key subplot is on the bottom left. The solid line is the nominal policy rate, which responds to the inflation of the boom phase and the disinflation of the crash in the usual fashion. The dotted line is the notional fixed-spread CBDC rate, which deducts the fixed spread of the CBDC price rule

(73) from the policy rate.[63] The dashed line is the actual CBDC rate under the countercyclical CBDC price rule. We observe that the policy and CBDC rates form a corridor whose width fluctuates over the business cycle, widening during the upturn and narrowing during the downturn. In other words, while the policy rate *rises* during the boom, the CBDC rate, relative to the policy rate, *drops* in order to withdraw CBDC from circulation and thereby help to prevent overheating. The opposite happens during the downturn. This is an illustration, highly stylized of course, of what monetary policy might look like under a CBDC regime.

2. Demand Shocks and Technology Shocks

Figure 10 studies the performance of countercyclical CBDC price rules following a contractionary investment demand shock. Investment contracts by around 18% on impact, and GDP by around 2.75%. The resulting drop in inflation, of slightly less than 2 percentage points, calls for a large countercyclical injection of CBDC, of around 15 percent of GDP in the case of the most countercyclical alternative considered. This reduces liquidity taxes, in an analogous fashion to tax cuts, and thereby contributes to a faster recovery in GDP. But even for this very large injection, the effect on GDP is at most equal to 0.25% of initial steady state GDP immediately following the shock. The reason is that, unlike the credit cycle shock, this shock is to the real fundamentals of the economy rather than to monetary conditions. Countercyclical CBDC policy most directly affects monetary conditions, and its ability to dampen economic fluctuations that originate on the real side of the economy is therefore more limited. The same can be shown for consumption demand shocks and for technology shocks.

3. Shocks to the Demand for Total Liquidity

Figure 11 studies the performance of countercyclical CBDC price rules following a shock to the demand for total liquidity. The solid line of Figure 11 corresponds to $i^m_{\pi^p} = 0$, and is identical to the solid line in Figure 3. The other two lines represent $i^m_{\pi^p} = 0.4$ and $i^m_{\pi^p} = 0.8$. An increase in demand for total liquidity has strong contractionary and disinflationary effects. Under a countercyclical rule lower inflation, which drops by well over 1 percentage point, signals to the policymaker to raise the CBDC interest rate, and thereby to permit the injection of sizeable quantities of CBDC into the economy. Because this mitigates the shortage of liquidity, the increase in liquidity taxes is reduced, and as a result the contraction in GDP is smaller. Figure 12 shows that using a countercyclical quantity rule instead of a price rule, with $m_{\pi^p} = 4$ and $m_{\pi^p} = 8$, again has very similar quantitative effects.

4. The Role of Substitutability between CBDC and Bank Deposits

Figures 13 and 14 simulate the same shock to the demand for total liquidity as Figures 11 and 12, but under the assumption that the elasticity of substitution between CBDC and bank deposits equals $\epsilon = 0.075$. All other calibration targets are kept at the values of Tables 1-4. This includes the calibration of the countercyclical coefficients of the CBDC rules, which are kept at

[63] Of course this is counterfactual, because if the CBDC interest rate did not react to inflation, the policy interest rate itself would evolve differently.

$i^m_{\pi^p} = 0.4/0.8$ and $m_{\pi^p} = 4/8$, respectively. Because of the different elasticity of substitution, the steady state and dynamic behaviour of the two model economies are different, so that Figures 13 and 14 cannot be directly compared to Figures 11 and 12 in terms of the magnitudes of the impulse responses. What we can compare however is the differential responses to countercyclical policies.

Figure 13 shows that, using the baseline calibration of the countercyclical CBDC price rule, the implied changes in the quantity of CBDC are far smaller than in Figure 11, and therefore so are the effects on liquidity taxes and on GDP. The reason is that with a low ϵ the demand for CBDC is highly unresponsive to its interest rate. The CBDC price rule would therefore have to be calibrated far more aggressively to deliver the same stimulus results as in Figure 11.

Figure 14 shows that the opposite is true for CBDC quantity rules. First of all, given the strong complementarity between CBDC and bank deposits, and the absence of an endogenous response of the quantity of CBDC, the baseline rule with $m_{\pi^p} = 0$ now shows a deeper and more protracted contraction in GDP. But for the same reason, the countercyclical injection of CBDC now becomes much more effective at lowering liquidity taxes and thereby dampening the contraction in GDP. This is amplified by the strong effects of CBDC injections on bank lending and thus bank deposit creation, which is again due to the strong complementarity between CBDC and bank deposits. The observed differences in the behaviour of the CBDC interest rate are of course much larger for this case, because with a low ϵ any quantity change implies a much larger price change.

The general lesson from Figures 13 and 14 is that the policy rule coefficients that equalize the degree of countercyclicality of CBDC price and quantity rules are contingent on the elasticity of substitution between CBDC and bank deposits. The implication is that under a CBDC regime policymakers need to anticipate technological, institutional or legal changes that might affect this elasticity, because these changes can materially change the countercyclicality of a policy rule, away from what may be desired by the policymaker.

D. Discretionary Monetary Stimulus Through CBDC

The ability of monetary policy to deliver additional stimulus near the zero lower bound on nominal policy rates has been the subject of a lively policy debate. Much of that debate has concentrated on the requirements for delivering additional stimulus through a negative policy rate, including the abolition of cash as a way around the zero lower bound. Figure 16 illustrates that CBDC offers an additional quantitative tool that does not require negative policy rates, and that will in fact leave the policy rate mostly unaffected while *raising* the interest rate on CBDC.

Figure 15 simulates the economy's response to a contractionary borrower riskiness shock[64] that, absent countercyclical policy, would reduce GDP by almost 1.5% on impact, with a protracted adjustment back to the original steady state thereafter. Because this shock increases the riskiness of bank borrowers, banks immediately cut back on lending by around 8% of GDP, with a somewhat smaller contraction in deposits.[65] This rapid loss of liquidity has a highly detrimental effect on monetary transaction costs, with the main liquidity taxes rising by around 2% on impact. This loss of liquidity, together with higher lending rates, explains the 1.5% contraction of GDP, which is accompanied by a 1.5 percentage point drop in the rate of inflation.

[64] This is the same type of shock as the one underlying the crash portion of our credit cycle shocks simulations.

[65] The difference between the contractions in loans and deposits represents sizeable losses on the part of banks.

The simulations in Figure 15 assume that monetary policy follows a strict CBDC quantity rule without a *systematic* countercyclical response, $m_{\pi^p} = 0$. Instead, monetary policy is assumed to respond to the shock through a *discretionary* injection of CBDC, represented by S_t^{ms} in the CBDC quantity rule(72). The stimulus is equal to around 10% of GDP under the most aggressive policy response shown in Figure 15. Its effect is to partly offset the loss of bank-created liquidity, with a sizeable countercyclical effect on liquidity taxes and on GDP. Because this injection reduces the non-pecuniary liquidity value of existing CBDC balances, it is accompanied by an increase in the interest rate on CBDC of almost 1.5 percentage points under the most aggressive policy response shown. In other words, this stimulus takes the CBDC interest rate away from its own zero lower bound. Furthermore, because the effect of the additional liquidity is inflationary relative to the baseline of no stimulus, the policy rate is also moved away from the zero lower bound, but this effect is quite small.

Figure 16 simulates the same shock as Figure 15, but under the assumption that $\epsilon = 0.075$, with all other calibration targets kept at the values of Tables 1-4. As in our comparison of Figures 13/14 with Figures 11/12, a direct comparison of the magnitudes of the impulse responses is not appropriate, because the steady state and dynamic behaviour of the two model economies are different. What we can compare however is the differential responses to monetary stimulus. The size of the countercyclical injection of CBDC in Figure 16 is about one third of that in Figure 15, relative to GDP. Yet the stimulus effects, especially after the first two quarters, are larger. This is mainly because, with the much stronger complementarity between CBDC and bank deposits, the stimulative effect on bank lending and deposit creation is much stronger. Similar to countercyclical CBDC policy rules, countercyclical discretionary CBDC stimulus is therefore much more effective when the elasticity of substitution between CBDC and bank deposits is low. In that case it therefore represents an even more promising alternative to negative policy rates.

E. Fiscal Policy Interactions with CBDC

Figure 17 illustrates how the overall countercyclicality of policy changes, under credit cycle shocks, when the fiscal instruments that are varied countercyclically are distortionary taxes and government spending rather than, as in Figures 2-16, lump-sum taxes. Because the transmission mechanism is more transparent for CBDC quantity rules, Figure 17 is based on such rules, with $m_{\pi^p} = 0/4/8$, and can therefore be directly compared to Figure 7. The difference to Figure 7 is that in the sixth row we replace the tax revenue-to-GDP and government deficit-to-GDP ratios with the labour income tax rate and with government spending. These two fiscal instruments, together with other tax rates, are now determined by the auxiliary fiscal rules (67) - (70).

Countercyclical distortionary taxes (and countercyclical government spending) now add to the overall countercyclical policy stance, by being increased (decreased) during the initial credit-driven expansion and decreased (increased) during the crash. As a result, under the baseline of $m_{\pi^p} = 0$, the boom itself is now considerably weaker, and the crash much smaller, than in Figure 7.

But in addition, countercyclical CBDC policies, with $m_{\pi^p} > 0$, now have a smaller effect. This is because countercyclical CBDC lowers the government's interest bill during the upturn and then increases it during the downturn, so that, in order to satisfy the fiscal rule, distortionary tax rates increase by less (and government spending decreases by less) during the upturn, and vice versa during the downturn. This partly offsets the countercyclical increase in liquidity taxes during the upturn, and their decrease during the downturn. Countercyclical CBDC policies and fiscal policies, for credit cycle shocks, are therefore substitutes. The offset is however not complete, and countercyclical CBDC policy remains very effective.

The broader lesson of Figure 17 is that CBDC policies have an intrinsic fiscal dimension that cannot be ignored when studying their effectiveness. More systematic work will be required to study the extent to which countercyclical CBDC and fiscal policies interact.

VII. Conclusions

Both central banks and private financial institutions are paying increasing attention to the emergence of digital currencies and the distributed ledgers on which they are based, as this technology may present an opportunity to improve the efficiency, resiliency and accessibility of systems that facilitate monetary and financial transactions. There are, however, serious problems with existing private versions of such currencies. These problems are not associated with the viability of distributed ledgers in general, but rather with their prohibitively high costs of transaction verification. Alternative implementations, such as "permissioned" systems, may potentially avoid these costs by stepping away from purely decentralised designs while still retaining many of the benefits. One possible application of such a permissioned system would be the issuance of a central bank digital currency (CBDC) – universal, electronic, 24x7, national-currency-denominated and interest-bearing access to a central bank's balance sheet.

Any attempt to study the macroeconomic consequences of adopting a CBDC regime faces the problem that there is no historical experience to draw on, and thus also a complete absence of data for empirical work. Our approach has therefore instead relied on using a monetary-financial DSGE model, calibrated to match the US economy in the pre-crisis period, and extended to add features related to CBDC, as a laboratory that allows us to study the issues that are most relevant for policy, including the efficiency gains and stabilisation effects of CBDC. The model also clarifies the open empirical questions that need to be answered to make quantitative estimates of the effects of CBDC more reliable.

Our model should be familiar to central bankers and academics, because it starts from a canonical New Keynesian monetary model with nominal and real frictions. There are however two departures from the standard model.

The first is that it incorporates the financing through money creation (FMC) banking model of Jakab and Kumhof (2015), where monetary bank deposits are created by the extension of loans. A realistic model of banks is essential for the exercise performed in this paper, because of the key role of banks as providers of the monetary transaction medium that would compete with CBDC in the real world. A critical ingredient in the FMC framework is the existence of private-sector demand functions for monetary transactions balances, with the supply of monetary transactions balances prior to the introduction of CBDC, in the form of bank deposits, determined by commercial banks. These demand functions for monetary transaction balances can ignore central bank money other than CBDC because, during normal economic times, central bank reserves are endogenous and are, in any event, not held by the non-bank private sector, and because cash is a very small and non-constitutive part of the financial system.

The second departure from the standard model is the incorporation of the minimal structure necessary to accommodate interest-bearing CBDC into our model. This amounts to the reintroduction of central bank-issued money into the money demand function, but with the difference that this money can (unlike reserves) be held by the non-bank private sector, that it is (unlike cash) interest-bearing, and that it competes with endogenously created private bank-issued money. Many critical questions concerning CBDC are therefore related to the functional form and calibration of the joint demand function for CBDC and bank deposits.

Using this model, we find that a system of CBDC offers a number of clear macroeconomic advantages, with few obvious large costs. The advantages that we identify in this paper include large steady state output gains of almost 3% for an injection of CBDC equal to 30% of GDP, and sizeable gains in the effectiveness of systematic or discretionary countercyclical monetary policy, particularly if a sizeable share of shocks is to private credit creation or to the demand for monetary transaction balances, and if the substitutability between CBDC and bank deposits in transaction technologies is low. Our analysis suggests that the only conditions needed to secure these gains are that a sufficiently large stock of CBDC is issued in steady state, and that the issuance mechanism for CBDC ensures that the central bank only trades CBDC against government debt instruments.

In addition to these results, we also find that the theoretical and empirical gaps in our knowledge concerning CBDC have become much clearer. We are hopeful that filling these gaps will form part of a multi-pronged research agenda across central banks, covering economic theory, empirical work, and a research program on the technological aspects of distributed ledgers that are relevant to CBDC.

Important theoretical questions include the following: What are the welfare properties of alternative CBDC policy rules, including their interaction with traditional monetary policy rules, with macroprudential policy rules, and with fiscal policy rules? Should CBDC policy rules also react to financial variables, rather than simply to inflation as in this paper? What are the advantages and disadvantages of introducing CBDC into the economy through spending (on goods/services and/or transfers), lending (directly or via the banking system), or the purchase of financial assets, including not only government bonds but also other financial assets? Which of these would best safeguard financial stability? How might the issuance of CBDC interact with the unwinding of Quantitative Easing? What could be the impact of CBDC on international liquidity and exchange rate dynamics? How might the introduction of CBDC affect the likelihood of a bank run when bank deposits carry default risk, or the dynamics of a run if one were to occur?

In order to more reliably answer many of these questions, we need better answers to four empirical questions. We use the remainder of the conclusions to summarise these.

The first empirical question concerns the appropriate calibration of the main sources of demand for bank liabilities. In our model, we split the overall demand into demands related to the scale variables consumption, investment, working capital and asset holdings. The question is whether data on different holders of bank liabilities could allow us to better identify the relative sizes of these different demands.

The second empirical question concerns the interest semi-elasticity of the demand for bank deposits. Estimates for this have been produced in the past, but most often for narrower monetary aggregates rather than for the very broad aggregates that are relevant for CBDC according to our model. The model also suggests that it is important to know how this elasticity varies over different holders of bank deposits. These questions could be answered using currently available data and techniques.

The third empirical question concerns the interest semi-elasticity of the demand for CBDC relative to bank deposits, and therefore, by implication, the elasticity of substitution between CBDC and bank deposits in household and firm portfolios of monetary transaction balances. Our discussion in Section V.C alluded to how one might attempt to answer this question.

The fourth empirical question concerns the appropriate calibration of the steady state spread between the interest rate paid on CBDC and that paid on bank deposits. This spread, which currently does not exist and therefore cannot be quantified using historical data, will be influenced by many different factors, including differences in perceived risk relative to bank deposits, differences in the convenience of use, and differences in operating costs due to differences in underlying technologies.

All of the foregoing represents essential building blocks in a research agenda that puts monetary quantity aggregates, and therefore financial sector balance sheets, back at the centre of macroeconomic analysis, with obvious relevance also to financial stability issues. We are hopeful that there will be a broad collaborative effort to study these questions.

References

Agarwal, R. and Kimball, M. (2015), "Breaking Through the Zero Lower Bound", *IMF Working Paper*, WP/15/224.

Ali, R., Barrdear, J., Clews, R. and Southgate, J. (2014a), "Innovations in Payment Technologies and the Emergence of Digital Currencies", *Bank of England Quarterly Bulletin*, **54(3)**, 262-275.

Ali, R., Barrdear, J., Clews, R. and Southgate, J. (2014b), "The Economics of Digital Currencies", *Bank of England Quarterly Bulletin*, **54(3)**, 276-286.

Ashcraft, A.B. and Steindel, C. (2008), "Measuring the Impact of Securitization on Imputed Bank Output", Working Paper, Federal Reserve Bank of New York.

Autonomous Research (2016), "Blockchain: Back-Office Block-Buster", available at: http://www.autonomous.com/fintech.html.

Ball, L. (2001), "Another Look at Long-Run Money Demand", *Journal of Monetary Economics*, **47**, 31-44.

Bank of England (2015), "One Bank Research Agenda", *Bank of England Discussion Paper*, available at: http://www.bankofengland.co.uk/research/Documents/onebank/discussion.pdf.

Bates, T., Kahle, K. and Stulz, R. (2008), "Why Do US Firms Hold So Much More Cash Than They Used To?", Working Paper, Ohio State University.

Benes, J. and Kumhof, M. (2012), "The Chicago Plan Revisited", IMF Working Paper, WP/12/202.

Bernanke, B., Gertler, M. and Gilchrist, S. (1999), "The Financial Accelerator in a Quantitative Business Cycle Framework", in: Taylor, J.B., Woodford, M. (eds.), Handbook of Macroeconomics, Volume 1C. Elsevier, Amsterdam, pp. 1341-1393.

Böhme, R., Christin, N., Edelman, B. and Moore, T. (2015), "Bitcoin: Economics, Technology, and Governance", *Journal of Economic Perspectives*, **29(2)**, 213-238.

Brock, W.A. (1975), "A Simple Perfect Foresight Monetary Model", *Journal of Monetary Economics*, **1**, 133-150.

Calvo, G.A. (1983), "Staggered Prices in a Utility-Maximizing Framework", *Journal of Monetary Economics*, **12**, 383-398.

Carabenciov, I., Freedman, C., Garcia-Saltos, R., Laxton, D., Kamenik, O. and Manchev, P. (2013), "GPM6 - The Global Projection Model with 6 Regions", IMF Working Papers, WP/13/87.

Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, **113(1)**, 1–45.

Christiano, L., Motto, R. and Rostagno, M. (2014), "Risk Shocks", *American Economic Review*, **104(1)**, 27-65.

Competition and Markets Authority (2015), "Retail Banking Market Investigation: Provisional Findings Report", pp. 191-193, available at: https://www.gov.uk/cma-cases/review-of-banking-for-small-and-medium-sized-businesses-smes-in-the-uk.

Danezis, G. and Meiklejohn, S. (2016), "Centrally Banked Cryptocurrencies", Working Paper, University College London, available at: http://www0.cs.ucl.ac.uk/staff/G.Danezis/papers/ndss16currencies.pdf.

Deetman, S. (2016), "Bitcoin Could Consume as Much Electricity as Denmark by 2020", Opinion Article, Leiden University, available at: http://motherboard.vice.com/read/bitcoin-could-consume-as-much-electricity-as-denmark-by-2020.

Diamond, D. and Dybvig, P. (1983), "Bank Runs, Deposit Insurance, and Liquidity", *Journal of Political Economy*, **91(3)**, 401-419.

Engen, E.M. and Hubbard, R.G. (2004), "Federal Government Debt and Interest Rates", *NBER Macroeconomics Annual*, **19**, 83-138.

Erceg, C., Guerrieri, L. and Gust, C. (2006), "SIGMA: A New Open Economy Model for Policy Analysis", Board of Governors of the Federal Reserve System, International Finance Discussion Papers, No. 835 (revised version, January 2006).

Federal Accounting Standards Advisory Board (2012), *Federal Accounting Standards Advisory Board Handbook*, Statements of Federal Financial Accounting Standards 7, pp. 106-107.

Federal Financial Institutions Examination Council (2007), Annual Report 2006, Washington, DC.

Fish, T. and Whymark, R. (2015), "How Has Cash Usage Evolved in Recent Decades? What Might Drive Demand in the Future", *Bank of England Quarterly Bulletin*, **55(3)**, 216-227.

Farrell, J. (1987), "Cheap Talk, Coordination, and Entry", *RAND Journal of Economics*, **18(1)**, 34-39.

Fernández-Villaverde, J. and Sanches, D. (2016), "Can Currency Competition Work?", *NBER Working Paper*, No. 22157.

Friedman, M. (1969), *The Optimum Quantity of Money*, London: Macmillan.

Furfine, C. (2001), "Bank Portfolio Allocation: The Impact of Capital Requirements, Regulatory Monitoring, and Economic Conditions", *Journal of Financial Services Research*, **20(1)**, 33-56.

Gale, W. and Orszag, P. (2004), “Budget Deficits, National Saving, and Interest Rates”, *Brookings Papers on Economic Activity*, **2**, 101-187.

Gertler, M. and Karadi, P. (2011), “A Model of Unconventional Monetary Policy”, *Journal of Monetary Economics*, **58(1)**, 17-34.

Gesell, S. (1916), *Die Natürliche Wirtschaftsordnung*, available in English as *The Natural Economic Order*, 1958, Peter Owen Ltd, London.

Girouard, N. and André, C. (2005), “Measuring Cyclically-Adjusted Budget Balances for OECD Countries”, OECD Economics Department Working Papers, No. 434.

Goodfriend, M. and McCallum, B. (2007), “Banking and Interest Rates in Monetary Policy Analysis: A Quantitative Exploration”, *Journal of Monetary Economics*, **54**, 1480-1507.

Goodhart, C.A.E. (1975), “Problems of Monetary Management: The U.K. Experience”, Papers in Monetary Economics, Reserve Bank of Australia.

Gorton, G. (2010), “Questions and Answers about the Financial Crisis”, Working Paper, Yale School of Management.

Gorton, G., Lewellen, S. and Metrick, A. (2012), “The Safe-Asset Share”, NBER Working Paper No. 17777.

Heaton, J. and Lucas, D. (1996), “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing”, *Journal of Political Economy*, **104(3)**, 443-487.

Hördahl, P. and King, M. (2008), “Developments in Repo Markets During the Financial Turmoil”, *BIS Quarterly Review*, 37-53.

Independent Commission on Banking (2011), *Final Report*, London: Domarn Group, available at: http://bankingcommission.independent.gov.uk.

Ireland, P. (2001), “Sticky-Price Models of the Business Cycle: Specification and Stability”, *Journal of Monetary Economics*, **47**, 3-18.

Ireland, P. (2007), “On the Welfare Cost of Inflation and the Recent Behaviour of Money Demand”, Working Paper, Boston College.

Jakab, Z. and Kumhof, M. (2015), “Banks Are Not Intermediaries of Loanable Funds - And Why This Matters”, Bank of England Staff Working Papers, No. 529.

Lagos, R. and Wright, R. (2003), “Dynamics, Cycles, and Sunspot Equilibria in ‘Genuinely Dynamic, Fundamentally Disaggregative’ Models of Money”, *Journal of Economic Theory*, **109(2)**, 423-453.

Lamport, L., Shostak, R. and Pease, M. (1982), “The Byzantine Generals Problem”, *ACM Transactions on Programming Languages and Systems*, **4(3)**, 382-401.

Laubach, T. (2009), “New Evidence on the Interest Rate Effects of Budget Deficits and Debt”, *Journal of the European Economic Association*, **7(4)**, 858-885.

Lucas, R. (1976), “Econometric Policy Evaluation: A Critique”, Carnegie-Rochester Conference Series on Public Policy, **1**, 19–46.

McLeay, M., Radia, A. and Thomas, R. (2014a), “Money in the Modern Economy: An Introduction”, *Bank of England Quarterly Bulletin*, **54(1)**, 4-13.

McLeay, M., Radia, A. and Thomas, R. (2014b), “Money Creation in the Modern Economy”, *Bank of England Quarterly Bulletin*, **54(1)**, 14-27.

Nakamoto, S.(2008), “Bitcoin: A Peer-to-Peer Electronic Cash System”, Working Paper, available at: http://bitcoin.org/bitcoin.pdf.

Neumeyer, P. and Perri, F. (2005), "Business Cycles in Emerging Economies: The Role of Interest Rates", *Journal of Monetary Economics*, **52**, 345-380.

O'Brien (2000), "Estimating the Value and Interest Rate Risk of Interest-Bearing Transactions Deposits", Working Paper, Board of Governors of the Federal Reserve System.

O'Dwyer, K. and Malone, D. (2014), "Bitcoin Mining and Its Energy Footprint", *Proceedings of Irish Signals and Systems Conference 2014*, 280-285.

Poole, William (1970), "Optimal Choice of the Monetary Policy Instrument in a Simple Stochastic Macro Model", *Quarterly Journal of Economics*, **84**, 197-216.

Pozsar, Z., Adrian, T., Ashcraft, A. and Boesky, H. (2010), "Shadow Banking", Federal Reserve Bank of New York Staff Reports, No. 458.

Rogoff, K. (2015), "Costs and Benefits to Phasing Out Paper Currency", *NBER Macroeconomic Annual 2014*, **29**, 445-456.

Schmitt-Grohé, S. and Uribe, M. (2003), "Closing Small Open Economy Models", *Journal of International Economics*, **61**, 163-185.

Schmitt-Grohé, S. and Uribe, M. (2004), "Optimal Fiscal and Monetary Policy under Sticky Prices", *Journal of Economic Theory*, **114**, 198-230.

Sidrauski, M. (1967), "Rational Choice and Patterns of Growth in a Monetary Economy", *American Economic Review Papers and Proceedings*, **47**, 534-544.

Ueda, K. and Brooks, R. (2011), "User Manual for the Corporate Vulnerability Utility: The 4th Edition", International Monetary Fund.

Tobin, J. (1987), "The Case for Preserving Regulatory Distinctions", in *Restructuring the Financial System*, Federal Reserve Bank of Kansas City, 167-183.

Visa (2015), "Visa Inc. at a Glance", available at: https://usa.visa.com/dam/VCOM/download/corporate/media/visa-fact-sheet-Jun2015.pdf.

Van den Heuvel, S. (2005), "The Bank Capital Channel of Monetary Policy", Working Paper, Wharton School, University of Pennsylvania.

White, M. (1998), "Why Don't More Households File for Bankruptcy?", *Journal of Law, Economics and Organization*, **14(2)**, 205-231.

Table 1. Directly Calibrated Parameters

Description	**Calibration Target**	**Parameter**	**Parameter Value**
Real Growth Rate	2% p.a.	x	1.005
Inflation Rate	3% p.a.	$\bar{\pi}^p$	1.00753
FI Population Share	5%	ω	0.05
TA Cost Resource Cost Share	25%	$\mathfrak{r}$	0.25
Consumption Habit		v	0.7
Labor Supply Elasticity		η	1
Investment Adjustment Cost		ϕ_I	2.5
Price Adjustment Cost		ϕ_p	200
Wage Adjustment Cost		ϕ_w	200
Steady State Price Mark-up	10%	μ_p	1.1
Steady State Wage Mark-up	10%	μ_w	1.1
Policy Rate Smoothing		m_i	0.7
Policy Rate Inflation Feedback		m_π	2.0
Fiscal Output Gap Feedback		d^{gdp}	0.34 (0.5 in Fig.1)
Basel MCAR	8%	Υ	0.08
Risk Weight, Consumer Loans	75%	ζ^c	0.75
Risk Weight, Mortgage Loans	50%	ζ^a	0.50
Risk Weight, WCAP Loans	100%	ζ^y	1.00
Risk Weight, Investment Loans	90%	ζ^k	0.90
Financial Collateral Coefficients	100%	κ^{cf}, κ^{af}, κ^{yf}, κ^{kf}	1.0
Loan Adjustment Costs		φ_c, φ_a, φ_y, φ_k	0.00001

Table 2. Calibrated Steady State Moments and Implied Parameters: Real Variables

Description	**Calibration Target**	**Parameter**	**Pre-CBDC Value**	**Post-CBDC Value**
	National Accounts Ratios			
Labor Income Share	61%	α	0.330	0.329
Investment/GDP	19%	Δ	0.0270	0.0271
Government Spending/GDP	18%	s_g	0.18	0.18
	Fiscal Accounts Ratios			
Government Financing/GDP	80%	$\overline{gd^{rat}}$ or $\overline{gdx^{rat}}$	3.9712	3.9712
Labor Income Taxes/GDP	17.6%	$\bar{\tau}_L$	0.3174	0.3174
Capital Income Taxes/GDP	3.2%	$\bar{\tau}_k$	0.2492	0.2510
Consumption Taxes/GDP	4.6%	$\bar{\tau}_c$	0.0730	0.0730
Lump-Sum Taxes/All Taxes	33.3%	$\overline{trf}/\overline{gdp}$	0.1928	0.1982
	Elasticity of Real Risk-Free Rate w.r.t. Gov.Debt/GDP			
Elasticity Ξ	2 bp per pp	ϕ_b	0.00005	0.00005
	Normalisations			
FI/HH Consumption	1:1	ι	0.9629	0.9466
FI Labor Supply	1	ψ_h	0.5798	0.5836
Price of Land	1	ψ_a	0.0491	0.0470

Table 3. Calibrated Steady State Moments and Implied Parameters: Financial Variables

Description	**Calibration Target**	**Parameter**	**Pre-CBDC Value**	**Post-CBDC Value**
	Bank Balance Sheets			
Capital Conservation Buffer	2.5%	δ^b	0.0573	0.0573
	Borrower Balance Sheets			
Investment Loans Leverage	100%	ψ_f	0.0114	0.0109
Mortgage Loans Leverage	200%	a	8.754	8.677
Consumer Loans/GDP	20%	κ^{cr}	0.8786	0.7289
Mortgage Loans/GDP	60%	κ^{ar}	1.1687	1.1724
WCAP Loans/GDP	20%	κ^{yr}	0.3490	0.3398
Investment Loans/GDP	80%	κ^{kr}	0.5241	0.5104
Consumer Deposits/GDP	50%	A_c	0.2425	0.7183
Real Estate Deposits/GDP	10%	A_a	0.0003	0.0004
Producer Deposits/GDP	10%	A_y	0.0052	0.0148
Investment Deposits/GDP	30%	A_k	1.0264	1.8652
	Real Interest Rates			
Real Policy Interest Rate	3.00% p.a.	β_u	0.9975	0.9960
Real Deposit Interest Rate	2.00% p.a.	β_c	0.9869	0.9873
Wholesale Spread over Deposit Rate	1.50% p.a.	χ	0.0117	0.0117
Retail Spread, Consumer Loans	5.00% p.a.	σ^c	0.6327	0.6396
Retail Spread, Mortgage Loans	2.00% p.a.	σ^a	0.3116	0.3159
Retail Spread, WCAP Loans	3.00% p.a.	σ^y	0.3112	0.3152
Retail Spread, Investment Loans	1.50% p.a.	σ^k	0.0996	0.1015
	Interest Semi-Elasticities of Deposit Demand			
ε_u^d	250	ϑ	2.7182	2.4994
ε_c^d	5	B_c	0.0161	0.0253
ε_a^d	5	B_a	0.0187	0.0206
ε_y^d	5	B_y	0.0145	0.0256
ε_k^d	5	B_k	0.0176	0.0224
	Failure Rates			
MCAR Violation Rate, Banks	2.5% p.q.	σ^b	0.0129	0.0129
Bankruptcy Rate, Consumer Loans	4.0% p.q.	ξ^c	0.0585	0.0563
Bankruptcy Rate, Mortgage Loans	2.5% p.q.	ξ^a	0.0486	0.0472
Bankruptcy Rate, WCAP Loans	3.0% p.q.	ξ^y	0.0322	0.0308
Bankruptcy Rate, Investment Loans	1.5% p.q.	ξ^k	0.0199	0.0192

Table 4. Calibrated Steady State Moments and Implied Parameters: CBDC Variables

Description	**Calibration Target**	**Parameter**	**Parameter Value**
CBDC/GDP	30%	$\overline{m^{rat}}$	30
CBDC Quantity Rule, Inflation Feedback	various	m_{π^p}	various
CBDC Price Rule, Inflation Feedback	various	$i^m_{\pi^p}$	various
Policy Rate minus CBDC Rate	0.80% p.a.	T^{fintec}	1.153
LGF Exponent, Separable LGF		θ	0.95
LGF EoS, CES LGF		ϵ	2.0
Consumer CBDC/GDP	19.31%	γ_c	0.3244
Real Estate CBDC/GDP	0.66%	γ_a	0.0894
WCAP CBDC/GDP	4.00%	γ_y	0.3314
Investment CBDC/GDP	6.03%	γ_k	0.2104

Table 5. Steady State Output Gains of Transition to CBDC

	Distortionary Taxes	Lump-Sum Taxes
$\Xi = 0.02$	2.94%	1.43%
$\Xi = 0$	0.72%	-0.04%

Figure 1. Transition to New Steady State with CBDC at 30 Percent of GDP

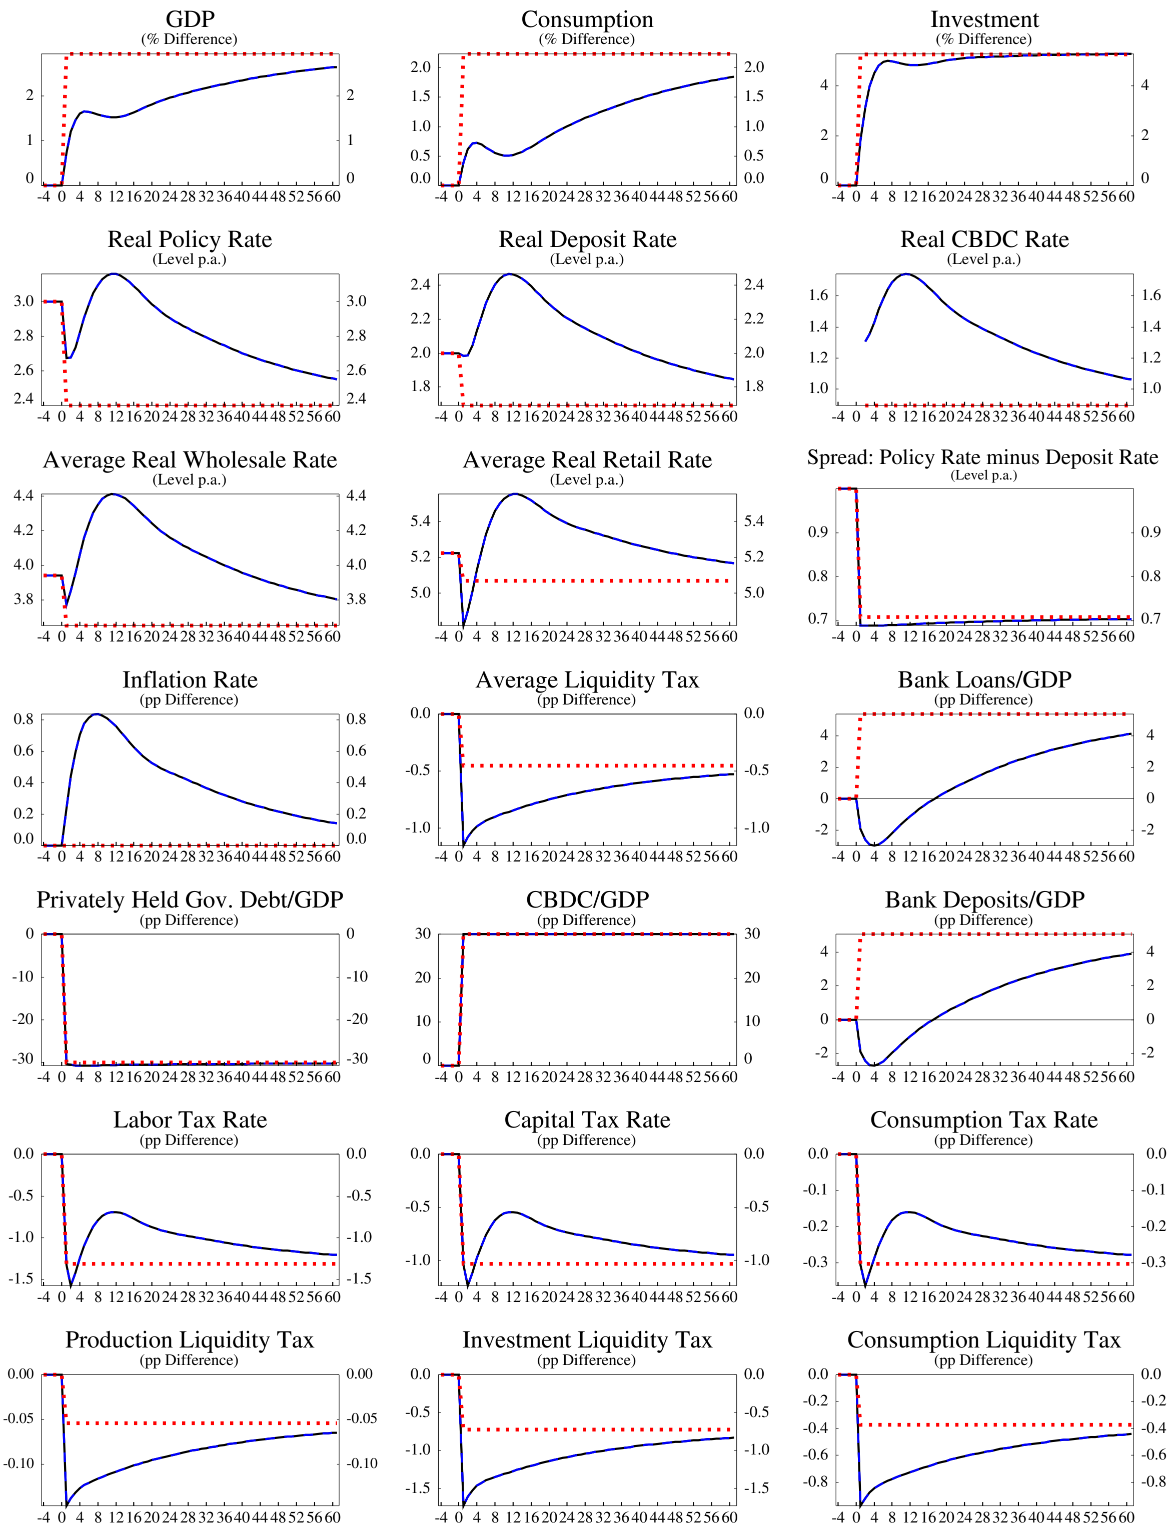

Solid Line = Actual Transition, Dotted Line = Change in Long-Run Steady State

Figure 2. Quantity versus Price Rules for CBDC - Credit Cyle Shocks

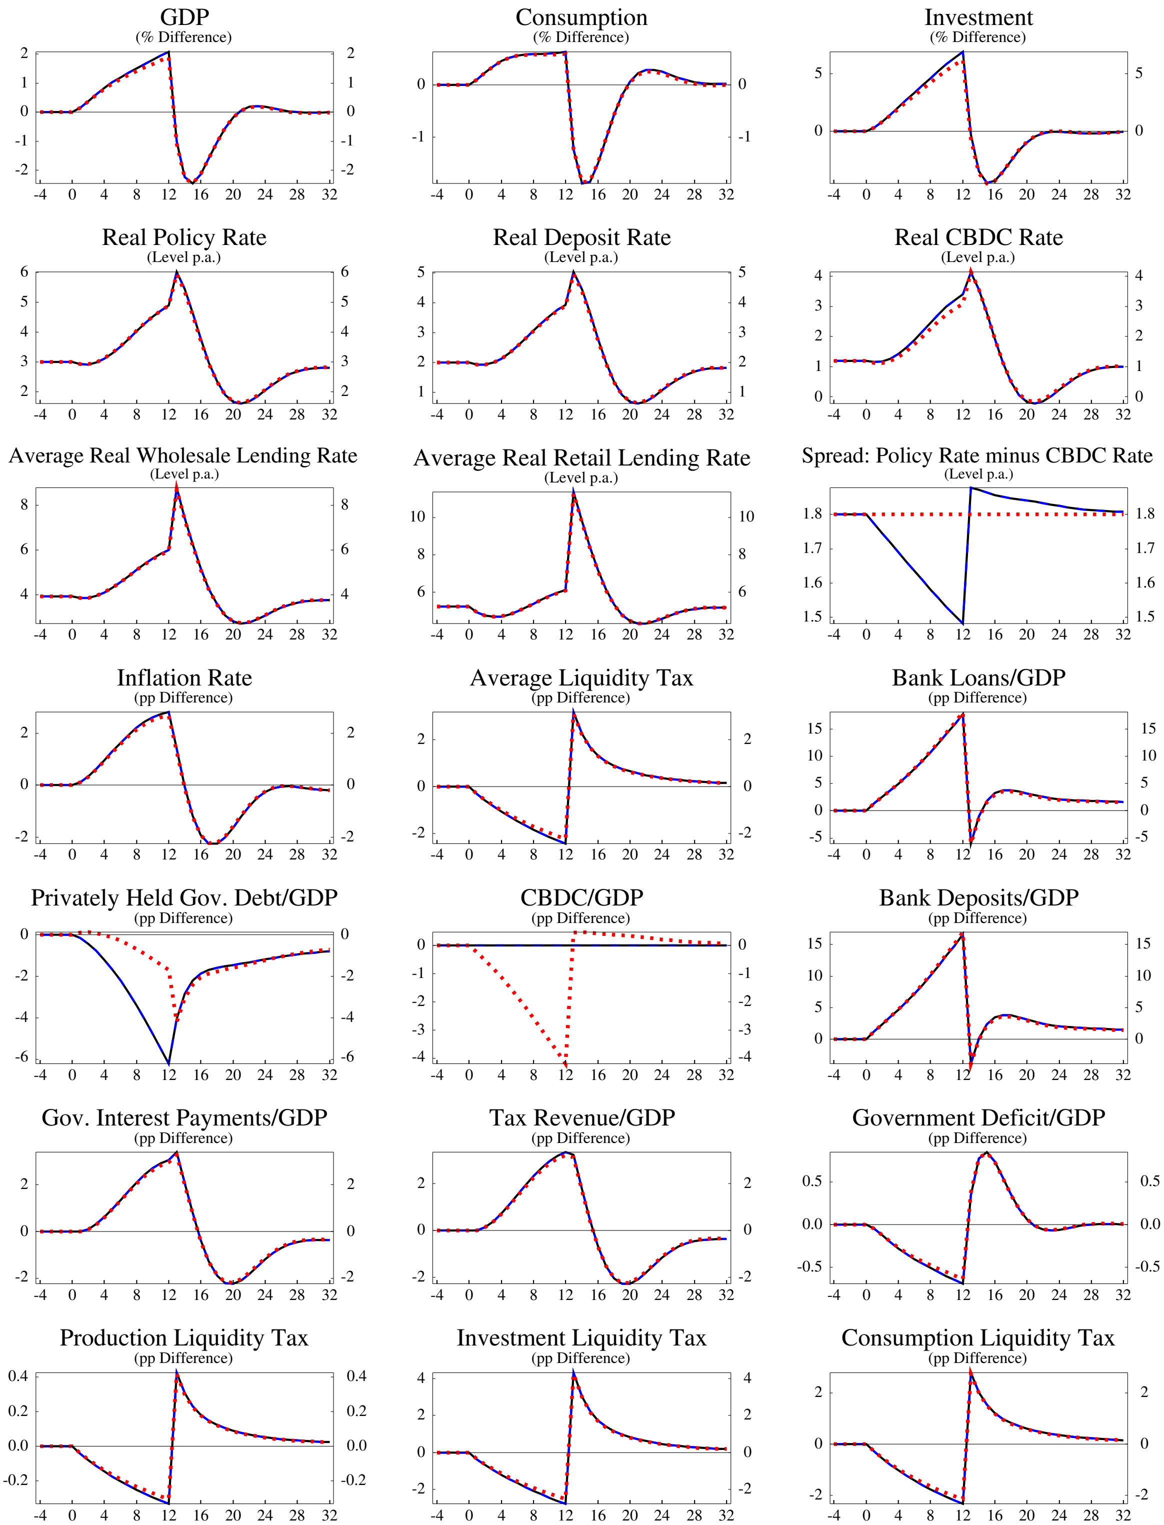

Solid Line = Quantity Rule (fixed m^{rat}), Dotted Line = Price Rule (fixed spread $i - i_m$)

Figure 3. Quantity versus Price Rules for CBDC - Higher Demand for Total Liquidity

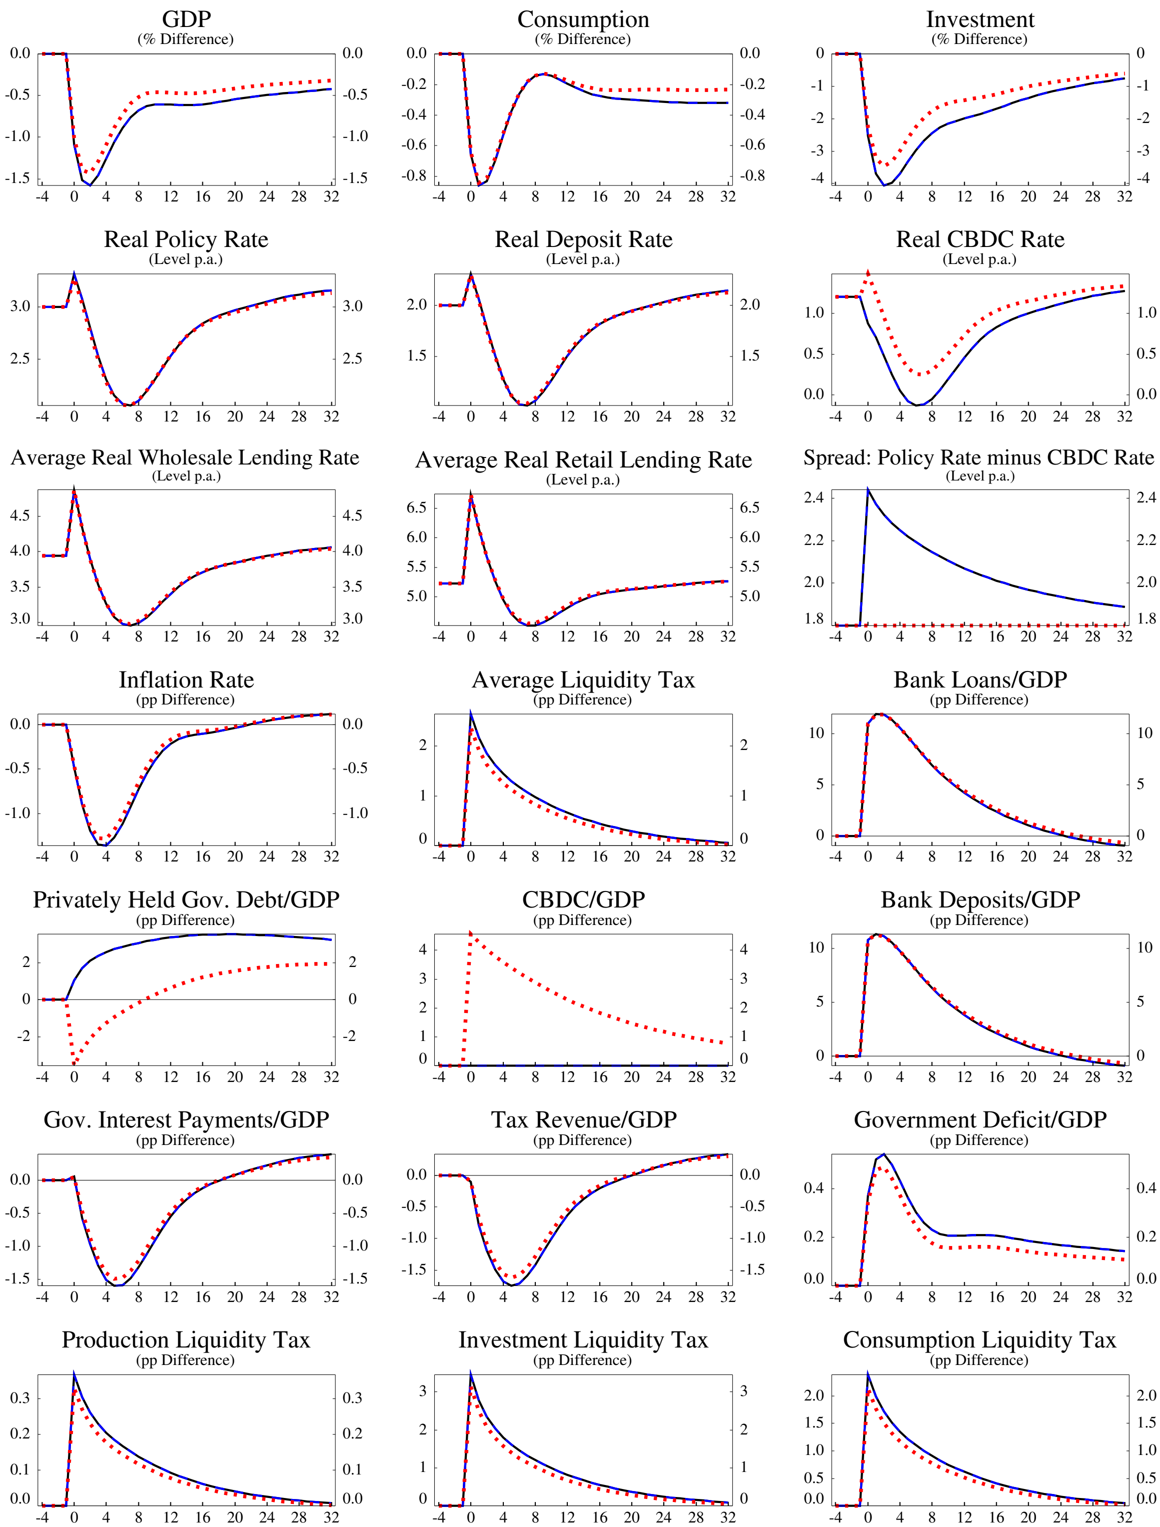

Solid Line = Quantity Rule (fixed m^{rat}), Dotted Line = Price Rule (fixed spread $i - i_m$)

Figure 4. Quantity versus Price Rules for CBDC - Higher Demand for CBDC Liquidity

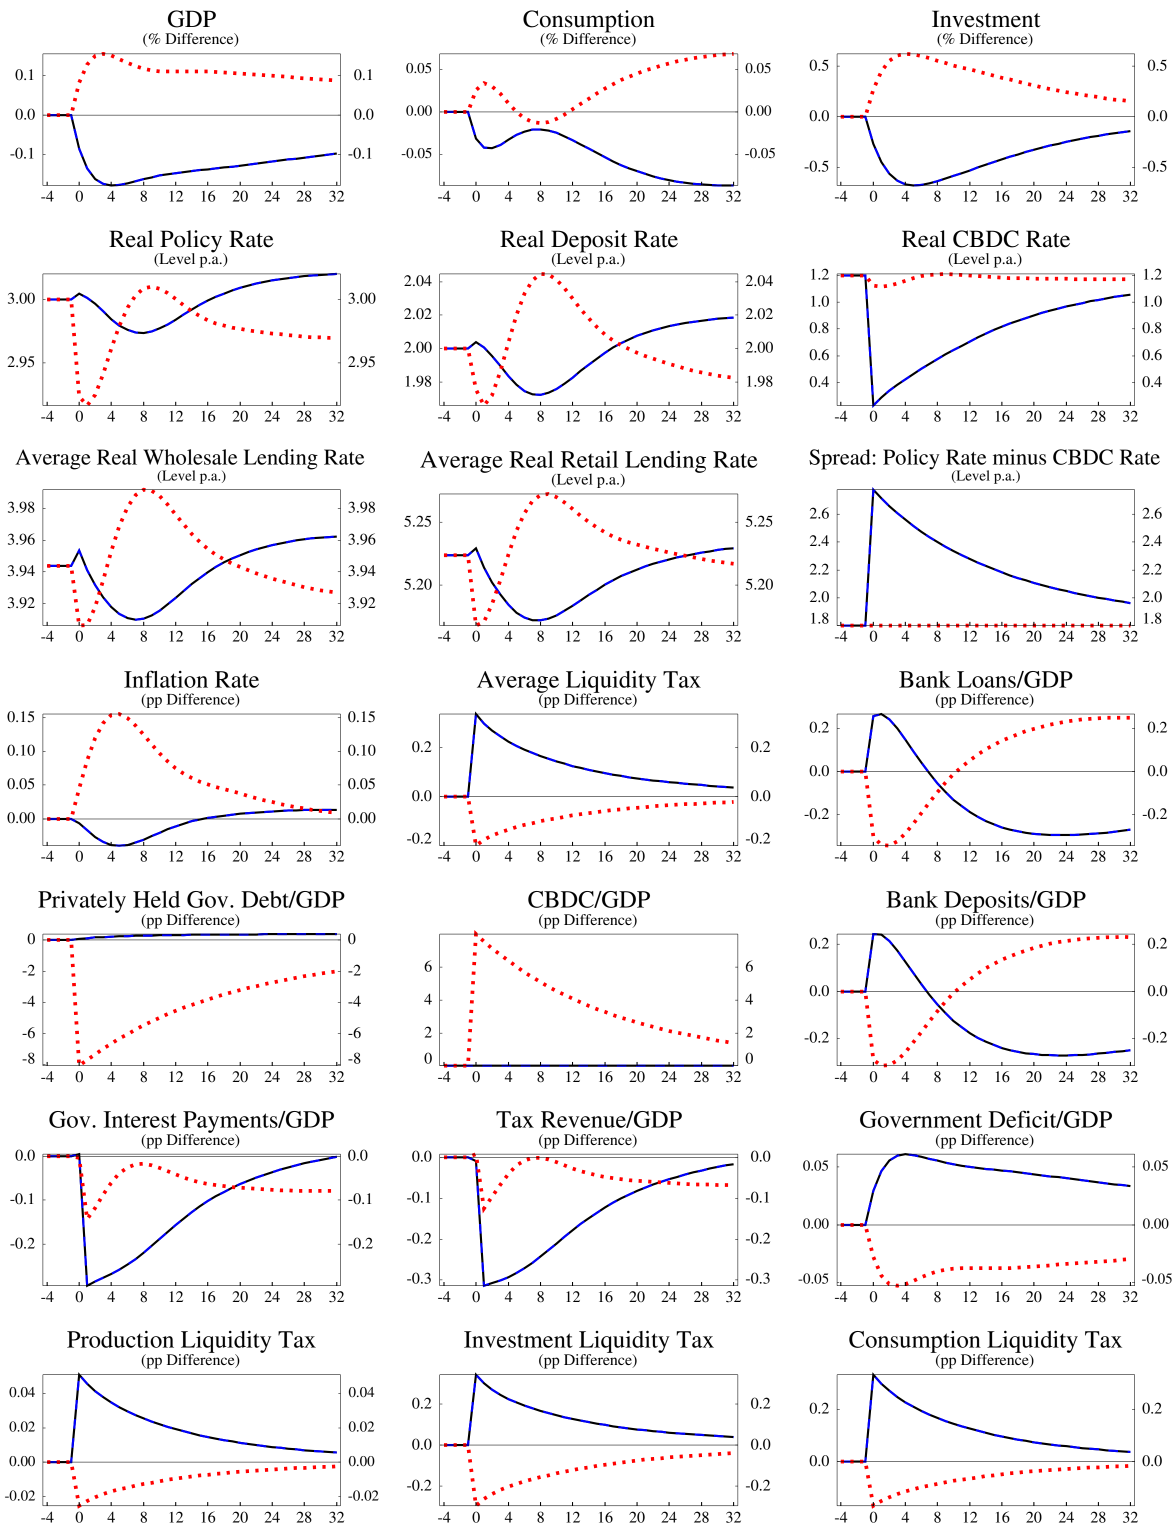

Solid Line = Quantity Rule (fixed m^{rat}), Dotted Line = Price Rule (fixed spread $i - i_m$)

Figure 5. Low EoS versus High EoS - Quantity Rule for CBDC - Credit Cycle Shocks

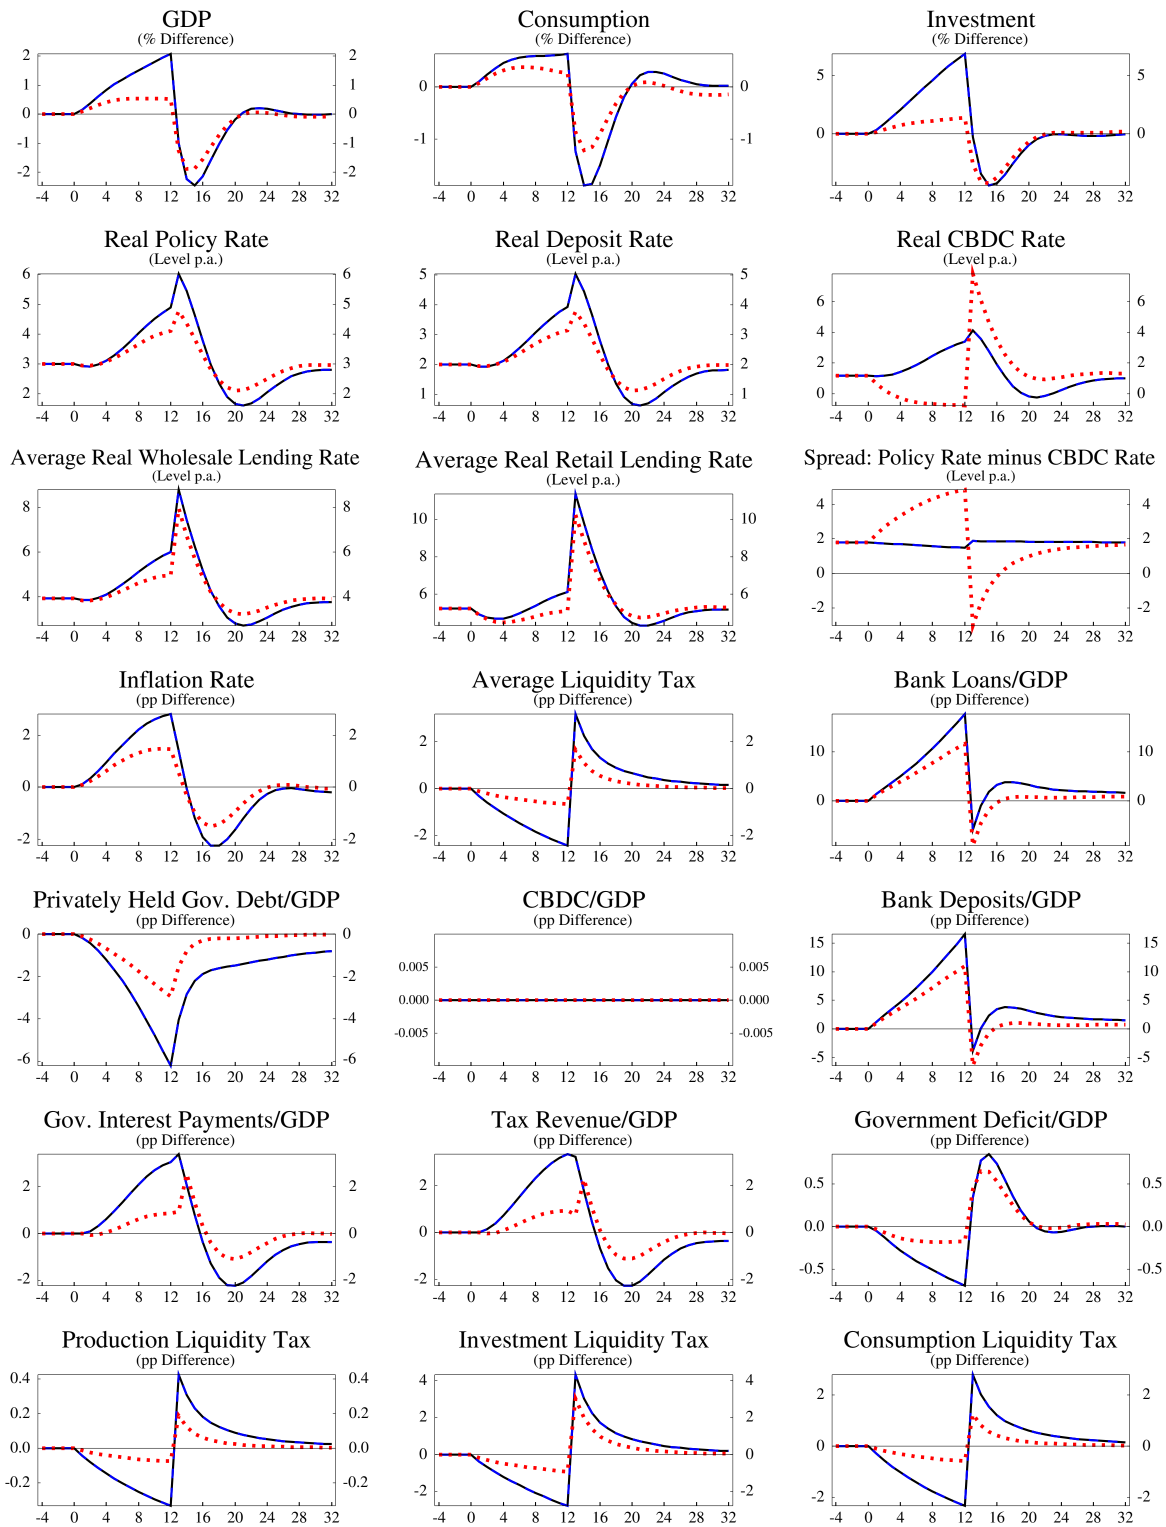

Solid Line = Baseline EoS (2), Dotted Line = Low EoS (0.075)

Figure 6. Low EoS - Quantity versus Price Rules - Higher Demand for Total Liquidity

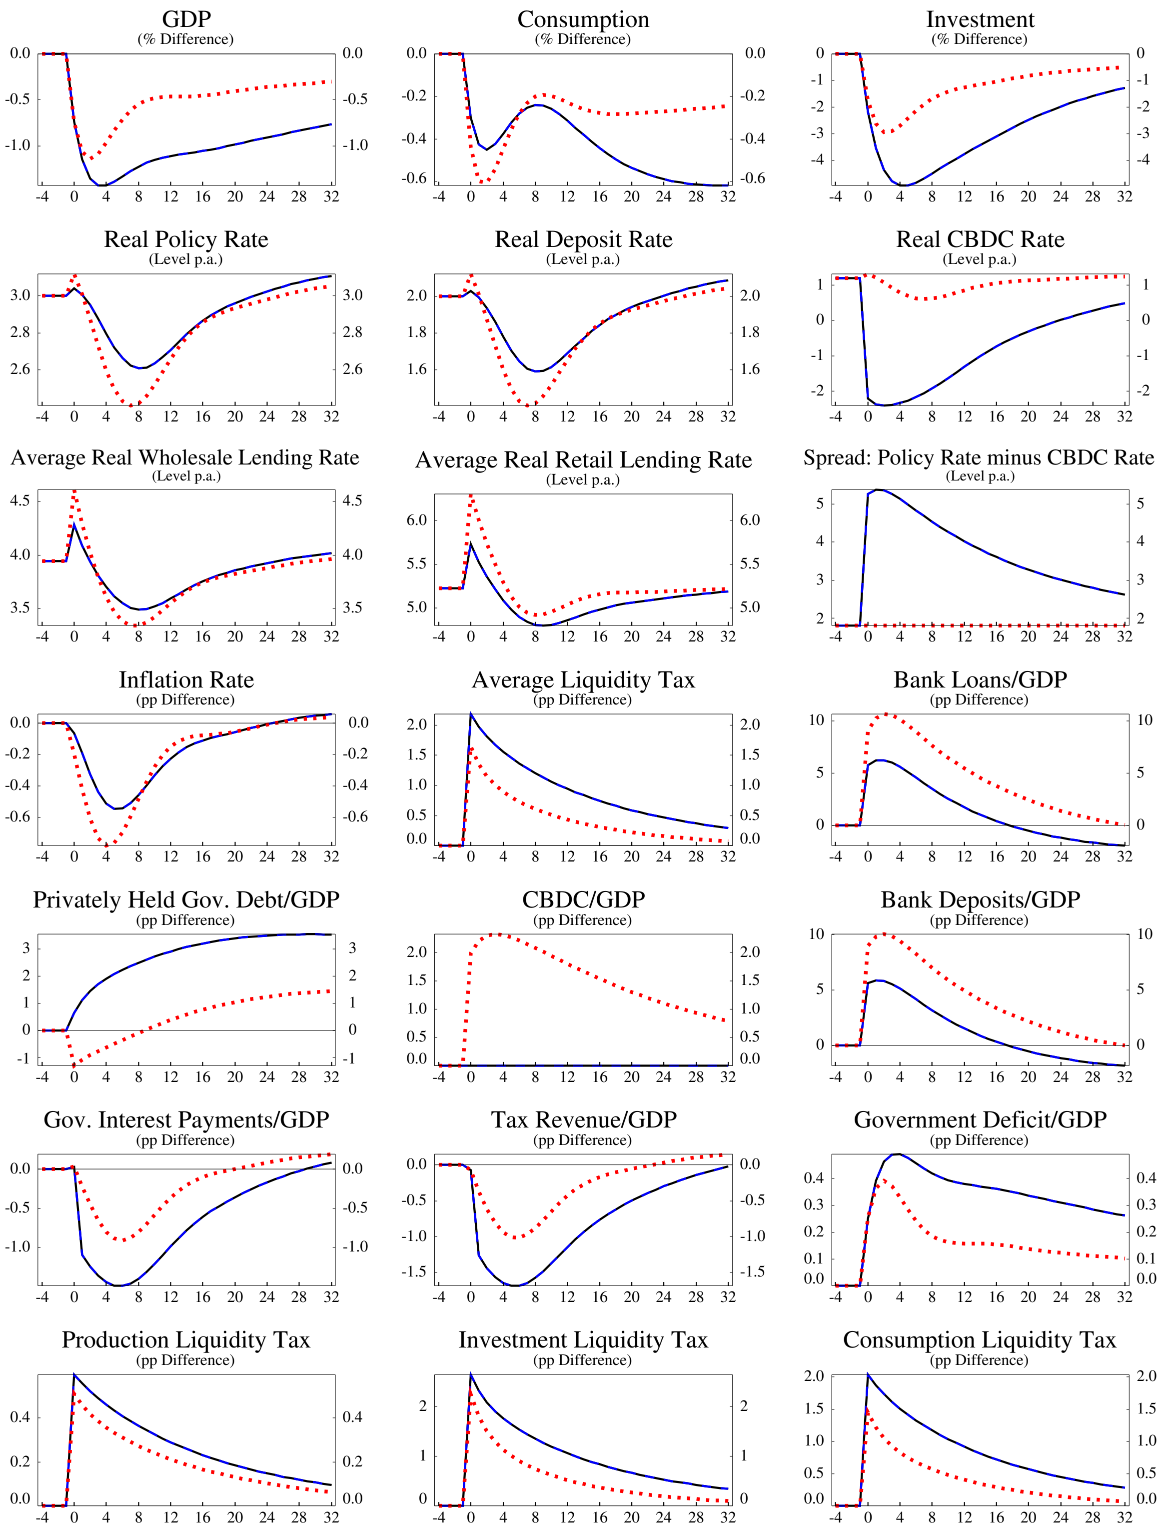

Solid Line = Quantity Rule (fixed m^{rat}), Dotted Line = Price Rule (fixed spread $i - i_m$)

Figure 7. Countercyclical CBDC Quantity Rules - Credit Cycle Shocks

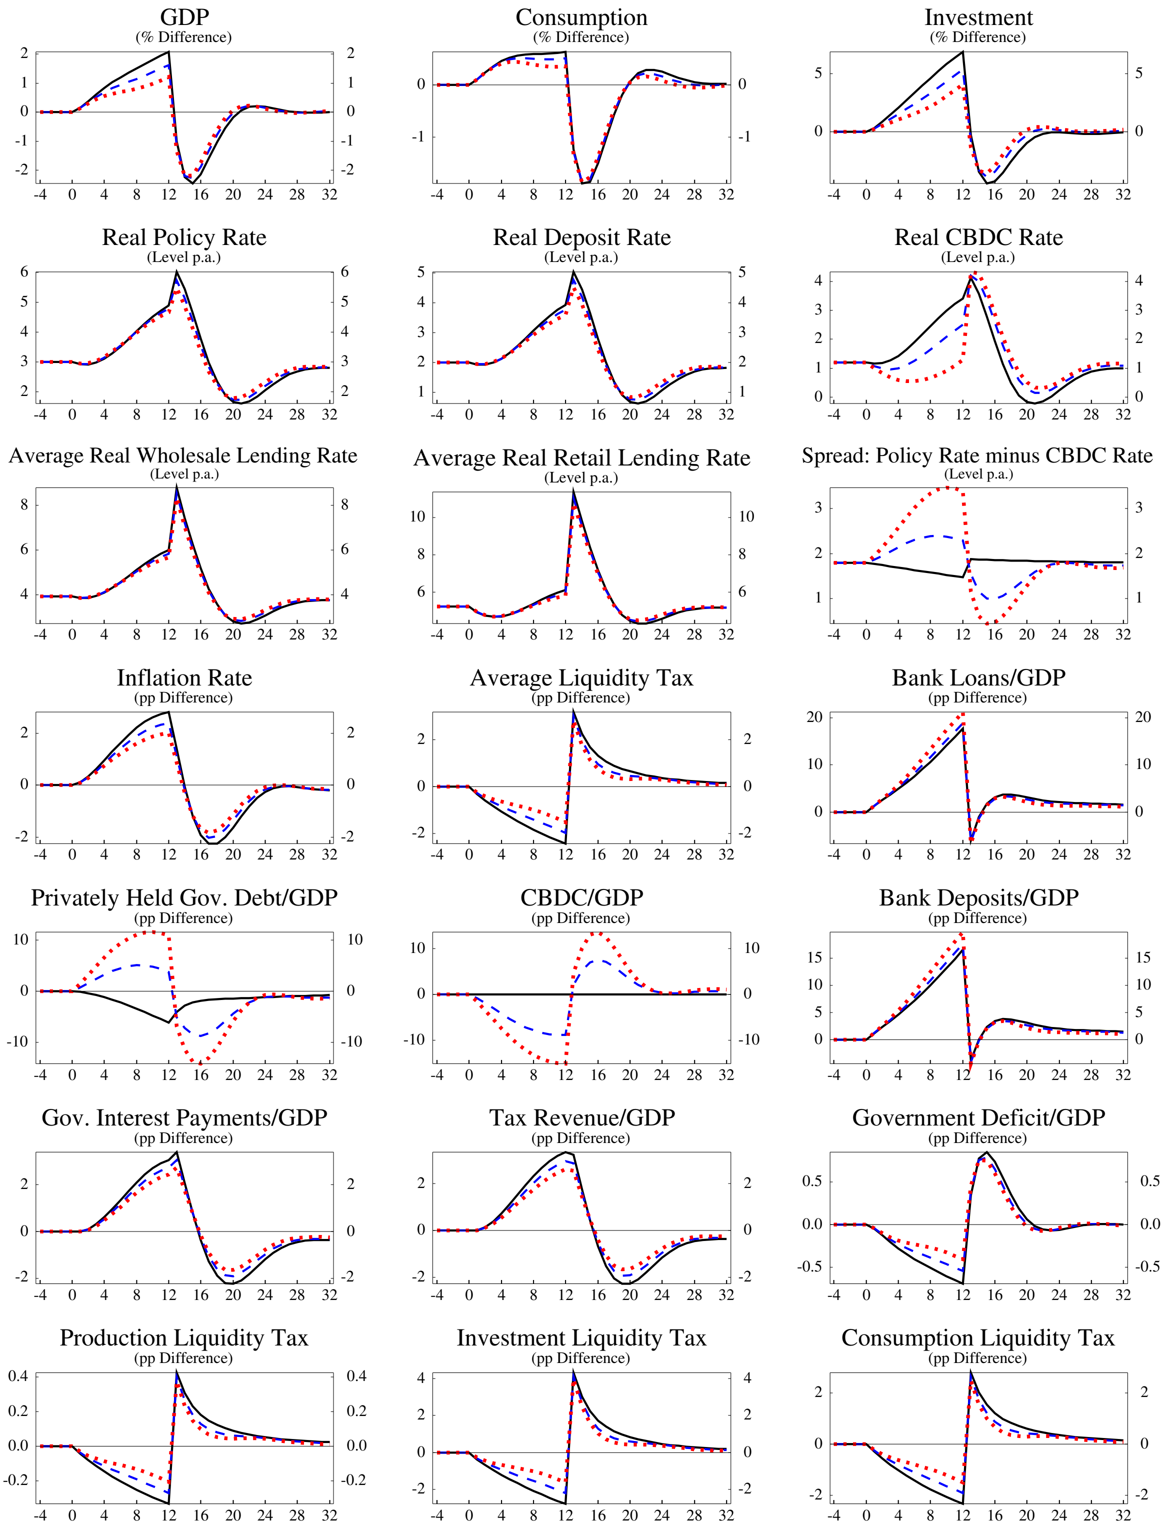

Solid Line = Baseline ($m_{\pi^p} = 0$), Dashed Line = Intermediate ($m_{\pi^p} = 4$), Dotted Line = Aggressive ($m_{\pi^p} = 8$)

Figure 8. Countercyclical CBDC Price Rules - Credit Cycle Shocks

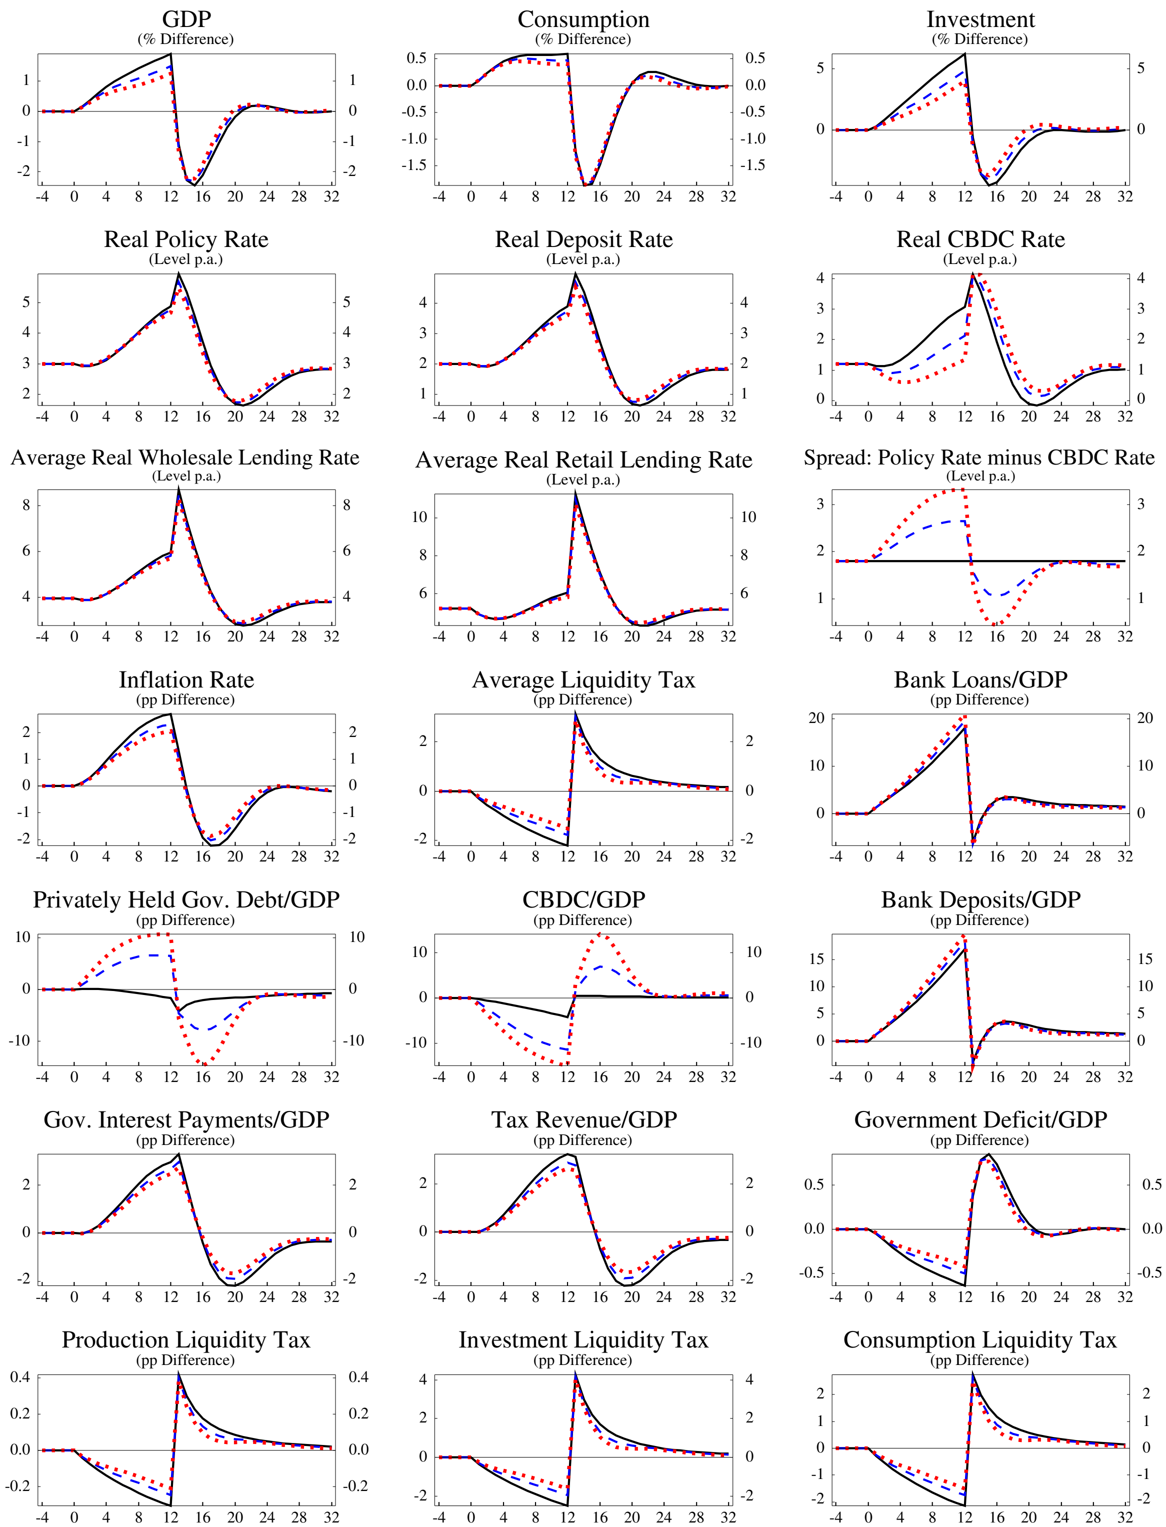

Solid Line = Baseline ($i^m_{\pi^p} = 0$), Dashed Line = Intermediate ($i^m_{\pi^p} = 0.4$), Dashed Line = Aggressive ($i^m_{\pi^p} = 0.8$)

Figure 9. Countercyclical CBDC Price Rules - Credit Cycle Shocks - Policy Rate Corridor

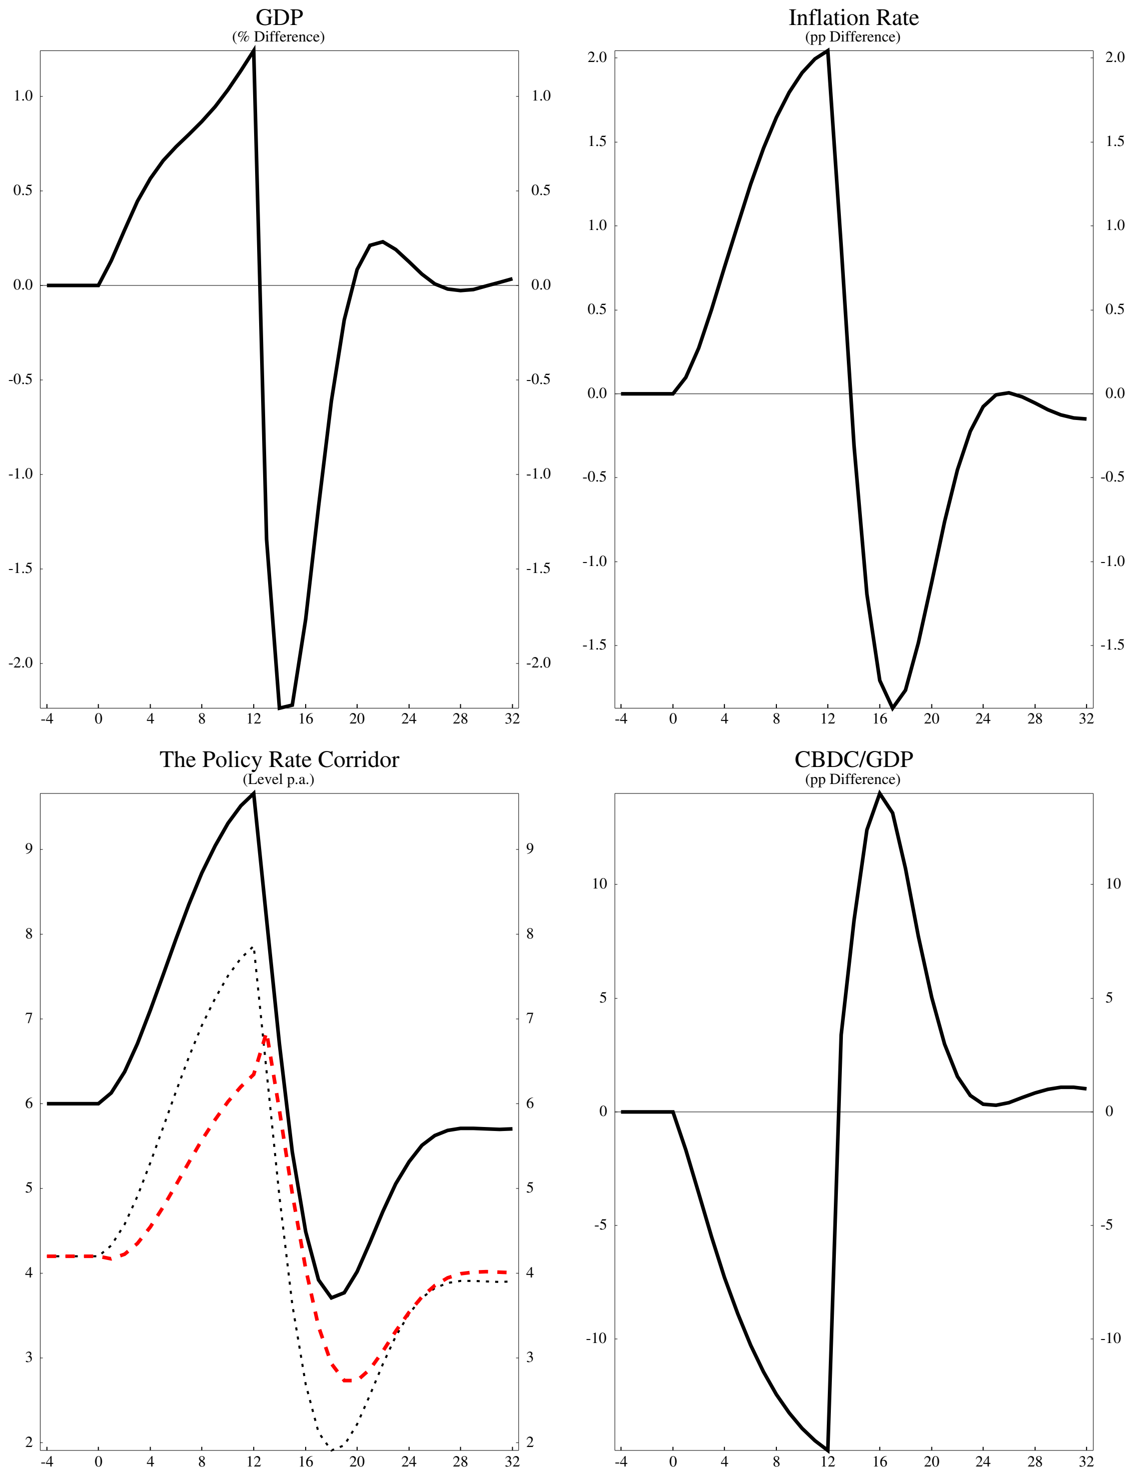

Bottom Left: Nominal Policy and CBDC Rates
Solid Line = Policy Rate, Dotted Line = Policy Rate minus Fixed Spread, Dashed Line = CBDC Rate

Figure 10. Countercyclical CBDC Price Rules - Lower Investment Demand

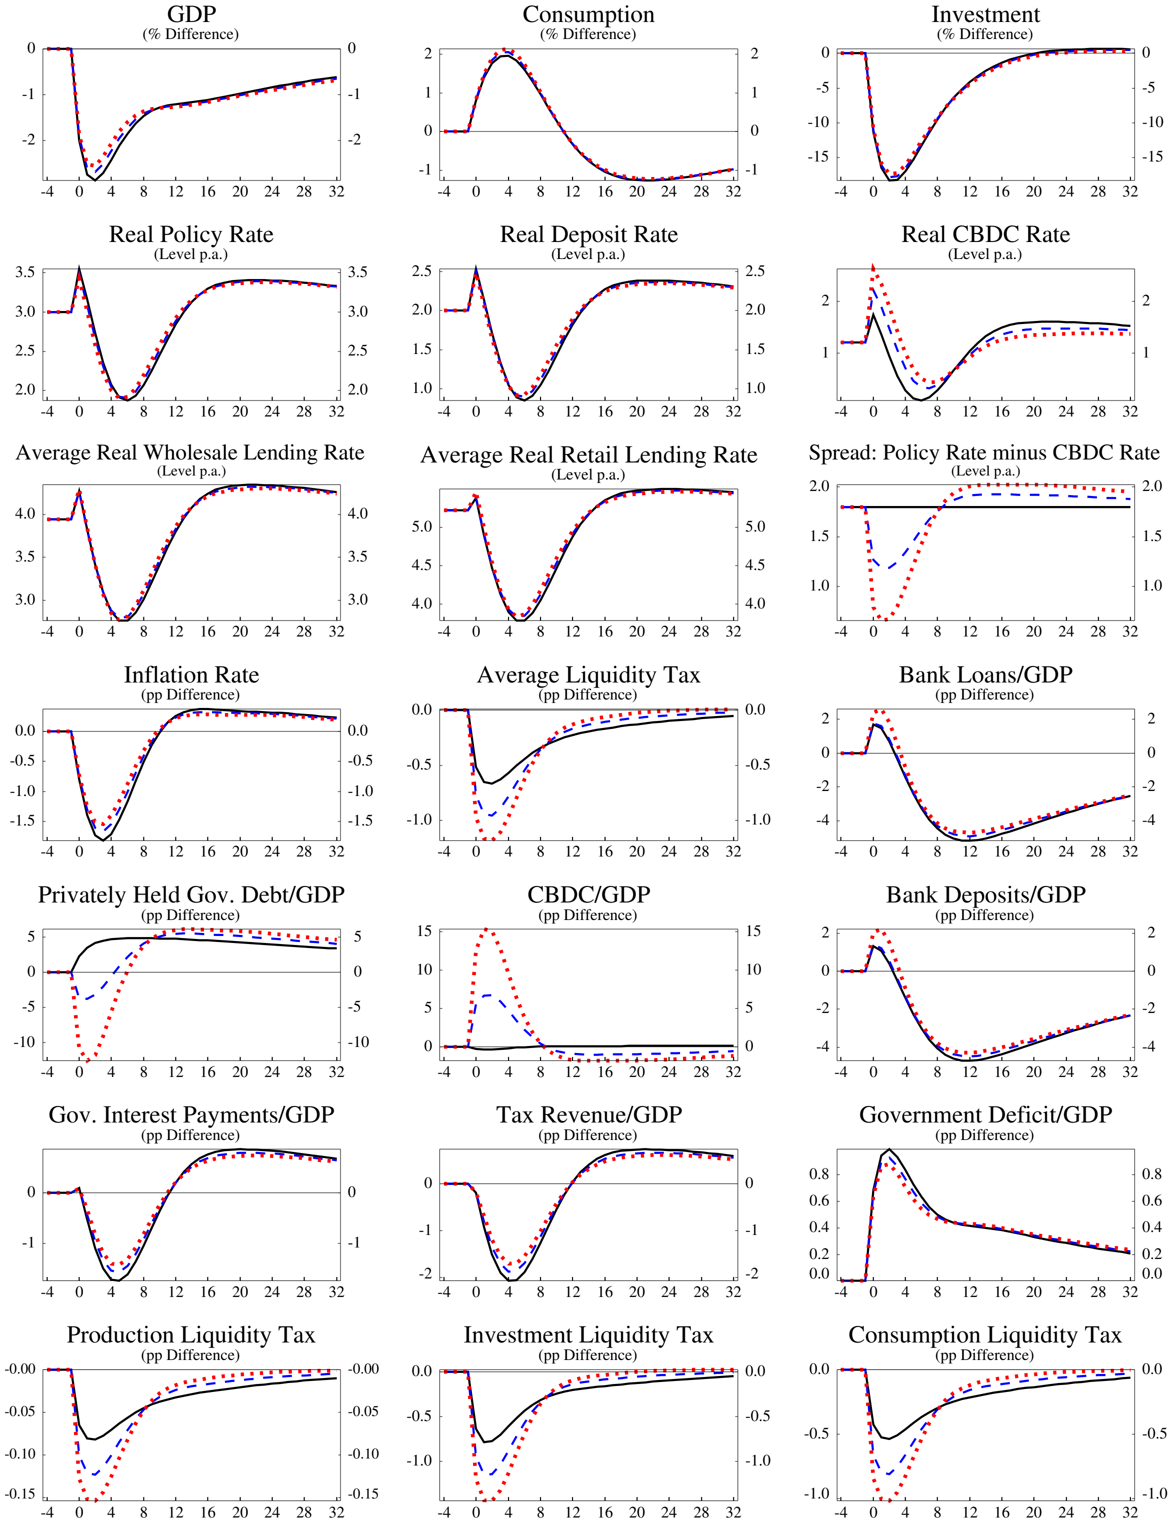

Solid Line = Baseline ($i^m_{\pi^p} = 0$), Dashed Line = Intermediate ($i^m_{\pi^p} = 0.4$), Dotted Line = Aggressive ($i^m_{\pi^p} = 0.8$)

Figure 11. Countercyclical CBDC Price Rules - Higher Demand for Total Liquidity

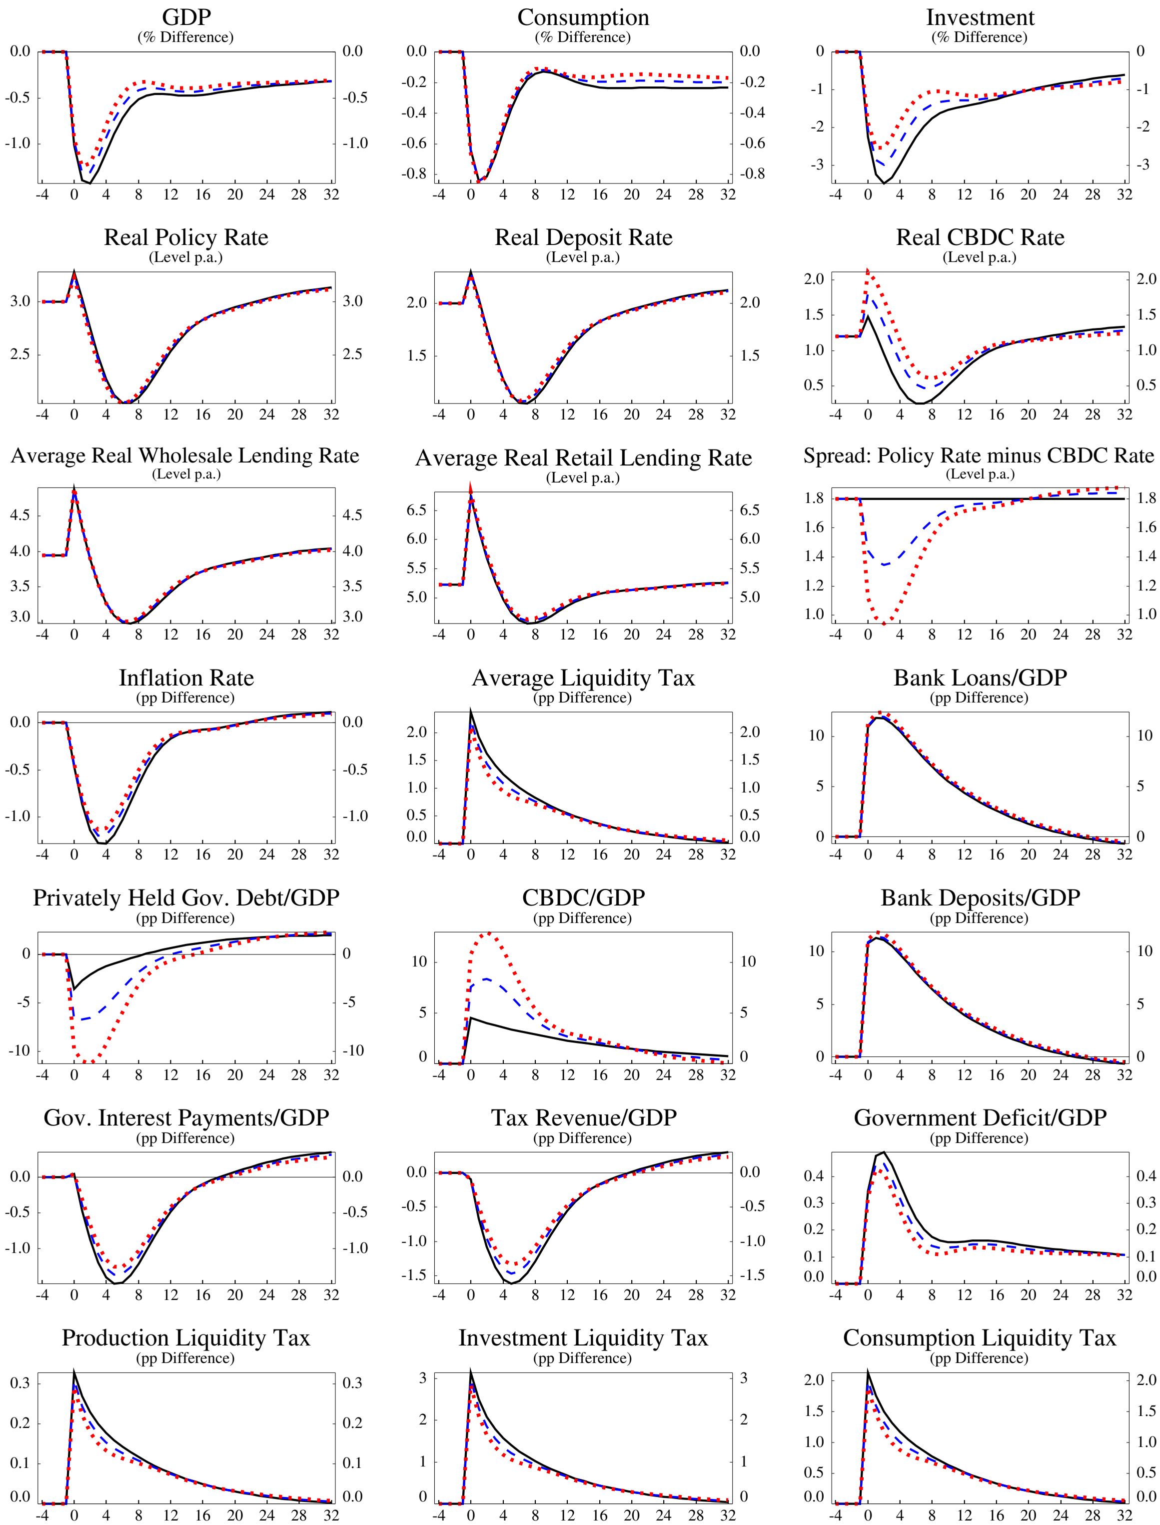

Solid Line = Baseline ($i^m_{\pi^p} = 0$), Dashed Line = Intermediate ($i^m_{\pi^p} = 0.4$), Dotted Line = Aggressive ($i^m_{\pi^p} = 0.8$)

Figure 12. Countercyclical CBDC Quantity Rules - Higher Demand for Total Liquidity

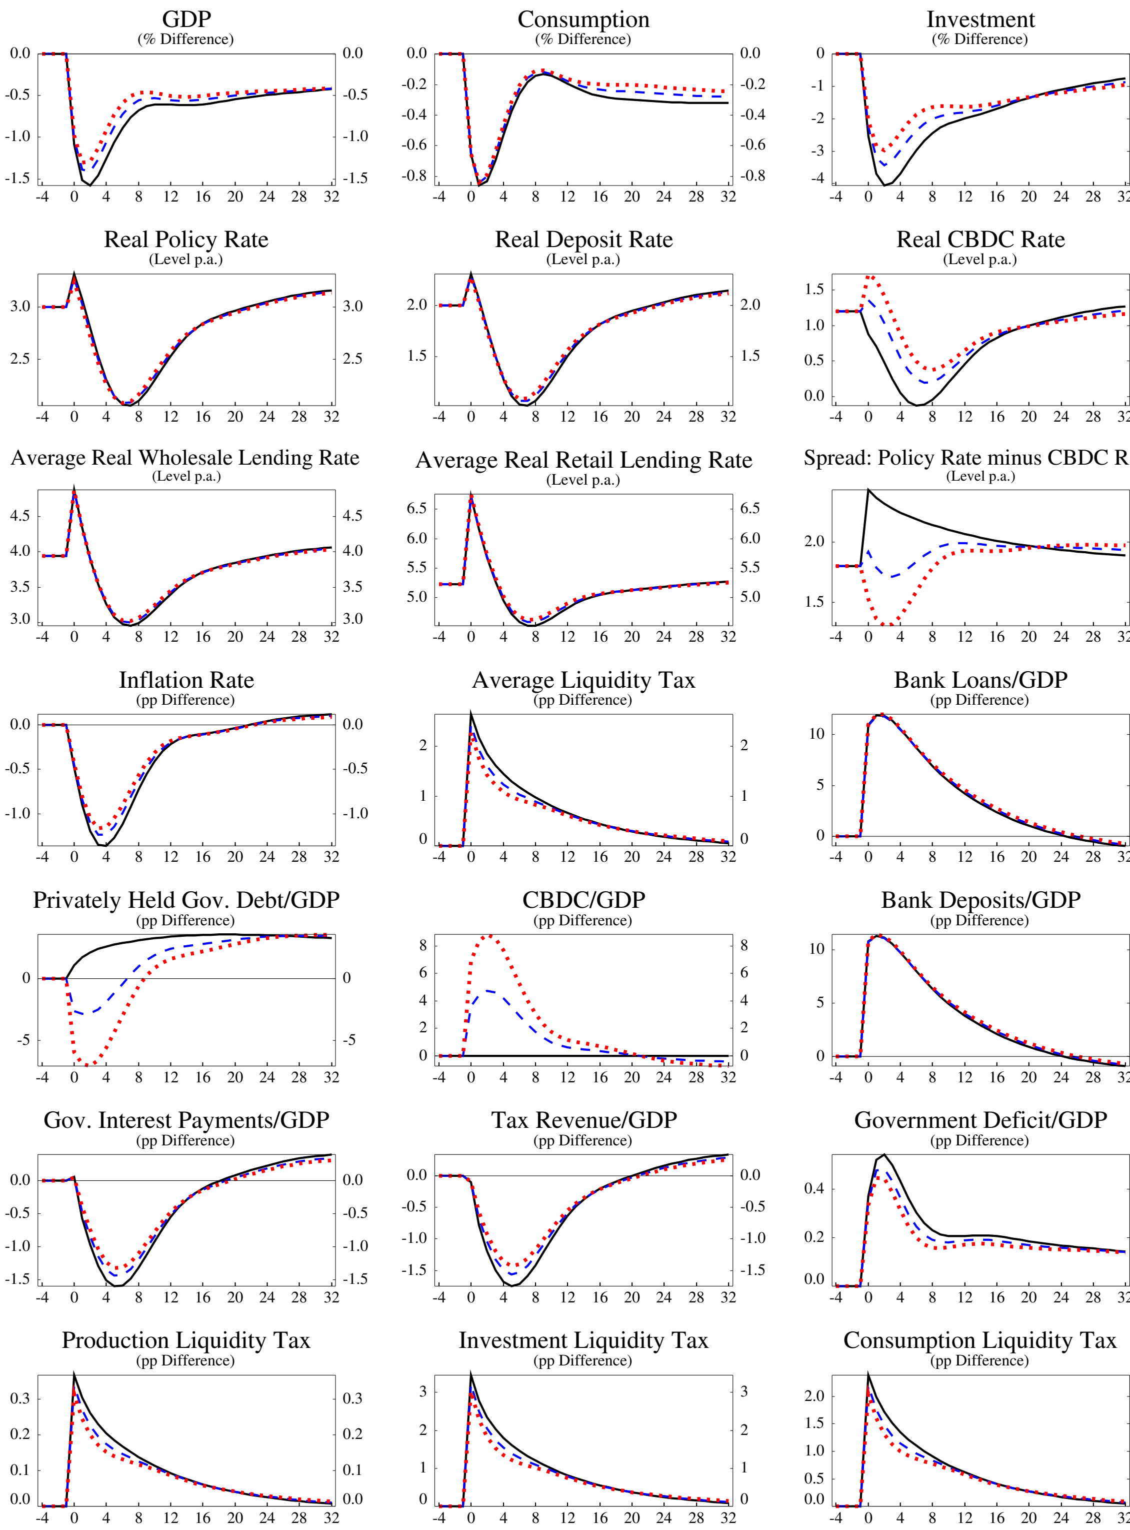

Solid Line = Baseline ($m_{\pi^p} = 0$), Dashed Line = Intermediate ($m_{\pi^p} = 4$), Dotted Line = Aggressive ($m_{\pi^p} = 8$)

Figure 13. Low EoS - Countercyclical CBDC Price Rules - Higher Demand for Total Liquidity

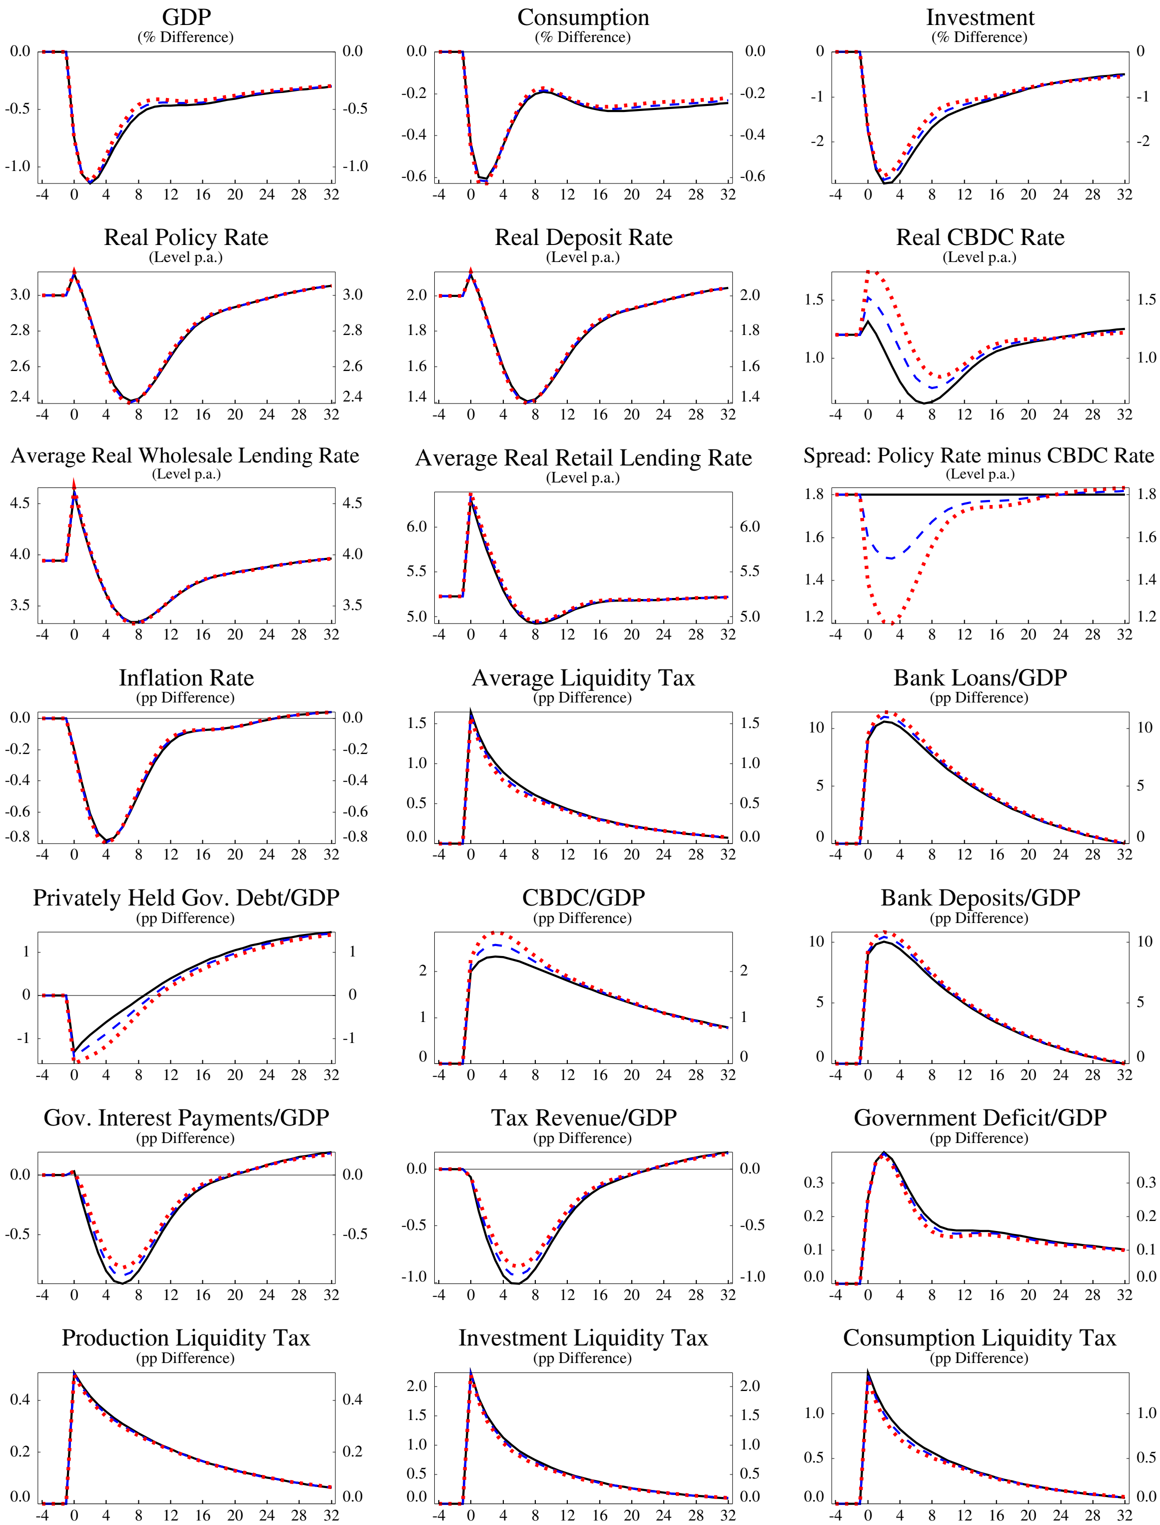

Solid Line = Baseline ($i^m_{\pi^p} = 0$), Dashed Line = Intermediate ($i^m_{\pi^p} = 0.4$), Dotted Line = Aggressive ($i^m_{\pi^p} = 0.8$)

Figure 14. Low EoS - Countercyclical CBDC Quantity Rules - Higher Demand for Total Liquidity

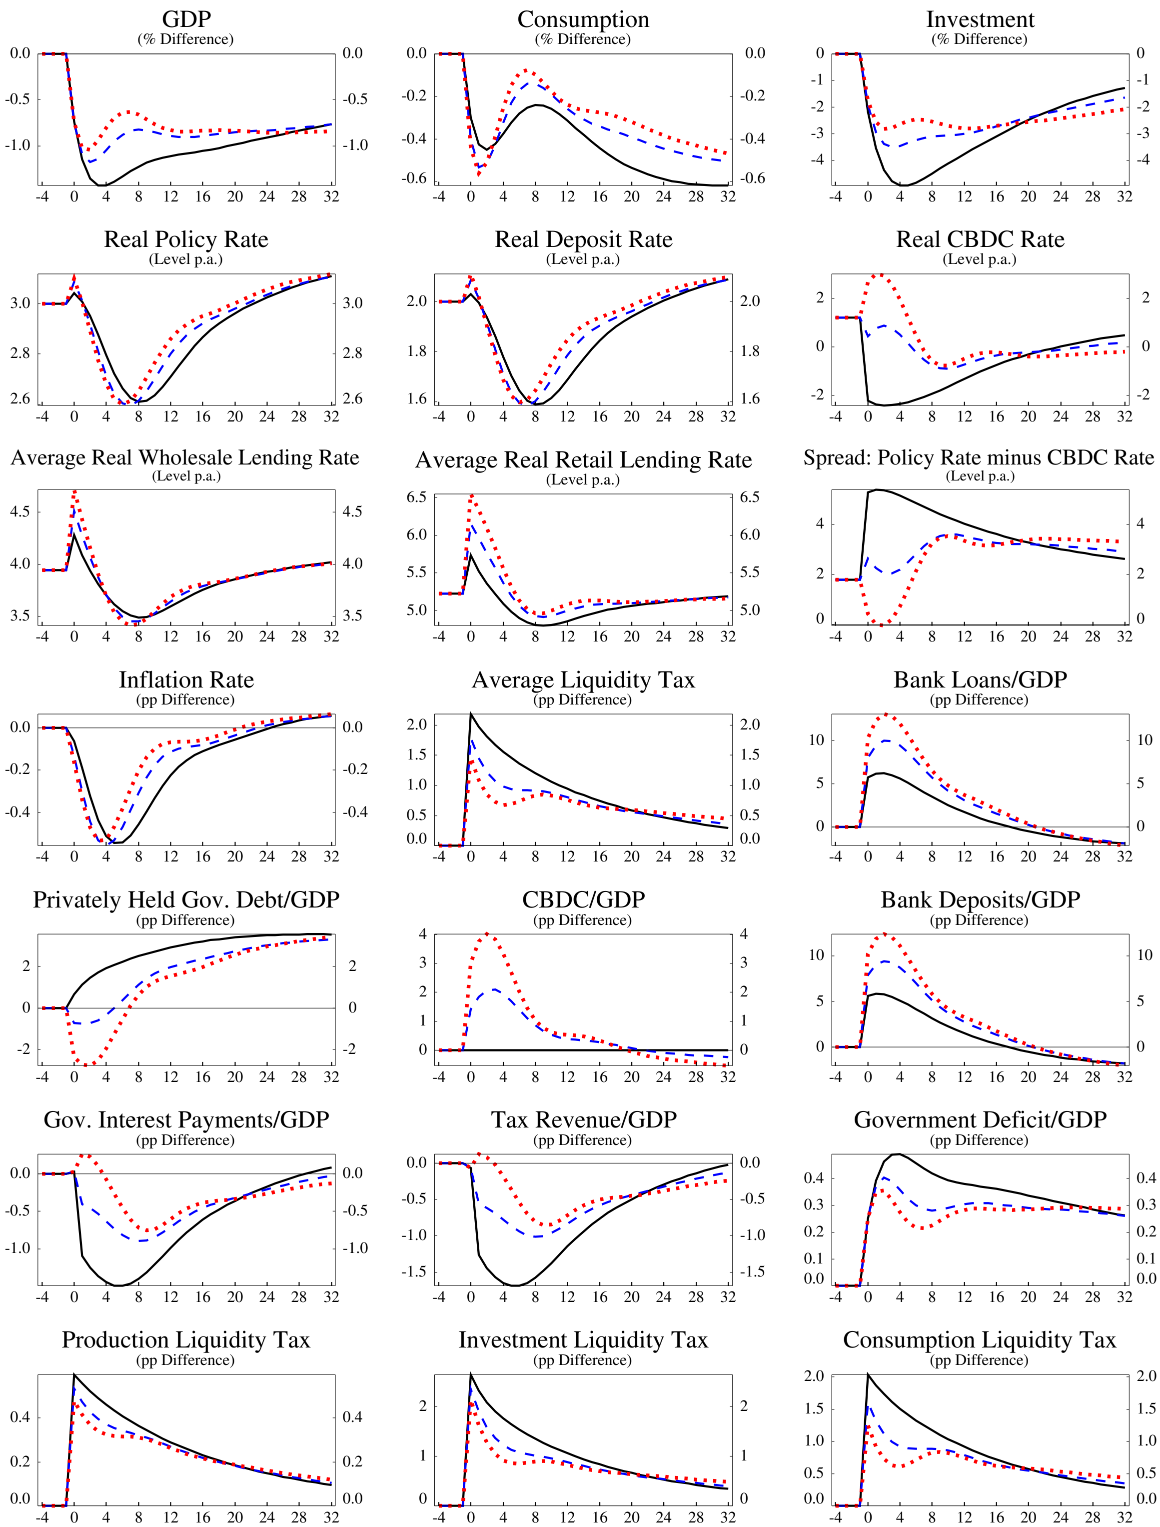

Solid Line = Baseline ($m_{\pi^p} = 0$), Dashed Line = Intermediate ($m_{\pi^p} = 4$), Dotted Line = Aggressive ($m_{\pi^p} = 8$)

Figure 15. CBDC-Based Discretionary Stimulus in Response to a Credit Risk Shock

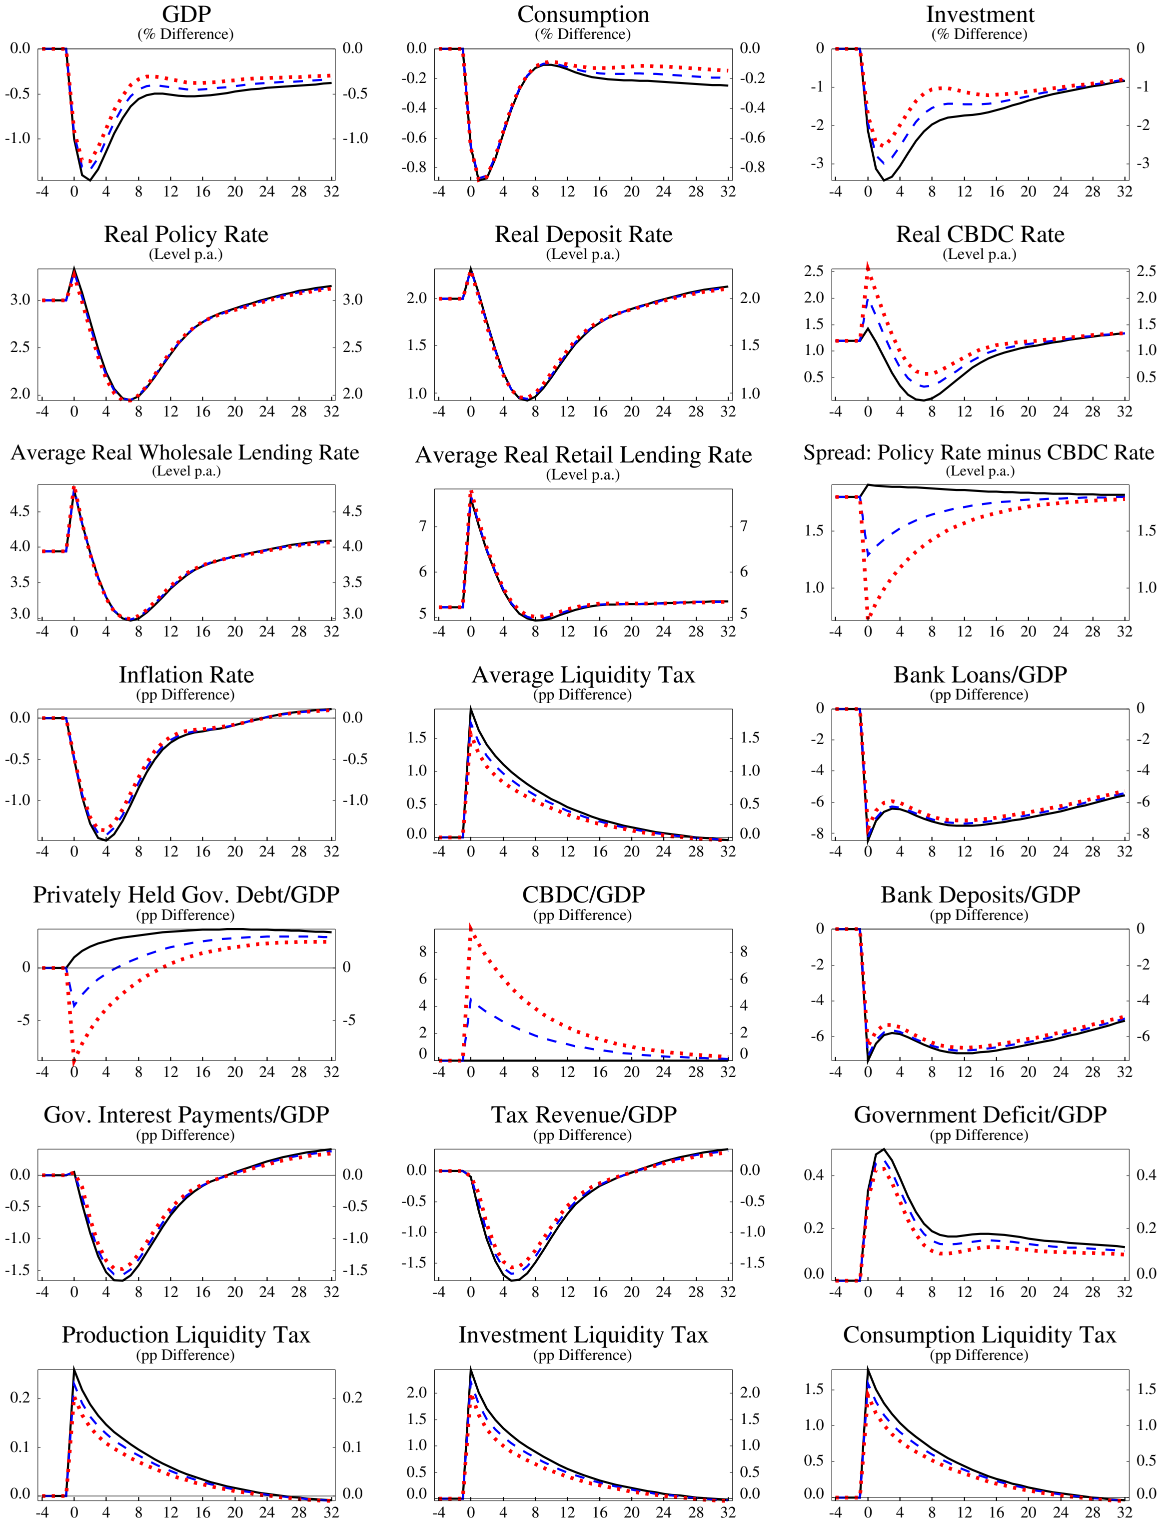

Solid Line = Baseline (no stimulus), Dashed Line = Intermediate Stimulus, Dotted Line = Aggressive Stimulus

Figure 16. Low EoS - CBDC-Based Discretionary Stimulus in Response to a Credit Risk Shock

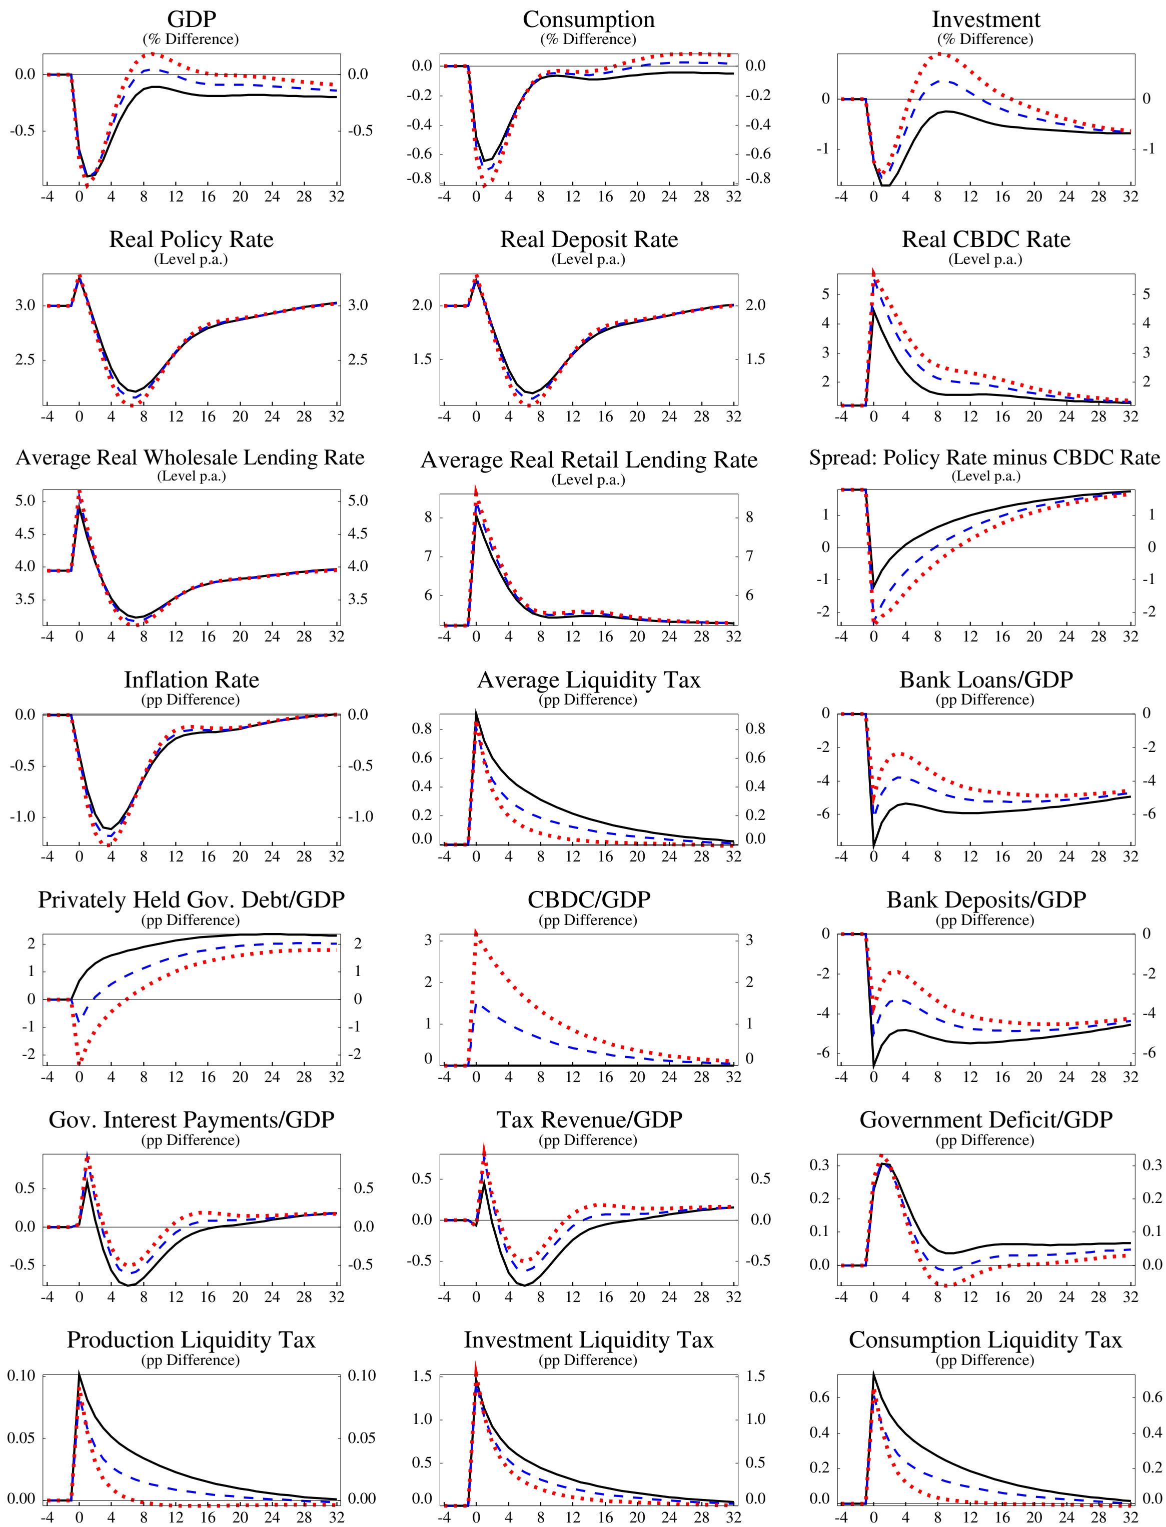

Solid Line = Baseline (no stimulus), Dashed Line = Intermediate Stimulus, Dotted Line = Aggressive Stimulus

Figure 17. Distortionary Taxes - Countercyclical CBDC Quantity Rules - Credit Cycle Shocks

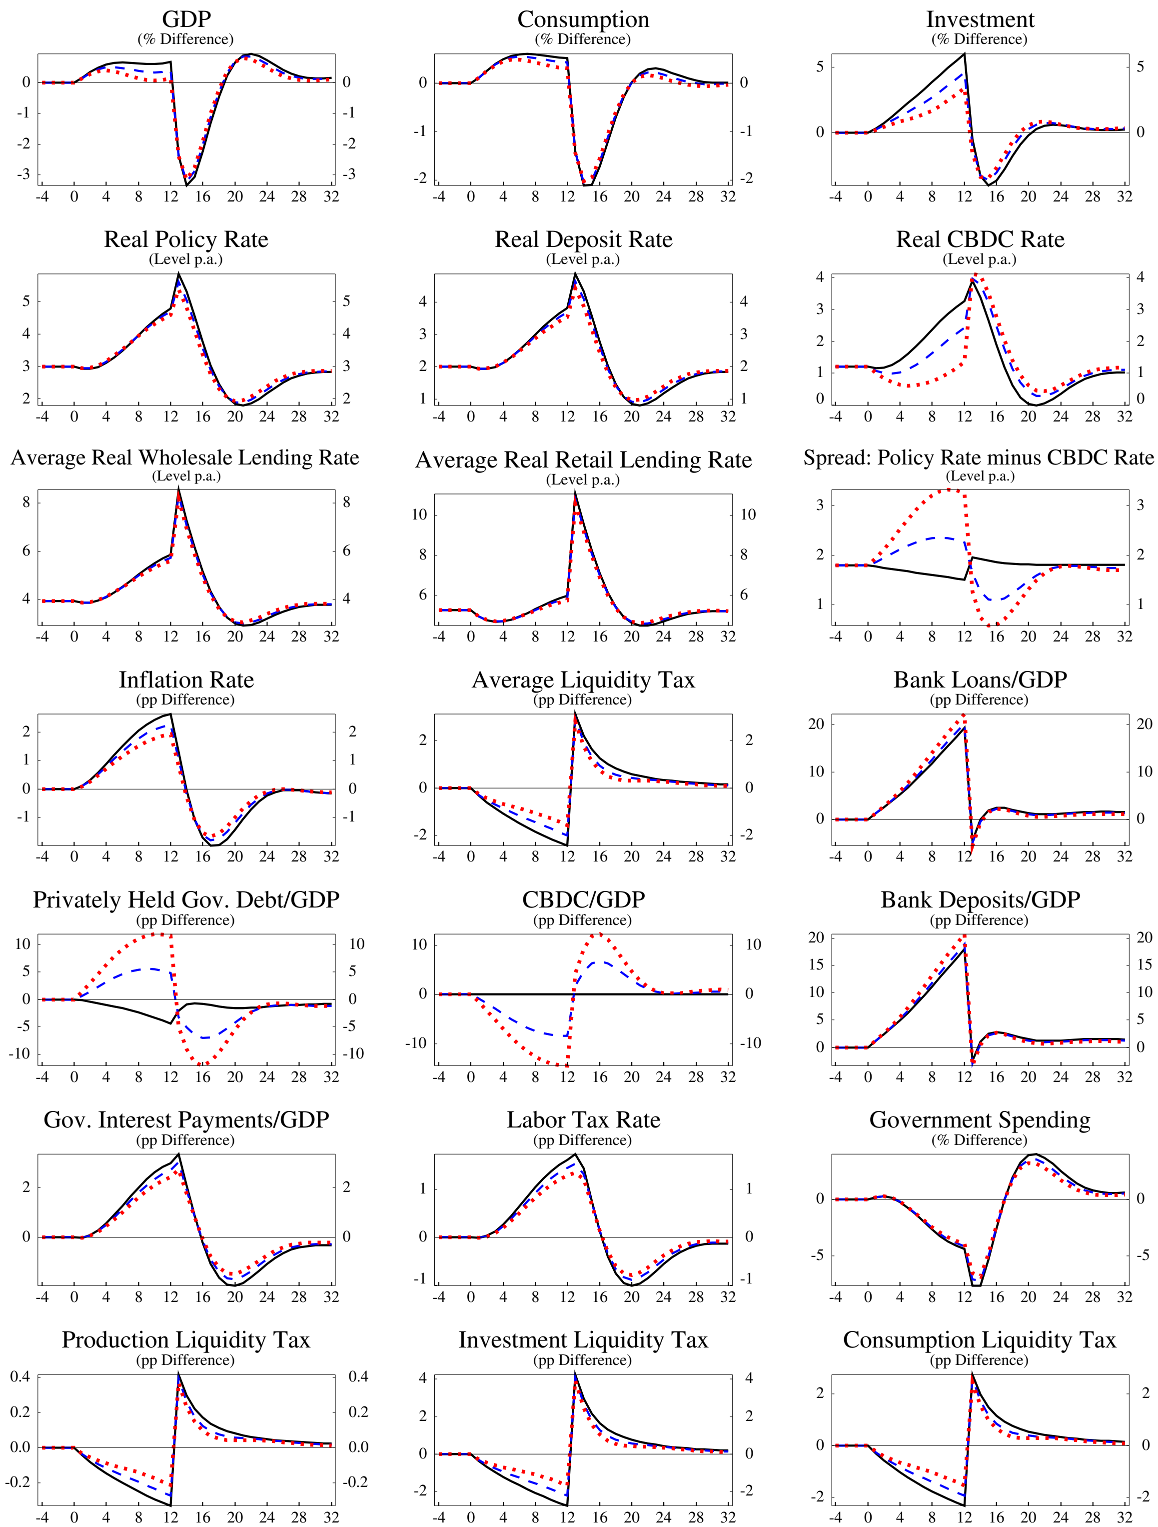

Solid Line = Baseline ($m_{\pi^p} = 0$), Dashed Line = Intermediate ($m_{\pi^p} = 4$), Dotted Line = Aggressive ($m_{\pi^p} = 8$)