Clouds, chips or chains? The three ways to implement a digital dollar: FedPesa, FedDex or FedCoin

Received: 23rd October, 2020

David G.W. Birch
Principal, 15Mb Ltd., UK

David G.W Birch *is Principal at 15Mb Ltd, Global Ambassador for Consult Hyperion and Technology Fellow at the Centre for the Study of Financial Innovation. He is an internationally-recognised thought leader in digital identity and digital money, and was named one of the global top 15 sources of business information by* Wired *magazine. His most recent book, 'The Currency Cold War', was published in May 2020. Dave graduated from the University of Southampton with a BSc (Hons) in physics.*

Abstract

This paper explores the 'evolutionary tree' of electronic cash using a series of case studies; presents a taxonomy to facilitate a discussion of the strategic options available for central bank digital currencies; and highlights the three practical alternatives for implementing such a currency: as balances maintained offline in tamper-resistant hardware ('FedDex'); as balances maintained online in a database ('FedPesa'); or as tokens managed by an online shared ledger ('FedCoin'). The paper concludes with the observation that a FedCoin solution may offer the most opportunities for innovation and economic growth. The idea of some form of FedCoin as a platform not just for smart money but very smart money is appealing in the central bank context as a way to implement the population-scale electronic cash system required to deliver a national digital currency. This is important because the renewed focus of central banks on the potential for central bank digital currencies should highlight the potential for electronic cash as a platform for new opportunities and a new generation of financial services, not simply a more efficient means to implement domestic retail payments. There is some urgency to this, as the beta testing of a Chinese digital currency is already underway.

Keywords: digital currency, electronic cash, evolution, digital money

INTRODUCTION

When people talk about a digital currency, they generally mean something more interesting than just another form of electronic money (e-money). After all, we already have e-money, and lots of it. In developed markets almost all money is already e-money and only a small fraction is in the form of notes and coins. If there is a point to creating digital currency, then it must be as a potential substitute for those notes and coins. This would imply that digital currency, then, is not a form of e-money but a form of electronic cash (e-cash).

The crucial distinction between e-money and e-cash is that if I want to pay you for something, I can send e-money to your bank account; people cannot use PingIt, Swish or Venmo to put money into someone else's pocket, only into an account somewhere. But e-cash is just like physical cash: if I send it to you, it becomes yours. You can put it into a bank account if you want to, but that is up to you. If you want to leave it in your smartphone, or your television or your car, then that is your choice. To simplify, then, there is a difference between e-money that is stored in *accounts* and e-cash that is stored in *wallets*.

Visa and MasterCard, Amex and Discover, Unionpay and others have delivered

David G.W. Birch

15Mb Ltd.
1 Armadale Road, Woking,
Surrey GU21 3LB,
UK
E-mail: dave@15mb.ltd

Journal of Payments Strategy & Systems
Vol. 14, No. 4 2020, pp. 339–352
© Henry Stewart Publications,
1750-1806

convenient e-money to everyone on Earth (well, everyone on Earth with a bank account plus some prepaid cards around the edge) and they have been astonishingly good at it. When I get off a plane in Sydney or Samarkand, I expect to stroll into my hotel and pay with my chip-and-PIN Amex card. This is a frankly astonishing achievement. However (to steal a phrase from Facebook CEO Mark Zuckerberg), as it stands, neither consumers nor businesses can send currency to their counterparties across the globe with the same ease as they might send them a photo. For that to happen, we need the e-cash that is not yet available but which could be provided most obviously either as a private claim against other currencies or as a public good in the form of a central bank digital currency (CBDC).

The Bank of England assumes that any CBDC will be introduced alongside, rather than as an immediate replacement for, both e-money and cash,[1] but developing the narrative from there on is complicated by the fact that despite the paucity of useful taxonomies for CBDCs, there would appear to be no useful shared taxonomy of implementation options.

One novel way to develop an implementation taxonomy is to consider what the 'genes' of a digital currency might be and then look at how these have 'mutated' over the years under the pressure of natural selection from the payment marketplace. This gives a way of understanding an evolutionary tree of electronic cash that serves to illustrate the recent history of the sector, provides for some interesting speculation on where it might develop further, and facilitates useful conversation between digital currency stakeholders.

To explore and experiment with this approach, take two genes, which one might broadly label *authentication* and *authenticity*. This very structure is set out in 'The Currency Cold War', which notes that these genes address the related but separate problems of how a value transfer system determines whether the payer is entitled to transfer value to a payee (authentication) and whether the value that they are sending is real or fake (authenticity).[2]

This evolutionary approach is explored in the present paper, leading to a discussion of the three most practical options for a CBDC. These are set out along with personal views on how the options can illuminate current digital currency discussions. The paper concludes with a suggestion as to how central banks in developed countries might bind their technology choices to the benefit of not only the financial sector but also the wider economy.

THE E-CASH GENES

To understand how e-cash will be used to create a digital currency that is likely to work, it is useful to step back and look at what one might call the evolutionary tree — not in exhausting detail, but in order to see why one might come to some conclusion about how a digital currency can be implemented. To do this, one must first identify the fundamentals that are responsible for the evolution of e-cash. These rest on the observation that creating a working means of exchange means solving two essential problems, as shown in Table 1. If you are going to transfer value to me, I need to know that the value is yours to give and I need to know that the value is real. Each of these involves solving different but related problems.

One can think of these two problems as the 'genes' and then look at evolution of e-cash as the story of new technologies coming along to mutate solutions to problems under the pressures of natural selection in the payment marketplace.

The authentication problem

The authentication problem comes down to proving the ownership of a cryptographic key. This is something we do whenever we

Table 1: The digital currency problem

	The identity issue	*The value issue*
What is the problem?	The authentication problem: How do I know that it's your money?	The authenticity problem: How do I know that the money is real?
What new technology can help?	Devices, biometrics, artificial intelligence	Clouds, chips, chains

wave a smartphone over a reader to get on a bus, punch in a PIN after inserting a bank card in a shop terminal or type in a verification code (when buying something online). Each of these works to different degrees of certainty. Access to a device, remembering a PIN, possession of a card, having a face, knowledge of the verification code — all of these are examples of authentication 'factors'. Generally speaking, people are happy with their counterparties demonstrating one factor for small transactions, two factors for medium-sized transactions and three or more for larger transactions.

This paper considers the authentication problem to be solved and will not dwell on it further. While the authentication journey for consumers is not yet optimised (to say the least), the appropriate technologies and standards are progressing, although many banks still rely on their own platforms for the time being. This is unsatisfactory, as may be illustrated using the recent example of when I called my bank to enquire about a new service — not to order anything, just to ask a question about business bank accounts. Now, as is normal when one calls a bank that one has been with for many years, the bank first requested authentication based on a selection of publicly available information (eg date of birth and mother's maiden name) before asking a series of questions, the answers to which they knew but I had forgotten. This seems an odd way to authenticate a customer who has the bank's app installed on their phone.

The bank app on my iPhone not only knows who I am and where I am, it also knows what I have been doing. Thus, the combination of the bank and the mobile operator really ought to deliver something special. For example, because of continuous passive authentication — the software running in the mobile phone that checks how I hold the phone, where I go, what I do, how I type and so on — the use of PINs and passwords is redundant: when the bank calls, there should be no question of asking for my mother's maiden name or my PIN because the phone will already know whether or not it is me.

There is no need to discuss authentication further, so the paper moves on to the other 'gene' for electronic cash: authenticity.

The authenticity problem

The authenticity problem is straightforward in mundane transactions: if you give me a physical banknote, I need to know that it is not counterfeit. In virtual transactions, however, there is an additional factor: if you give me a digital bank note, I not only need to know that it is not counterfeit but also that you have not given it to someone else already. Thus, authenticity is a relatively easy problem to solve in the physical world (because I can see if you have put the banknote into my hand) but not in the digital world. This is why Mr Zuckerberg's goal of making sending money as easy as sending photos is difficult. When you send your

friend a photograph of you on holiday, you are sending them a copy of the photograph and not the original. This is all very well for photographs, but not so good for money.

As has been understood from the very earliest days of electronic money, the crucial problem in the world of electronic cash is this *double-spending* problem. This is the architectural foundation of e-cash schemes and it must be solved at population scale because moving from a form of electronic money that is held in the accounts of financial institutions to a digital currency that is a form of money that can be held, well, anywhere, means moving from e-money to e-cash for everyone.

So how does one stop people from making up fraudulent e-cash or giving the same e-cash to more than one person? Cryptography works very well with regard to counterfeit prevention and detection because it is to all intents and purposes impossible to forge a digital signature (exactly how digital signatures work is beyond the scope of the present discussion, suffice to say that they do). In the digital money world, this is a simple problem to solve.

Double spending, however, is a tougher nut to crack. When it comes down to it, there are only two ways of preventing double spending: online, by having some sort of database to prevent unauthorised copying of value, or offline, by storing value in tamper-resistant hardware from where it cannot be copied.

E-CASH EVOLUTION

Together, the authentication and authenticity genes give shape to e-cash systems so it is worth exploring how these have mutated since the early experiments in the field. The first e-cash pioneers began by having a wallet on a PC with password authentication and a central database, which these days would be off in the cloud somewhere. This approach has an obvious disadvantage if one is trying to implement software e-cash: no privacy, because the database operator can always see exactly what is going on. As the paper will discuss in due course, however, cryptography can deliver some counter-intuitive services and it was one of the pioneers of e-cash who found a way to use cryptography to have a centralised database yet provide privacy to users. This was David Chaum's DigiCash.

DigiCash

The cryptographer David Chaum launched DigiCash b.v. in Amsterdam back in 1990 with a contract from the Dutch government to build and test technology to support anonymous road-toll payments.[3] Along with a number of other inventions in the field, Chaum came up with something called 'blinding'. Blinding means that someone issuing e-cash (eg a bank) can use cryptography to know with absolute certainly that some e-cash is real but find it cryptographically impossible to know who the e-cash was issued to.

Enter the internet. It was apparent that DigiCash's eCash, as the product was then called, could work very well in this new online environment and deliver an anonymous equivalent of cash into the new space. To guard against double-spending, eCash maintained a database of used coins. So, if someone sent you some e-cash, you could check the database to make sure that they had not previously given it to someone else.

The system never crossed the chasm. A few banks signed up to experiment with the new system, but it never gained any traction and DigiCash eventually filed for bankruptcy in 1998. It turned out that consumers were not interested in anonymity and were wholly content to use their credit cards to buy things online. If you had to get an account at a bank to use eCash, then you might as well just use the debit card they gave you.

DigiCash was a valuable experiment — an attempt to use wallets and the web to create an alternative payments network built around the individual. It contained many ideas that were influential in later developments in cryptography, but it was just as important in providing lessons about how to make a wholly new online form of money work.

Around the same time that DigiCash launched, NatWest in the UK decided to try a really radical experiment in digital money by creating a scheme that allowed true person-to-person value transfer — a genuine attempt to deliver e-cash into the mass market.

Mondex

Mondex was invented by Tim Jones and Graham Higgins at the National Westminster Bank (NatWest) in 1990. Instead of a centralised database, it used smart cards to store electronic value. Although pilot schemes were launched around the world, it never took off. While the technology worked, thanks to the work by Consult Hyperion among others, the banking processes did not. I can remember the first time I walked into a bank to get a card: I wandered in with £50 and expected to wander out with £50 loaded onto a card, but that was not how it worked. Customers had to fill out some forms to set up an account and then wait for the card to be posted to them. The hassle was too much for most people, so ultimately only around 14,000 cards were issued.

When the card eventually arrived, it had to be taken to an automated teller machine (ATM) in order to load it. The process of loading the card was especially crazy. As customers had to have a bank account to have one of these cards, this meant that they also had an ATM card. Anyone who wanted to load money onto their Mondex card had to take their ATM card to the ATM, pop it in, enter their PIN, select 'Mondex value' (or whatever the menu said), and then insert their Mondex card. Most people never bothered. Anyone who takes their ATM card to the ATM might as well get cash — and they did.

Mondex used its secure hardware to allow offline, card-to-card transfers — a radical decision at the time. A variety of other smart-card based schemes were introduced around the world at the time (eg Visacash) but these were account-based and did not facilitate such peer-to-peer value transfer. Similarly, DigiCash was not the only database-based software e-cash experiment of the time. There were many experiments using different kinds of cryptography, such as DEC's Millicent, QPass and eCharge,[4] as well as experiments in using different forms of value (eg e-gold, Beenz, Flooz and so on). An early example was CyberCoin, manifest in the UK as 'Barclaycoin'.

Barclaycoin

In 1997, my Consult Hyperion colleagues and I took part in an experiment with paid internet content.[5] For this experiment, we chose Barclays' BarclayCoin scheme, which was based on the CyberCoin scheme developed by the US internet payments pioneers CyberCash. This was a software-based scheme, which required consumers to download a free wallet to their computer. This wallet was then charged up from a credit card to hold digital money that could be spent on the web. The wallet was downloaded free from Barclays.

Once a customer had the wallet they could then go online and start buying. When they clicked on a link to the digital product they wanted to buy, they were sent an encrypted version of the product that triggered the wallet to ask if the customer wanted to confirm a payment. If the customer agreed, the balance in their wallet was reduced, and the merchant's balance

was credited. Once the customer confirmed payment, the key for decrypting the product was made available to the customer's wallet and the product could be used. As consumers spent money, it accumulated in the merchant's BarclayCoin account, less Barclays' commission, which was 25 per cent. That charge may seem high, but it was typical for micropayment schemes involving digital goods and, indeed, similar to the Apple App Store's 30 per cent today.

The combination of technological limitations and the rise of subscription business models (which removed the need for micropayments — one of the key uses cases for e-cash at the time) meant that none of these first attempts to create something truly new went anywhere. The cryptographic solutions were too complex and web browsers did not support encryption or authentication anyway.[6] It looked as if no person-to-person solution would gain traction.

Then along came Paypal and, eventually, Venmo.

Venmo

PayPal was established in 1998 (as Confinity) to develop security software for handheld devices such as the Palm Pilot. PayPal itself was a money transfer service developed within Confinity and, as *The Economist* observed at the time was 'more like real digital money, because it allows consumers to pay each other as well as merchants'[7] by transferring between customers' wallets. In 2000, Confinity merged with Elon Musk's online banking startup X.Com, which soon terminated other operations to focus on the money transfer business. The company, then renamed PayPal, went public in 2002. In 2011, PayPal acquired the mobile payment player Zong (led by David Marcus, of whom more later) before adding Braintree and Venmo to its portfolio in 2013.

Venmo is a particularly interesting example of a peer-to-peer scheme because of its pro-social functions, which, while not the subject of this paper, were a key factor in its growth.[8] Venmo now handles almost US$30bn per quarter and is still growing, as shown in Figure 1. Venmo (and its rival Square Cash) are now embedded in US day-to-day experiences.

In one sense, Venmo, while a breakthrough, is not a revolution: underneath its façade of convenience, everything is running on the bank rails of debit networks and automated clearinghouse transfers. The failure of genuine alternatives such as DigiCash and the continued growth and development of international card schemes and domestic payment networks meant that these rails carried more and traffic (successfully) and by the mid-2000s some observers began to question whether any alternative to the bank-led payments infrastructure would ever get off the ground. However, even before PayPal purchased Venmo, this question had already been answered by the success of M-Pesa, the first mobile-centric person-to-person payment system.

M-Pesa

Back in 2003, Safaricom was the market leading mobile operator in Kenya, with just over half the market. The company had the idea of using mobile phones to make the distribution of microfinance loans in Africa more efficient and submitted a proposal to the UK Department for International Development for matching funding. This was granted and M-Pesa was born in a feasibility study supported by Consult Hyperion. The pilot launched in 2005 and within a year the scheme had 2 million subscribers and was handling US$1.5m per day. The scheme, which allows people to deposit and withdraw cash from wallets associated with their mobile phone numbers, has been an incredible success, with more than two-thirds of Kenya's adult population using it and tens of thousands

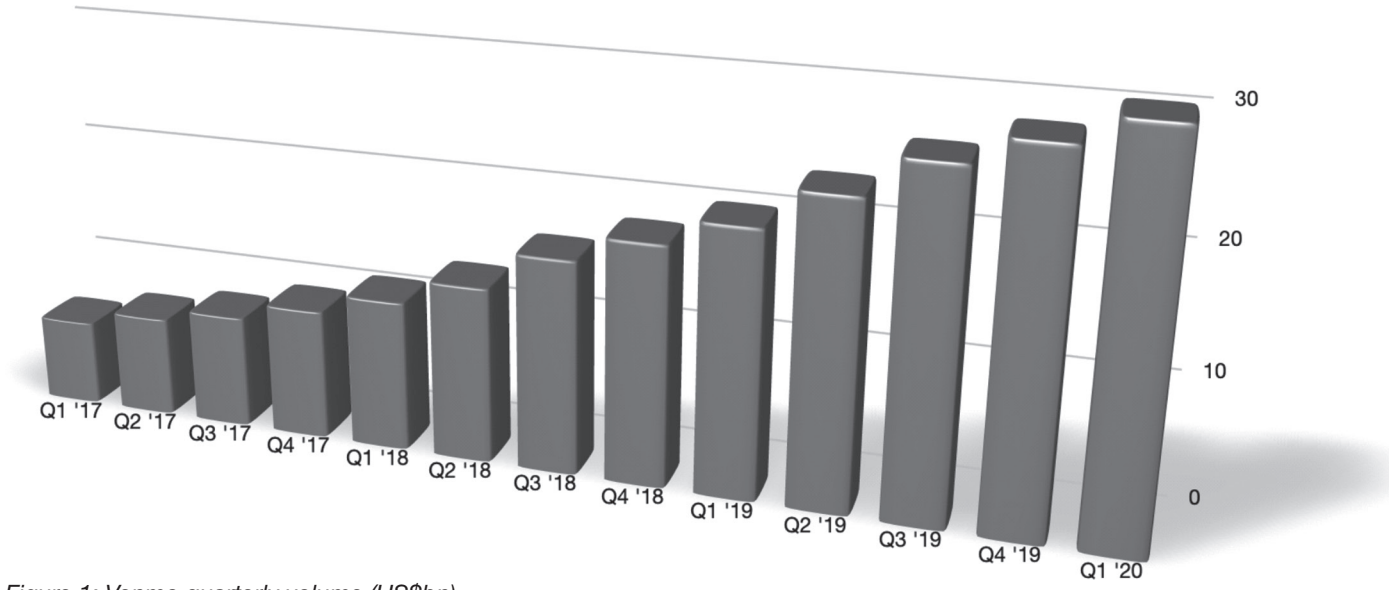

Figure 1: Venmo quarterly volume (US$bn)

of agents allowing consumers to pay cash into the system or take cash out of it. To put these numbers in context, it took banks in Kenya a century to create a mere 1,000 bank branches, 1,500 ATMs and 100,000 credit card customers. A non-bank payment system founded on new technology rather than legacy infrastructure changed people's lives in ways that could not have been envisaged by the people who created it, creating countless new services that would not have been possible without the safe, fast and instant peer-to-peer money transfer service.[9]

M-Pesa was electronic money that used tamper-resistant hardware to authenticate access to wallets that hold the prepaid value (conceptually — the actual value is held in commercial banks). PayPay was electronic money that used software and some heavy-duty anti-fraud backend systems to authenticate access to the data representing value. DigiCash was electronic cash that used software: someone sent value to you, and you could do what you wanted with it. Mondex, on the other hand, was electronic cash that used tamper-resistant hardware to store the data representing value as well as to authenticate access.

What these and other e-cash schemes had in common was that they needed someone in the middle to make everything work. Clever software people, building on ideas from all of these schemes and more, were keen to find a way to avoid this. However, they kept running into the same problem about maintaining the overall integrity of the system. How to solve the authenticity problem without the tamper-resistant chips round the edges or a database in the cloud in the middle?

In 2008, someone came up with a genuinely new way to put together the technologies of digital money to give a decentralised software solution: the blockchain. This brings the discussion, of course, to Bitcoin.

Bitcoin

The Bitcoin story is, by now, well known. Person or persons unknown, under the pseudonym 'Satoshi Nakamoto', published a white paper setting out how to create a

person-to-person e-cash system without a central system operator or database.[10]

At the core of the Bitcoin system were three main concepts: to replace a central database with a shared ledger, to use a new consensus technique to build that ledger absent a central coordinator (the 'Nakamoto consensus' that uses 'proof-of-work' to determine which version of the ledger is correct) and to use a particular kind of mathematical puzzle to incentivise the proof-of-work. The technologies used to implement each of these concepts were all well known, but Nakamoto made a crucial breakthrough by combining them to embed the incentive mechanism into the consensus-forming process ('mining') giving an energy to the ecosystem.[11] This combination — the blockchain — was a revolution.

One aspect of Bitcoin that sowed the seeds for new ways of working is the introduction of programmability. For some observers, the invention of what one might call smart money — money that has its own apps — is actually much more interesting than the peer-to-peer payment system — a point this paper will return to in due course.

In practice, Bitcoin has struggled to find mainstream adoption beyond speculation. The Financial Conduct Authority's mid-2020 survey found that some 2.6 million UK consumers had bought cryptocurrencies at some point, and that a surprisingly high 1.9 million that still owned some, with half owning more than £260 worth. These purchases are presumably highly speculative as cryptocurrencies are little used for 'real-world' transactions. One of the reasons for the lack of retail interest is the instability of such digital transaction intermediaries, hence the interest from many directions in what have become generically known as 'stablecoins' that could in principle form currencies for general use.

In account-based cryptocurrencies (eg Ethereum), each account has an associated public key for verifying the validity of transactions signed by the account holder using a corresponding private key. Only validly signed transactions are included on the shared ledger, thus providing a mechanism for secure value transfer.[12] One can then build a crypto-asset layer on top of that to link the values to something in the real world. Note, of course, that this crypto-asset layer could be null and the digital value itself be the value traded, as in the case of Bitcoin.[13] Either way, one then has some form of digital asset that can be traded without clearing or settlement and e-cash is one particular kind of asset that one can use to make a stablecoin.

USDC

A good example of a digital currency built using a token-based digital asset is USD Coin (USDC). This is a token pegged to the US dollar. A collaboration between Circle and Coinbase, it is intended to provide the equivalent of the fiat currency for use over the internet and public blockchains. USDC tokens can be redeemed for US dollars at any time.

The reason for wanting to do this is clear. In the absence of an actual digital dollar of some kind, then bringing the current prime currency to the blockchain means currency can be moved anywhere in the world in minutes, bringing much-needed stability to the cryptocurrency world to facilitate use for payments and opening up new opportunities for trading, lending, risk-hedging and more.

The issuing and redeeming of USDC tokens is executed through an ERC20 smart contract on the public Ethereum blockchain.

The tree of e-cash, as shown in Figure 2, has therefore evolved to the point where, in essence, the 'meta-technology' means that anyone can now use cryptocurrency create their own e-cash of any kind. In that 'branch' of the evolutionary tree

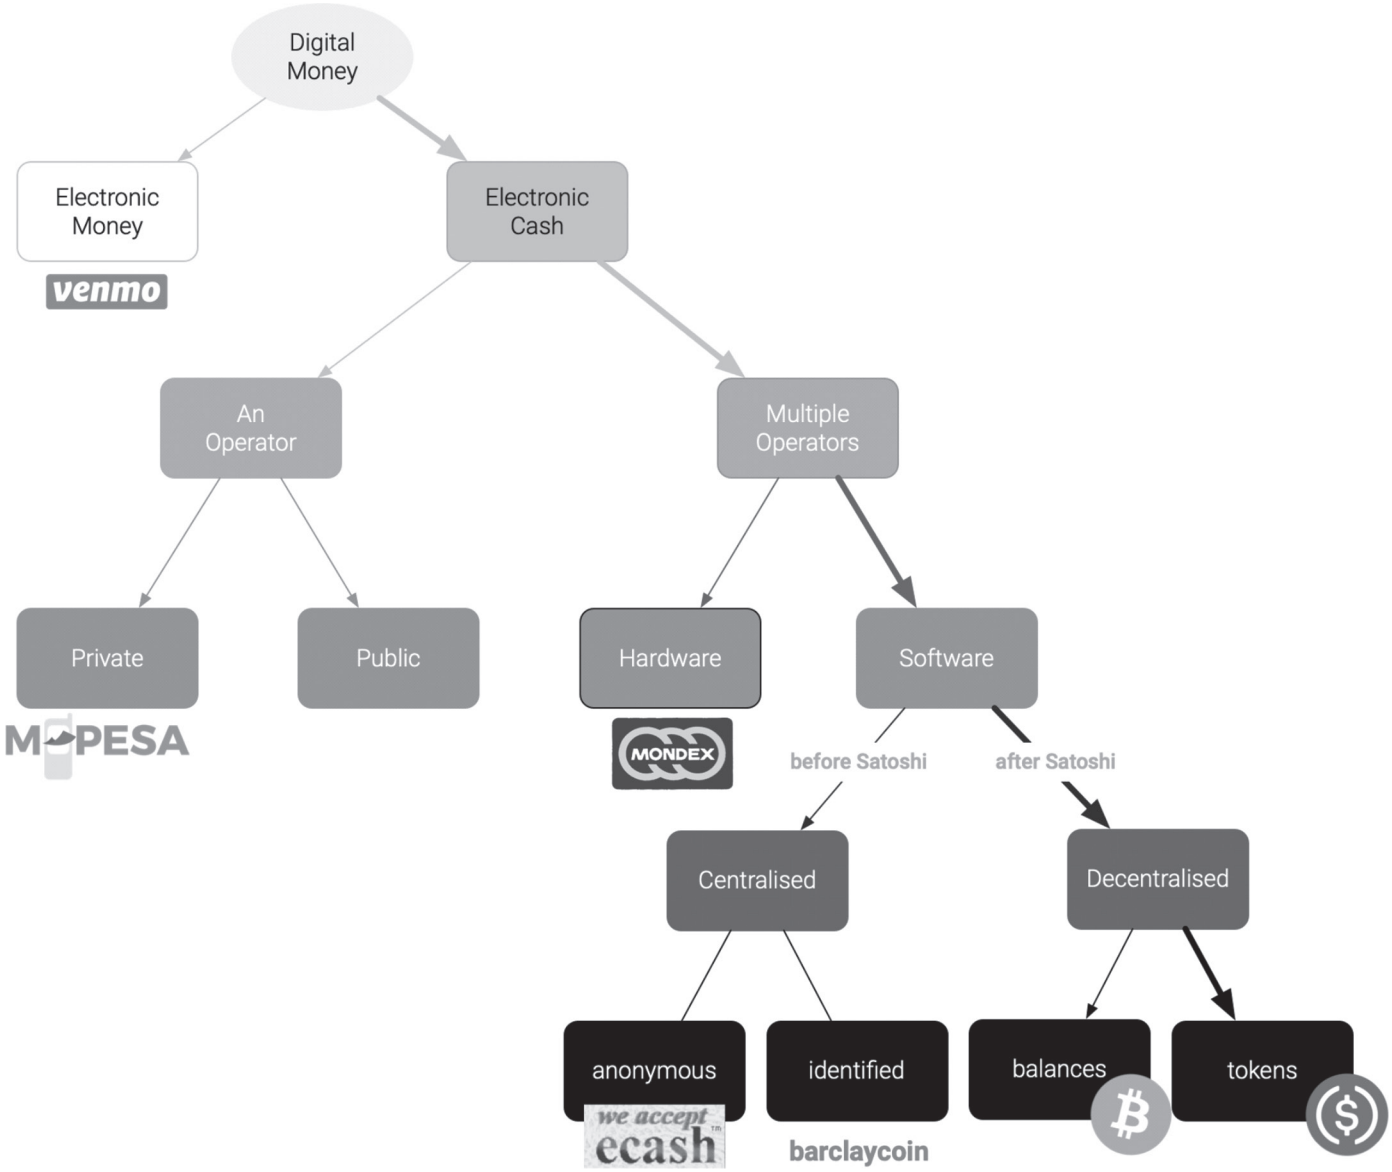

Figure 2: The digital currency tree

there is a value transfer layer that may or may not be implemented using some form of shared ledger that makes for the secure transfer of digital values from one storage area (ie the wallets mentioned earlier) to another.

In this formulation, then, e-cash is a digital bearer instrument (whoever holds the cryptographic key holds the value, whether that value is a dollar, one-thousandth of the Mona Lisa or gold in a depository somewhere). Digital bearer instruments can be exchanged by what the blockchain fraternity insist on calling 'smart contracts' (the term 'consensus applications' is more apposite as they are neither smart nor contracts). The general term for these bearer assets is 'tokens'. We have the technology to make a central bank digital currency.

BRITPESA, BRITCOIN OR BRITDEX?

There are various ways that a central bank could implement a digital currency as part of a strategy to move to a cashless society (ie a society where cash is irrelevant rather than illegal). Way back in the 1990s, the model that was chosen for the Mondex experiment that began in the UK was to have the central bank control the creation of the digital currency, but to have it distributed by the commercial banks through their existing channels. This is what is known as the 'two-tier' implementation. These days, however, it would be implemented using mobile phones rather than smart cards — a kind of BritDex.

A cheaper alternative is to have the central bank create accounts for all citizens, businesses and other organisations. Imagine something like M-Pesa but on population scale — BritPesa if you like, in the UK example. This would be cheaper because it would be completely centralised and the marginal cost of transferring value from the control of one personal organisation to another through such a system would be negligible. Central banks do not really want to implement such a 'one-tier' solution, however, because it would mean having to manage millions of accounts, and they would prefer somebody else to do this and deal with everything else that goes with interacting with the general public. The commercial banks and plenty of other non-bank players (for example, Alipay in China) already have the apps, the infrastructure and the innovative approach that would not only bring the digital currency to the mass market but would also open up the potential for the digital currency as a platform for innovation and development.

Alternatively, there could be something like USDC — a digital asset backed by central bank reserves. This BritCoin would still be a two-tier solution distributed to the public by the commercial banks, but it would remain under the control of the central bank. Of the three practical CBDC options that are presented in context in Figure 3, this is arguably the most interesting because it offers the potential for significant innovation.

The UK way

Given these three alternatives, what should a central bank do? To use the specific example of the UK, what should the Bank of England do? If one imagines that the economy of UK might demand an e-cash infrastructure capable of handling 100,000 transactions per second, what are the options?

There is no problem managing this volume using BritDex for the obvious reason that the transactions are peer-to-peer. Similarly, there is no problem with some form of BritPesa. In China right now, Alipay handles more than 1 billion unique customers, with peak volumes in the region of 250,000 transactions per second. When it comes to BritCoin though, the kind of public shared ledger used for mining and value exchange simply cannot support the necessary volumes, which means either building a 'level 2' network on top of it or using an alternative structure. The Bank for International Settlements' report on the foundational principles of CBDC notes that research into scalability has shown that performance problems associated such public solutions can be overcome with permissioned shared ledger networks.[14] It further observes that estimating future volumes and throughput requirements is complicated by ongoing industry developments (eg payment requests generated by smart devices and the potential for high-volume micro-transactions) — an interesting point beyond the scope of this paper.

The Bank of England's excellent recent discussion paper on CBDC[15] discusses a 'platform approach' and quite rightly notes that one of the key advantages of such an approach is that it will help innovation throughout the 'stack'.

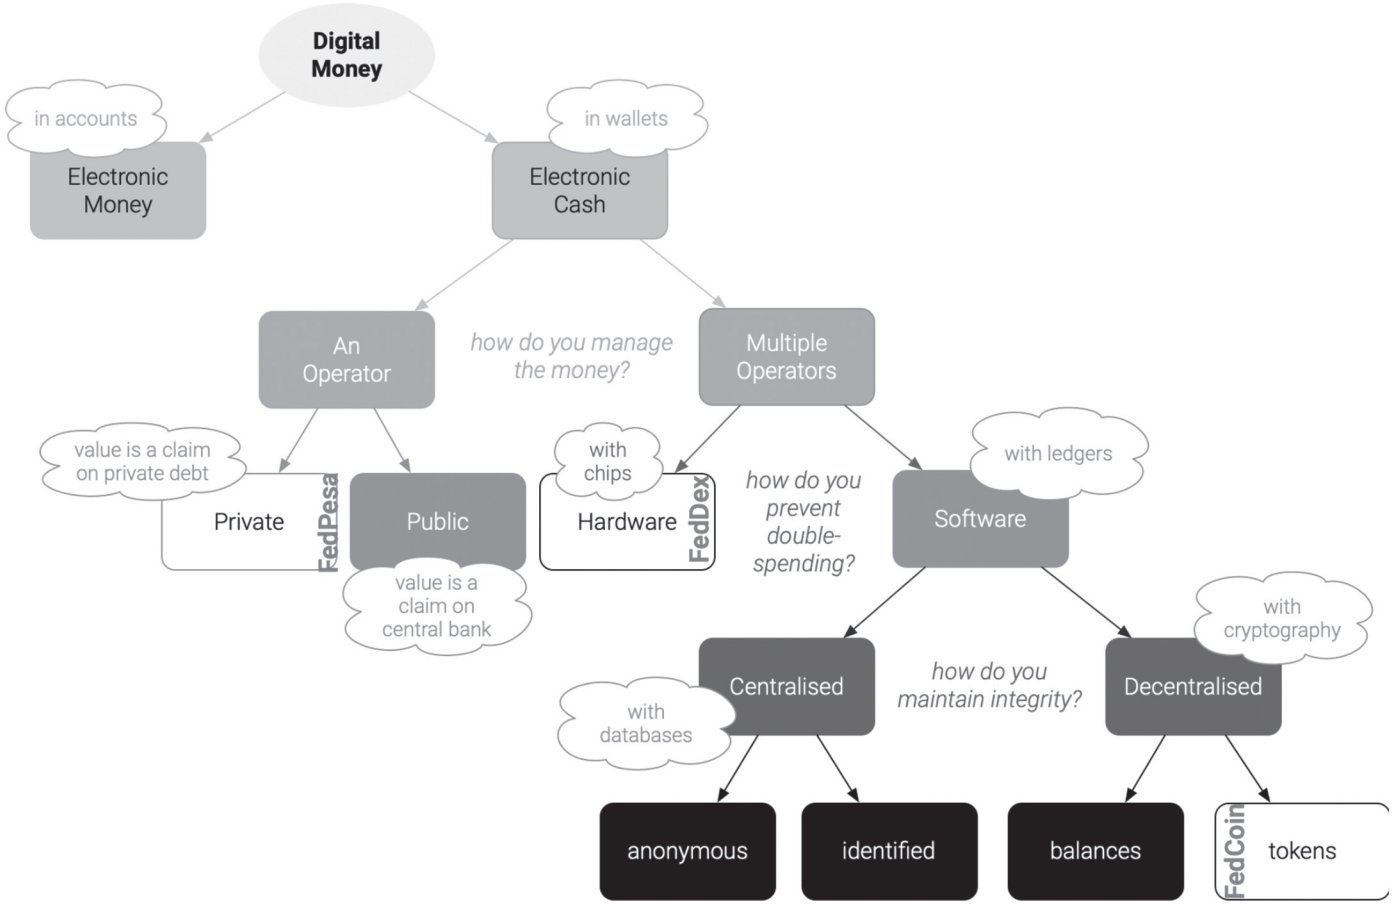

Figure 3: Options for a CBDC

How would this work? Imagine a merging of something like India's UPI, M-Pesa, social media and the lifestyle apps coming from the Far East and one can begin to develop a picture of just how powerful such an implementation might be in all markets. The Bank of England uses some specific terminology which makes sense and will allow for constructive discussions between regulators, businesses and innovators in the payments space. In the Bank of England's platform model, it is assumed that the central bank runs the platform and provides what the Bank of England calls 'API access' to this platform. The people allowed to access the platform are labelled 'payment interface providers' (PIPs) and it is these providers (banks among them) who interact with users.

Were post-Brexit Britain to create a British CBDC issued and managed by commercial banks (as some form of digital sterling should be a priority in the country's financial rebuilding post Brexit), this would not use either the smart cards of the Mondex days or the basic SIM toolkit and SMS technology of M-Pesa. Rather, it would use smartphones, chatbots, artificial intelligence, fingerprints, voice recognition and all that jazz (or 'biometrics, blockchains and bots' as I often unhelpfully and incorrectly paraphrase for management).

It would not be that difficult or complicated to create a basic centralised CBDC:

there would be a system shared by the commercial banks with the funds held in a central account. Regardless of whether digital fiat is the long-term future of money (and it probably is not), it seems prudent to get on with digital sterling, whether BritPesa or BritCoin or BritDex, and give everyone access to payment accounts without credit risk. Aside from re-forging the UK financial system in the white heat of new technology, there is a very good reason for doing so. According to the Bank of England Staff Working Paper No. 605:

> 'CBDC issuance of 30 per cent of GDP, against government bonds, could permanently raise GDP by as much as 3 per cent, due to reductions in real interest rates, distortionary taxes, and monetary transaction costs. Countercyclical CBDC price or quantity rules, as a second monetary policy instrument, could substantially improve the central bank's ability to stabilise the business cycle'.[16]

Aside from increasing gross domestic product, there is another excellent reason for taking this step, which is that cash has no API. Writing in the Bank of England's 'Bank Underground' blog, Simon Scorer from the Digital Currencies Division made a number of very interesting points about the requirement for some form of digital fiat. He remarked on the transition from dumb money to smart money, and the consequent potential for the implementation of digital fiat to become a platform for innovation, saying that:

> 'other possible areas of innovation relate to the potential programmability of payments; for instance, it might be possible to automate some tax payments (eg when buying a coffee, the net amount could be paid directly to the coffee shop, with a 20 per cent VAT payment routed directly to HMRC), or parents may be able to set limits on their children's spending or restrict them to trusted stores or websites'.[17]

CONCLUSION

If digital fiat were to be managed via some form of shared ledger to create a BritCoin, then the Bank of England's comments suggest that it is not the shared ledger but the consensus applications that will become the platform for radical innovation as they are used to implement new digital currencies. The Bank of England is clear that it does not envisage the ledger as a cryptocurrency platform, but it does say that the technologies of shared ledgers might be the best way to implement it. Were such a system to come into existence, its resilience and availability would become matters of vital national interest. It therefore makes complete sense to take advantage of the new technologies and construct a decentralised and robust solution. Each PIP would have a gateway and the option of maintaining its own copy of the shared ledger or accessing somebody else's. All banks above a certain size would be mandated to keep a copy of the ledger and the PIP gateways would simply talk to each other (through the normal protocols of consensus chosen for the particular architecture), but there would be no central system in the middle — no equivalent to real-time gross settlement that could fail either because of management failings (as is usually the case), unforeseen technical problems or subversion by foreign powers.

As the Bank says, the most game-changing aspect of such an implementation would be what it calls 'programmable money' and what I have previously labelled 'smart money'.[18] This is where the real innovation will take place that will make the money of the future so very different from

the money of today, and it will be very interesting to see thinking develop in this area. There are obviously overheads associated with overloading the ledger with the consensus applications, but on the other hand it may be that there are some truly revolutionary features that can only be delivered through such applications. The Bank suggests a compromise whereby certain distributed applications are provided for the use of the PIPs in order to give them infrastructure that they can then use to develop innovative end-user services and this seems a good place to start.

It would be good to see the Bank of England take the global lead in the race to create money for the digital future, rather than continue with digitised versions of money from the analogue past, and I for one would bet on them to succeed were this smart money to be built from translucent transactions that deliver ambient accountability for the new economy.[19] Huw van Steenis of UBS predicts a 'three-horse race'[20] around the future of money, with private tokens and CBDCs developing in parallel with efforts to improve the current system (see, for example, SWIFT gpi and the UK's new payments architecture). This is wise counsel, and there is indeed every possibility of competition between these approaches stimulating innovation in the short term but then a longer-term convergence as the platforms for exchanging digital asset tokens are used to implement both private and public tokens (including CBDCs).

Thus, the long-term competition will not be between public and private versions of digital dollars but between public central bank money tokens and private tokens that are backed by other assets. In this appealing vision of the future, there is nothing technological to distinguish dollar bills from Bill's dollars: one will be backed by risk-free central bank money, the other by Microsoft revenues.

References

(1) Bank of England (2020) 'Central Bank Digital Currency — Opportunities, Challenges and Design', Bank of England, London.

(2) Birch, D. (2020) 'The Currency Cold War', London Publishing Partnership, London.

(3) Levy, S. (1994) 'E-money (that's what I want)', *Wired*, December, available at: https://www.wired.com/1994/12/emoney/ (accessed 3rd November, 2020).

(4) Essex, D. (1999) 'Big dreams for tiny money', *Computerworld*, 13th December, p. 66.

(5) Birch, D. (1998) 'An experiment in micropayments', *Financial Times Virtual Finance Report*, Vol. 12, No. 2, p. 2.

(6) Green, M. (2018) 'Thirty years of digital currency: from digicash to the blockchain', paper presented at IACR Eurocrypt Conference, Tel Aviv, 29th April.

(7) *Economist* (2000) 'E-cash 2.0', *Economist*, Vol. 354, No. 8158, 19th February, pp. 67–71.

(8) Acker, A. and Murthy, D. (2020) 'What is Venmo? A descriptive analysis of social features in the mobile payment platform', *Journal of Telematics and Informatics*, Vol. 52, September, DOI: 10.1016/j.tele.2020.101429.

(9) Donovan, K. (2012) 'Mobile money, more freedom? The impact of M-pesa's network power on development as freedom', *International Journal of Communications*, Vol. 6, pp. 2647–2669.

(10) Nakamoto, S. (2008) 'Bitcoin: a peer-to-peer electronic cash system', available at: https://bitcoin.org/bitcoin.pdf (accessed 3rd November, 2020).

(11) Brunton, F. (2019) 'Emergency money', in 'Digital Cash — The Unknown History of the Anarchists, Utopians and Technologists Who Created Cryptocurrency', Princeton University Press: Princeton, NJ, pp. 153–170.

(12) Allen, S., Capkun, S., Eyal, I., Fanti, G., Ford, B., Grimmelmann, J., Juels, A., Kostiainen, K., Meiklejohn, S., Miller A., Prasad, E., Wüst, K. and Zhang, F. (2020) 'Design Choices for Central Bank Digital Currency: Policy and Technical Considerations', Global Economy and Development, Brookings Institution, Washington, DC.

(13) Birch, D. (2018) 'Who will make money? Tokens and the "5Cs" of future currency', *Journal of Payments Strategy and Systems*, Vol. 12, No. 2, pp. 111–121.

(14) Bank for International Settlements (2020) 'Central Bank Digital Currencies: Foundational Principles and Core Features', Bank for International Settlements, Geneva.

(15) Bank of England (2020) 'Central Bank Digital Currency: Opportunities, challenges and design', available at https://www.bankofengland.co.uk/paper/2020/central-bank-digital-currency-opportunities-challenges-and-design-discussion-paper (accessed 3rd November, 2020).

(16) Barrdear, J. and Kumhof, M. (2016) 'The Macroeconomics of Central Bank Issued Digital Currencies', Bank of England, London.
(17) Scorer, S. (2017) 'Beyond blockchain: what are the technology requirements for a central bank digital currency?', available at https://bankunderground.co.uk/2017/09/13/beyond-blockchain-what-are-the-technology-requirements-for-a-central-bank-digital-currency/ (accessed 3rd November, 2020).
(18) Birch, D. (2017) 'Before Babylon, Beyond Blockhain', London Publishing Partnership, London.
(19) Birch, D., Brown, R. and Parulava, S. (2016) 'Towards ambient accountability in financial services: shared ledgers, translucent transactions and the legacy of the great financial crisis', *Journal of Payment Strategy and Systems*, Vol. 10, No. 2, pp. 118–131.
(20) Steenis, H. (2020) 'The new digital payments race', available at: https://www.project-syndicate.org/onpoint/central-banks-digital-payments-by-huw-van-steenis-2020-04 (accessed 3rd November, 2020).