BANK OF ENGLAND

Staff Working Paper No. 724

Broadening narrow money: monetary policy with a central bank digital currency

Jack Meaning, Ben Dyson, James Barker and Emily Clayton

May 2018

BANK OF ENGLAND

Staff Working Paper No. 724

Broadening narrow money: monetary policy with a central bank digital currency

Jack Meaning,[1] Ben Dyson,[2] James Barker[3] and Emily Clayton[4]

Abstract

This paper discusses central bank digital currency (CBDC) and its potential impact on the monetary transmission mechanism. We first offer a general definition of CBDC which should make the concept accessible to a wide range of economists and policy practitioners. We then investigate how CBDC could affect the various stages of transmission, from markets for central bank money to the real economy. We conclude that monetary policy would be able to operate much as it does now, by varying the price or quantity of central bank money, and that transmission may even strengthen for a given change in policy instruments.

Key words: Central bank digital currency, money, monetary policy, cryptocurrency.

JEL classification: E42, E52, E58.

(1) Bank of England. Email: jack.meaning@bankofengland.co.uk (corresponding author)
(2) Bank of England. Email: ben.dyson@bankofengland.co.uk
(3) University of Exeter. Email: jab255@exeter.ac.uk
(4) Bank of England. Email: emily.clayton@bankofengland.co.uk

We would like to thank John Barrdear, Tim Boobier, Rashmi Harimohan, Andrew Hauser, Michael Kumhof, Becky Maule and Matthew Trott for a number of informative discussions, Jamie Walton for help in proofing this draft, and participants at the CEBRA Annual Conference 2017, particularly Charles Kahn for his insightful discussion, and the ECB workshop on Money Markets, Monetary Policy Implementation and Central Bank Balance Sheets 2017. We also thank Andrew Levin for his insightful comments and suggestions as referee. The views here represent those of the authors only and should not be interpreted as an official position of the Bank of England in any way.

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

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

ISSN 1749-9135 (on-line)

1 Introduction

The feasibility and desirability of central banks issuing their own fiat versions of digital currencies has been the focus of a growing debate in recent years.[1] Numerous central banks around the world are researching the topic, including the Bank of Canada (2017), the European Central Bank (Mersch, 2017), the People's Bank of China (Qian, 2017), the Sveriges Riksbank (2017) and the Bank of England (2017b). The policy community is being aided by the academic community, with projects such as MIT's *Digital Currency Initiative*, and by collaboration with the private sector (Monetary Authority of Singapore, 2017). Despite this, the literature remains in its relative infancy with consensus around some fundamental issues only slowly beginning to form. What is more, the very concept of a central bank digital currency can feel inaccessible and otherworldly to economists who are not familiar with blockchain, distributed ledger technology or similar terminology.

A first contribution of this paper is therefore to offer a general definition of central bank digital currency (CBDC) around which consensus can build and which, importantly, should feel familiar to anyone with a basic understanding of the monetary process. Central bank digital currency is *any electronic, fiat liability of a central bank that can be used to settle payments, or as a store of value.* As such, CBDC can be viewed as electronic narrow money and in some senses already exists in the form of central bank reserves. Within our general definition there are a wide range of subsets and characteristics that may vary for any given design of CBDC. We will discuss these further in Section 2.

The nascent literature on CBDC covers an array of topics, from computer science to monetary economics.[2] For instance, discussion of the appropriate technological framework to underpin CBDC can be found in Danezis and Meiklejohn (2016). While they advocate the use of distributed ledger technology, similar to that behind Bitcoin and other private digital currencies, Scorer (2017) argues that this need not be the case. Benos, Garratt, and Gurrola-Perez (2017) evaluate distributed ledger technology in relation to securities settlement and suggest that it has the potential to improve efficiency and lower the costs of trading in these markets, but highlight the current immaturity of the technology.

Ketterer and Andrade (2016), Gouveia et al. (2017), FirstRand Bank (2016) and others focus on the implications of CBDC for the financial system, incumbent banks and their business models. A significant strand of the literature concerns itself with CBDC and the payments system, and relatedly, financial inclusion (Dyson & Hodgson, 2016).

Within the more traditionally monetary and macroeconomic sphere, Barrdear and Kumhof (2016) offer one of the first assessments of the potential impact CBDC may have on the real economy. In the context of a large-scale DSGE[3] model they suggest that the introduction of a CBDC via purchases of government bonds could increase real GDP by as much as 3%. Bordo and Levin (2017b) provide a comprehensive analysis of the design of CBDC and how it would affect the monetary framework on a deep level.

[1] Bordo and Levin (2017a) define a digital currency as "an asset stored in electronic form that can serve essentially the same function as physical currency, namely facilitating payments transactions".

[2] This is another reason that having consensus on a general definition is useful as it provides a common language to the diverse range of fields involved in research on CBDC.

[3] Dynamic Stochastic General Equilibrium.

They conclude that CBDC could act as a highly effective form of money and promote true price stability, as the real value of CBDC could be easily held stable over time.

In contrast to the existing literature, this paper approaches the topic of CBDC from the perspective of its impact on the monetary transmission mechanism (MTM). More specifically, we consider how CBDC might affect the implementation of monetary policy decisions and the channels, speed and strength with which those decisions are transmitted to the real economy. For the most part we focus on a specific form of central bank digital currency: universally accessible, account-based and interest-bearing. Conceptually, this CBDC variant could be thought of as a widening of access to the existing system of central bank reserve accounts, which are currently only accessible to a limited subset of economic agents.[4] In Section 6 we consider a scenario in which different types of CBDC account holder are paid different interest rates and show how this is largely analogous to reserves existing alongside a second CBDC. In Section 7 we look at an e-cash variant of CBDC.

It is difficult to draw definitive or quantitatively-robust conclusions about the impact of CBDC on the monetary transmission mechanism, due to the large degree of uncertainty around the ultimate design of CBDC, the economic environment it will be introduced into, and the structural changes that may accompany it. Rather, this paper aims to provide a framework for thinking about these issues and begins to populate that framework with broad brush conclusions based on existing theory and reasonable, transparent assumptions. These conclusions would need to be reviewed were a CBDC to be actively introduced. What is more, the impact on the monetary transmission mechanism will need to be balanced by considerations of the other dimensions of CBDC, such as its effect on payments systems and financial stability. In Section 8 we lay out some of the most pertinent questions that still need to be addressed and which would be fundamental to the effective implementation of CBDC.

Despite these uncertainties, our analysis leads us to the broad conclusion that a universally accessible, interest-bearing, account-based CBDC could be used for monetary policy purposes in much the same way that central bank reserves are now. On the margin, there may even be reason to believe that the monetary transmission mechanism would be stronger for a given change in policy instruments.

The rest of the paper is laid out as follows. We begin by elaborating on our general definition of CBDC and the various sub-characteristics it may exhibit. We then provide a balance sheet exposition of an economy with universally accessible CBDC and use it to work through some of the key CBDC-related transactions in the economy. In Section 4 we discuss the likely impact of CBDC on the monetary transmission mechanism, building from the control of CBDC rates by the central bank, to the channels through which this impacts the real economy. Section 5 considers quantitative easing under CBDC, while sections 6 and 7 look at alternative designs for CBDC, respectively a single CBDC that pays different interest rates dependent on the account holder type, and e-cash. Finally, we outline some outstanding questions and areas of interest that follow from our research, before offering some concluding remarks.

[4] The Bank of England has recently widened access to central bank reserves to include non-bank Payment Service Providers and Electronic Money Issuers (Bank of England, 2017a).

2 Defining central bank digital currency

The term *central bank digital currency* is currently used to refer to a wide range of potential designs and policy choices, with no single commonly agreed definition. This is, in part, due to the fact that the concept sits at the nexus of a number of different areas of research and brings together many complex elements, covering topics as diverse as computer science, cryptography, payments systems, banking, monetary policy and financial stability. As a result, the debate around CBDC may at times appear otherworldly to economists who are not familiar with cryptocurrencies, blockchain and distributed ledger technology (DLT).

As a response, we offer a general definition that is both an accurate representation of CBDC, and is stated in terms that will be familiar to any economist with a basic understanding of the monetary process. It is our hope that this will remove any perceived mystique around CBDC and promote wider discussion of CBDC among economists.

To that end, we define a central bank digital currency simply as *an electronic, fiat liability of a central bank that can be used to settle payments or as a store of value.* It is in essence electronic central bank, or 'narrow', money.

Within our general definition there exists a wide range of sub-characteristics and parameters that could be set or varied and which would more precisely define any given CBDC.

One of the key parameters relates to access to CBDC. A CBDC may be universally accessible — in other words it can be held by anyone for any purpose — or access may be restricted to a limited subset of economic agents, or for a limited range of purposes. This may seem counterintuitive to some monetary economists who take *currency* to be something that can be held by anyone i.e. is universally accessible[5]. Indeed, Fung and Halaburda (2016) and Bjerg (2017) identify universal access as a fundamental characteristic of any CBDC. However, it seems possible that central banks may issue a CBDC that is available only to a sub-sector of the economy, such as a 'retail CBDC' for households and non-financial businesses only, or a 'wholesale CBDC' which can be used as a settlement asset in financial markets by firms that do not currently have access to central bank reserves (Bech & Garratt, 2017). The question is then a semantic one about whether a CBDC that is *not* universally accessible could still be considered a central bank digital *currency*. In fact, the ECB has chosen to use the broader term "digital base *money*" (Mersch, 2017), which we find to be more accurate. For consistency we will stick with the more widely-used term 'CBDC', but acknowledge the potential mild misnomer.

A second key parameter relates to whether CBDC is interest bearing. An interest-bearing CBDC might pay positive, zero or even negative rates at various points in the economic cycle. As we will discuss in Section 2.1, this interest rate could be used to pursue a number of objectives, although not simultaneously. For instance, it could be used to stabilise inflation and output, as the primary instrument of monetary policy, or it could be used to regulate demand for CBDC. Alternatively, a non-interest-bearing CBDC could be considered closer in spirit to central bank notes, and so is often referred

[5]Making a CBDC universally accessible to non-residents as well as residents could have implications on the exchange rate and capital flows. These are complex issues that we do not address in this paper, except very briefly in Section 4.

to as ‘e-cash’.

An important question relates to whether a CBDC trades at par with other central bank liabilities. In most existing monetary frameworks, different types of central bank liabilities can be exchanged with one another 1:1, e.g. one unit of central bank notes can be exchanged for one unit of reserves. However, a number of authors have suggested breaking with this convention, particularly in the context of CBDC. Kimball and Agarwal (2015) for instance outline a framework in which a flexible exchange rate can be operated between cash and electronic central bank money in order to facilitate a negative interest rate on cash and overcome the effective lower bound. What such a system would mean though is that the economy would be operating with two distinct fiat currencies simultaneously, albeit with a managed exchange rate. We doubt that this would be plausible in practice. For instance, it would raise significant questions around which of the two currencies, physical cash or CBDC, would be the unit of account in the economy. Also, were both currencies to become widely used, the prices of goods and services would need to be quoted in both, adding an administrative cost that could be significant. In general, we believe managing two fiat currencies simultaneously would pose a significant risk for monetary stability, so while we acknowledge the theoretical possibility of a CBDC that does not trade at par, for this paper we assume that all forms of central bank money can be exchanged at par.

Some design choices relate to the technology used to power a CBDC. For instance, a CBDC could be *token-based* currency, such as the proposed Fedcoin of Koning (2014) or BitDollar, proposed by Motamedi (2014). This would mean that once issued, units of CBDC could be transferred from one agent to another independently of the central bank, in much the same way that central bank notes are currently. The alternative would be an *account-based* CBDC in which agents had an account recorded by the central bank and transactions were made by the central bank debiting one account and crediting another.

Perhaps the characteristic that causes the most confusion is whether or not the CBDC is a cryptocurrency. Cryptocurrencies, such as Bitcoin, Litecoin or Ether, make use of distributed ledgers that rely on cryptographic techniques to maintain their veracity. Much of the discussion around CBDC implies, either explicitly or implicitly, that it would also be a cryptocurrency, but this need not be the case. CBDC could equally be based on a more mature and established technology, such as that which powers existing central bank real-time gross-settlement systems.[6] This kind of CBDC would not be a cryptocurrency, but would remain a central bank digital currency.

The optimal setting of each of these parameters will depend on the reason for which the CBDC is being introduced; a CBDC designed to provide a secure payments service may look very different to one that is primarily used as an instrument of monetary policy. However, our general definition provides a framework by which any future research can define the subset of CBDC to which it refers, and assess how its conclusions may vary were those parameters to be set differently. We therefore see this as an important first contribution of our paper. There is a small existing literature seeking to clarify the terminology around CBDC. Bjerg (2017) offers a taxonomy that is a subset of our

[6] Scorer (2017) explains that distributed ledger technology may have other technological advantages over centralised ledgers — in particular, a higher level of resilience. However, the technology is considered still too immature to power a critical national payment system such as the Bank of England’s Real-Time Gross Settlement system (Bank of England, 2017a).

more general definition, additionally requiring universal accessibility. Bech and Garratt (2017) present a taxonomy broadly consistent with our own. Within their *money flower* representation, our definition of CBDC is depicted by the four core segments incorporating both retail and whole central bank crypto-currencies, settlement or reserve accounts with the central bank, and deposited currency accounts. The analysis within Bech and Garratt (2017) goes on to focus on the subset of money that is based on cryptographic technology, so discusses central bank *crypto*currencies - CBCC rather than CBDC.

To provide some context to our definition, Table 1 compares CBDC to some common money-like assets and shows the characteristics they each have. As shown by the first column, under our general definition, the first two characteristics are both necessary and sufficient to class an asset as CBDC, while the others may or may not hold. Before any real-world CBDC could be launched, careful consideration would need to be given to the setting of these parameters.

Table 1: Characteristics of CBDC and other money-like assets

	CBDC	Reserves	Central Bank Notes	Deposits	Bitcoin	Ether[a]
Liability of the central bank	✓	✓	✓	x	x	x
Electronic	✓	✓	x	✓	✓	✓
Universally accessible	?	x	✓	✓	✓	✓
Interest bearing	?	?	x	?	x	x
Trades at par[b]	?	✓	✓	✓	x	x
Cryptocurrency	?	x	x	x	✓	✓
Token or account based	?	A	T	A	T	A

[a]We have taken Bitcoin and Ether as the best known examples of privately-issued cryptocurrency. The characteristics shown are also accurate for the majority of cryptocurrencies, although the economic and technological design of different cryptocurrencies can vary significantly.

[b]Trades at par with other central bank liabilities

To give a clear example, under our general definition, one form of central bank digital currency already exists, and plays a fundamental role in the conduct of monetary policy: reserves. Reserves are electronic and account-based. They are *not* a cryptocurrency as they are not based on distributed ledger or other cryptographic system of account. In many cases they pay a rate of return, either on the total balance held, or on those parts of the balance deemed *excess*. However, they are only accessible to a limited number of participants, mainly banks and other select monetary financial institutions (MFIs). Their primary objective is as an active instrument for monetary policy, however, for a small subset of the economy they also provide secure, efficient payments services, and they can serve as an instrument of financial stability policy. It should be noted here that under the stricter definition of CBDC suggested by Bjerg (2017), reserves would not qualify as a CBDC as they fail to meet the universal accessibility requirement. However, a number of studies have either implicitly or explicitly made the assumption that reserves should be thought of as CBDC, including He et al. (2017) at the IMF.

2.1 CBDC in this paper

For the majority of this paper we will focus on a particular form of CBDC. Following the conclusions of both Scorer (2017) and Bordo and Levin (2017b) we concentrate on an account-based CBDC, in which the identity of CBDC account holders is known (in contrast to a token-based CBDC which could offer truly anonymous payments). We also begin from an assumption of universal accessibility such that everyone can hold a CBDC account with the central bank.

Beyond assuming that a CBDC is technologically feasible, we abstract away from questions over the precise nature of the underlying technology, remaining agnostic over whether the accounts are managed on a distributed ledger, a form of RTGS or some other means.[7] However, we feel that the most plausible method of implementation would be that the central bank would provide the underlying payments platform for CBDC, but would not deal directly with members of the public. Instead, it would allow private sector firms, such as financial technology firms, payment institutions or banks, to identify customers, register CBDC accounts on their behalf and provide the customer interface and customer services related to CBDC accounts. These firms would be responsible for administering CBDC accounts, but would not take custody of the CBDC itself, ensuring that CBDC remained a liability of the central bank to the end user rather than to an intermediary. We assume that the payment services available to CBDC account holders would be comparable to those available to holders of bank deposits, and that the CBDC and deposit-based payment systems would be interoperable (so that any deposit account could make a payment to any CBDC account and vice versa). This means that CBDC would serve as a close — but not perfect — substitute for bank deposits. This point is crucial: as discussed by Broadbent (2016), if CBDC only serves as a substitute for central bank notes, then the monetary policy implications are negligible, but once CBDC starts to offer payment services similar to bank deposits, there will be an impact on the quantity and price of bank funding, with interesting implications for monetary policy. Importantly, we assume that CBDC accounts would not offer credit facilities, such as overdrafts, to the vast majority of users.[8]

Lastly, in the main we will be working on the assumption that CBDC pays a rate of interest and, more specifically, that this interest rate is used as an instrument of monetary policy. The arguments in favour of paying interest on short-term central bank liabilities have their origins in Friedman (1960), who argued that they should pay a rate of interest equivalent to the risk-free rate. As we will expand upon in Section 4, in our benchmark framework the rate of interest paid on CBDC balances is the main instrument of monetary policy, and would be set in order to guide the macroeconomy. However, an alternative approach would be to set the interest rate paid on CBDC with a different objective in mind, for instance to guide the quantity of CBDC in response to changes in demand. This would require a different formulation of CBDC to the one considered in our benchmark. Primarily it would require a distinction between CBDC held by banks, and CBDC held by non-banks, including the public more broadly. The rate paid on the CBDC held by MFIs could then be used to set the stance of monetary policy (just as the rate on reserves is used today), while the

[7] See Scorer (2017) on the relative merits of DLT or other delivery methods.

[8] As discussed in Section 4, we do assume CBDC credit can be lent to certain financial institutions at a discount window facility.

rate paid on the broader CBDC could be used to regulate demand for CBDC through the price mechanism. In Section 6 we will discuss a variant of our benchmark that approximates this by allowing the central bank to pay differentiated rates of interest to different account holder types. Barrdear and Kumhof (2016) go further and explicitly differentiate reserves and a broader CBDC as distinct assets. But as we will show, the two frameworks are largely equivalent except under specific conditions.

Conceptually, one way to think about the CBDC discussed in this paper is simply as a widening of access to the existing system of reserve accounts offered by central banks. Since the financial crisis of 2007-09 the Bank of England, like a numberof other central banks, has undertaken projects to widen access to central bank money to broker-dealers and central counterparties (CCPs), although these firms still make up a relatively small subset of the economy. The universally accessible CBDC discussed below takes that process to its extreme conclusion of access being granted to all agents within the economy.

3 A balance sheet perspective on CBDC

As a first step in considering how CBDC could affect monetary policy, Figure 1 shows a series of stylised balance sheets for three key sectors of the economy — the central bank, banks, and the non-banks private sector — in a world without CBDC, and a world with a CBDC of the type described above.[9]

The left column, showing the situation without CBDC, should be familiar to many. Central bank money exists in two forms: *electronic* reserves, and *physical* central bank notes.[10] Central bank notes can be held by both banks and non-banks, while reserves can only be held by banks. These two types of central bank liability are backed by the central bank's assets. In this simplified exposition, these assets consist entirely of government bonds, although in practice a wider range of assets may be held.

Banks make loans to the non-bank private sector, simultaneously issuing deposits on the liability side of their balance sheet.[11] Non-banks hold these deposits as assets, alongside government securities, equity in the banking sector and physical central bank notes. On the liabilities side of their balance sheet, non-banks have the loans they owe to banks, and positive 'non-bank equity' (since their assets exceed their liabilities). The right-hand panel shows how this might be different in a world with a universally accessible CBDC. The two forms of central bank money remain, one electronic and one physical, although for simplicity, we have renamed electronic central bank money as CBDC. It should be stressed that this is not a model in which reserves exist separately from a universally accessible CBDC, although such a scenario is considered in Section 6 and more fully by Barrdear and Kumhof (2016).

Note that we do not assume, nor advocate, the removal of physical central bank money (i.e. notes). Such an innovation is often a feature of debates around central bank digital currencies, and in some cases is even used to motivate their introduction.

[9]Kumhof and Noone (2018) provide a more comprehensive exploration of the first-order balance sheet impacts of introducing CBDC under different design scenarios.

[10]In the UK coins are issued by the Royal Mint and recorded as liabilities of the Treasury, so are not 'central bank money' despite trading at par with central bank money and being considered as another form of 'cash'.

[11]For a clear exposition of the money creation process, see McLeay, Radia, and Thomas (2014).

Figure 1: Balance sheets without and with CBDC

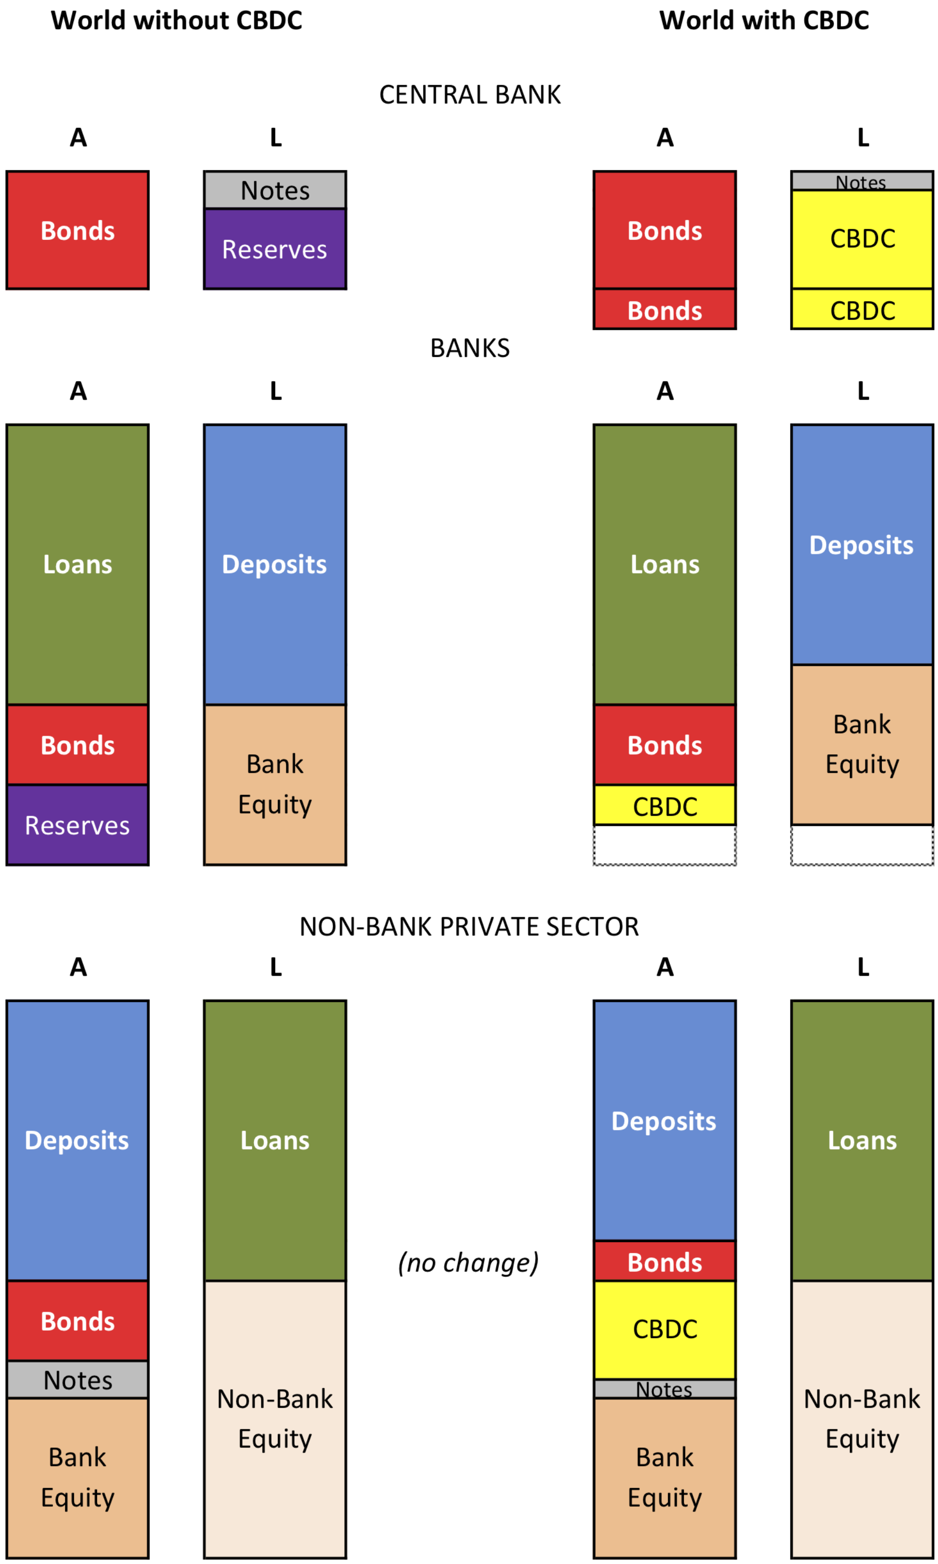

Goodfriend (2016), Kimball and Agarwal (2015) and Rogoff (2016) all suggest that replacing cash with a CBDC could make it easier to set a negative rate on central bank money and thus alleviate the lower bound on interest rates. However, as recognised by these authors, the removal of cash is not a necessary consequence of CBDC and while one may have a bearing on the other, the two policies should be assessed on their own merits.

The most obvious difference between the left and right panels is that non-banks can now hold electronic central bank money, in the form of CBDC. Our example assumes that non-banks partly substitute from bonds, deposits and banknotes into CBDC. The impact on the size and composition of the different sectors' balance sheets depends on the type of substitution from each asset to CBDC.

Substitution by non-banks from central bank notes into CBDC simply changes the composition of central bank liabilities and private sector assets. This has limited impact beyond some minor implications for seigniorage.[12] However, substitution by non-banks from either deposits or bonds into CBDC will have more significant impacts on the balance sheets, through the transactions described below.

3.1 Varying the aggregate supply of CBDC

A central bank could increase the aggregate supply of CBDC in the economy by purchasing financial assets from the non-bank private sector (or indeed from the banking sector), and paying for these assets with newly created CBDC. This is shown in Figure 2. The non-bank sells an asset — the bond — to the central bank, transferring the bond across to the central bank.[13] The central bank pays for this asset by issuing new CBDC, which is credited to the CBDC account of the non-bank. The results are shown in the right hand column of Figure 2: the central bank balance sheet has expanded, since the new asset (the bond) is matched by a new liability (CBDC). For the non-bank private sector, one asset (the bond) has been replaced by another (CBDC), and so while the composition of its assets has changed, there is no change in the total quantity of assets. This transaction is very similar in balance sheet terms to quantitative easing, and we will consider its economic effects in Section 5.

A key question concerns the assets that should be used to back CBDC. If CBDC were to be backed exclusively by government bonds, similar very safe assets, or even repos against these safe assets, complications could arise during a period of financial turmoil, or a banking crisis, when a surge in demand for CBDC could coincide with a surge in demand for these safe assets. To avoid leaving the market significantly short of government bonds in such a moment, the central bank may need to accept a wider range of assets to back CBDC — indeed, many central banks already accept a wider range of assets. The decision on which specific assets would be eligible for asset purchases to control the supply of CBDC would depend on a range of factors including the risk tolerance and other balance sheet objectives of the central bank, and the objectives of the monetary expansion itself. For instance, in an economy where the

[12] The impact on seigniorage could be positive or negative depending on (a) the interest rate paid on CBDC (if any) and (b) the relative cost of producing CBDC versus the costs of producing the equivalent amount of cash (including fixed and marginal costs).

[13] The central bank could alternatively repo bonds or other eligible assets, essentially making a secured loan of CBDC. This would have different balance sheet impacts, not shown here.

Figure 2: Varying the aggregate supply of CBDC through asset purchases

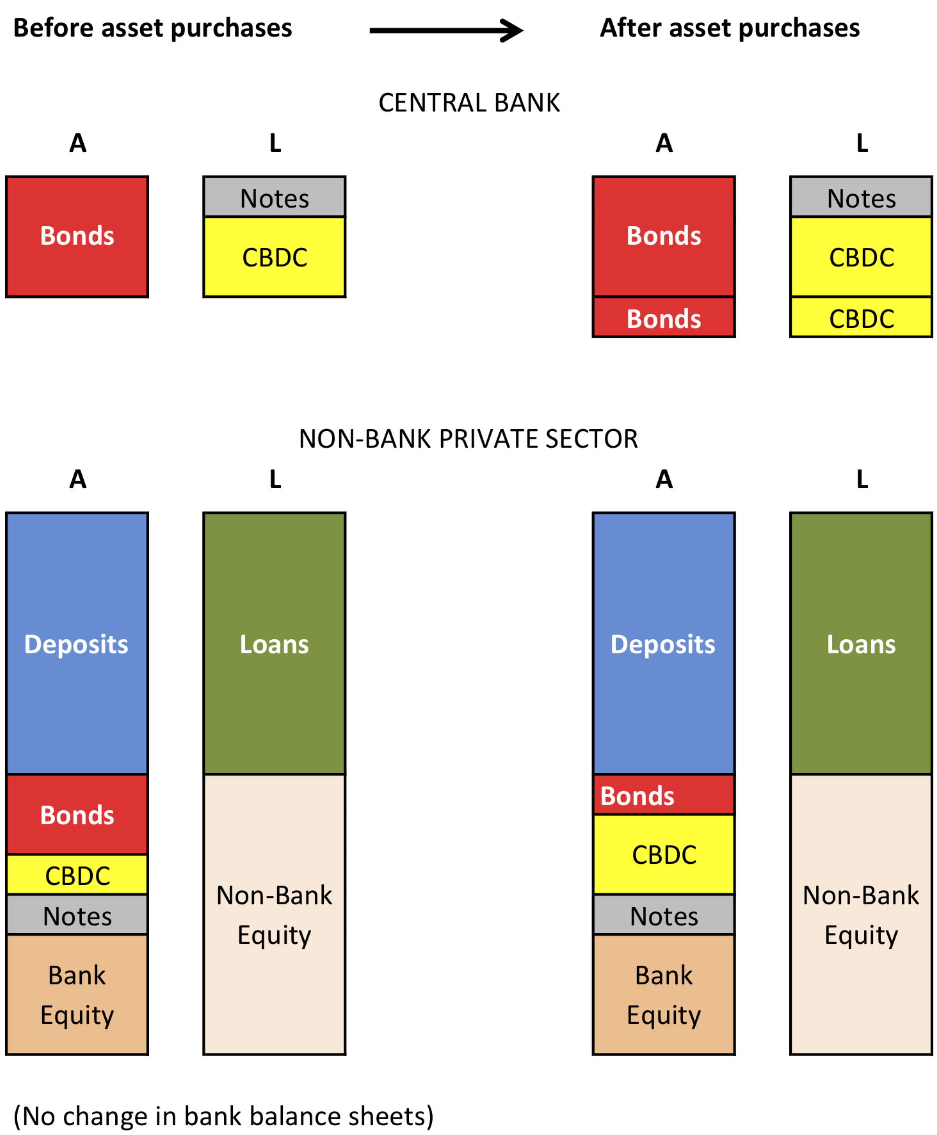

level of government debt is very low but the majority of transactions are carried out using electronic payments, it may be necessary to accept a wider range of collateral than in an economy with a large government debt market and little latent demand for CBDC.

3.2 Exchanging CBDC and deposits

A crucial design choice around CBDC would be how it interacted with the most prevalent form of money in the current system, bank deposits. We assume that depositors are able to withdraw CBDC on demand in the same way that they can currently withdraw central bank notes. This process is shown in Figure 3. When the depositor requests to withdraw CBDC from their deposit account, the commercial bank would reduce the balance of the depositor's account, and pay CBDC across the central bank's balance sheet to the depositor's CBDC account. The results are shown in the second row of balance sheets: the banking sector's balance sheet has shrunk by the amount of the withdrawal, whilst the depositor has simply swapped one asset (deposits) for another (CBDC), with no overall change in the size of their balance sheet. The overall impact on the banking sector's balance sheet is identical to that of a cash withdrawal.

3.2.1 Risk of disintermediation of banks

Allowing depositors to withdraw CBDC on demand has implications for bank funding and liquidity, since banks would lose both deposits, and CBDC, which would be transferred across to non-banks' CBDC accounts at the central bank (Broadbent, 2016). If this movement from deposits to CBDC is not offset by other agents moving from CBDC to deposits, then the overall reduction in demand for commercial bank deposits has the effect of shrinking the aggregate size of the banking sector's balance sheet. A slow but large substitution from deposits into CBDC could pose a growing challenge to the sustainability of current bank business models.

However, such an extreme scenario is not considered likely, especially in the short to medium run, because the central bank has considerable control over the attractiveness of CBDC relative to deposits through its design of CBDC itself. For instance, as discussed in Section 2.1, CBDC would not be a perfect substitute for bank deposits, and would not offer many of the services that deposit accounts with commercial banks do, such as overdrafts. The choice over how attractive to make CBDC vis–a–vis deposits will depend on the extent to which the central bank is willing to encroach on the role that the banking sector has traditionally played in the provision of an electronic means of payment.

What is more, banks are not passive in the face of the introduction of CBDC, and could respond to the competition from CBDC on price terms, offering a higher return in order to maintain their deposit base, a dynamic that is central to the monetary policy implications discussed in Section 4. They could also respond by evolving their offer to depositors, a process that is likely to be necessary anyway in light of technological developments in the payments field.

Figure 3: Balance sheets without and with CBDC

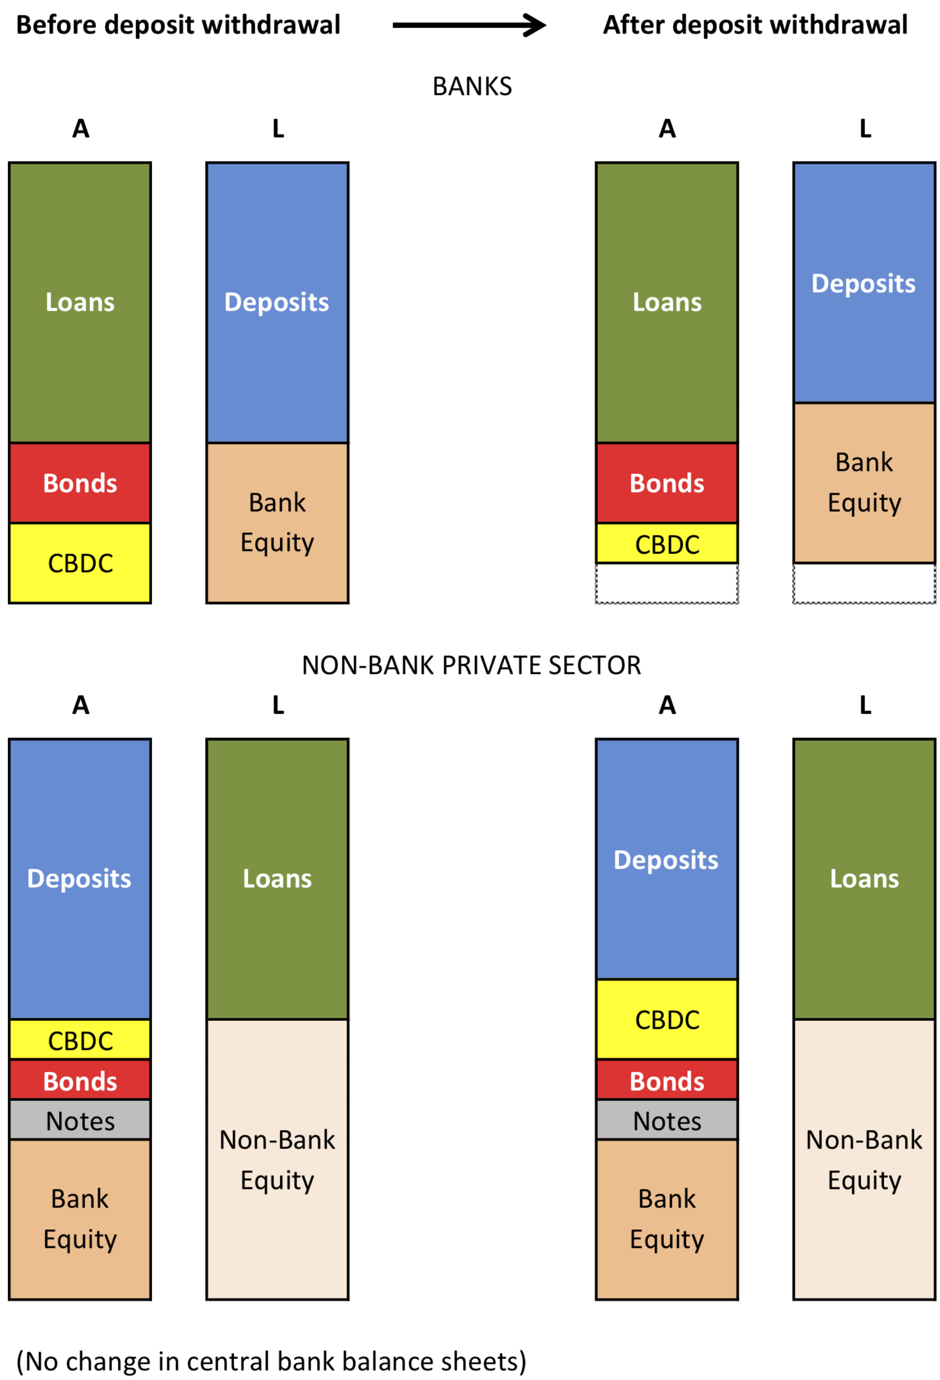

3.2.2 Risk of runs to CBDC

A related concern is the possibility of a very rapid substitution from deposits to CBDC — in other words, a bank run. In this extreme case, the contraction in bank funding and the withdrawal of liquidity (in the form of CBDC) may put the central bank in the position of having to replace the funding that commercial banks had lost.[14]

These risks have been highlighted by policymakers (Broadbent, 2016; Carney, 2018), and pose a genuine challenge to the feasibility of CBDC. However, we suggest that such risks can be managed. One approach would be to artificially introduce some of the frictions that currently discourage large-scale runs to cash, for example by introducing a notice period for large CBDC withdrawals, not paying interest on balances held above a given limit, or imposing fees on unusually large balances that could approximate the storage costs of cash. Alternatively, CBDC accounts might have a daily transfer limit. Such design features are within the control of the central bank, and would affect the attractiveness of CBDC, limiting the extent to which it was a substitute for bank deposits. A key challenge is to find the optimal trade-off between encouraging uptake of CBDC in order to have a monetary policy impact, and limiting the creation of new financial stability risks.

It is also not certain that run-risk would necessarily be increased under a CBDC. It is certainly true that the speed of a run could be quicker, and that the cost and frictions of running *could* be less. However, it is also possible that those who have the highest sensitivity to the (real or perceived) credit risk of bank deposits are those most likely to substitute from deposits to CBDC over a period of time after CBDC has first been introduced. Therefore, the marginal impact of a change in the perceived level of risk in the banking sector could conceivably be less, as the most 'flighty' depositors would have already substituted into a safe asset. Thus the probability of a run, for a given level of risk, may be lower when a safe outside option such as CBDC has already been provided.

Theoretically, run risks could also be mitigated by removing the requirement for banks to convert deposits to CBDC.[15] In our view, the ability of depositors to exchange commercial bank money for central bank money on demand is fundamental in maintaining the confidence in bank deposits, and many of the activities of the monetary authority (such as lender of last resort, liquidity regulations, and deposit insurance) are geared towards ensuring that this is always possible. In fact, many would consider it a necessary part of establishing a stable monetary framework in which CBDC and bank deposits coexisted and exchanged at parity (i.e. 1:1). Preventing depositors from withdrawing CBDC as a fundamental design feature of the system would therefore seem to be a significant and potentially risky departure from usual central bank practice.

4 The monetary transmission mechanism

We next look at how the introduction of a CBDC could affect the transmission of monetary policy decisions, that is to say, how a change in the policy instrument leads

[14]Initially this re-funding would likely be against collateral as part of its normal liquidity operations, minimising the risk to the central bank, although this would depend on the availability of eligible collateral for the banking sector and the willingness of the central bank to accept a wide range of collateral.

[15]Such a proposal is presented in Kumhof and Noone (2018).

to a change in the path of the real economy. In the analysis that follows we will discuss three broad stages of the monetary transmission mechanism. The first is the setting of the policy instrument - either the interest rate on, or quantity of, electronic central bank money in the secondary market. Second is the pass-through of changes in the price and interest rate on CBDC to the interest rates and prices of other assets in the economy. Lastly there is the pass-through from these financial market movements to the real economy. This final stage can itself be subdivided into a range of transmission channels including the real interest rate channel, the bank lending channel and the expectations/signalling channel, among others.

For the most part we will focus on the marginal impact of a change in a given policy instrument. CBDC would likely bring with it a host of changes to the steady state structure of the economy, such as changes to the equilibrium interest rate or steady state credit spreads. For a fuller analysis of such issues, we refer the reader to Barrdear and Kumhof (2016) and Bordo and Levin (2017b). We will also focus on the most common existing instruments of policy, namely (1) the short-term nominal interest rate and (2) quantitative easing. This provides a clearer lens through which to view the marginal change to the policy transmission mechanism that is purely a consequence of a universally accessible CBDC, rather than a discussion of more unconventional policies which would represent an innovation to the monetary framework in and of themselves.

4.1 MTM Stage 1

4.1.1 The overnight rate on central bank money

A fundamental element of monetary policy is the central bank's ability to set the overnight rate for central bank money. Operationally there are a number of ways in which this can be achieved, but in a general form it can be described by Figure 4. Prior to the financial crisis of 2008/09, most advanced economy central banks targeted the interest rate in the secondary market for central bank money (reserves) and set supply such that the market cleared at that rate.[16] In order to insulate this rate from large shocks to demand, a central bank could set a floor to the market, by offering to remunerate at least some balances at a deposit rate (which would be lower than the target rate). Similarly , a ceiling could be set by offering to lend central bank money perfectly elastically[17] at a rate above the target rate via a lending facility. This structure is known as a corridor system.

Since the financial crisis of 2008/09 many central banks have expanded the quantity of reserves significantly as part of their asset purchase programmes. This has led to a shift in monetary frameworks to a floor system, also in Figure 4. In a floor system the quantity of reserves are expanded such that the market clears at the rate paid on reserve balances and all reserve balances are remunerated at that rate. Thus the rate of relevance for the stance of policy is no longer the secondary market rate for central bank money, but rather the rate of interest on reserves. In fact, as the supply of reserves is now at least sufficient to essentially satiate demand for central bank money,

[16]Other mechanisms could also be used to incentivise the market to clear at that rate. For example, the Bank of England remunerated reserves at the policy rate and penalised banks who held reserves in quantities that differed on average from their voluntary targets over a given period, to anchor banks' reserves demand close to supply in the short-term.

[17]Subject to the central bank's risk tolerance.

Figure 4: Secondary market for central bank money under corridor and floor systems

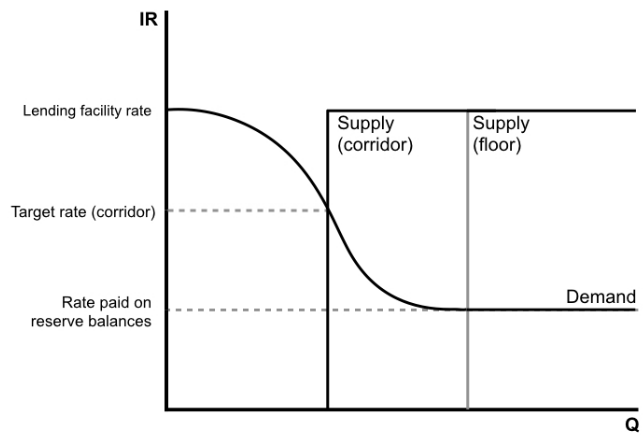

Figure 5: Secondary market for central bank money with a CBDC

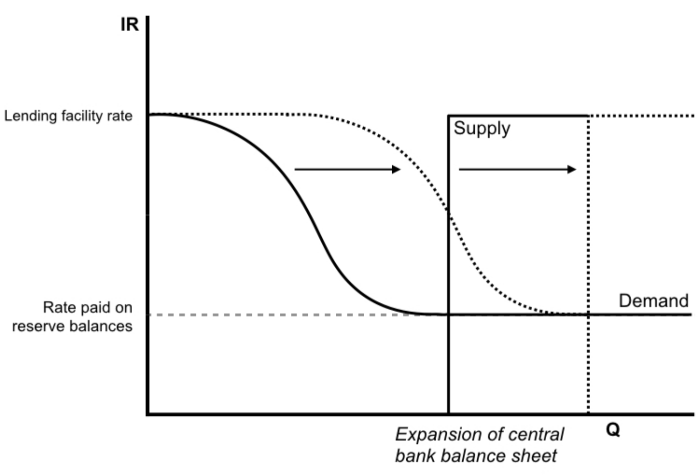

secondary market activity has notably diminished in frameworks where a floor system has been introduced. In practice it has been found that the floor created by paying interest on reserves is far from fully binding. This is largely attributed to the tiered nature of access to electronic central bank money.

How would this change with the existence of a universally-accessible CBDC? First, the demand curve for electronic central bank money would shift to the right, as in Figure 5, as there would be increased demand from non-banks who previously could not access CBDC. The extent of this rightward shift would depend on the attractiveness of CBDC and ultimately the whole range of design features it offered, as well as other determinants including whether a credible deposit guarantee scheme applied to bank deposits, and the balance between electronic and cash payments in the economy. The impact on interest rates would depend on the response of the central bank. If they accomodated the increase in demand by expanding supply sufficiently then they would be able to keep the prevailing interest rate at the floor. However, if they were to fail to fully accommodate the increase in demand, then the interest rate in the secondary market would rise in a fashion more akin to a corridor. If the central bank were to maintain the initial solid supply curve in Figure 5, not only would the interest rate rise, but there would also be a redistribution of CBDC within the economy. Banks would divest themselves of CBDC to the point of inflection in the original demand curve, something they are unable to do in the current framework. Non-banks would acquire these excess CBDC holdings from the banks to partially satiate their demand. Any increases in supply would lead to a larger central bank balance sheet, consistent with the balance sheet diagrams in Section 1.

Beyond this shift in the demand curve there would be a fundamental change in the nature of the overnight market for central bank money. Currently that market is an *interbank* market as only certain banks can participate. This means that there are a relatively small, homogenous set of agents trading for a relatively homogenous set of purposes, predominantly to smooth liquidity and payments shocks. With universally-accessible CBDC, banks could still trade central bank money with each other in this way, but this interbank lending would now only represent one possible segment of the market for CBDC. A bank requiring central bank money to meet liquidity needs at the end of the day could potentially borrow those funds from a non-bank as much as it could borrow them from a bank. This increased access to the market would be likely to increase both liquidity and competition.[18] Similarly, banks with excess CBDC at the end of the day would be able to lend it to a wider range of agents than they can currently, as they could now lend it outside of the banking sector, although demand for such short-term funding is likely to be limited in the real economy.

Non-banks would now have the option to hold funds at the central bank, in the form of CBDC, and earn the CBDC rate. This would increase the efficacy of the floor and make it more binding. If anyone can earn the central bank policy rate, then there is no incentive for them to put their money on deposit or lend it to someone else for less than the rate they can earn risk-free from the central bank.

Crucially, the central bank could control the rate on CBDC — the first stage of the monetary transmission mechanism — in much the same way as it does now. By

[18] In practice, it may not be feasible for banks to borrow the amount of funds they require at such short notice from multiple non-banks, and so they may stick to existing relationships with other banks, which may limit any impact of CBDC on the interbank market.

expanding the quantity of CBDC such that the market clears at the floor, the central bank could use the rate of interest paid on CBDC balances, and the expectation thereof, to guide rates in the rest of the economy. It could also vary the aggregate quantity of CBDC as an (operationally) independent instrument to stimulate the economy. Goodfriend (2002) shows how both the quantity and price of central bank money can be varied independently in a floor system based on the existing reserves framework.

It would also be possible for the central bank to operate a corridor system with a universal CBDC, by reducing the supply such that a more developed secondary market evolves that clears at a rate above the central bank deposit rate. This would, however, require more active management of the central bank's balance sheet, perhaps with significant and regular open market operations, and it would not allow for the independent adjustment of both the price and supply of CBDC.[19]

4.1.2 The term structure of CBDC

As well as being lent or deposited overnight, CBDC could also be lent or deposited at term. Currently, term lending of reserves is limited to relatively short horizons. However, the broader nature of the market for a universally accessible CBDC would result in a more varied range of motives for lending and borrowing CBDC, and thus it may be more likely that CBDC is lent at longer terms than traditional reserves, and for a wider range of purposes. These CBDC loans would give rise to a term structure on central bank money that was based on expectations of the overnight rate over the term of the loan, and a premium that compensated for characteristics such as default risk. Similarly, there would likely be a market for contracts that traded on the expectations of overnight rates, like overnight index swap rates and Fed Funds Futures do currently, for which the rate for CBDC in the secondary market could become the point of reference.

All of this would mean that by guiding market expectations for the future path of the policy rate, central banks could shift the term structure of CBDC, which would then directly affect the interest rates on any contracts specified in CBDC.

4.1.3 Non-CBDC rates and the wider term structure

As is currently the case, the process of arbitrage would mean that changes in the rates on CBDC would pass-through to rates on other assets in the economy. In Appendix A we present a stylised model that provides an intuitive formalisation of this process under a universally-accessible CBDC. In essence what one would expect to see is that the interest rate that clears the market for CBDC would be the theoretical risk-free rate, minus a premium derived from the non-pecuniary transactional utility (or convenience yield) provided by CBDC.[20]

$$R^C = R - \phi^C \tag{1}$$

[19] Advanced economy central banks have not yet specified whether they will maintain a floor system once their post-crisis asset purchases are unwound.

[20] It may also be a positive function of the cost of administering accounts, or the underlying payments system. ϕ^C could therefore be thought of as a more complex, composite premium of factors that drive a wedge between the risk-free rate and the rate on CBDC.

This transactional service can be motivated in a number of ways and need not be constant, but could vary through time and, as posited by Friedman (1960), it could be a function of the stock of CBDC itself.[21]

The rates of interest on other assets in the economy can then be written in the general form by

$$R^x = R^c + \phi^x \tag{2}$$

where ϕ^x is a premium for asset x that could be a function of a wide range of factors. These could include the relative probability of default compared with CBDC,[22] the relative transactional utility provided by the asset and the relative utility provided by the asset for liquidity management or regulatory purposes. For instance, government bonds would have little to no default risk, so that element of the premium would be low, but they also provide little to no transactional utility, which would cause a positive spread against the CBDC rate to occur. Deposits with commercial banks could provide transactional services similar to CBDC, and so would see an equivalent premium along that dimension, but do inherently contain risk of default, and so would see a spread occur over the CBDC rate.

An initial consequence of a universally-accessible CBDC is likely to be that deposit rates offered by commercial banks would be higher than the interest rate paid on CBDC. Prior to the 2008 financial crisis, deposit rates consistently cleared at below the policy rate (Figure 6). This could be explained by two phenomena. First, banks had a monopoly on the creation of *electronic* money that could be used to make transactions in the wider economy, and non-banks had no outside alternative. Banks could make use of this monopoly to lower the rate that they paid on deposits. In terms of equation 2 this meant that the transactional utility for non-banks of reserves was zero, while the transactional utility of commercial bank deposits was high, and so the latter rate naturally fell below the former. What is more, state backstops, implicit or otherwise, essentially gave the perception that bank deposits were risk-free and so there was no default premium built into deposit rates. Second, banks offered a bundling of services, such as overdrafts and preferential mortgage or lending rates, which provided a utility to holding deposits with the banking sector and which were reflected in a negative premium, lowering the rate at which the deposit market cleared. With a universally-accessible CBDC, central banks would be making available to non-banks an outside, competitive option for the provision of a means of payment (although unlike some bank deposits, this would not be bundled with other services such as overdrafts). With such an outside option, were a commercial bank to try and pay less than the policy rate on funds, the non-bank depositor could withdraw CBDC in return for a reduction in their deposit balances. This would constitute a large structural change for the banking sector which could have consequences for credit provision and banks' funding models. In Section 6 we will suggest an extension of the framework outlined here that could ameliorate any negative impact of such a change.

It may be argued that deposit guarantee schemes make bank deposits essentially risk-free, so that there is no significant difference in default risk between deposits and

[21] Bordo and Levin (2017a) develop the ideas of Friedman in the context of CBDC. They suggest that if $\phi^c = f(Q^c)$ and the creation of CBDC is costless then the supply of CBDC should be expanded to the point where $\phi^c = 0$ and the rate of return on CBDC equals the risk-free rate, $R = R^c$.

[22] We assume CBDC has zero default risk.

Figure 6: Sight deposit rates vs policy rates in the UK

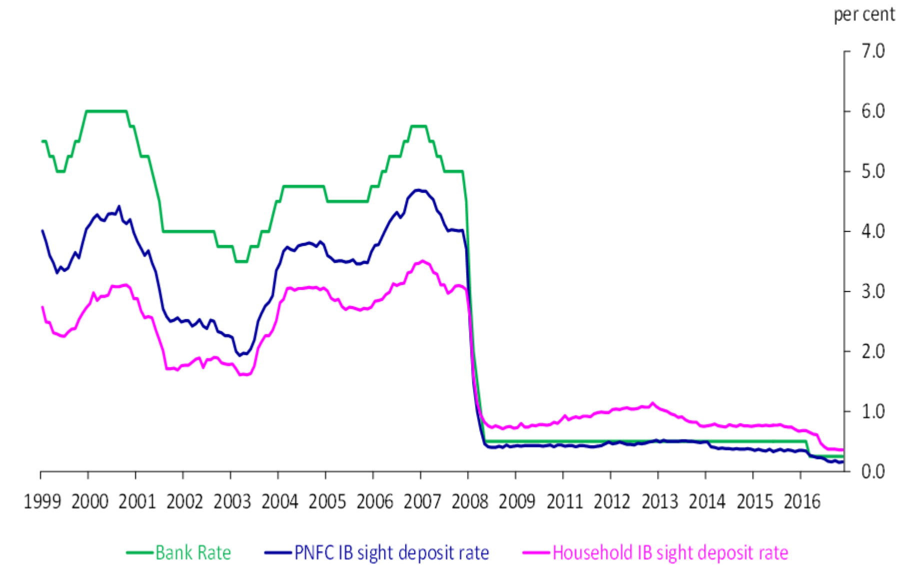

Source: Bank of England.

CBDC. However, coverage of deposit guarantee schemes is not exhaustive, especially for business account holders and accounts with large balances. What is more, even those that are covered face the risk of their money being frozen or tied up, even if only for a few days, while the process of unwinding a failing bank is undertaken and the deposit guarantee scheme enacted. Thus there are grounds to believe that deposits have some inherent risk, and they are likely seen this way by many depositors.

It is worth noting that the financial crisis of 2007-09 revealed the extent of the 'implicit state guarantee' that applied at the time even to bank liabilities which were not explicitly covered by deposit guarantees. A CBDC could potentially further reduce the importance of such implicit and explicit guarantees, by providing universal access to a truly risk-free asset that can be used to make payments and which does not depend on the underlying health of the banking system. Crucially, the two methods of providing a secure payments system — either providing risk-free CBDC, or trying to make deposits risk-free through guarantees — have different implications for the incentives of the banking sector. Deposit guarantee schemes may give rise to distortions of moral hazard, since banks can take risks without facing a commensurate increase in their funding costs. This may in turn lead to the over-allocation of credit due to low funding costs. Essentially, banks can obtain funding at a cost below a true reflection of the risk they take. Conversely, meeting the same need to provide a risk-free means of payment in the form of CBDC could act as a disciplining influence on commercial banks as they would now have to show themselves as safe to compete with the service offered by CBDC, or face higher costs of funding.

4.2 MTM Stage 2: transmission to financial markets

As well as potentially changing the structure of interest rates in the economy, there are a number of features of the model of CBDC outlined earlier that could affect the speed and extent of pass-through between changes in the policy rate and other rates. Most notably, the central bank is now offering an outside option to those who want to hold electronic money balances and who previously could only do so via deposits with a commercial bank. To explore the effect of this, we abstract from the steady state shift that might occur as a result of this outside option, and also from the transitional period as any new steady state is achieved. The existence of a competitive money alternative to bank deposits is likely to mean that if the interest rate on that alternative changes, but deposit rates do not move by an equal amount, then people may reallocate their portfolio to take advantage of the relatively higher return that has opened up. This would create flows between the two assets. If the policy rate which is paid on CBDC is increased then this could result in a fall in demand for bank deposits, while if the policy rate is cut, this could drive demand from CBDC into bank deposits. This will be particularly acute when it is easier to convert between CBDC and deposits. To the extent that pass-through from policy rates to deposit and wholesale rates has been estimated to be less than one, CBDC is likely to strengthen this stage of transmission.

We agree with Bordo and Levin (2017b) that deposit-taking institutions that engage in customer-focussed relationship banking are likely to be less vulnerable to the outflows of deposits than the areas of the market that compete purely on price terms. We also acknowledge that evidence for the UK shows that deposits are sticky — depositors do not often switch banks — and so have low price elasticity (Chiu & Hill, 2015). The extent to which this remains the case once the central bank offers an outside option, and in the face of structural changes that will make it easier to switch accounts (such as the EU's Second Payment Services Directive), remains to be seen.

Ultimately, whether there is a change in the speed of pass-through from changes in the policy rate to changes in deposit rates will depend on how banks react. We could assume that technology is likely to mean that depositors can more easily and costlessly move between deposits and CBDC, so banks would have to respond quickly to stem deposit flows. However, banks may react to the increased flightiness of deposits by changing their funding models to rely more on term funding, to 'lock in' deposits. This would mean that pass-through to rates paid on deposits would ultimately be slower.

The strength of pass-through to lending rates for commercial bank money would likely be affected along two avenues. First, larger changes in deposits and wholesale rates for a given change in the policy rate would also mean a larger impact on banks' funding costs, all else equal. For a given mark-up on these funding costs, this would have a larger impact on loan rates. Second, a universally accessible CBDC might also enable greater competition in the provision of credit. CBDC could allow non-banks to make loans more easily. For instance, peer-to-peer lenders would no longer have to clear settlements through their competitors in the banking sector, as is currently necessary in the existing system of tiered access to central bank money. This process incurs a cost which CBDC could potentially eliminate, putting non-bank credit providers on a more equal footing with their banking sector counterparts and would limit the extent to which banks could vary margins in light of changes in funding costs. This increased sensitivity of both funding costs and lending rates to changes in the policy rate could act to strengthen the *bank lending channel* which we discuss further in the next section.

4.3 MTM stage 3: transmission to the real economy

Being able to influence interest rates is only an intermediary step in the monetary transmission mechanism; the ultimate goal of changes in policy is to guide the real economy. The process by which this is achieved can be characterised by various channels, for which there is a wide range of taxonomies.[23]

A number of the channels by which the monetary transmission mechanism leads to the real economy begin with the change in interest rates we have discussed previously. The most obvious of these is the real interest rate channel, but the cash-flow channel and bank lending channel are also dependent on changes in either nominal or real interest rates.

Given that our previous analysis suggests that universally accessible CBDC is likely to increase the extent of changes in interest rates for a given change in the policy rate, all else equal, this will serve to amplify the strength of all of these channels, for a given change in the policy rate. Beyond this however, it is hard to pin down any structural reason why a CBDC would influence the strength of, for example, the real interest rate channel. At its heart, this channel is based on the deep parameters determining the intertemporal preferences and decision making of economic agents. These are unlikely to be influenced by whether or not agents are fulfilling their intertemporal allocation through a CBDC, bank deposits, or some other asset.[24] Similarly, the cash-flow channel is based on the marginal propensities to consume of various agents, which there is little reason to believe would change as a result of CBDC becoming universally accessible.

The process by which agents form expectations would also appear to be largely independent of the introduction or otherwise of a universally accessible CBDC. Therefore one may expect the expectations/signalling channel to be affected only insofar as CBDC adds or detracts from the clarity and credibility of central bank communications of the policy stance. This effect seems unlikely to be large, but if the pass-through from policy rates to wider rates were to become fuller then the potential for policy becoming constrained would be reduced, all else equal

The exchange rate channel could become more sensitive to a change in the policy rate, as international differentials in market interest rates would widen by more. For the logical argument behind this, think of an uncovered interest rate parity (UIP) condition in which the two rates of interest that matter are the bond rate in each currency. If a given change in the domestic policy rate leads to a greater change in the bond rate, then the exchange rate will have to move by more to clear the market consistent with UIP.

Perhaps the channel of transmission that would be most affected by the introduction of a universally accessible CBDC would be the bank lending channel. As discussed previously, the funding costs of banks would likely become more sensitive to changes in policy rates. This should strengthen the bank lending channel. What is more, the additional competition in credit provision may make pass-through to lending rates more complete. However, this picture has the potential to be complicated by a number of features of an economy with universally accessible CBDC. First, the fact that deposit

[23] The interested reader should see Mishkin (1995) which provides a broad coverage and forms the basis of the characterisation in this paper.

[24] The caveat to this may be that if CBDC is a truly risk-free asset, then this may lower the risk of transferring wealth between periods and increase the amount of intertemporal allocation that occurs, but this would predominantly be a level change rather than a change in the extent of transmission.

rates are now above policy rates could squeeze the net interest margins of the banking sector, which, among other things could result in lower profits; this would mean bank capital grows at a slower rate, constraining banks' ability to lend and therefore weakening the bank lending channel.[25]

Perhaps more fundamentally, were a CBDC to disintermediate the banking sector and significantly reduce the size of its aggregate balance sheet, as discussed in Section 3, this would reduce the importance and traction of the bank lending channel. This disintermediation is more likely when a CBDC is a close substitute for bank deposits that fully or partially dominates them in some aspects. This would be precisely when the benefits to the monetary transmission mechanism discussed above would be at their greatest. It is therefore apparent that a trade-off exists between the two forces CBDC would exert on transmission through the banking sector, which would need to be managed.

What is more, banks are crucially involved in money creation, a key part of the impact of the bank lending channel. Currently banks lend by issuing new deposits, in effect creating new money and purchasing power (McLeay et al., 2014; Jakab & Kumhof, 2015). In contrast, non-bank lenders transfer existing purchasing power (either deposits or CBDC) from savers to borrowers, but do not create any new purchasing power in the process. With the introduction of CBDC, banks can continue to lend by issuing deposits, but they could alternatively choose to 'lend CBDC'. In practical terms, this would require that they ask the borrower to nominate a CBDC account into which the lent funds could be transferred.

In practice, there are a number of reasons why banks would continue to prefer to lend by issuing new deposits. Firstly, there is unlikely to be demand for borrowing in CBDC specifically: as long as the sellers of goods, services or assets are able to substitute freely between deposits and CBDC (which we assume they can), then they should be neutral between receiving payments in CBDC or deposits, and should not offer any incentive for buyers to pay by one medium over the other. This means borrowers themselves should be neutral between borrowing CBDC and borrowing deposits — they will borrow from whichever lender offers them the best borrowing rate. Consequently, the interest rate on loans for a given level of risk and term should be the same whether it is CBDC or deposits that are borrowed. For this reason, banks are unlikely to receive requests to borrow in CBDC specifically. Secondly, for a bank, lending CBDC will have a much more negative impact on current regulatory ratios (specifically the Liquidity Coverage Ratio and Net Stable Funding Ratio) than lending via issuing deposits, because lending CBDC ensures that the bank will lose £100 of liquidity for every £100 lent. In contrast, while lending by issuing deposits could still lead to some outflow of CBDC, it is likely to be less, on average, than 100% of the amount lent.

All taken together, this analysis suggests that a universally accessible CBDC would most likely strengthen the impact of changes in the policy rate on the real economy, predominantly through increased pass-through from policy rates to other interest rates in the economy. To the extent that this were the case, a CBDC would imply that policy rates needed to vary by less over the cycle to stabilise the economy, conditioned on the same shocks afflicting the economy.

[25] We suggest a partial solution to this in Section 6 based on paying different rates of interest on CBDC held by banks and on CBDC held by other agents.

5 Quantitative Easing with CBDC

In recent years central banks have purchased assets fom the private sector and funded these purchases with newly created central bank money, a policy known as quantitative easing. However, as non-banks cannot currently hold electronic central bank money, and purchasing large amounts of assets by printing central bank notes is neither practical nor desirable, non-banks must use commercial banks as intermediaries to sell assets to the central bank. The commercial bank sells the asset to the central bank on behalf of the non-bank and receives an increased balance of electronic central bank money in its reserve account. It simultaneously generates a new deposit on the liability side of its own balance sheet, which it credits to the ultimate seller – the non-bank. Crucially, a universally accessible CBDC would remove the need for this intermediated interaction: QE could be carried out directly with non-bank participants. This process is shown in Figure 2. The central bank can purchase an asset from a non-bank and simply increase the balance on the seller's CBDC account. In this way, QE can be more targeted, as the central bank can choose to alter the balance sheet of the non-bank or banking sector independently, as it sees fit. It also has consequences for the onward transmission of QE.

One of the main theoretical transmission mechanisms of asset purchases is the portfolio rebalancing channel. This has its basis in work by Tobin (1969) and others, who show that if assets are imperfect substitutes for one another, and there are market frictions, then changes in the relative supplies of assets in the privately held portfolio will have consequences for prices and yields. Empirical literature on the recent asset purchase programmes of the Federal Reserve, Bank of England, European Central Bank and Bank of Japan shows significant incidences of this channel of transmission. See Bhattarai and Neely (2016) for a survey of the literature on the U.S. experience, or Joyce, Liu, and Tonks (2014) for the UK, as well as Meaning and Zhu (2011, 2012).

As stated above, a consequence of the current system is that central bank asset purchases from the non-bank private sector increase in the size of the banking sector's aggregate balance sheet, and change its composition, for instance average duration and liquidity ratios. This will lead banks to rebalance their portfolios. They can do this in two ways, each of which would have different consequences for the transmission of QE, and for the impact of a universally accessible CBDC. If banks react to the enforced additional holding of reserves by divesting themselves of similar, safe, liquid assets, they may sell government bonds back to the non-bank private sector. This would offset the reduction in the publicly-available supply of government bonds and counter some of the intended policy impact. In this case a universally accessible CBDC would strengthen the impact and monetary transmission of QE. Alternatively, banks could choose to rebalance their portfolios by increasing the size of their balance sheets by extending loans and thus arrive back at the same relative portfolio mix. Christensen and Krogstrup (2017) develop a theoretical model along these lines in which QE induces an unwanted change in the duration of banks' portfolios and leads to an amplification of the traditional portfolio rebalancing channel. Were this to be the dominant reaction of banks then a universally accessible CBDC would remove this secondary rebalancing effect and reduce the strength of QE. In the case of the UK, there is some empirical evidence that banks behaved more in the former fashion than the latter (Butt, Churm, & McMahon, 2015), and so we believe that the most likely case is that CBDC would lead to more effective QE.

6 Differentiated rates of interest on CBDC

CBDC as envisioned in the majority of this paper pays a single interest rate, regardless of the type of account holder. However, this need not be the case and it may even be beneficial to pay different rates on CBDC depending on who holds it. For instance, perhaps the most logical way to differentiate CBDC holdings would be between those held by banks[26] and those held by non-banks. This could be motivated by the special role that banks play in monetary transmission and the economy more widely. For instance, banks create money and purchasing power in the economy (McLeay et al., 2014), and we would expect them to continue to create the marginal unit of money even with a universally accessible CBDC. Paying differentiated rates would allow monetary policy to influence banks — and therefore credit and money creation — differently to non-banks.

One constraint faced by banks is that, due to regulation and as a result of the maturity transformation which they undertake, they must hold some fraction of their balance sheet in the form of liquid assets, such as central bank money (currently as reserves) or similar. There is an opportunity cost to holding these assets, as liquid assets usually pay a lower rate of return than other, less liquid alternatives such as loans. Banks could increase their interest income by minimising the amount of liquid assets they hold, but this increases the risk that they will run into liquidity problems in the future.

Paying a single rate on a universally-accessible CBDC may exacerbate this problem. Historically, the return on central bank money has at least been in excess of the rates banks have to pay on short-term liabilities such as sight deposits (see Figure 6). However, as discussed above, CBDC would mean that deposit rates were higher than the rate paid on CBDC, so banks would be making a net loss on all CBDC held. This could reduce their overall net interest margins and profitability, and strengthen the incentive to hold the minimum amount of central bank money.

Differentiating the interest rates paid to banks and non-banks offers a solution to this problem. If the rate paid on CBDC balances held by banks were greater than the rate paid to non-banks, then this would increase the return banks receive on the asset side of their balance sheet without necessarily increasing the cost of deposits on the liability side, thus lowering the overall cost of holding liquidity. From the depositor side, if the return on CBDC (the risk-free rate) has not changed, then the return they require to hold deposits, adjusting for risk and other factors would also not need to change.

The spread between the rates paid to banks and non-banks could be set as a positive, but fixed, level with the two rates moving together. In this case there would be a steady-state impact on banks' balance sheets, but the consequences for monetary policy beyond our benchmark case would be limited. Alternatively, the spread itself could be varied through time; this would have implications for monetary policy, and potentially financial stability. An increase in the spread could lower the cost of holding a given level of liquidity for a bank, which in turn would improve its balance sheet position. There would therefore be a monetary stimulus through the bank balance sheet channel, or through the bank lending channel. Even if banks passed some or all of the increase

[26] This would also include building societies and other firms that are classified as Monetary Financial Institutions (MFIs).

in the return on their assets to depositors in the form of higher deposit rates, this would be stimulative through a cash-flow channel.[27] The spread between the banking sector CBDC rate and the non-bank CBDC rate would therefore provide an extra dimension to the stance of monetary policy.

It should be noted that banks already benefit from a form of differentiation similar to that described here: they have access to the central bank's balance sheet on different terms to those of the general public. In fact, the difference is even starker under the current paradigm: not only do banks receive interest on their central bank money (reserves) while non-banks do not (on cash), but the functionality and utility of the two forms of central bank money is different. Although paying different rates to banks and non-banks may assist in implementing monetary policy, there is the potential that a system that explicitly pays more to banks than to the public at large could be objectionable due to concerns around equity and fairness, These debates may be especially acute in times of financial crisis.

6.1 Reserves and CBDC coexisting as distinct assets

In the discussion above, we have one electronic central bank liability, CBDC, which pays different rates depending on whether the holder is a bank or non-bank. An alternative scenario is to consider the introduction of a universally accessible CBDC *alongside* the existing system of restricted-access central bank reserves.[28] For example, in the model of Barrdear and Kumhof (2016) and Kumhof and Noone (2018), banks would be able to hold both reserves and CBDC, whilst non-banks would only be able to hold CBDC. The primary rate of monetary policy would be the rate paid on reserves, while the rate paid on CBDC would be used to control demand for CBDC relative to bank deposits.

At first glance this scenario seems to offer extra possibilities for monetary policy under a CBDC. However, under some sensible assumptions, we believe the outcome would be practically identical to the scenario above in which different rates are paid to banks and non-banks on a single liability in the form of CBDC.

To see why this is the case, let us begin from the point of two distinct electronic central bank liabilities — reserves and CBDC — following Barrdear and Kumhof (2016). Banks can hold both reserves and CBDC as assets, but non-banks can only hold CBDC. For our assertion that this would be equivalent to a single liability with differentiated rates, we require two assumptions to hold:

1. Both forms of electronic money offer the same, or highly similar, non-pecuniary utility. For instance, both provide the same transactional services to banks, meaning that it is just as easy for banks to clear balances between themselves with CBDC as with reserves. Similarly, both reserves and CBDC would be treated equivalently as regards regulations such as the liquidity coverage ratio.[29]

2. The central bank stands willing to convert any form of central bank money into another form, on demand, at no cost.

[27] Banks may choose to pass on the higher return in order to strategically gain market share, or because competition is such that if they didn't, another bank would.

[28] If the universally accessible CBDC were non-interest bearing, then it would constitute e-cash, which we discuss in Section 7.

[29] Existing capital and liquidity regulations do not tend to distinguish between different types of central bank liability, and so would treat reserves and CBDC as identical in calculating key regulatory ratios.

Under these conditions the only difference between the reserves and CBDC from a bank's point of view would be the rate of interest each paid. Banks would not need to simultaneously hold buffer stocks of both reserves and CBDC, since they could freely convert between them on demand, for example converting reserves to CBDC at the point at which depositors request to withdraw CBDC. This means banks would exclusively hold whichever of the two assets paid them the highest rate of interest. If the rate paid on reserves is greater than the rate paid on CBDC, this results in banks only holding reserves (at the higher rate) and non-banks only holding CBDC. This outcome is practically identical to the scenario above in which banks are paid one rate on the CBDC they hold, whilst non-banks pay a different rate. If the rate paid on reserves is *lower* than the rate paid on CBDC, banks and non-banks alike will only hold CBDC, which is equivalent to our earlier scenario in which a single rate is paid to banks and non-banks alike.

Consequently, in this dual asset regime, assuming free convertibility for different types of central bank money, then for a non-zero quantity of reserves to exist, the rate on reserves must be greater than the rate on CBDC. This limitation would not exist in a framework of differentiated rates on a single CBDC, in which banks receive the rate determined by the central bank, and have no option to earn the rate paid to non-banks.

Under what circumstances would the two assumptions above not hold? One plausible possibility is if reserves and CBDC had different functionality. For instance, in the UK, reserves are currently the accepted means of settlement for most interbank and national payment systems, and we assume they would continue to serve that purpose. However, perhaps CBDC is introduced as a cryptographic token on a distributed ledger platform with a specific usage, such as settling securities trades in specific financial markets. In this case, banks would have an incentive to hold both reserves *and* CBDC. However, if the rate on reserves were higher than the rate on CBDC, banks would still prefer to hold the minimum possible stock of CBDC, and convert reserves to CBDC at the point at which they are needed for settlement of securities. Consequently, the second assumption, that reserves and CBDC are freely convertible at the central bank, must also be false in order to make it worthwhile considering CBDC and reserves as two distinct assets. Kumhof and Noone (2018) explore scenarios where neither of these assumptions hold.

7 E-cash

Although the primary focus of this paper is a single, universally accessible CBDC used for monetary policy purposes, an alternative formulation is to introduce multiple forms of electronic central bank money and use them for different objectives. The simplest of these is the introduction of a digital, non-interest bearing, 'e-cash' alongside central bank reserves. To be clear, such a framework amounts to the central bank issuing two forms of CBDC simultaneously and differentiating them in terms of access. In this world, reserves (the first form of CBDC) would continue to function as they currently do, being used to settle between banks but could not be used to pay for goods, services and assets in the wider economy. They would also continue to be at the heart of setting monetary policy (Boel, 2016). E-cash (the second form of CBDC) would not be used in setting monetary policy, but rather as a means of establishing an efficient and safe payments system.

Consideration needs to be given to how the two assets would interact. As we have noted before, it is common practice for central banks to allow free convertibility, at par, between their different types of liability (e.g. £10 of reserves for £10 of cash). It seems unlikely that they would break from this practice if they were to introduce a third type of liability — CBDC — to their balance sheet. Therefore, an e-cash CBDC would need to be provided perfectly elastically to meet demand at an interest rate that is fixed at zero. This means that agents must be able to freely substitute between cash, CBDC, and reserves (for banks). In this set-up the central bank can set the overall size of the monetary base, but not directly the composition of liabilities within it. This is a central feature of the par-convertibility of central bank liabilities even now: the central bank can choose to inject reserves into the banking system, but if there is a surge in demand for central bank notes then they must allow the stock of notes to increase while the stock of reserves falls, or else inject additional reserves, expanding the monetary base. The alternative is that the increased demand for notes goes unmet and the price of a note changes relative to the price of reserves, breaking parity between different types of central bank money and undermining monetary stability.

The demand for e-cash is likely to be a function of, inter alia, other interest rates in the economy as changes in the rate paid on other assets change the relative return on e-cash. These shifts in demand could drive flows between deposits and e-cash. In this way demand for e-cash is likely to be counter-cyclical to policy rates. In the case of a tightening of monetary policy, an increase in the rate paid on reserves would lead to an increase in rates paid on bank deposits and thus a decrease in the relative return of e-cash. This would make e-cash less attractive to hold and lead to a substitution from CBDC into bank deposits. In contrast, a monetary policy expansion would have the opposite effect, making e-cash relatively more attractive and leading to a substitution from deposits into CBDC. These flows in and out of an e-cash CBDC could be a source of instability in the banking sector. They could also affect the level of liquidity on banks' balance sheets and the cost/availability of bank funding, which may in turn act as an unintentional dampener of the desired change in policy, loosening credit conditions, all else equal, when policy is aiming to tighten, and vice versa.

8 Outstanding questions and future research

The literature on CBDC is growing, but this is still a new topic of investigation, and there are a wide range of questions that still need to be explored.

This paper has investigated the possible benefits to monetary policy of CBDC, but also highlighted some of the dangers CBDC may pose for financial stability. Better understanding the determinants and the nature of this trade-off is crucial. For instance it will have implications for the optimal design of CBDC. Should a CBDC be a very close substitute for commercial bank deposits then that may maximise the improvement in monetary transmission but, all else equal, it also increases the risk of runs from the banking sector. Could many of the benefits of CBDC be achieved by increasing access to electronic central bank money to the wholesale market, without the need to extend it to retail customers? Even abstracting from bank runs, at what point do the benefits of a new competitive force for the banking sector get outweighed by the negative consequences of the central bank disintermediating a large part of banks business models?

Relatedly, it is important to find ways to mitigate against any financial stability risks of CBDC. Critically, research should explore how CBDC can be introduced in an orderly manner, without having a destabilising impact on the banking sector. Mechanisms need to be designed to allow interoperability between deposits and CBDC whilst creating incentives to discourage a run from deposits to CBDC in a panic scenario. If the financial stability risks of CBDC can be reduced, this may tip the cost-benefit balance in favour of CBDC, from the point of view of central banks.

Another area in which progress might tip the balance on the cost-benefit analysis of CBDC is technology. Questions on this topic are far from resolved, but, currently, many consider that distributed ledger technology is too immature and energy intensive to represent a viable replacement to existing large-scale payments systems. However, developments in this field are moving at such a rapid pace that what may currently seem infeasible may be perfectly possible, or even obsolete, by the time a widely accessible CBDC is introduced.

As well as influencing the functioning of the existing monetary toolkit, CBDC also has the potential to enable more significant change in the toolkit itself. This paper has deliberately considered the impact of CBDC on widely-used monetary policy, namely changes in the central bank's policy rate, and asset purchases (quantitative easing). However, CBDC has also been discussed in the context of less conventional monetary policy, such as the use of negative rates (Kimball & Agarwal, 2015) or direct distributions of newly issued CBDC to citizens — so-called 'helicopter money' (Turner, 2015). These policies do not necessarily require a CBDC to be implemented, but the existence of CBDC may affect their feasibility and impact. It may even be that CBDC and future technological progress give rise to monetary instruments that have not yet been considered.

Spanning all of these areas is the need for robust quantitative insights into CBDC. Obviously this is complicated by a lack of data and a lack of good proxies for CBDC — a problem that would be resolved with time once a CBDC had been introduced, but that is hard to deal with ex-ante. High on the list of empirical questions that will need to be answered are: What would be the size of demand for a CBDC? How volatile would this demand be? What would be the interest rate elasticity of substitution between CBDC and bank deposits? And what would CBDC do to steady state of interest rates and credit spreads?

9 Concluding remarks

At its simplest, a central bank digital currency can be thought of as electronic narrow money and so in many ways it should feel familiar to economists and policymakers. Within this general definition there exist a number of dimensions along which any specific CBDC could be varied: access, whether it is interest-bearing, the objective it is designed to achieve, and the underlying technology on which it is based, to name a few. Careful consideration will need to be given to these parameters both in future research and also in order for a central bank to effectively introduce CBDC.

The main conclusion of this paper is that under a universally accessible, account-based CBDC, monetary policy could operate in much the same way as it currently does, guiding the economy through varying the rate of interest paid on balances of electronic central bank money and the aggregate quantity of that money. The untested nature of such a CBDC means that the impact on the monetary transmission mechanism is uncertain, but we believe the most likely consequence would be that CBDC would strengthen the monetary transmission mechanism, for a given change in policy instruments.

This paper is intended as an early step in the development of the literature on central bank digital currencies, and many fundamental questions remain unanswered. These relate to, among other things, the impact on the wider financial sector, the implications for financial stability, steady state changes in the economy resulting from CBDC, and how CBDC would affect the balance sheet management of central banks. Significant work, both theoretical and eventually empirical, will be required to inform any policy decision to introduce CBDC.

References

Bank of Canada. (2017). *Digital Currencies and Fintech.* Retrieved from `www.bankofcanada.ca/research/digital-currencies-and-fintech/`

Bank of England. (2017a). *A blueprint for a new RTGS service for the United Kingdom.*

Bank of England. (2017b). *Digital currencies.* Retrieved from `http://www.bankofengland.co.uk/research/Pages/onebank/cbdc.aspx`

Barrdear, J., & Kumhof, M. (2016). *Staff Working Paper No. 605 The macroeconomics of central bank issued digital currencies.* Bank of England.

Bech, M. L., & Garratt, R. (2017). *Central Bank Cryptocurrencies.* Bank for International Settlements.

Benos, E., Garratt, R., & Gurrola-Perez, P. (2017). The Economics of Distributed Ledger Technology for Securities Settlement. *Bank of England Staff Working Paper No 670.*

Bhattarai, S., & Neely, C. J. (2016). *A Survey of the Empirical Literature on US Unconventional Monetary Policy.* Federal Reserve Bank of St Louis.

Bjerg, O. (2017). *Designing New Money: the policy trilemma of central bank digital currency.* Copenhagen Business School.

Boel, P. (2016). *Thinking about the Future of Money and Potential Implications for Central Banks.* Sveriges Riksbank.

Bordo, M. D., & Levin, A. T. (2017a). *Central Bank Digital Currency and the Future of Monetary Policy.* Hoover Institution.

Bordo, M. D., & Levin, A. T. (2017b). *Central Bank Digital Currency and the Future of Monetary Policy.* National Bureau of Economic Research.

Broadbent, B. (2016). *Central banks and digital currencies (speech).* Bank of England.

Butt, N., Churm, R., & McMahon, M. (2015). *Did Quantitative Easing boost bank lending?* Bank Underground.

Carney, M. (2018). *The future of money.* Speech to the Inaugual Scottish Economics Conference.

Chiu, J., & Hill, J. (2015). *Staff Working Paper No. 540 The rate elasticity of retail deposits in the United Kingdom: a macroeconomic investigation.* Bank of England.

Christensen, J., & Krogstrup, S. (2017). A Portfolio Model of Quantiative Easing. *Presented at the European Central Bank Workshop on money markets, monetary policy implementation and central bank balance sheets*.

Danezis, G., & Meiklejohn, S. (2016). *Centrally Banked Cryptocurrencies.* University College London.

Dyson, B., & Hodgson, G. (2016). *Digital Cash: Why Central Banks Should Start Issuing Electronic Money.* Positive Money.

FirstRand Bank. (2016). *The Advent of Crypto Banking: A New Paradigm for Central and Commercial Banking.* FirstRand Bank.

Friedman, M. (1960). *A Program for Monetary Stability.* Fordham University Press.

Fung, B. S., & Halaburda, H. (2016). Central Bank Digital Currencies: A Framework for Assessing Why and How. *Bank of Canada Staff Discussion Paper*.

Gagnon, J. E., & Sack, B. (2014). *Monetary Policy with Abundant Liquidity: A New Operating Framework for the Fed.* Peterson Institute for International Economics.

Goodfriend, M. (2002). Interest on Reserves and Monetary Policy. *Federal Reserve Bank of New York Economic Policy Review*, *8*(1), 77–84.

Goodfriend, M. (2016). *The Case for Unencumbering Interest Rate Policy at the Zero Bound* (Vol. 26).

Goodfriend, M., & McCallum, B. T. (2007). Banking and Interest Rates in Monetary Policy Analysis: a quantitative exploration. *Journal of Monetary Economics*, *54*(5), 1480–1507.

Gouveia, O., Dos Santos, E., de Lis, S. F., Neut, A., Sebastián, J., et al. (2017). *Central Bank Digital Currencies: assessing implementation possibilities and impacts.* BBVA.

He, D., Leckow, R., Haksar, V., Mancini-Griffoli, T., Jenkinson, N., Kashima, M., . . . Tourpe, H. (2017). Fintech and Financial Services: Initial Considerations. *International Monetary Fund Staff Discussion Note*, *17*(05).

Jakab, Z., & Kumhof, M. (2015). *Staff Working Paper No. 529 Banks are not intermediaries of loanable funds - and why this matters.* Bank of England.

Joyce, M. A., Liu, Z., & Tonks, I. (2014). *Institutional Investor Portfolio Allocation, Quantitative Easing and the Global Financial Crisis.*

Ketterer, J. A., & Andrade, G. (2016). *Digital Central Bank Money and the Unbundling of the Banking Function.* Inter-American Development Bank.

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

Koning, J. (2014). *Fedcoin.* `https://jpkoning.blogspot.com/2014/10/fedcoin.html`.

Kumhof, M., & Noone, C. (2018). *Central Bank Digital Currencies design principles and balance sheet implications.* Bank of England.

McLeay, M., Radia, A., & Thomas, R. (2014). Money Creation in the Modern Economy. *Bank of England Quarterly Bulletin Q1*.
Meaning, J., & Zhu, F. (2011). The Impact of Recent Central Bank Asset Purchase Programmes.
Meaning, J., & Zhu, F. (2012). *The Impact of Federal Reserve Asset Purchase Programmes: another twist.*
Mersch, Y. (2017). Digital Base Money: an assessment from the ECBs perspective. *Speech at the Farewell ceremony for Pentti Hakkarainen, Deputy Governor of Suomen Pankki–Finlands Bank. Helsinki*, *16*.
Mishkin, F. S. (1995). Symposium on the Monetary Transmission Mechanism. *The Journal of Economic Perspectives*, *9*(4), 3–10.
Monetary Authority of Singapore. (2017). *Project Ubin Phase 2.* Monetary Authority of Singapore.
Motamedi, S. (2014). *Will Bitcoin Ever Become Money? A path to decentralized central banking.* http://tannutuva.org/blog/2014/7/21/will-bitcoins-ever-become-money-a-path-to-decentralized-central-banking.
Qian, Y. (2017). Digital Currency and Central Bank Bank Accounts [translated]. *Tsinghua Financial Review*.
Rogoff, K. (2016). *The Curse of Cash.* Princeton.
Scorer, S. (2017). *Central Bank Digital Currency: DLT, or not DLT? That is the question.* Bank Underground blog https://bankunderground.co.uk/2017/06/05/central-bank-digital-currency-dlt-or-not-dlt-that-is-the-question/.
Svierges Riksbank. (2017). *Does Sweden need the e-krona?* Retrieved from www.riksbank.se/en/Financial-stability/Payments/Does-Sweden-need-the-e-krona/
Tobin, J. (1969). A General Equilibrium Approach to Monetary Theory. *Journal of money, credit and banking*, *1*(1), 15–29.
Turner, A. (2015). *The case for monetary finance: an essentially political issue.* IMF 16th Jacques Polak Annual Research Conference.

Appendix A: The structure of interest rates and arbitrage conditions

To aid thinking about the structure of interest rates in an economy with a universally-accessible CBDC we present a stylised model of the key arbitrage conditions that might prevail. While this model is partial rather than general equilibrium, and appeals to sensible assumptions rather than strict microfoundations to motivate the range of premia, it still serves as a useful expositional tool and offers some initial, intuitive insights.

We begin from a theoretical risk-free interest rate, R, that represents the return on a pure store of value asset with no associated premia. There is no risk of default, no illiquidity and no term. CBDC is assumed to be risk-free in the same way, but alongside its store of value function it also provides an additional service as a means of exchange, for instance, lowering transaction costs, ϕ^C. The total expected return from a unit of CBDC is therefore[30]

$$R^C + \phi^C \tag{3}$$

and no-arbitrage would imply

$$R = R^C + \phi^C \tag{4}$$

meaning that

$$R^C = R - \phi^C \tag{5}$$

This means that CBDC would clear at a rate below the theoretical risk-free rate by a spread determined by the transactional utility supplied by CBDC. In the exposition here we assume this transactional utility is fixed per unit of CBDC, independent of the quantity of CBDC. This is purely for clarity of presentation and a credible alternative assumption is that the degree of transactional utility is a function of the quantity of CBDC held. Were transactional utility to be a negative function of quantity, but to a decreasing extent (the implied function has a negative first derivative and positive second derivative) then Friedman (1960) argues that, as the creation of central bank money is costless, the supply should be expanded to the point where ϕ^C is zero and $R^C = R$. This would imply that the rate paid on CBDC could be considered a reflection of the true risk-free rate.

Building from CBDC to a wider array of assets we look at each sector of the economy in turn and work through the no-arbitrage conditions that their balance sheets would imply.

Beginning with the non-bank private sector, consistent with our stylised balance sheets (see Figure 1), they can hold their wealth as a combination of three assets: CBDC (denoted by C), bank deposits (D) and government bonds (B).[31] For simplicity of exposition we assume that each of these assets is one period in term, but can be differentiated by other characteristics. As discussed above, each unit of CBDC held provides a non-pecuniary benefit to the holder, ϕ^C, for instance as a result of lowering transactional costs. Similarly, the interest rate paid on bank deposits is R^D and bank deposits offer a similar but not necessarily equivalent non-pecuniary return emanating from its role as a means of exchange, ϕ^D. However, unlike CBDC, there is a probability, γ, that banks will default on their deposits, in which case the depositor gets neither the pecuniary return nor the non-pecuniary benefit. Lastly, government bonds offer an interest rate R^B. They are assumed to offer no transactional services, but are defaultable with probability δ. Taken all together this means that the non-bank's end of period objective function can be written as

$$R^C C + \phi^C C + (1 - \gamma)[R^D + \phi^D]D + (1 - \delta)R^B B \tag{6}$$

And so is maximised where

$$\frac{dU}{dC} = R^C + \phi^C = \tag{7}$$

[30] With no default yet in the model, the expected returns are equal to the agreed returns.

[31] We abstract from equity without any loss of insight for the themes with which we are concerned.

$$\frac{dU}{dD} = (1-\gamma)[R^D + \phi^D] = \tag{8}$$

$$\frac{dU}{dB} = (1-\delta)R^B = \tag{9}$$

For the non-bank sector then, assuming that the rate on CBDC is set by the central bank

$$R^D = \frac{R^C + \phi^C - (1-\gamma)\phi^D}{(1-\gamma)} \tag{10}$$

and

$$R^B = \frac{R^C + \phi^C}{(1-\delta)} \tag{11}$$

This gives rise to two spreads. The spread of the deposit rate over the CBDC rate is a positive function of the relative transactional services of CBDC compared with deposits, and a positive function of the default rate. The spread of bond rates over the CBDC rate is a positive function of the transactional service of CBDC money, and a negative function of the default risk in the government bond. This all occurs in a one period setting. In practice there is likely to be another significant premium driving a wedge between the two rates, which is the term premium. This could be derived in a multi-period model by the additional risk of locking funds away when you are subject to unknown payments or liquidity shocks, and would appear as a positive function of the term of the bond.[32] As discussed previously, were we to assume that the transactional utility of CBDC were a decreasing and concave function of the quantity of CBDC, then the supply can, and arguably should, be expanded to the point at which $\phi^C = 0$. This would mean that the only differences between government bonds and CBDC were default inherent in government bonds and term. For short-term government bonds in stable economies both of these elements could be expected to be negligible, and so we would expect the short-term government bond rate to be extremely close to the interest rate on CBDC. As noted by Bordo and Levin (2017b), were the central bank to engage in open market operations between treasury bills and CBDC, they could ensure that this would be the case in practice.

We follow the same process for the banking sector. Banks can hold assets in the form of CBDC (C), loans (L) or government bonds (B). Again, this is consistent with our balance sheet diagram. As with non-banks, the banking sector receives both a pecuniary return on its CBDC holdings, R^C, and a non-pecuniary benefit from CBDC's transactional services, ϕ^B. Unlike non-banks, they receive an additional non-pecuniary benefit, η, from CBDC through its role as a High Quality Liquid Asset (HQLA). This could be due to a regulatory need to hold HQLA, or a portfolio preference of the bank itself. In a system of mandated reserve requirements, this benefit could be very significant. Government bonds also provide a benefit as HQLAs, but provide no transactional services, and default with probability δ.

[32] As another point of reference, bond rates will differ from the theoretical risk-free rate discussed above to the extent of inherent default risk only. In practice there would also be other premia, such as term, which would mean that the bond rate clears at a spread above the risk-free rate.

The last asset that banks can hold on their balance sheet is loans. The pecuniary return is R^L with a default probability of μ. We assume a cost of producing and monitoring each loan, M. For a more developed model which includes monitoring costs of this type, see Goodfriend and McCallum (2007). We further assume loans offer no transactional services, nor liquidity services. Lastly, banks must finance the asset side of their balance sheet with liabilities, meaning that they must pay R^D on all deposits owed to the non-bank sector.

Taken together we can write the bank's optimisation problem as

$$R^C C + \phi^B C + \eta C + (1-\delta)[R^B + \epsilon]B + (1-\mu)R^L L - ML - R^D[C+B+L](1-\gamma) \quad (12)$$

the first order conditions of which give

$$\frac{dU}{dC} = R^C + \phi^B + \eta - (1-\gamma)R^D = 0 \quad (13)$$

$$\frac{dU}{dB} = (1-\delta)R^B + (1-\delta)\epsilon - (1-\gamma)R^D = 0 \quad (14)$$

$$\frac{dU}{dL} = (1-\mu)R^L - M - (1-\gamma)R^D = 0 \quad (15)$$

and which optimises to give

$$R^B = \frac{R^C + \phi^B + \eta - (1-\delta)\epsilon}{(1-\delta)} \quad (16)$$

$$R^L = \frac{R^C + \phi^B + \eta + M}{(1-\mu)} \quad (17)$$

This shows that from the viewpoint of the banking sector, the spread of bond rates over the CBDC rate is a positive function of the transactional benefit of CBDC, of the default risk in bonds, and the relative benefits of CBDC as a HQLA when compared with bonds. When combined with the non-bank condition for bonds, this implies that the relative transactional services received by banks compared with non-banks must equal the additional benefit that banks receive from holding bonds as HQLA. The loan rate spread is a positive function of the cost of producing a loan and the probability of default.