Central bank digital currency with asymmetric privacy

Katrin Tinn, McGill University, Desautels Faculty of Management (Finance), katrin.tinn@mcgill.ca

Christophe Dubach, McGill University, Department of Electrical and Computer Engineering, and School of Computer Science, christophe.dubach@mcgill.ca

Executive summary

1. Background

There are increasingly many digital alternatives to physical cash as basic means of payment. In addition to card payments that have been offered by the banking sector for a while, further means of payments are being created by non-traditional players such as technology firms (e.g., Amazon Cash, Apple Pay) along with many different cryptocurrencies (e.g., Bitcoin, Z-Cash) that emerged outside the traditional financial system.

As a result, Central Banks, who are responsible for overseeing prominent payment systems to ensure that they are safe, available to everyone and efficient, have become increasingly interested in central bank digital currency (CBDC) design and potential implementation. For example, the Bank of Canada's contingency plan highlights that the need for a CBDC may arise in an environment with multiple alternatives provided by the private sector to ensure the existence of a safe, convenient and reliable digital option alongside physical cash. Namely, they emphasize: 1) Safety; 2) Universal accessibility, 3) Privacy, 4) Resilience, 5) Competition and efficiency, and 6) Monetary sovereignty, as key desirable features that a Canadian CBDC should have.

Similar considerations and desirable features have been emphasized by the European Central Bank and the Bank of International Settlements, among others.

2. Token money, account-based money and a proposal of P-hybrid CBDC

A useful conceptual distinction is that of **account-based money** *versus* **token money**. An example of account-based money is credit and debit cards, where the identity of account holders (both sender and receiver) is known and verifiable. An example of token money is traditional physical cash, where the validity of the banknote itself is verifiable, while the identities of the sender and of the receiver are private. Many other emerging forms of money fall somewhere in between. For example, different crypto-currencies resemble token money, as the identities of senders and receivers are not easily known, while Amazon Pay or PayPal are closer to account-based money.

The advantage of token money is that consumers can make their purchasing decisions without needing to worry about outside parties collecting data on their purchasing choices, which may lead to unwanted advertising, adverse impact and potential discrimination when it comes to their credit score, or other unintended consequences. The weakness of token money is that it limits the core benefits of digital records (e.g., the possibility to create innovative financing products and other services built on these records) and makes regulatory compliance (e.g., anti-money laundering regulation) harder. This is because accounting records with sales in token money (e.g., physical cash) require voluntary disclosure and auditing, which can take time and can be imperfect. The disadvantages and advantages of account-based money are just the opposite: while users lack privacy, their digital records can be used by regulators to ensure compliance, but can also be used by innovative financial institutions and technology firms to offer new and better services. One could then ask: how much should privacy be compromised for regulatory compliance and to encourage innovation? Should a CBDC be more token money like, or account-based money like? **Our key contribution in this paper is to show that there is a solution that mitigates, if not overcomes, this inconvenient trade-off, and we propose a Privacy-Hybrid (we call it P-Hybrid) CBDC.**

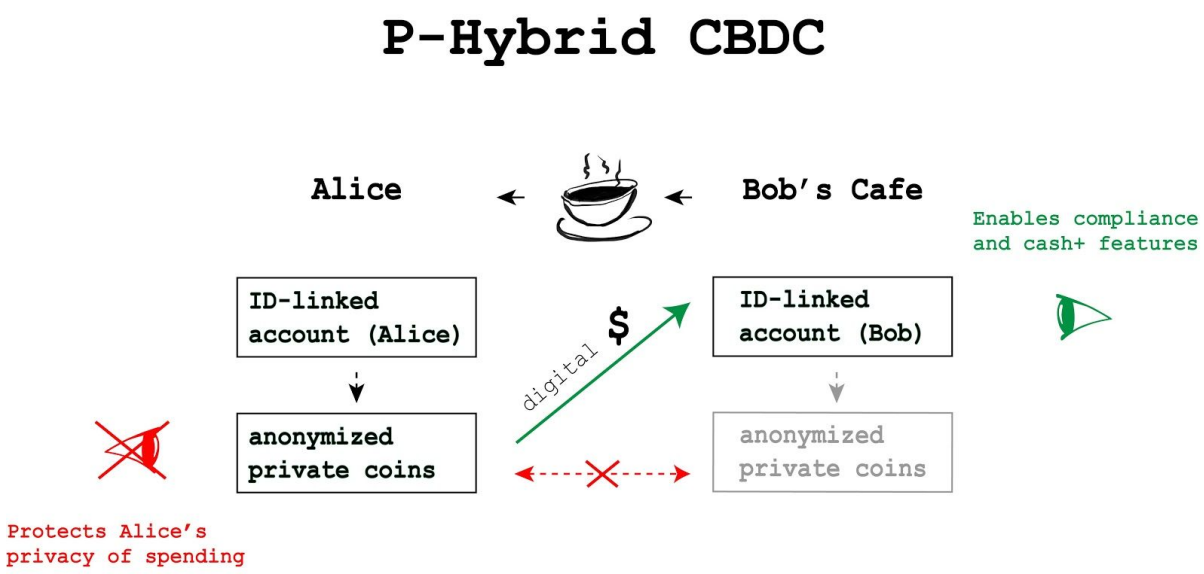

Figure 1: Illustration of a transaction using P-Hybrid CBDC

The key feature of our P-Hybrid CBDC is the **intentional asymmetry** between receiving and sending money. We argue that the central bank's digital currency should offer a better alternative regarding privacy than account-based money when it comes to protecting the anonymity of individual spenders. Namely, the proposed P-Hybrid CBDC aims to make it impossible (or more precisely, statistically close to impossible) to associate an individual with their purchases of coffee, groceries, alcohol, entertainment, medicine, etc. At the same time, we argue that there is little reason to protect the identity of the receivers of money. For example, even if someone uses cash to buy coffee, the cafeteria still needs to keep accounting records of these sales and pay taxes. If an individual receives income, it should also be traceable by parties that need to know, such as tax authorities, or the institutions responsible for detecting incidences of money laundering or big donations to political parties. We further argue that digital records of incoming money, such as sales records, not only simplify accounting and facilitate compliance, but also bring Cash+ benefits and encourage innovation, as mentioned above and discussed in greater detail in the paper.

Figure 1 illustrates the basic functioning of the P-Hybrid CBDC by considering a transaction between Alice and Bob's Cafe for the purchase of a coffee. Alice and Bob have both gained access to the P-Hybrid CBDC system and its functionalities by creating beforehand an ID-linked account which is uniquely associated with their identity. However, before Alice buys her coffee, she has transferred some of her money to her "anonymized wallet" (more technically, she has created private coins). No outside parties (i.e., not Bob, not her bank, no outside eavesdropper, and not even the system actors) are able to digitally prove and observe that it was Alice who bought the coffee. At the same all those who need (e.g., system actors, tax authorities, anti-money laundering authorities, authorized firms who provide services to Bob at his consent) will observe that Bob received the digital payment from someone who used their private coins/"anonymous wallet".

As explained further below and in our paper, private coins can only be sent to ID-linked accounts; this overcomes one weakness of physical cash and anonymous digital money, which allow the introduction of illegal money into the system. ID-linked accounts are also useful for making it harder to use the P-Hybrid CBDC as means of payment for illegal goods and services.

3. Economic case for P-Hybrid CBDC

We propose a model of privacy concerns to highlight the distortions that these concerns impose on the consumer-producer relationship. We build on tools from the Differential Privacy literature and show that even if the producer is not benefitting from knowing the identity of the buyer of its goods or

services, these concerns affect the preferences of rational privacy concerned consumers and make them less willing to pay, alter their optimal choices, or both. All this negatively affects the profits of the firm, and the joint surplus of the consumer and of the producer. We also highlight that observing the total spending on more goods (possibly over longer time) mitigates these frictions.

We then incorporate these insights to a model of money demand, where consumers face a classical trade-off between earning interest income (or more broadly returns) on their savings and some potential costs of transferring funds from their interest bearing account to their spending account. We introduce some new twists. First, financial assets can also be used as means of payments (i.e., there is no Clower constraint), and non-interest bearing fiat money (be it physical cash or CBDC) has a privacy benefit, more so when consumers withdraw less frequently.

We show that without privacy concerns and with low enough costs of transfers between accounts, there is no demand for a non-interest bearing asset. In contrast, privacy concerns re-establish a classical looking money demand function, where consumers have an incentive to convert their interest bearing assets to privacy preserving cash-like money relatively infrequently. We further show that rational consumers are more likely to adopt cash-like money when:

1. Privacy concerns are important enough;
2. Consumers are wealthier;
3. Inflation is lower.

We also analyse in detail the benefits of the ID-linked part of our P-Hybrid CBDC for

1. Regulatory compliance and taxation;
2. Facilitating "smart contracts" and financial innovation;
3. Enabling new tools for fiscal and monetary policy.

4. System Architecture and choices of technology

To implement our vision of P-Hybrid CDBC (see Figure 1), we propose a system architecture based on a blockchain technology with Proof-of-Authority and Zero-Knowledge Proofs. These design features have been chosen to ensure:

1. **Compliance** - which is achieved by utilizing ID-linked accounts, and by allowing private coins to be only sent to ID-linked accounts, rather than any unauthorized private account.
2. **Privacy** - which is achieved by allowing users to transform part of funds in their ID-linked account into private coins. Disassociating the individual's ID from their coins is achieved using Zero-Knowledge Proofs, which implies that no one (not even system actors and authorities) can identify the sender of digital cash.
3. **Transparency** - which is achieved by utilizing a blockchain such that all users can observe and validate that the transactions recorded do not violate the rules of the system.

We propose Proof-of-Authority as a consensus method, as added complexities associated with methods like Proof-of-Work (used for Bitcoin among others) would slow down the system, and make little sense in cases where the Central Bank maintains control over digital CBDC issuance. The transparency of the blockchain technology enables users to build trust in the system as they are able to see when a new amount of digital CBDC has been created. It is worth emphasising that the transparency of a blockchain system does not imply that the users' identities are at risk, but it simply implies that outsiders can verify that the manifested rules of the system are obeyed.

The paper also compares our preferred privacy generation system of Zero-Knowledge Proofs with alternatives such as anonymous blockchains (where no cryptographic address is used twice), and "shuffling services" (which mixes users' coins and creates new bundled coins for transactions), and highlights the benefits and workings of our chosen approach.

The implementation of our proposed system architecture requires **four sets of system actors:**

1. Identity Certifiers - who enable creation of ID-linked accounts by certifying users;
2. ID-linked Database keepers - who maintain ID-linked accounts;
3. Block Producers - who manage the blockchain system and ensure that double-spend is not possible;
4. CBDC Issuers - who issue and withdraw CBDC from circulation (e.g., manage money supply, and ensure convertibility between physical and digital cash)

We further analyze the scalability of our proposed system, and highlight that under conservative estimates such a system would easily manage at least 2,000 transactions per second using today's compute technology, and noticeably more transactions when using systems with faster computational speed, larger storage capacity and higher internet bandwidth, which is likely to be commonplace by the time a P-Hybric CBDC could be introduced in practice.

5. System actors, economic agents, policy choices and adoption

As highlighted above, our system design identifies four types of necessary system actors. It has been set up to enable the Central Bank to either become responsible for all these separate roles, or delegate some of these to other authorized and monitored entities. From an economics perspective, the key role that the central bank needs to maintain is control over high-powered money supply (mainly the sum of physical and digital cash in circulation). We discuss a number of considerations and trade-offs when delegating these roles to other authorized entities. We also discuss the role of governmental authorities and private sector actors who should or could have the right to view information in ID-linked accounts to ensure regulatory compliance or to enable innovation. Finally, we discuss the adoption incentives of users, the means to encourage the design of consumer interfaces and cash+ features, and potential sources of income for outside developers.

Central bank digital currency with asymmetric privacy

Katrin Tinn
McGill University
katrin.tinn@mcgill.ca

Christophe Dubach
McGill University
christophe.dubach@mcgill.ca

11 February 2021

Abstract

There are increasingly many digital alternatives to physical cash as basic means of payment, many of these offered by non-financial private sector actors. While such a trend towards a cashless society brings the convenience of digital payments to users and can encourage innovation in financial contracts, there are also increasing concerns about users' privacy as well as the regulators' ability to ensure compliance with anti-money laundering prevention and other regulations. As a result, Central Banks, who are responsible for overseeing prominent payment systems to ensure that they are safe, available to everyone and efficient, have become increasingly interested in a CBDC (Central Bank Digital Currency) design and potential implementation.

In this paper we analyze the economic rationale for, and the technical feasibility of, a new form of CBDC. The key feature of our proposed Privacy-Hybrid CBDC is intentional **asymmetric privacy** between the receiver and the sender of money. We show that protecting the privacy of consumer spending is necessary to avoid distortions between producers and consumers. Furthermore, we model the demand of money and show that privacy protection is necessary for a classical money demand function in a framework where interest bearing financial assets can be used as means of payments. We also analyze how adoption incentives and the frequency of transactions depend on income and other factors. At the same time, limiting the privacy of money received enables *compliance* with tax and anti-money laundering regulations, and can enable new value adding services by financial institutions and technology firms.

To achieve asymmetric privacy, we propose a CBDC system architecture that relies on a central registry of ID-linked accounts to ensure *compliance* and the use of zero-knowledge proofs to mint private tokens to achieve *privacy*. The system ensures that private tokens cannot be exchanged between users in a trust-less manner and that private tokens can only be used to send digital cash to the holder of an ID-linked account. This paper discusses how this system could be implemented with today's technology to support a throughput in excess of $2,000$ transactions per second.

Keywords: blockchain, central bank digital currency, money demand, privacy, proof-of-authority, regulatory compliance, zero-knowledge proofs

JEL Codes: C70, D18, D83, E41, E42, E58, G21, G23, L86

1 Introduction

What added value can and should a CBDC bring? As pointed out by Duffie et. al. [DSR+20], there is currently a range of possible ambitions, with their related challenges. While the simplest forms of CBDC may aim to focus on settlement efficiency, the more ambitious ones may aim to introduce a form of retail digital currency as a substitute to fiat money and physical cash (see [ACF+20] for an overview).

We analyze the rationale for, and the implementable design features of, a retail/cash-like CBDC. This cash-like digital currency is not only of theoretical interest to central banks. It may also become a practical necessity, as highlighted by the Bank of Canada's background paper for a CBDC contingency plan [Ban20];

cash might soon no longer be accepted in many businesses or some private digital alternatives might start to replace the Canadian dollar as a method of payment.[1]

The question of whether a CBDC would be needed inevitably leads to the broader question of why there is a need for any form of fiat money (cash) at all. In seminal models of money, fiat money is either imposed as sole means of payment (see *e.g.* Chapter 4 in [BF89] for a review), or the role of fiat money endogenously stems as a means to overcome the double-coincidence of wants across goods or time (see *e.g.* [Wic13], [KW93], [KM02], and [KM18]).

There are nowadays many digital alternatives to physical cash as basic means of payment. A useful conceptual distinction is "account-based money" vs "token money [BJL19], [KRW18], [AB20]. An example of account-based money is credit and debit cards, where the identities of the account holders (both sender and receiver) are known and verifiable. An example of token money is traditional physical cash, where the validity of the banknote itself is verifiable, while the identities of the sender and the receiver are private. Many other emerging forms of money fall somewhere in between. For example, many different crypto-currencies (e.g., Bitcoin, Z-Cash) resemble token money, as the identities of senders and receivers are not easily known,[2] while many means of payment provided by technology firms (*e.g.* Amazon Cash, or PayPal) enable to trace individuals' spending similarly to account money.

In this paper we propose a hybrid form of CBDC, a fiat cash scheme that embodies an **intentional asymmetry regarding privacy**, between receiving and sending money. We call it P-hybrid CBDC (abbreviated from Privacy-hybrid CBDC). [3] We argue that the central bank's digital currency (CBDC) should offer a better alternative with respect to privacy than account-based money when it comes to protecting the anonymity of individual spenders. Namely, it should be impossible (or more precisely, statistically close to impossible) to associate an individual with their purchases of coffee, groceries, alcohol, entertainment, medicine, etc.

We present a set of theoretical arguments and a privacy-centered model of money to highlight why this feature is value adding compared to forms of account-based money (or more broadly, to using financial assets as means of payment), and how it affects money demand. In our proposal, outgoing flows should hence bear no information on the identity of the sender. At the same time, we argue that full privacy regarding incoming money is a less necessary feature, as most flows of incoming money are subject to taxation and at least some institutions in the economy are entitled to have verifiable information on these flows. Linking incoming flows to individual identities, for instance via ID-linked accounts, facilitates the prevention of fraud, money laundering, tax evasion, and, perhaps more importantly, it can facilitate new forms of financial contracts, and enable new fiscal and monetary policy tools.

Overall, we will argue that such a P-hybrid CBDC has many similarities to physical cash, but enables further desirable features, sometimes referred as "Cash+" features. While such a P-hybrid CBDC does not need to pay interest, the ID-linked accounts can be associated with financial assets that do give returns to the account/asset holders.

From the system's architecture side, designing a hybrid system is more challenging than either of the extremes: a system where no transactions are private; or where all transactions are private. This paper shows how the ZeroCash [BsCG$^+$14a] approach, which relies on ZK (Zero-Knowledge) proofs, could be leveraged to achieve our main objective: a digital cash system that offers privacy for the spender while guaranteeing it is always possible to identify the receiver of a transaction. The paper also provides some estimation of

[1] We further note that the Bank of Canada's top desirable features for a cash-like CBDC are: 1) Safety, 2) Universal accessibility, 3) Privacy, 4) Resilience, 5) Competition and efficiency, and 6) Monetary sovereignty, and we discuss all these aspects. We tackle all these issues separately, pointing out those that we consider most essential from a theoretical/philosophical perspective (from the perspective of computer science, economics and finance), those that are ultimately a matter of policy choice, and those that necessitate practical trade-offs for the system's architecture, e.g. between safety and universal access

[2] Note that crypto-currencies are not as private and anonymous as payments in physical cash, and could be considered pseudonymous rather, as the sender's identity may be traceable with enough data. There also exists crypto-currencies and crypto-tokens (such as Ethereum), the role of which expands beyond fiat money and means of payments, *e.g.* by enabling smart contracts. Lessons learned from different cryptocurrencies motivate many of the technological solutions that facilitate the implementation of the P-hybrid CBDC we propose.

[3] Our term P-hybrid CBDC should not be confused with the term "hybrid CBDC" in *e.g.* [AB20], where it means that the central bank delegates retail payments to intermediaries while remaining the issuer of CBDC. In the case of P-hybrid CBDC the digital currency itself has a hybrid nature and treats money received and paid differently. As we discuss in Section 5, our proposed system enables both "direct" and "hybrid" CBDC as defined in [AB20], as the central bank can manage all functions the system architecture requires, or delegate some of these functions.

computing system requirements. It shows in particular that using today's technology, the system could potentially support in excess of $2,000$ transactions per second.

The rest of this paper is structured as follows. Section 2 presents theoretical models and arguments for both key aspects of our proposal — privacy as a key driver of demand for cash (or CBDC) and the key benefits of records on incoming money being visible. Section 3 provides the technical details of the system architecture from the Computer Science perspective, and section 4 analyzes the system scalability. Section 5 discusses the choices, roles and policy trade-offs of the system actors and economic agents interacting with the system. Section 6 is about promoting adoption and covering initial investment and finally, section 7 concludes.

2 Economic arguments for the P-hybrid CBDC

This section has three parts. In Section 2.1. we model the value of spending privacy, and the resulting demand for cash-like (fiat) money. Section 2.2. focuses on cash+ and regulatory compliance benefits that the ID-linked accounts and records of incoming payments can bring. Section 2.3. brings both arguments together to make a more detailed case for a hybrid digital cash initiated and monitored by a central bank.

In this section, we will focus on stylized theoretical arguments using tools from economics, and will discuss the feasible technological solutions and policy choices that the implementation of such a P-hybrid CBDC involves, and the interactions with the private sector (be it traditional financial institutions, FinTechs, non-financial technology firms and retailers) in subsequent sections.

2.1 Value of privacy and resulting demand for cash

The recent emergence of multiple new forms of means of payment poses a challenge to traditional models of fiat money and warrants revisiting the question of what the key sources of demand for cash-like fiat money are (be it physical cash, CBDC, and fiat-money-like digital currencies). As mentioned in our introduction, we find it informative to separate two classes of traditional models of money: those that directly or indirectly impose the assumption of some form of Clower constraint - money as the only means of payment; and those that focus on money as means to overcome the double coincidence of want across goods and time.[4]

Taking these considerations at face value, there should be no demand for any form of cash-like fiat money, because there exist many other forms of highly liquid means of payment that offer greater than zero nominal returns (e.g., most bank cards satisfy this criterion). Continued demand for cash-like fiat money is even more surprising, as using one's bank card involves a lower cost and is less trouble than withdrawing cash from an ATM or figuring out how to acquire and use bitcoins or other cryptocurrencies.[5]

This section emphasizes privacy concerns as a new key separating factor between using cash-like assets (that offer a degree of spending privacy) and equally liquid financial assets (that offer less privacy,[6] but offer interest/returns on idle money) as means of payment. Privacy concerns[7] are traditionally not integrated in models of money and monetary economics, and we propose an approach that can be used both in micro-economic and macroeconomic contexts. See also related models by [KRW18], [GV19]. Privacy concerns also feature in many recent debates around privately created forms of digital money, *e.g.* crypto-currencies,

[4]The seminal models are by Baumol and Tobin [Bau52], [Tob56], many other papers, including Romer (1986) [Rom86] we build on have formulated this idea in general equilibrium context. These models have further been used to rationalize cash-in-advance constraints and money in utility (see [BF89] for a review) which ultimately feed into models of monetary policy; and those that emphasize the liquidity aspect of money as means to overcome the double coincidence of wants, trust and commitment issues (see e.g., [KW93], [KM02], and [KM18]

[5]Of course there are other reasons why individuals may use cash, such as habit, them not having a bank account etc. In our model we will focus on rational consumers who can use financial assets as means on payment, and discuss unbanked further in later parts of this paper. It is also important to note that the current driving rationale for demand for cryptocurrencies has less to do with their instantaneous value as means of payment, and is more to do with benefits of early adoptions and bets on future currency/payment value of these currencies, somewhat similar to any foreign exchange/currency speculation.

[6]Private sector players may lack incentives to offer privacy preserving means of payment because the can monetize individual information. Financial institutions may use it to assess credit worthiness [BBGP20], and technology firms selling goods and services may have an incentive to price discriminate (see *e.g.* [GL20] and [AV05]). As another benchmark, [BN19] consider a model without privacy concerns, and highlight the conditions for equivalence of publicly and privately generated money.

[7]A recent speech by Augustin Carstens, General Manager, Bank of International Settlements [Car21] includes recent survey results on privacy preferences across countries.

but these come with an additional and different set of system considerations that are specific to prominent verification systems (*e.g.* Proof-of-work).

For this reason, we first suggest a unifying modeling framework of privacy concerns, highlight the real effects that privacy preferences imply, and then analyze the transaction demand and the key drivers of choice between different forms of liquid means of payment.

2.1.1 Modeling privacy preferences and the real effects of privacy preferences

When modeling privacy preferences, we take inspiration from the recent literature on differential privacy (see [DR$^+$14] for a review)[8], and adopt the functional form of privacy concern where ϵ is a parameter that captures how important the privacy concern is.

$$\epsilon \ln \left(\frac{\Pr\left(action|true\_type\right)}{\Pr\left(action|another\_type\right)} \right). \tag{1}$$

A broad message from this literature is that hiding the identity of individuals does not protect them from the possibility of statistical inferences about personal characteristics that individuals may rather wish to keep private (*e.g.* whether they have a disease, a socially condemned habit, or belong to a disadvantaged racial or gender group).[9].This consideration is increasingly important in an environment where data analytic capabilities are improving. Furthermore, empirical findings by [BBGP20] show that even seemingly innocent data, such as the time of the day an individual interacts with a retail platform, or the device used for the transaction (iPhone or Android phone), can reveal information about a consumer's credit worthiness. A related insight from the differential privacy literature is that a degree of randomization may be needed for efficiency[10] We apply these insights to our particular question of the privacy effects on money demand and firm-consumer interactions.

Consider a stylized setting, where there is a risk neutral firm which has zero marginal cost and aims to maximize profits/sales. Without loss of generality, the firm has monopolistic power and is able to set the price. There are N potential consumers, and each consumer j is also risk neutral and characterized by two type parameters: their utility of consuming the product which they value either highly or not at all, i.e., $v_j \in \{0, 1\}$, and another personal characteristic $T_j \in \{A, B\}$, about which the consumer has privacy preferences. While the characteristic T_j has no direct relevance for the firm, all participants are aware of the joint distribution. For example, consider this numerical example:

	$v_j = 1$	$v_j = 0$
$T_j = A$	2/6	1/6
$T_j = B$	1/6	2/6

,

which implies that type A is more likely to value the good highly compared to type B, i.e., $\Pr\left(v_j = 1|T_j = A\right) = \frac{2/6}{2/6+1/6} = \frac{2}{3}$ and $\Pr\left(v_j = 1|T_j = B\right) = \frac{1}{3}$.

Further assume that both types are privacy concerned and obtain disutility from revealing their type, in the spirit of 1, i.e., consumer j's utility from buying the product is

$$v_j - p - \epsilon \ln \left(\frac{\Pr\left(action = buy|T_j; information\_revealed\_in\_PBE\right)}{\Pr\left(action = buy|T_{-j}; information\_revealed\_in\_PBE\right)} \right),$$

where p is the price the firm sets, T_j is the true character type (e.g., $T_j = A$, when the consumer has true character A), and T_{-j} is another possible, non-true, character type (e.g., $T_{-j} = B$, when the consumer has

[8]The differential privacy literature focuses on a different question: how to best structure queries from a database such that it is difficult enough to identify an individual, but such that it still enables robust statistical inference

[9]The literature cites classical examples of a particular Netflix user being identified despite all data on identities being anonymized. Similar ideas feature in papers that analyze indirect discrimination by mainstream banks and Fintech, and compare the two, see e.g., [BMSW19] and [FGPRW20]

[10]One classical example of randomization as means to ensure some privacy involves a "yes-no" question and instructions of coin flips to answer a survey question about a characteristic that survey participants may feel uncomfortable to reveal (e.g. do they have a particular disease, bad habit, religious belief). To make the participants more relaxed one can use the following instructions: flip the first coin - if it is "heads" answer truthfully, if it is "tails" flip another coin and answer "yes" if it is "heads" and "no" if it is tails. Under these instructions, neither "yes" or "no" fully reveals whether the individual has the surveyed characteristic, yet on avarage there is a possibility of statistical inference.

true character A). PBE refers to Perfect Bayesian Equilibrium, which requires that all agents act rationally, and the beliefs about the agents' actions are consistent with equilibrium actions and Bayes' rule.

Using the same notation, the decision by consumer j to not buy the good gives them utility

$$-\epsilon \ln \left(\frac{\Pr\left(action = dont\_buy | T_j; information\_revealed\_in\_PBE\right)}{\Pr\left(action = dont\_buy | T_{-j}; information\_revealed\_in\_PBE\right)} \right).$$

Technical details aside, all agents are aware that the decision of privacy concerned consumers to buy or not buy the product may reveal information about their character type T_j, depending on equilibrium beliefs. As we prove in the associated Appendix (see Appendix A)[11], there are three types of interesting Perfect Bayesian Equilibria:

1. The one where the firm maintains all its customers (i.e., all consumers with valuation $v_j = 1$ buy the good) but the firm needs to sell its good at a discount, which in this example is

$$p = 1 - \epsilon \ln(4)$$

   and leads to the firm's profit $\frac{N}{2}(1 - \epsilon \ln(4))$.

2. The one where type A consumers use a mixed strategy to buy and not buy the product, such that their decision to buy or not buy reveals nothing about their type. In such equilibrium, the firm charges price $p = 1$, but loses consumers, as only $\frac{1}{2}$ type A consumers buy the product. The firm's profit is $\frac{1}{3}N$.

3. The one where type A consumers use a mixed strategy to buy and not buy the product, such that their decision to buy or not buy reveals something about their type, and the firms offers some discount (see further details in Appendix A.1)

A key message from this analysis is that interacting with privacy concerned individuals, whose purchase of a particular product is public information, implies losses for a firm, either via offering large discounts, or via accepting to lose some consumers who may have an incentive to alter their consumption behavior. This may be sub-optimal for two reasons: first, some innovative firms may anticipate lower profits due to privacy concerns and may not produce the good to start with; second, as we prove in Appendix A, the joint surplus of consumers and producers is always lower.

There are two further more subtle effects worth emphasizing. First, when comparing the outcomes under equilibria with full discount and full consumer driven privacy, i.e., case 1 and case 2 above, the firm's profit is higher under full privacy preserving randomization, whenever $\frac{1}{3}N > \frac{N}{2}(1 - \epsilon \ln(4)) \Leftrightarrow \epsilon > \frac{1}{3\ln(4)}$. That is, consumers altering their behavior leads to better equilibrium outcomes whenever privacy concerns are important enough.[12] Second, privacy concerns are fundamentally different from consumer decisions being correlated with "good" or "bad" types analyzed in many adverse selection models. If in the above example, being type A was "bad", we would reach the same equilibrium, while if being A was "good", type B would benefit from mimicking type A and the firm could sell their good at a premium, rather than at a discount.

Preserving privacy is hardest when purchases of each individual good are observed, as in the example above. When there are more than two goods, and we maintain our two character setting, what is gained by observing only the total money spent (as we propose) rather than individual purchases, depends on the statistical dependence structure of preferences.

To give some examples, consider that there are two goods, as described above, and one or two producers. As an extreme example, suppose that consumers' preferences for these goods are perfectly correlated, i.e., the conditional probability that type A (or B) values both goods highly is the same, whether we observe exactly what they bought or whether the total amount spent on identical goods is equivalent. Similar holds when consumer preferences are uncorrelated across time (or good). The latter follows from the fact that a sum is a sufficient statistic for independent Bernoulli trials (and whether or not the conditional distribution $\Pr(v_j = 1 | T_j)$ is different for $T_j = A$ and $T_j = B$ does not alter this argument).

[11] The model in our appendix considers a more general joint distribution matrix, and presents further details of the framework and proofs.

[12] Similar argument, albeit at a higher threshold, applies for joint surplus.

As another extreme example, suppose consumer preferences are perfectly negatively correlated, and the products have the same price. More formally, suppose that $\Pr\left(v_j^1 = 1|T_j\right) = \Pr\left(v_j^2 = 0|T_j\right)$, for all T_j, where v_j^1 reflects the consumer valuation for the first good and v_j^2 reflects the consumer valuation for the second good. In such case, observing the total amount spent does not reveal T_j as

$$\frac{\Pr\left(\left\{v_j^1, v_j^2\right\} = \{1,0\} \text{ or } \left\{v_j^1, v_j^2\right\} = \{0,1\} \ |T_j = A\right)}{\Pr\left(\left\{v_j^1, v_j^2\right\} = \{1,0\} \text{ or } \left\{v_j^1, v_j^2\right\} = \{0,1\} \ |T_j = B\right)} = 1$$

In this case the price of both products is one, and there are no distortions.

The last example could offset one credit score effect highlighted in [BBGP20], who find that those buying a product using credit in the evening are considered less credit-worthy than those buying it in the morning. When consumers internalize these effects, those preferring to buy in the evening may worry that their credit score may suffer as they may be seen as disorganized. On the opposite, there is also a view that those preferring to buy in the morning may have nothing more productive to do than shopping at that time, and this can also affect their credit scores. Indeed, there are examples of Fintechs which consider those who spend more on entertainment as more credit worthy, as their spending pattern is more flexible.

2.1.2 Privacy driven transaction demand of money and choice between means of payment

The key takeaway from Section 2.1.1 is that, due to privacy concerns, consumers may have an incentive to alter their purchasing behavior. We now analyze the behavior of such consumers in a stylized setting in the spirit of Romer (1986)[Rom86], which can be more directly related to familiar macroeconomic settings relevant for monetary policy. We focus on the determinant of their transaction demand of money, and the key determinants of their choice to use cash-like assets (instead, or in parallel) to using interest paying financial assets as means of payment. These questions are interesting from a theoretical perspective (*i.e.* why there is demand for cash-like assets in an environment where account-money is at least as convenient as means of payment), from an adoption perspective (i.e., what forces make the adoption of a private CBDC such as our P-hybrid CBDC more likely), and from the system's architecture perspective (i.e., what determines the frequency of conversions of financial assets to privacy preserving forms of cash).

We delegate the technical details of our economic framework and proofs to Appendix A, and highlight here only the key features of the setting and our main findings. Similarly to [Rom86], we consider a stylized continuous time model, where each consumer is endowed with a nominal amount $W = P_0^P Y$ (where Y is the real value and P_0^P is the price of the consumption basket) at date 0, and fully spends this amount within an interval $[0, T]$. For example, we can think of a consumer who obtains a monthly wage at the beginning of a month, spends a nominal amount W of this on consumption over the full month. While cash-like means of payment do not offer any nominal return (denoted with i), financial assets do. The real interest rate, r, and inflation, π, are exogenous and constant, Fisher's equation holds (i.e., $i = r + \pi$), and inflation affects all goods and baskets of goods similarly. If consumers use cash, they will choose how many times (N) and when to convert their financial assets to cash.

There are two key differences. First, we do not impose a Clower constraint and interest bearing financial assets can be used as means of payment at least as easily as cash. Second, we consider that consumers are worried that their purchasing decisions reveal information about their personal characteristics, which they want to keep private. Similarly to Section 2.1.1., we consider the simplest informative case with two consumer types who have a personal characteristic (A or B), and their personal characteristic is correlated with preferences for their favorite basket of goods. It is further worth noting that unlike the one-firm one-consumer interaction considered in Section 2.1.1, modeling aggregate effects requires considering the full set of goods available in the economy, which implies that when a consumer chooses not to buy something, (s)he will buy something else.

In Appendix A.2, we provide a formal theoretical rationale for modeling privacy concerned groups of individual consumers as a representative consumer whose instantaneous utility from consuming an endogenous privacy sensitive basket of goods at some specific point in time $t \in [0, T]$ is

$$u\left(c_t\right) = (1 - \varepsilon) \ln\left(c_t\right), \tag{2}$$

where $\varepsilon \in [0, 1)$ captures the instantaneous loss consumers incur, if their individual purchasing decisions are not private and observed (at least by some parties, e.g., their bank or technology firms offering means of

payments). We show that parameter ε must be positive, whenever ϵ (the degree of privacy concern in section 2.1.1.) is positive, and there is a utility loss associated with replacing one's favorite consumption basket with another, less favored consumption basket. Consumers suffer a utility loss at date t, because it is optimal to sometimes replace their favorite consumption basket with another to hide their personal characteristic.[13] To give some intuitive examples, when no purchasing decision is private and individuals worry about their choices being stigmatized, someone with diabetes has an incentive to sometimes buy a less healthy version of bread, and some other individuals may alter their preferred consumption behavior in multiple ways if they expect their choices will affect their credit score chances to get a favorable mortgage arrangement, etc. At the limit case where $\varepsilon = 0$, i.e., privacy effects are mute, the instantaneous utility of consumption (2) is the same as in Romer (1986) [Rom86].

Building on our analysis and discussion in Section 2.1.1 we further argue that ε is a parameter which captures the maximum privacy driven instantaneous utility loss that a representative privacy concerned consumer obtains. Namely, **the maximum utility loss consumers suffer in our framework due to the lack of spending privacy is obtained in one of these two circumstances:**

1. Consumers only use a non-private financial asset as means of payment.
2. Consumers convert their financial assets to private cash at arbitrarily short frequencies.[14]

While both of these strategies imply a utility loss due to privacy, these also maximize the interest rate (or more broadly any form of returns) obtained from holding financial assets. Case 2 above is more subtle, but intuitive when considering our analysis in Section 2.1.1., where it does not make any informational difference whether an outside party observes that a consumer bought a particular product, or that (s)he spent an observable amount of money to buy this good. More broadly, converting financial assets into private cash just before buying an item may give the illusion of spending privacy, but when the frequency of such conversions is high, it is not very difficult (from a statistics perspective) to figure out what the consumer bought.[15]

Frequency of conversions Returning to the issue of the transaction demand of cash-like money, we consider that converting financial assets to private cash comes at a utility cost of privacy which is lower when consumers make these conversions less frequently. Consequently, we assume that the relationship between the privacy cost and the number N of conversions in the interval of length T is a function $\xi(N)$, where $\xi(0) = 0$, $\lim_{N\to\infty} \xi(N) = \varepsilon > 0$ and $\xi'(N) > 0$.[16] Considering all the above, a representative consumer who uses private cash as means of payment has utility

$$U_{TD}^{M} = \int_0^T (1 - \xi(N)) \ln(c_t)\, dt - Nb,$$

where b is the cost of taking the trouble to convert a financial asset to private cash (e.g. a trip to an ATM or bank in the case of physical cash, arguably zero when turning a digital asset into another one). Consumers choose how much to consume at each point in time, when to convert, and how many times to convert financial assets into money.

[13] We show that everyone always buying their favorite basket is not sustainable as equilibrium strategy when purchasing decisions are never private. Therefore, even if ϵ is very small, consumers must mix between their favorite and less favorite baskets of goods in a contemporaneous Bayes-Nash Equilibrium. We consider symmetric mixed strategies and derive (2).

[14] At the extreme, just before any specific purchase of a grocery item: be it a piece of bread, a bottle of alcohol, or a cinema ticket.

[15] Here is one example: someone buys a ticket for a show. If a statistician knows when exactly they bought "something", when shows usually start and what they cost, and what exact amount the consumer spent, the statistician can assign a high probability of it being a ticket for a show. There are many less obvious cases highlighted in the differential privacy literature (e.g., a specific Netflix user identified based on limited information on their movie choices, and despite anonymized data), see D[DR+14]

[16] The assumed functional property that $\xi'(N) > 0$ stems from the observation that under most general statistical structures, observing the total amount spent over a longer interval reveals less information about what a consumer bought, and thus what his/her personal characteristic type is. The limit case of $N \to \infty$ was discussed above, and the limit case $\xi(0) = 0$ corresponds to the situation where the consumer does not convert assets in the middle of interval T. We consider a general functional form of $\xi(.)$ with mentioned properties, because we neither need nor want to take a precise stance about the details of the statistical learning process.

We find that a consumer either always uses cash and never converts (after date zero), or optimizes the trade-off between interest bearing savings in financial assets and privacy loss, and converts financial assets to cash at regular intervals, where the optimal time between conversions, μ, is[17]

$$\mu \equiv \frac{T}{N+1} = \sqrt{\frac{2b + 2\frac{\partial \xi(N)}{\partial N}\left(\ln\left(\frac{Y}{T}\right) + \frac{rT^2}{2}\right)}{i \cdot \left(1 - \xi(N) + \frac{\partial \xi(N)}{\partial N}(N+1)\right)}}. \tag{3}$$

Our benchmark model [Rom86], is a special case where privacy concerns do not matter and the optimal time between withdrawals is

$$\mu = \sqrt{\frac{2b}{i}}.$$

Notice that when privacy concerns are mute (i.e., $\xi(N) = 0$), and it is (becomes) easy and costless to convert interest bearing financial assets to cash, i.e., $b \to 0$, the optimal time between withdrawals becomes zero. This would render any form of non interest bearing cash useless, unless cash must be used as means of payment.

In contrast (3) shows that rational privacy concerned individuals, using physical cash or the P-hybrid CBDC as means of payment, will maintain infrequent withdrawals even if $b \to 0$. This provides a potential rationale why physical cash demand has not disappeared, and why there is still a negative relationship between (high powered) money demand and interest rate. Stretching this argument further (and possibly too far), it can explain why there is demand for fiat-money-like cryptocurrencies (e.g., bitcoin).[18]

Another observation from (3) is that when consumers are wealthier, i.e., have higher Y, they convert financial assets to cash less frequently. Relatedly, privacy considerations generate a money demand function that not only incorporates the negative interest rate effect, but naturally generates a positive income effect on money demand.

A key finding relevant to the system's architecture is that rational privacy concerned consumers would not transfer their financial assets to private cash very frequently, which puts less pressure on the system. Things would be different if consumers would not care about privacy, but as we will show next such individuals would not demand cash-like money. Furthermore, (3) shows that more wealthy individuals have less incentives to convert their financial assets into cash frequently even if cash is their preferred means of payment. We can then speculate based on this stylized setting that when the economy grows, the average (rational) consumer who becomes wealthier will have incentives to convert their financial assets into cash less frequently over time, while the computer power necessary to manage transactions is likely to improve over time. Both of these factors may make the P-hybrid CBDC or any other form of digital cash-like CBDC increasingly sustainable over time.

Choice between means of payment Let us now consider the choice of consumers to either always use an interest bearing, but not privacy preserving, financial asset as means of payment vs. using non-interest bearing but private cash-like money to buy goods. For the sake of clarity, we abstract from the nuances of the timing of conversions of one form of assets into another discussed above, but re-introduce the possibility of distortions for producers, by considering that in addition to baskets that involve privacy considerations, there also exist baskets of goods that are privacy neutral (e.g., buying light bulbs).

A consumer who always uses financial assets as means of payment during the interval $[0, T]$ solves

$$\max_{c_t} U^{FA} = \int_0^T \left\{\beta(1-\varepsilon)\ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right)\right\} dt, \text{ s.t.}$$
$$P_t^P c_t^P + P_t^N c_t^N + \dot{b}_t = i b_t;\ b_0 = P_0^P Y \text{ and } b_T = 0,$$

where c_t^P and c_t^N are quantities of privacy sensitive basket and privacy non-sensitive basket consumed; ε is a privacy parameter as earlier; β is a preference parameter between the baskets of goods; b_t is the nominal

[17] See details in Appendix A.3

[18] These forms of money were created outside the traditional financial system, and have at least two appealing characteristics relevant to our setting: being digital and protecting privacy more than digital account money. Whether or not these currencies become a liquid means of payments is a different question.

amount of financial asset held at date t, and P_t^P and P_t^N are the prices of the privacy sensitive and the non-sensitive basket, respectively.

A consumer who always uses cash-like assets as means of payment solves

$$\max_{c_t} U^M = \int_0^T \left\{\beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right)\right\} dt, \text{ s.t.}$$
$$\int_0^T \left(P_t^P c_t^P + P_t^N c_t^N\right) dt = P_0^P Y.$$

and a consumer who uses cash-like assets to buy the privacy sensitive basket and financial assets to buy the privacy non-sensitive basket solves

$$\max_{c_t} U^{M//FA} = \int_0^T \left\{\beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right)\right\} dt, \text{ s.t.}$$
$$P_t^N c_t^N + \dot{b}_t = i b_t;\ b_0 + \int_0^T P_t^P c_t^P dt = P_0^P Y; b_T = 0.$$

Without loss of generality, all the above settings assume that the representative consumer's initial endowment is represented in units of the privacy sensitive basket, and the details of the solution are in Appendix A.4.

The most expected finding is that $U^{M//FA} - U^M = (1-\beta)\frac{iT^2}{2} > 0$, i.e., it is better to use privacy protecting means of payment to buy privacy sensitive goods, and interest rate maximizing assets to buy privacy non-sensitive goods. This is not surprising as the consumer is shielded from privacy concerns in either case. Whether or not consumers would pursue such nuanced behavior is subject to considerations outside the model, e.g., there may be attention costs associated with optimally managing multiple means of payments. More importantly, the choice between using cash for all baskets or some baskets does not depend on the privacy parameter, ε.

The choice between using only financial assets or cash is more important for the question of the sources of value of a CBDC and its adoption incentives (be it comparisons between U^{FA} and U^M, or $U^{M//FA}$ and U^M)

We find (see Appendix A.4 for details) that **consumers prefer cash to financial assets as means of payment** when

1. Privacy concerns are important enough;
2. Consumers are wealthier;
3. Inflation is lower.[19]

The dynamic effects are best illustrated graphically (see Figure 2).[20]

Panel A shows that privacy concerned individuals prefer cash if their privacy preference (ϵ) is high and the share of privacy sensitive goods (β) in their optimal basket is high. Perhaps surprisingly, we find that wealthier consumers are more likely to prefer cash, and Panel B illustrates this effect. It is surprising because common wisdom suggests that poorer individuals are more cash dependent. A natural explanation of this seeming discrepancy is that these individuals may find it too costly to have a bank account. This implies that they use cash not because they prefer it, but because they cannot use financial assets as means of payment.

The observation that wealthier consumers are more likely to prefer private cash as means of payment echoes our earlier finding regarding the frequency of transactions. Note that this outcome is purely driven by the interplay between privacy preferences and wealth, and has nothing to do with the common wisdom

[19] Also the real interest rate affects choices not necessarily in the same way, depending on whether the alternative to using financial assets as means of payment is using cash always or using cash only to buy privacy sensitive baskets.

[20] These figures assume that inflation is 2% (standard target), real interest rate is 3% (close to historical averages in Canada and the USA), T=30 days (*i.e.* a month). The baseline scenario in panel A, assumes CAD$4000 monthly income (close to average net monthly income in Canada). These figures have been constructed under the assumption that the representative consumer uses cash-like money to buy privacy sensitive basket and financial assets to buy privacy non-sensitive basket. The case of using cash-like assets to buy all goods is qualitatively similar. The only difference is that the indifference frontier is flatter.

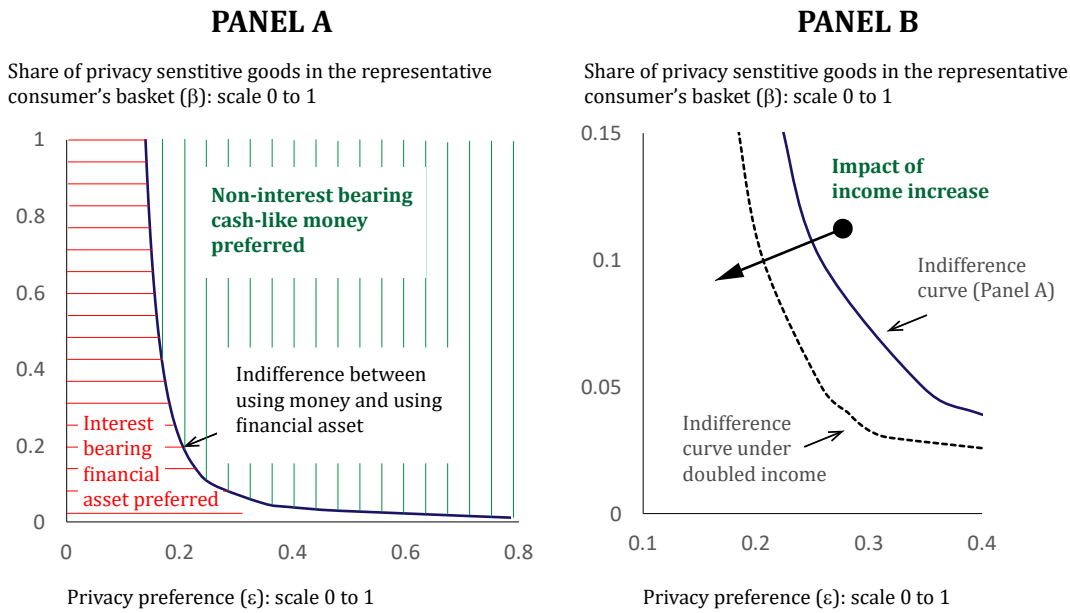

Figure 1: Consumers' optimal choice between cash-like money (privacy preserving and non-interest bearing) and liquid financial assets (interest bearing and not privacy preserving) as means of payment.

that richer individuals demand cash to avoid or optimize taxes (we discuss this issue in section 2.2). It has further consequences for the likely pattern of adoption along income distribution and over time.

Finally, similarly to the setting in 2.1.1., there are potential negative real effects to production whenever consumers' privacy concerns play a role. Suppose that it is optimal for consumers to use financial assets. From Appendix A.4 their optimal consumption is

$$c_t^P = \left(\frac{Y}{T}\right)\frac{(\beta-\beta\varepsilon)}{(1-\beta\varepsilon)}\exp(rt) < \left(\frac{Y}{T}\right)\beta\exp(rt),$$
$$c_t^N = \left(\frac{Y}{T}\right)\frac{P_0^P}{P_0^N}\frac{(1-\beta)}{(1-\beta\varepsilon)}\exp(rt) > \left(\frac{Y}{T}\right)\frac{P_0^P}{P_0^N}(1-\beta)\exp(rt).$$

and consumers optimally tilt towards buying the non-privacy sensitive basket whenever $\varepsilon > 0$. (The right hand side of the inequalities above reflects the case where $\varepsilon = 0$). However, this can be costly at an aggregate level as a privacy sensitive basket is likely to include more innovative goods, and investing in their equity may offer higher returns.

2.2 Benefits of keeping digital records of incoming money for regulatory compliance and innovation

While the previous sub-section argued that maintaining spending privacy is important for demand for cash (and other privacy preserving forms of token money), and that it helps to mitigate the distortions that privacy preferences can generate for the production side of the economy, we now argue that there are important benefits of maintaining verifiable records of incoming money. These records are a characteristic of different forms of account money (be it individual and business accounts held in banks or our P-hybrid CBDC). It is important to highlight that non-private digital accounts may but do not need to be visible to all economic agents all the time. Instead, there can be a limited set of economic agents, institutions and regulatory bodies who have the right to make (potentially aggregated) queries only when it is either necessary or value generating. We discuss the trade-off associated with the degree of privacy further in Section 5.2.

2.2.1 Regulatory compliance and taxation

One of most compelling reasons for having verifiable records of incoming money is that almost all sources of incoming money (be it wages, sales records of a company, gains from financial investments etc.) are anyway subject to taxation and need to be declared. Consequently, **economic agents who do not aim to avoid taxes do not need full privacy of funds arriving to their ID-linked accounts** (at least in front of tax authorities). Furthermore, as these records are digital, we can envision many technical solutions provided either by private firms or the public sector which automate the calculation of tax obligations associated with these incoming funds and ease the associated administrative burden that individuals and firms face. Such services are already provided to a degree (e.g., by banks who offer individuals access to calculated tax obligations on interest rate earned), and can be expanded.

The second important benefit is that the existence of digital records of incoming money can greatly facilitate compliance with anti-money laundering and terrorist financing regulations [Fin19] and facilitate the prevention of money laundering incidences. The inability to trace incoming funds is a known weakness of physical cash and to some degree also a weakness of known private sector generated cryptocurrencies.[21]

In our proposed P-hybrid CBDC system, all individuals can open an account only when they provide some form of identification (similarly to opening a bank account), and it should be very difficult for the same individual to open multiple ID-linked accounts. Similarly, corporate entities could open one or a limited number of ID-linked accounts associated with records of their corporation and owners' identities. This has several immediate benefits for compliance and prevention, such as

1. It can be made impossible (or very hard) for an illegal business (e.g., seller of illegal drugs or a terrorist organization) to open a valid ID-linked account and immediately benefit from the system.

2. Compliance with KYC (Know-your-customer) is easy to achieve as opening an account involves a form of identity verification. Furthermore, there are also efficiency benefits as individuals do not need to verify their identity multiple times, unless they need to gain superior rights.

3. Strategies like "smurfing"[22] are impossible or difficult as no entity can have an unlimited number of ID-linked accounts.

4. Many abnormalities (e.g., winning in casinos too frequently, or receiving an abnormal number of "gifts") which may suggest money laundering could be detected using data analytics and machine learning tools.

All this makes the placement stage of introducing the proceeds of crime into the financial system more difficult. Furthermore, as we will discuss in section 3, since the system's architecture will not allow private transfers of tokens between users of our P-hybrid CBDC, it is likely to make using such a CBDC relatively unappealing as a tool of layering to disguise illegal proceeds.

Of course, as long as physical cash remains in circulation, digital currencies which offer privacy for incoming money exist, and as long as not all countries adopt a form of CBDC (as sole means of payment) which prevents the reception of payments for illegal activities, no system can fully prevent money laundering. However, we argue that ID-linked accounts for incoming money can be at least as efficient for detecting money laundering incidences as the provisions of the current financial system, and are arguably more efficient because of the four reasons highlighted above.

2.2.2 Facilitating "smart(er) contracts" and financial innovation

One important traditional friction which prevents individuals and firms that do not have high enough net worth to obtain external financing is that their income is not verifiable or is expensive to verify (see [Tir10] and especially the chapter 3 in [Tir10] for a literature review). As digital account-based money (including

[21] It should be noted that digital currencies are pseudonymous rather than fully private and anonymous, and thus it is possible to analyze data which is available in the associated public blockchain, and find the individuals who use it for money laundering. However, while the more naive criminal users could be found, there are also mechanisms that make the tracing of individuals harder.

[22] Smurfing involves opening many accounts involving small sums that are below extra regulatory requirements of reporting to the government. .

our P-hybrid CBDC) provides verifiable records, it is therefore likely to facilitate more efficient external financing contracts.

Namely, when there are digital records of incoming flows, the verification costs are greatly reduced, which alone can make external financing cheaper for borrowers, while not reducing the returns of investors.[23] Furthermore, reliable digital records of incoming funds come with the associated possibility to write "smart contracts", i.e., contracts that specify pre-agreed contractual terms (e.g., cash flow rights) as a computer code which is a function of future recorded incoming flows. As the existence of verifiable records makes commitment and the execution of contractual terms easier, it facilitates the creation of new innovative financing contracts that are superior to existing loans and other forms of external financing contracts (see e.g., [Tin18]).

To give a few specific examples, consider first a student loan. If an individual's salary always arrives to his/her ID-linked account, it is easy for a financial institution which manages this student loan (be it a mainstream bank or a FinTech) to automate their student loan re-payments as a function of the student's future income (e.g. wage income), and to automatically trigger any pre-agreed conditions for loan breaks.[24] Furthermore, many student loan schemes (*e.g.* in the United Kingdom, already incorporate some insurance against negative income shocks such that those earning more have to pay relatively more, and those who earn a very low salary have their loan forgiven. Such a loan/insurance scheme is more efficient ex-ante as, at the time of their studies, all students are selected as having the ability to earn a good salary, but various shocks are likely to make some of them to earn more, and some of them to earn less, ex-post. Any such insurance features are more easy to manage when incoming flows are verifiable. One social benefit of this is that it can avoid the inefficient widening of income inequality due to the student loan's burden being too heavy on unlucky students and too easy to repay by the luckiest ones.

Consider a business which needs funds for investing in a new product. Financing the investment with new debt may be difficult if the firm already has a substantial debt, and may lead to foregoing profitable investment opportunities and to inefficient liquidations (i.e., debt overhang problem). Similarly, it is often difficult for producers of innovative products to obtain external financing using debt, as the risk of failure is high, and investors only obtain a small fraction of returns in case of success. However, if the firm's new project can be associated with the sales success of a particular product, which is recorded in their business-line specific ID-linked account, credible product-specific and equity-like financing contracts become possible on a wider scale. Equity-like financing contracts are often more efficient for both firms or individuals (in need of funds), and investors alike. This is because both the risk of failure and the possibility of extraordinary success is shared.

From a theoretical perspective, there could be many new forms of financing contracts that are even more efficient and can be built on verifiable records of a product sales [Tin18]. From a practical perspective, there exist FinTech initiatives that propose novel and potentially more efficient external financing schemes which do not fall easily into the traditional categories of debt or equity.

Existing FinTech initiatives and crypto-tokens already offer some "smart contract" functionalities. For example, via the Ethereum platform it is possible to use their inbuilt "smart contract" functionalities and their native Ethereum cryptocurrency to create new tokens that record investors' contractual rights, as well as to use Ethereum and other crypto-tokens associated with this system as means of payment for a particular service. There was a mini-boom of ICOs (initial coin offerings) which peaked around the late 2017 and early 2018, and attracted both "scams" and legitimate innovative producers and platform initiatives. More recently, STOs (Security Token Offerings) refer to a somewhat similar (to ICOs) private sector effort to enable novel financing contracts, but with a greater attention to compliance with KYC and anti-money laundering regulations. Overall, these and many other (e.g. crowdfunding) private sector initiatives indicate that there is demand for innovative financing arrangements and that many innovative FinTechs are eager to fill gaps in the market.

We further argue that many of these legitimate FinTech initiatives could flourish more if there were to be a (more) universal system of digital records of incoming flows, which a CBDC (as well as other forms of account-based-money) can provide. One key concern which limits innovation in financial contracting is that there are currently many parallel systems of physical and digital means of payment and systems of

[23] Indeed, the recent literature on the benefits of blockchain technology highlights the benefits of having verifiable records (see e.g., [CG16] and [Tin19] for reviews)

[24] A similar argument applies to individual mortgages.

recording incoming flows (be it based on bank records or more novel forms of blockchain based crypto-tokens). For example, an innovative firm with an ex-ante valuable investment project may wish to use an intelligent contracting scheme based on the Ethereum system to raise funds against its future sales of their innovative physical product (e.g., a domestic robot or a modernized bicycle). While the Ethereum platform enables a variety of contracting arrangements, it cannot ensure that the sellers of a physical product will have incentives or an obligation to report their sales of their product on the Ethereum ledger. Instead, they may sell their product for physical cash, accept transfers to their bank account, or accept any non-Ethereum cryptocurrency payment. This reintroduces the traditional verification problem - investors betting on the success of this firm's product, cannot be sure that the sales reported on the Ethereum ledger capture all sales (or a substantial part of them).

A complementary concern from the firm's side is that using crypto-tokens as means to raise funds may put them at odds with regulators. As a CBDC naturally comes with an official guarantee that accepting this means of payment is legal, contracts built on ID-linked accounts should be legitimate as well. This frees innovative firms and innovative financial institutions and Fintech startups from many regulatory worries that are not directly related to the contracts they aim to offer, and are to do with using specific privately generated crypto-tokens to implement their vision.

We note that welcoming innovation in financial contracting necessitates some more open-minded approaches on the regulator's side, as the optimal regulatory supervision relevant for welcoming financial contracting innovation is likely to be different from maintaining the status quo.

2.2.3 New tools for fiscal and monetary policy

Individuals having ID-linked accounts greatly facilitates the transfer of government subsidies to any target population. For example, imagine that a CBDC with ID-linked accounts would have existed when the Covid-19 pandemic emerged, and the government wanted to help the ones most affected (as it did). The system of transfers would have been easier, as the government could have transferred funds quickly and directly to target individuals' and firms' ID-linked accounts. The same applies to subsidies of any kind in normal times.

When it comes to monetary policy, theoretical models indicate that there is sometimes value in making "helicopter drops of money", i.e., issuing new physical or digital notes, and sharing new notes among the population. The benefits of helicopter drops of money emerge when there are deflationary pressures and economic agents have an incentive to save too much rather than spend. It is a problem because of the zero-nominal-interest bound, which may be binding in recessions. While the theoretical benefit is that it directly transfers new money to consumers and firms, generating helicopter drops of money is difficult under traditional financial structures [25] that rely on the central bank interacting with financial intermediaries and interacting with the rest of the economy only via these intermediaries. The problem is that financial institutions may rather benefit from absorbing the new high-powered money generated, and the monetary boost does not get passed on to the rest of the economy.[26]

During and following the financial crisis of 2007-2008, as well as during the Covid-19 crisis, many central banks opted for "unconventional" monetary policy tools, such as quantitative easing strategies, to inject money into the economy more directly. While these actions are closer to helicopter drops of money, they are still not fully equivalent and somewhat more complicated.

If all (or most) economic agents were to have ID-linked accounts associated with a CBDC, helicopter drops of money would become rather easy - the central bank could simply issue a new high powered cash-like digital currency and spread it across ID-linked account holders in Canada. Furthermore, the central bank could make targeted helicopter drops of money and spread them across a narrower subset of ID-linked account holders in Canada.

We do not take a stance on whether or when a central bank (e.g., the Bank of Canada) should use this policy tool, but just note that this additional monetary policy tool exists.

[25] In theory, helicopter drops of money could also be sent to individuals' bank accounts, but it is not common, perhaps due to the current privacy framework. Furthermore, such an approach may discriminate against individuals who do not have a bank account, and lead to too much money being sent to individuals who have multiple bank accounts. Our approach with one ID-linked account per person would facilitate such a process.

[26] It is rational for financial institutions to do so, as deflationary pressures tend to be associated with recessions that coincide with banks' assets being under a distress.

2.3 P-hybrid CBDC: a solution to consolidate the benefits of token money and account-based money

In section 2.1 we highlighted the reasons why spending privacy is necessary for cash-like token money demand, and in section 2.2 we indicated the reasons why account-based money which maintains digital records of incoming money is needed for regulatory compliance, and enables "cash+" features which facilitate financial contracting and expand fiscal and monetary policy tools.

Instead of envisioning a CBDC which needs to compromise one aspect for another, we argue that a CBDC that adds the most value can achieve both. Our proposed P-hybrid CBDC incorporates an intentional asymmetry regarding privacy and a radical separation between money spent and money received. Namely under our P-hybrid CBDC

1. individual spendings are physical cash-like fully private, i.e., not traceable even by those managing the CBDC system.

2. incoming money arriving to an individual or firm's account is not fully private - it is observable at least by tax and regulatory authorities, and there is value in allowing financial institutions and other private sector participants to access parts of this data to build novel interfaces and financing contracts.

We argue that, without the cash-like privacy feature in point 1, there would be no demand for CBDC, and without point 2, a CBDC would not be different from account-based money, which can be provided by financial institutions or technology firms. However, consolidating these two features can make the P-hybrid CBDC more appealing than:

1. Physical cash - because it is digital and is as convenient to use as a bank card; and enables better monitoring and contracting.

2. Account-based money (offered by banks or technology firms) - because it is better in preserving spending privacy.

3. Private sector generated pseudonymous cryptocurrencies - because the central bank's endorsement and commitment to convertibility to physical cash makes holding it less risky, and more liquid;

4. Private sector generated pseudonymous crypto-tokens - because the official backing by the central bank can incentivize a more widespread use of CBDC, which facilitates financial contracting.

Our take on privacy emphasizes the need for protecting individuals from statistical inferences based on their spending behavior (in the spirit of the differential privacy literature [DR+14], which may turn out to be discriminatory, and can cause negative real effects. And we argue against the need to protect privacy when it comes with possibilities to engage in tax evasion, money laundering and other criminal activities. The degree to which ID-linked accounts are not private is a policy and cultural choice. It could be that everyone's income is public (as in Sweden), or that it is observed only by governmental institutions. When it comes to firms' incoming money, there are less arguments for it not to be observable by the general public.[27]

3 System Architecture

3.1 Desired System Properties

In order to fulfill the vision of *asymmetric privacy* exposed above, the following properties should always be guaranteed by the system:

1. **Compliance**: The identity of the receiver(s) of digital cash and the transacted amount should be observable by an authority[28];

[27] As caveat, there is a possibility of collusion or industrial espionage (see [CH19] when everyone observes cash flows when these arrive. This problem can be resolved when the firm's cash flows are made public with a delay. The latter would not be of worry for investors, as long as the code specifying the contractual terms is public.

[28] See section 5 for more information about the role of various authorities.

2. **Privacy**:
    - Only an authority should be able to establish the identity of the receiver(s) of digital cash.
    - No one — not even an authority — should be able to identify the sender(s) of digital cash;
3. **Transparency**: Any user should be able to observe that no transactions break the rules of the system, without affecting privacy.

Compliance As exposed in the previous section, our system should allow some economic actors, such as the tax authorities, to observe any incoming transactions and their amount, and identify the receiver. With this property, compliance can be achieved since it becomes impossible for users to hide their income.

Privacy We want to ensure that full privacy is achieved when it comes to spending digital cash while guaranteeing some level of privacy when it comes to receiving digital cash. In our system design, spending is completely private: although anyone can observe that transactions are taking place, no one should ever be able to establish who is actually spending digital cash.

On the other hand, receiving digital cash is semi-private: only a certain set of actors are able to determine who is receiving digital cash to satisfy compliance. As explained in the previous sections, this is a good compromise as it is already the case with today's system that any income is known to the authority (or at least it ought to be).

Transparency The final property we wish to uphold is transparency, following the "Trust, but verify" approach. Special actors in the system, known as block-creators, are responsible for enforcing the rules. However, the important point is that any user should be able to verify for themselves that the rules are indeed enforced. While the system does not need to be open, having a transparent and public system is desirable to increase user confidence and trust in the system. This gives reassurance to the users that big actors such as the government or financial institutions are kept in check.

We argue that to satisfy this property, the system should be built using *open* protocols and all transactions should be visible to everyone (while ensuring it is not possible to identify a spender). As we will see later, the use of an open block-chain provides transparency while the use of zero-knowledge proofs guarantee privacy. Zero-knowledge proofs are mathematical proofs that allow someone to prove that he possesses some information, without revealing the information itself. In the context of our system, zero-knowledge proofs are used by a user when spending his digital cash to prove:

- the user owns the digital cash they wish to transfer, without revealing who the user is;
- their digital cash they wish to spend has been fully checked for compliance by a specific authority, again without giving any details on when it was checked and who the user is.

3.2 Ledger technology choices

This section highlights the choice of technology to implement a system that satisfies the desired properties.

3.2.1 Transaction record: public Blockchain

The first choice we consider is about the type of ledgers used to record transactions. We advocate the use of a public blockchain ledger, with an open standard.

Having a public ledger is necessary to satisfy the *transparency* property of the system. Any user should be able to observe that the rules of the system are enforced by everyone, holding any authority accountable. The drawback of a public ledger is that it could potentially increase the attack vector since the system specifications are here for everyone to see. However, we argue that security by obfuscation is a bad idea and that having open specifications and an open system that can be scrutinized by anyone offers a strong security model in the long term.

The use of a blockchain, similar to Bitcoin [Nak08] or Guardtime which is based on [BLLV98], is very suitable for a CBDC. All transactions produced within a certain time frame are batched into a single block.

This ensures an efficient validation process since a block is validated all at once and then distributed over a peer to peer network between the users of the system, achieving *transparency*. Furthermore, the use of a Merkle tree to store transactions in a block allows users to verify specific transactions without having to download the full blockchain. We target a block production time of 2-3 seconds to ensure a fast settlement time similar to that of physical cash. Section 4.2 discusses the expected scalability of this approach.

Network fees Similar to other blockchain systems, we advocate the use of a small, sub-cent, *fixed* fee when creating transactions on the system. This fee would serve two purposes. First, it would act as a flood gate to prevent a user from bringing down the system by producing too many irrelevant transactions. Secondly, this fee could be collected by the block creators — who decide which transactions are accepted by the system — to cover their operating cost. For instance, assuming a small fee of \$0.001 per transaction, with an average of $2,000$ transactions/s, this would produce $\$172,800$ of revenue per day. Section 6 discusses further means to recover the initial costs of the system development and maintenance.

3.2.2 Consensus method: PoA

The major contribution of Bitcoin [Nak08] is the use of PoW (Proof-of-Work) for achieving consensus on a blockchain. Under this approach, the main role of the miners is to define what the new state of the ledger is, by adding a block of valid transactions to the blockchain. PoW is used to select which of the miners will get to make the next decision. Interestingly, the definition of the rules of the system can be influenced by the miners, as long as a consensus is reached by at least 50% of miners. Since miners have to expend resources (*e.g.* the initial equipment cost and energy) to have the chance to become the next decision maker, they are naturally inclined to play by the rule of the system and reach consensus. Otherwise they risk seeing their costly blocks being orphans. This is what makes blockchains like bitcoin trustless and decentralized.

In the case of a centrally controlled CBDC, by definition, the rule of the system is only decided by the central authority (*e.g.* the central bank). Therefore, the use of PoW in the context of CBDC would not make any sense. Instead, such a system should use PoA (Proof-of-Authority), where an appointed set of publicly recognized block-producers are used. Their role consists in deciding which transactions are valid, batch them into a block, and add it to the blockchain.

The decision on which block-producer ought to produce the next block could be done using a simple round-robin approach. However, a more robust approach could consist of emulating PoW by having each block-producer sample a value from a random distribution as soon as a new block is added to the block chain. The parameters of the distribution would depend on the number of block-producers in the system and be set to guarantee a specific target block time. After sampling, a block-producer would "sleep" with a timer duration set to the value sampled. If a new block has not been added to the blockchain during the sleep, the block-producer would then go ahead and produce a new block. This approach would add robustness to the system since block-producers could be added or removed from the system. This would also provide some protection against an attack such as DDoS (Distributed Denial-of-Service) as long as enough block-producers remain to perform their task.

Given the *transparency* property of the system, this ensures that the block-producers play by the rules. Any deviance from the rules would be detected by any system user (or more precisely by the software used by the users). In such case, the system could simply come to an emergency halt, until the bad block-producer is removed from the system. Alternatively, the block from a bad producer could be orphaned by the other block-producers to avoid stalling the system.

It might sound surprising that a block-producer might be willing to miss-behave. Especially given the fact that block-producers are not anonymous since they would be appointed or licensed by the central bank. A more likely scenario would be that a bad block is the result of a software bug. This is why multiple block producing node implementations should be encouraged to add robustness to the system.

3.2.3 Balance model: UTXO

Blockchains are typically implemented using either a UTXO (Unspent TransaXion Output) model as in Bitcoin [Nak08], or using an account model as in Ethereum [But14]. Both models have various pros and cons [Eth18] summarized below.

The UTXO model is simpler to scale, since verifying transactions *within* a given block is an embarrassingly parallel problem. In addition, there are added privacy benefits if new addresses are used for every received payment.

On the other hand, the account model is more space efficient with transactions and storage. Compared to Bitcoin, a typical Ethereum transaction is 2–3× smaller resulting in less network bandwidth. The total amount of storage required is also smaller, currently 2.6× smaller as of Dec 31, 2020. Finally, it also offers greater fungibility and makes it easier to implement smart contracts.

We argue that for a CBDC, the UTXO model is preferred for the following reasons. First, the main focus of a CBDC should be on supporting cash-like transactions, without adding unnecessary complexity to support advanced smart contracts features that are easier to implement with an account model. Secondly, the ability to scale by parallelizing the validation of transactions within each block is crucial in order to achieve a high system throughout. We also argue that it is more desirable than a mere 2–3× space and network bandwidth savings. Thirdly, the fungibility benefit of an account based system is irrelevant, since our proposed approach achieves fungibility through a different mechanism as we will see later.

So in the rest of this paper, when we refer to an ID-linked account, this account is in fact composed of many pairs of public/secret keys. Each pair of keys is controlled by the user associated with the ID-linked account. The balance on the ID-linked account is determined by aggregating the information contained in the UTXO as is the case with blockchain implementations such as Bitcoin.

Note that the funds stored in a given address is known to everyone using the system. However, except for the ID-linked Database Keeper(s), discussed below, users do not know the identity of the owner of an address, nor can they identify that a given set of addresses belong to the same user.

3.3 System actors

The system is composed of two main classes of actors. The first class, the *users*, refers to either an individual (*e.g.* John Doe) or an institution (*e.g.* McGill University). Users of the system are the ones producing transactions to exchange their digital cash in the form of CBDC.

The second class of actors are responsible to ensure the overall running of the system. To maximize privacy and separate as much as possible the different concerns, we identify four major system actors. We argue that it is best to keep these actors separated from each other to avoid concentrating too much power in the hand of a single entity.

Identity Certifiers These actors, appointed by the central authority, are responsible for certifying the identity of the users. Their role is discussed in more details in section 3.4.2.

ID-linked Database Keeper(s) These actors are responsible for maintaining a database that links users to their identity. These actors should be the only ones that are able to identify the receiving of any transactions. Note that, they are not able to identify how and when users spend their digital cash. More details about their role is given in section 3.4.3.

Block-Producers The role of these actors is to determine which transactions to accept, and add these valid transactions into new blocks. In particular, they ensure that no double-spend is possible and that a consensus is achieved on the ledger. Their role has already been discussed in section 3.2.2.

CBDC Issuers This set of actors is responsible for issuing digital cash and withdrawing it out of circulation. This role could be fulfilled by the central bank only, or delegated partially to commercial banks and other authorized financial institutions (see further details in section 5.1.3). New digital cash would be issued or withdrawn from circulation using dedicated transactions.

3.4 Identity System

To satisfy the *compliance* property of the system, it must always be possible to identify the receiver of a transaction. However, to ensure privacy, the identification should be restricted to a particular set of actors. The simplest solution would be for a central authority to be solely responsible for certifying the identity of

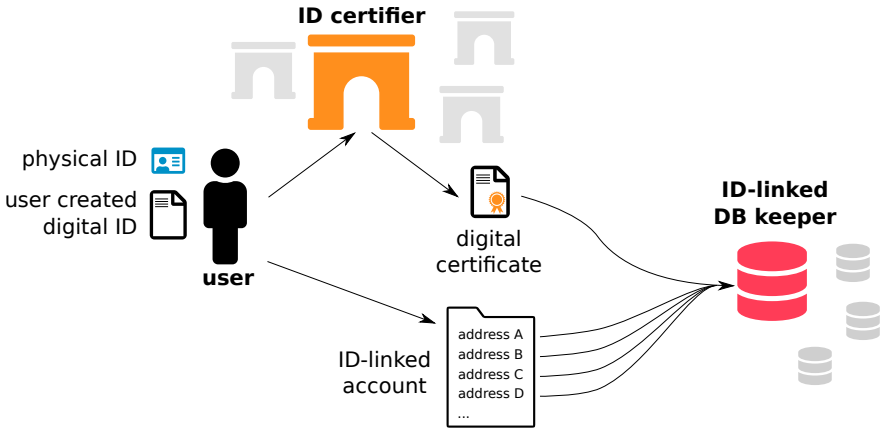

Figure 2: The identity system involves three actors: a user wishing to use the system; an identity certifier which checks the physical identity of the user and produces a digital certificate; and an ID-linked database keeper which associates the user addresses with the digital certificate.

any users and identify a receiver. However, this would involve setting up a new institution for that purpose and incur non-negligible economic overhead.

Instead, we propose to separate the actor(s) responsible for *certifying* the identity of a user from the actor(s) dealing with *identifying* the receiver of any transaction. We suggest leveraging the existing identity verification mechanisms currently available in the economy, such as — but not limited to — post offices or banks.

To achieve this goal, we propose a three steps approach illustrated in figure 2:

1. The user first creates a digital identity (*i.e.* public/secret key-pair);
2. An *identity certifier* checks the digital identify matches with some physical evidence and emits a digital certificate;
3. The user registers addresses that form the ID-linked account on the system with an *ID-linked database keeper* before receiving any transaction using the digital certificate received.

3.4.1 Digital Identity

Before using the system, a user starts by creating a digital identity. Such a digital identity simply consists of a Secret/Public key pair: (ID_{SK}, ID_{PK}). The generation of the key pair is done completely offline by the user on their own device which could be a mobile phone. We opt for this approach to guarantee that no one else, besides the user, has knowledge of the private key ID_{PK}.

3.4.2 Digital identity certification

The role of the identity certifier is to check the identity of a physical person, or institution, and establish a connection with the digital identity. This connection is established by producing a digital certificate that certifies that a given digital identity ID has been verified against some provided proof of identity. With this approach, we envision that only certain entities can operate as identity verifiers. These entities could include post offices, mobile telecom operators, or banks which already have processes in place to accomplish these tasks.

We propose to use a standard similar to `X.509` for public key certificates and a chain of trust, forming a hierarchy. The top level authority, *e.g.* the central bank, controls the root certificate. Then, it emits digital certificates to a set of identity certifiers such as banks or post offices. In turn, these identity certifiers would emit digital certificates to users of the system. As we will see later, these certificates are included in special transactions that allow a user to register a receiving address in the system.

Certificate management The verification of the validity of a certificate is done automatically by the system when they are used by a user. This process is performed by the block-producers, and, thanks to the *transparency* property, anyone else can check that invalid (*e.g.* expired) certificates are never used. In case a certificate should be canceled (*e.g.* stolen identity), the certificate issuers could simply revoke the certificate using a special system transaction recorded on the blockchain.

This approach facilitates the management of certificates, since a certifier's only actions are to emit a certificate when a user wishes to be certified, or revoke it when a user requests so. Except for its own certificate, the certifier does not need to store or manage the certificates it produces.

Receiving limits Each certificate could include certain constraints, such as limiting the total value that can be received within a given time period by a user. For instance, a lightweight certification of identity using just a phone number, might allow the user to receive up to $100 per month. We could also imagine that a user's digital identity could be certified by more than one identity certifier. In such a case, their overall limit would be raised cumulatively, reflecting the added trust the system has in the identity of the user.

Users certifying other users Another interesting aspect of using a hierarchy of certifiers is that users could certify the identity of other users. This might become useful for instance in the case of children, where a parent could certify the identity of their own children. In a similar style, this could be used to on-board deprived users that may not necessarily have a form of identification. A charity organization, for instance, could proceed to certify such users.

In the case where a user is certified by another parent user, any income received by the user would be counted towards the limit of the parent. This is necessary to ensure that a user does not certify nonexistent physical users for the purpose of increasing the total amount of income received. However, note that even if the system would not enforce a limit, the ID-linked database keeper would become aware of this. This is possible since the digital identity is linked to the parent and can be traced back all the way to the root.

3.4.3 Registering receiving addresses

To maintain *compliance*, it must be possible to match any destination address(s) in a transaction to a verified digital identity. This means that before being able to receive a payment, a user first needs to register their receiving addresses with one of the ID-linked database keepers. The system could have a single central keeper or several licensed ones.

While users would register their receiving addresses on the blockchain using a special transaction, the ID-linked database itself would exists as a separate system, off the blockchain To preserve *privacy*, we wish to ensure that only a specific set of authorities knows the identity linked to the receiving address, and therefore has control over the ID-linked database. The choice of which institutions (public or private) have access to this database is a matter of policy and will be discussed in section 5.2. Obvious ones would include tax authorities or the judicial system for instance.

The simplest technical solution could consist in using a public key from the ID-linked database keeper to encrypt the certified identity when sending the list of receiving addresses to register. A transaction would then be broadcasted to the network that contains the addresses to be registered together with the encrypted digital certificate. The keeper would sign it to certify that it has verified the digital certificate is valid. This signed transaction could then be added to the blockchain to inform the system that these receiving addresses have been linked to an ID. This would ensure that only the keeper is aware of the digital identity associated with a receiving address.

However, one potential weakness of this approach is that it does not fully satisfy the *transparency* property of the system. This would require trusting that the keeper would not accept any non-verified digital identity. Furthermore, it would not be possible for everyone to check that the limits associated with a given user are enforced by the system. A user may indeed want to register each address individually to improve privacy to prevent other users from knowing that a given set of addresses belong to the same user.

While it is probably not a deal breaker, more research would be needed to enable full transparency. One idea could be to use zero-knowledge proofs when registering the addresses. The proof would demonstrate that the encrypted digital certificate has been produced by a certifier that is itself certified, without revealing

any information about the digital certificate. The difficulty would be to ensure that anyone can check that all the limits linked to a given user are enforced.

3.5 Approaches to Privacy

Before presenting our proposed solution to provide privacy, we first briefly review the technological approaches to privacy.

"Anonymous" blockchains Blockchain systems, such as Bitcoin [Nak08], are in fact pseudonymous. The identity of the users is not directly visible on the network, since users use anonymous addresses — the pseudonym — to send or receive transactions. To maximize anonymity, users are encouraged to never reuse more than once the same receiving address and avoid transacting in whole amount. However, despite these precautions, many techniques have been proposed to effectively de-anonymized Bitcoin users [MPJ+16].

Shuffling services To increase privacy, it is now recommended to use shuffling services. These services connect groups of users wishing to mix their coins together. These users then coordinate to produce a single transaction that mixes their coins together and produces a set of new fresh coins.

The very first generation of such services were not trustless since it was possible for the server to steal the users' funds. The second generation of approaches, such as CoinShuffle [RMSK14] are built in a trustless manner. This makes it impossible for any user — including the shuffling server — to steal any funds.

The drawback of using shuffling services is that the level of privacy achieved is dependent on the number of users involved in the shuffling. More importantly, it is far from clear how compliance could be enforced and prevent users from colluding with each other to transact privately during the mixing.

Zero knowledge proofs The ultimate approach to privacy is the use of Non-interactive ZK proofs. ZeroCoin [MGGR13] and ZeroCash [BsCG+14a] were among the first Zero-Knowledge based approaches applied to blockchain. The main idea is that users *mint* their own coin by adding them to a list of minted coins. When they want to spend their coin, they simply provide a ZK proof that the coin they want to spend has been minted, without having to reveal which one it is. The end result is that when a coin is used, it could be anyone's coin. This contrasts with shuffling services where it is possible to infer that the coin must belong to the set of users involved in the shuffling at a given time.

However, the drawback of ZK proof is that building the proof might take several minutes [BsCG+14a]. This makes it difficult to use as a payment system if every time a user spends their digital cash, they have to wait several minutes for the transfer to occur. The next section presents a high-level sketch of how ZK proofs could be used for near instantaneous payment, while ensuring *compliance*.

3.6 Private Transactions Support

This section presents a sketch of how a system could be built to satisfy the three major properties presented earlier: *compliance*, *privacy* and *transparency*. Note that the system described here has not been built and any formal proofs that it satisfies the desired properties are left for future work. Also note that we assume that clients are connected to the system via a Tor-like network [DMS04] when broadcasting transactions to increase privacy.

To support private transactions, we propose the use of a protocol similar to ZeroCoin [MGGR13] and ZeroCash [BsCG+14a]. The main idea behind the protocol consists of minting private coins of specific values. When a user wants to spend them, they present a non-interactive zero-knowledge proof of ownership of such private coins. The use of ZK proofs ensures that no information about who mints a coin is revealed. In contrast to ZeroCoin/ZeroCash, our system aims to ensure that it is always possible for an ID-linked database keeper to identify the receiver of a transaction.

Proof generation Generating a zero-knowledge proof is the most time consuming part of the protocol [BsCG+14a]. As explained earlier, this makes it unsuitable for use at the time of making a payment. To get around this problem, we propose an approach where the proofs are built offline, well ahead of when payment takes place.

Exact change policy Given that the amount to spend is unknown at minting time, and at proof generation time, we adopt an *exact change* policy. When sending digital cash to a receiver, the sender must build a transaction made out of a set of private coins whose value matches exactly the amount to be transferred. To achieve this goal, the private coins can only be created with predetermined fixed values (*e.g.* \$0.01, \$0.02, \$0.05, ...). This is similar to physical cash which has a fixed set of possible values. We expect that any software managing the user fund (*i.e.* the wallet) will automatically break down any amount into smaller denominations.

Process To achieve our objective, the process of privately spending digital cash is broken down into three steps:

1. Minting private coins from an ID-linked registered address
2. Constructing the zero-knowledge proofs of private coin ownership
3. Spending the private coins

Each of these steps happens at different points in time. To maximize *privacy*, the minting of private coins should occur well ahead of the spending. Otherwise, it might be possible to establish the identity of the spender by looking at who recently minted private coins.

3.6.1 Minting private coins

Minting a private coin consists of transferring a certain amount from an ID-linked registered key-pair R (public/secret key R_{PK}, R_{SK}) into a private coin. Minting from an ID-linked registered address ensures that any private coin has been produced from funds that have satisfied compliance. However, once minted, it is no longer possible to link the private coin to the minter, hence achieving privacy.

From a *privacy* point of view, a user should use a single registered key-pair when minting a private coin. Otherwise, if multiple registered key-pairs are used during the minting of a private coin, it would be possible to infer for any observer that these key-pairs belong to the same user.

On the other hand, there is no problem reusing multiple times the same registered key-pair to mint new private coins. The system described below should ensure that it is not possible to trace-back a private coin to the registered key-pair used to mint it. Once minted, each private coin is indistinguishable from any other resulting in *fungible* coins.

Minting a private coin of a fixed denomination involves the following:

1. Randomly generate a new key-pair X (public/secret key X_{PK}, X_{SK});
2. Randomly generate a serial number sn;
3. Create commitment $cm_{IDX} = COMM_{r_{IDX}}(X_{SK}||R_{SK})$ for random r_{IDX};
4. Create commitment $cm = COMM_r(sn||cm_{IDX})$ for random r;
5. Produce a ZK-proof π_{IDX} that, given cm, cm_{IDX} and R_{PK} proves:
   - $cm == COMM_r(sn||cm_{IDX})$
   - $cm_{IDX} == COMM_{r_{IDX}}(X_{SK}||R_{SK})$
   - The secret key R_{SK} matches the public key R_{PK}
6. Broadcast a signed (using R_{SK}) mint transaction to the network which contains cm, cm_{IDX}, R_{PK} and π_{IDX}.

Key-pair X The key-pair X is used to claim, at a later stage, ownership of the private coin, without revealing the identity of the owner.

Serial number sn The purpose of sn is to show, again at a later stage, that a private coin is only used once.

Commitment cm_{IDX} This commitment is necessary to prove that whoever creates cm_{IDX} knows both X_{SK} and R_{SK}. This is important to ensure that the owner of the registered key-pair R is the same user as the owner of the private key-pair X. Otherwise, a different user could have created X, hence allowing completely private transfer of digital cash between two users.

Commitment cm This commitment is added to the list of commitments in the system when the mint transaction is broadcasted. The system contains a list of coin commitments, one such list for each possible coin denomination. The commitment will be used later on when spending the private coin with zero-knowledge proof, to prove that a coin has been minted, without revealing who minted the coin (*i.e.* without revealing R).

Similarly to [BsCG+14a], this list of commitments could be built using an append-only list of commitments based on a Merkle tree using a collision-resistant hash function. As shown in prior work [BsCG+14a], a tree with a fixed depth of 64, supports up to 2^{64} minted coins which is more than enough. The scalability of this approach is discussed in more details in section 4.5.

ZK-proof π_{IDX} The proof is here to ensure that the commitment contains the secret key R_{SK} associated with the public key R_{PK}. This is important to ensure that whoever signs the minting transaction has control of the registered key-pair R, and also has control of the newly minted coin by knowing X_{SK}. Without this, it would be possible to break *compliance* and transfer a coin during the minting process to another user.

Mint transaction broadcast Finally, the last step is to put together all the information in a transaction and broadcast it to the network. This includes the commitments cm and cm_IDX, the public key R_{PK} and the proof π_{IDX} that binds everything together. This transaction is signed using the secret key R_{SK}.

3.6.2 Building ZK-proof of private coin ownership

Once a coin has been minted, it is ready to be spent privately, without revealing the identity of the minter. To achieve this goal, we rely on another ZK-proof, π_{minted}, which proves that the coin has been minted. This proof is built by the user on their own device.

Given X_{PK}, sn, the proof π_{minted} proves the following:

- The list of commitments contains $cm = COMM_r(sn||cm_{IDX})$, where $cm_{IDX} = COMM_{r_{IDX}}(X_{SK}||R_{SK})$
- The secret key X_{SK} matches the public key X_{PK}

X_{PK} will be used later when spending the coin to make sure that whoever wants to spend the private coin is the same user that built the proof. Without this, there is a risk that someone could quickly reuse the proof when it is broadcasted. In such a case, an attacker might be able to spend the private coin (and in effect steal it) before the original transaction has been settled (appear in the next block).

3.6.3 Spending Private Coins

When a user wishes to spend one or more private coins, they build a spend transaction which contains:

- the destination ID-linked key-pair(s) and associated amount(s);
- for each private coin used:
  - the serial number sn
  - the ZK-proof π_{minted} associated with the coin which proves that the coin with serial number sn has been minted by X
  - a hash of the destination addresses and amount, signed with the key-pair X

The system ensures that the sum of all the private coins used matches the total amount spent.

Signing the hash of the destination registered address(es) and amount with X ensures that no one else besides X can use the proof to spend the minted private coins. In terms of performance, using ZeroCash [BsCG+14a] as a reference, verifying the proof can be done in the order of a few milliseconds.

3.7 Managing Private Keys

As with any system based on public-secret keys, it is important that the users manage their private keys properly. We expect that the heavy lifting regarding private key management will be handled by a digital wallet.

Digital wallets The use of HD (Hierarchical Deterministic) digital wallets is highly encouraged to avoid losing access to private keys. Such wallet software are already hugely popular among users of crypto-currencies. The main idea behind such wallets is that a seed is generated using a random selection of a fixed number of words (*e.g.* 12) from a dictionary (*e.g.* English). This seed is then used to produce — deterministic — all the private keys that a wallet will control. If the wallet is lost, the private keys can be recovered using the set of words that the user is encouraged to store somewhere secure. To limit the risk of having the private keys stolen, we advocate the use of secure hardware wallets, such as Ledger, that store the actual private keys on a dedicated physical device protected by a PIN code, and are used to sign transactions.

Id-registered fund recovery If a person is unable to recover control of their private keys, or the person is incapacitated, an authority could recover the funds of the user. This is possible since we assume that most of the user's funds would be stored in the ID-registered addresses and not stored as private coins.

To prevent any abuses or risk of a malicious actor impersonating the user in question to steal their funds, we could imagine that when such recovery operation is in progress, a special transaction could be broadcasted on the blockchain and only become effective after a certain lapse of time. If the user has not lost access to its private keys and notice this transaction, he could send a special transaction to cancel the fund recovery which would automatically alert authorities.

Private coins recovery Recovering private coins is much more difficult, since the system's goal is to ensure no one can ever know who is the owner of a private coin. In the case where a user wishes to store a large amount in private coins, we could imagine adding the ability in the system to support transactions that are only valid under certain conditions. For instance, a transaction could only be valid after a certain amount of time has passed since the private coins have been minted — or since they have been "refreshed" with a special transaction.

Using this feature, a user could prepare such transactions in advance and give it to a third party (*e.g.* family member). Every so often, the user would produce a refresh transaction, to indicate to the system that someone is still in control of the private coins. If one day, the private coins have lost their "freshness", the third party could broadcast the transaction he has received from the user and claim ownership of the private coins.

3.8 Access Requirements

Having presented the main key technical details of the system, we now discuss briefly three main issues related to accessing the system.

3.8.1 Need for identification

To satisfy compliance, our system requires users to certify their identity before being able to use the system. However, many users may not have a form of identification (*e.g.* passport). To cater to these users, we advocate the use of a hierarchy of certifiers together with the ability to enforce limits, in terms of the amount that a user is able to receive within a given time frame. As seen in section 3.4.2, this would allow a charity, for instance, to certify the identity of ID-less users with a very low limit set. We could imagine that such an organization might for instance have a lightweight identification process, such as taking a photo of a user when they register, or even issue their own form of identification document. (See further on the drawbacks of this in Section 5.1.1).

3.8.2 Need for electronic devices

One of the main requirements for using a digital cash system is the need for an electronic device for participating in the system. We argue that this is actually not an issue as it looks like we are heading for a society where every individual already has, or will soon have a mobile device.

In Canada, the percentage of over 15 years old owning a smartphone is at 88%, and is increasing every year [RtC18]. Furthermore, even the most deprived members of society have, in large numbers, access to mobile phones in North America. A not so recent study from 2012 conducted in three hospitals in Connecticut, USA, shows that 70% of the homeless people visiting the emergency department had mobile phones [PVD+13].

3.8.3 Need for internet connectivity

The final issue that we discuss is the need for an internet connection in order to receive payment. In our system, the spender can be completely offline, once they have minted their coin. However, the receiver must be connected to the peer-to-peer network to be able to verify any received transactions. Otherwise, they risk being the victim of a double-spent attack, where the spender, who is completely anonymous, could spend multiple times the same private coin.

We are of the opinion that internet connectivity is not going to be a problem in the near future. New satellite based internet services such as Starlink constellation of thousands of satellites promise to deliver low-cost ubiquitous internet access. Nonetheless, we offer some solutions to the problem of poor internet connection that might affect remote communities.

In the case of an intermittent internet connection, the problem becomes that the receiver must trust that the sender is not going to double-spent for the duration of the loss of internet connection. We could imagine that for small amounts, this might be a risk the receiver is willing to take. In such a case, the sender would send his transaction directly to the receiver device (the software would of course take care of this).

One possible approach to decrease the risk of fraud would be for the sender to give up his anonymity and reveal information about the registered address he used to mint the private coins. In such case, the receiver could always contact the authorities later on and present the following information:

- The proof of the double-spent (basically the transaction he has received, showing that the private coins involved have already been spent);
- The information about the source registered address used to mint the private coins

The authorities would then contact the ID-linked database keepers to reveal the identity of the fraudster who could be prosecuted.

4 System Scalability

4.1 Network topology

The network peer-to-peer topology of our system is illustrated in figure 3. As explained in section 3.2.2, we assume that the core of the network is composed of the *block-producing* nodes. These nodes are expected to be connected to each other with a high bandwidth connection to guarantee that their block propagates to one another as fast as possible.

The second layers of nodes are the *normal* nodes which do not create blocks but still verify blocks, propagate them and store the blockchain information. These nodes are used for instance by payment processors or by "power users" who wish to have their own copy of the blockchain.

Finally, the last layer is composed of light clients that do not need to store the full blockchain. Instead, such nodes only need to store the block headers. Under the assumption that the block creators are not cheating (in a PoA system this requires trusting the block creators), a light client can verify that any given transaction is valid, using just the block headers. This technique is explained in the original Bitcoin paper [Nak08] and is known as "Simplified Payment Verification". We discuss in more details scalability concerns for light clients in section 4.3.

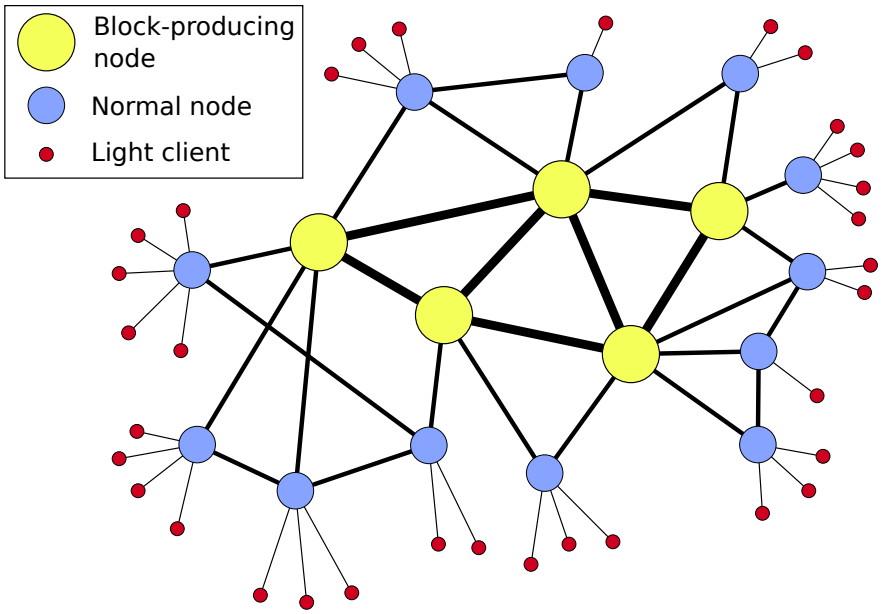

Figure 3: Network topology of the system

4.2 Requirements for verifying all transactions

We now estimate the various requirements necessary for a node that wishes to collect and verify *all* the transactions on the system. Note that the vast majority of the users should not be required to run a full node and should only require the use of a light client. Running a light client does not affect privacy or security since it is still possible to verify all incoming payments to a particular ID-linked address or and make private payments.

Compute A study [SCS+17] from 2017 has shown that an existing BitcoinCash implementation can already scale to $\sim 2,000$ transactions/sec using a desktop class machine to run a block-producing node.The same study shows that the number of transactions per second scales linearly with computing power and network bandwidth: a 500-cores machine (*e.g.* a GPU) with a 1.5Gbps network bandwidth could handle $50,000$ Bitcoin transactions/sec. The authors of the study [SCS+17] highlight that it might be even possible to increase throughput by a factor 2–4×, by further optimizing the software implementation.

Bandwidth The largest type of transactions in our system are produced when spending the private coins. Each such transaction includes a zero-knowledge proof which is expected to be around 300B [BsCG+14a]. Let us assume that a transaction without the proofs, but with the signatures, is similar in size to a Bitcoin transaction, which is around 500 B. Furthermore, let us assume that a typical spent transaction involves 10 private coins. Together with the proofs, the size of a typical spent transaction in our system would end up being around $300\text{B} \times 10 + 500\text{B} = 3,500$ B

Let us first estimate the bandwidth required to receive all the transactions not yet included in a block. If we wish to support $2,000$ payments/s, the incoming bandwidth required for a full node to receive all the private coins *spent* transactions, as well as the corresponding *mint* transactions (produced ahead of time), is around $2,000$ T/s $\times 3,500$ B/T $\times 2 \approx 14$ MB/s. This assumes that a node receives a transaction only once.

In the worst case scenario, each transaction might be received multiple times, up to the number of connected peers. To prevent this, we could easily imagine that transactions would initially be transmitted without the proofs, reducing their size by 10×. Only in the case where a node has not seen the transaction before, would the proofs be requested. Therefore, the bandwidth requirements would only grow by around $2,000$ T/s $\times 300$ B/T $\times 2 \approx 1$ MB/s per peer connected. Assuming that a node is connected to ten other full nodes, the total bandwidth required would be around 14 MB/s $+10 \times 1$ MB/s $= 24$ MB/s.

The second need for network bandwidth is for propagating newly created blocks to the rest of the network. Fortunately, there are practical solutions that drastically reduce the required bandwidth. This is based on the observations that all nodes on the network will have already seen most of the transactions produced since the last block was created. In such a case, it is possible to use probabilistic data structures, such as an invertible bloom lookup table, to propagate information about which transactions a given node has not yet seen. Graphene [OAL+19], which has been deployed on the Bitcoin Cash network, shows that encoding this information for $2,000$ transactions requires less than 10KB of data for a failure rate under 3% (*i.e.* only 3% of the transactions will need to be re-transmitted). Therefore, the bandwidth required to propagate blocks is negligible compared to the bandwidth required to broadcast the individual transactions on the network.

We argue that these bandwidth requirements are totally reasonable, especially since it is only needed by block creating and normal nodes which verify every single transaction. The majority of the users would be running light clients which would only need to verify their own transactions, and not *all* the transactions on the network. Light clients are discussed in more details in section 4.3.

Storage We now turn our attention to the storage requirements for the proposed system. Using the same numbers as above, we have seen that to support a target of $2,000$ payments/sec, each payment would generate $2 \times 3,500$B of data, leading to a production rate of 14 MB/s. At this rate, storing one year of transaction records would require approximately 400 TB of storage.

However, a significant portion of this data can actually be thrown away without impacting the security of the system. For instance, the zero-knowledge proofs do not need to be kept forever and could slowly be "forgotten" and *pruned*. These proofs are important to guarantee the *transparency* property of the system; we want to ensure any user of the system has the opportunity to check that the block-producers only accept transactions with valid proofs. However, it is perfectly reasonable to expect that after a certain amount of time has passed (*e.g.* a day, a week, or a month), a node in the system might no longer be required to produce these proofs when old blocks are requested from them. The rationale is that any observer would have had enough time to notice that a block-producer has broken the rules of the system. Similarly, further savings could be achieved by dropping the signatures attached to each transaction in the blockchain after a certain lapse of time. Such techniques requires to modify slightly the internal blockchain data structures. One such approach is SegWit [LLW15] which has been implemented on the Bitcoin blockchain.

By discarding the proofs alone, the storage required for a single transaction would drop to 500 B (the size of an average bitcoin transaction). Storing one year of blockchain data, with such a system supporting $2,000$ payments/s, would only require 57 TB. At the time of writing, the price of a single 4 TB SSD (Solid-State Drive) is less than $\$1,000$. Storing one year of blockchain data would cost less than $\$15,000$ at today's price.

We argue that storage cost is totally affordable today for any block-producer nodes or any institution wishing to have their own nodes verifying (and storing) *all* transactions on the network. Furthermore, historical trends [Kom14] show that the cost of storage per dollar is decreasing exponentially over time (roughly halving every 2 years).

4.3 User Payment Verification with Light Clients

Having seen the requirements for nodes that produce blocks and nodes that verify *all* transactions on the system, we now turn our attention to lightweight nodes, which would be used by almost all users of the system. The lightweight nodes are able to verify specific transactions and create new ones, without having to store, exchange or request information about the whole blockchain.

As explained in the first Bitcoin paper [Nak08], the use of a Merkle tree enables very fast payment verification on the receiver side. When a user verifies a payment is valid, it is sufficient to check that the transaction has been included in a block. This is done using only the block headers information together with the path in the Merkle tree of the block that contains the transaction.

Using Bitcoin as an example, and assuming less than $65,535$ transactions are produced per block, the size of the Merkle path would only be 512 B. The Bitcoin block header is less than 100 Bytes of data and does not depend on the number of transactions included in the block. Assuming a block production time of 3 seconds, this only represent 2.75 MB of data produced per day.

Let's assume that a merchant/receiver is running a light node and is out of synchronization with the network over a weekend for instance (2 days). In this case, it would only require 5.50 MB of download

when re-opening for business on a Monday. Even with a very slow internet connection, the time required to synchronize at the beginning of the day is totally negligible and only depends on the amount of time the merchant has been offline.

4.4 List of commitments and coin serial numbers

The discussions in the previous sections have not considered the list of serial numbers of coins already spent, and the list of commitments. As discussed in the extended version of the ZeroCash paper [BsCG+14b], it is possible to apply several techniques to reduce the need to hold the entire list of coin serial numbers and commitments by light clients. One of them, for instance, is similar to sharding. The list of serial numbers or commitments, could be subdivided into sub-lists, with new sub-lists created at regular interval for instance. These sub-lists could be of any size; however, the larger the list, the more privacy is achieved since the private coins minted are mixed in a larger pool. The ideal size would be determined by using the average storage available on mobile devices and could grow as technology allows.

4.5 Maximum number of private coins

As mentioned earlier, the depth of the merkle tree of commitments is fixed to 64, limiting the total number of minted coins, per denomination, to 2^{64}. Let's assume that each transaction uses n private coin commitments for each denomination. This means that for each such transaction, n private coins must have been minted per denomination, whose commitments are stored in a single tree corresponding to the specific denomination.

With a goal of supporting $2,000$ transactions/s, this means that $n \times 2,000$ coins are minted every second on average. A reasonable value for n might be around 10, meaning $20,000$ coins are minted per second. Given that the system supports up to 2^{64} minted coins, at this rate, it would take about 29 millions years to fill up the tree. And even in the case where the tree would reach its limit a few years from its initial launch, the system could be updated with an extra new larger tree.

4.6 Proof construction time

Building the non-interactive ZK proofs is expected to be one of the most computational intensive tasks in the system. As seen in section 3.6, our system relies on two ZK proofs, before a user can send digital cash. In our design, the proofs can be built at any time prior to spending, since they do not contain any information about the receiver of the transaction. Therefore, the amount of time required to generate the proofs is not on the critical path. Furthermore, the process of building the proofs for each individual minted coin can be completely parallelized.

Nonetheless, as a reference point, the ZeroCash paper from 2014 [BsCG+14a], reports a proof construction time of only a couple of minutes on a standard desktop. Since our approach is heavily based on this work, we expect construction time to be in the same order of magnitude.

However, given that computing power increases exponentially over time, it is very likely that proof construction time will be in the order of just a few seconds in the near future.

5 System actors and economic agents engaging with the system - discussion of roles, choices and policy trade-offs

Our analysis throughout the previous sections maintained a stance as broad as possible on which institutions could or should take up the role of system actors , as well as on how to define the economic agents who could benefit from the existence of a P-hybrid CBDC (see section 2). We now discuss a number of potential choices regarding the roles and rights of more specific economic agents, taking into consideration the constraints and priorities identified in the Bank of Canada's vision of a digital Canadian Dollar. Many of these priorities are shared by other central banks, *e.g.* the European Central Bank [Eur20]. Similarly, while many of our points of discussion below will be focusing on Canada, most of these also apply for similar countries.

5.1 The central bank and necessary system actors

Our proposal of a P-hybrid CBDC considers the central bank as the main institution driving the design and introduction of the digital currency. Section 3.3 identified four distinct roles for the system actors needed for the implementation of our vision of a P-hybrid CBDC. One possible way to implement this is to have the central bank directly taking up all these roles and directly managing the systems of Identity Certification, ID-linked Database Keeping, Transactions Verification, and CBDC issuance. Such a **possibility of having the central bank as the sole central authority is consistent with all the technological features discussed** in sections 3 and 4.

However, most of these roles can also be delegated to other public and private sector agents and actors in cooperation with the central bank and under a sufficient degree of supervision by the central bank. Under the scenario of maximum delegation, the central bank's involvement is mainly needed at the design and implementation stage of the P-hybric CBDC, but the central bank does not need to accumulate many new extra obligations in the long term (e.g., take up the full responsibility of detecting and reporting incidences of money laundering, which currently largely relies on the financial intermediation sector)[29] **Out of the four system actor roles that need to be undertaken by broadly defined central authorities, the only one that requires a long-term engagement by the central bank is the maintaining of a sufficient control over CBDC issuances and withdrawals.** This is to satisfy the policy goal of maintaining the central bank as sole institution who can conduct an independent monetary policy via sufficient control over CBDC issuances and withdrawals. Given the stated policy preferences, it is also necessary to incorporate the constraint that the central bank wishes to maintain central bank issued physical cash equivalent to CBDC, and to ensure the convertibility of one into the other. Consequently, control over the high powered money supply implies that after the introduction of CBDC the central bank fulfills the **new obligation to control the sum of physical and digital cash-like money in circulation.**

We now discuss which agents the central bank could delegate the roles of different system actors to (in Section 3.3), in a somewhat greater detail.

5.1.1 Identity Verifiers

The sole role of Identity Verifiers in the system is limited to verifying that a person (or a firm) is who it claims to be, and issuing a digital certificate.

This task would be trivial in countries that have a compulsory national digital ID-card system (e.g., Estonia) and firm registries that are integrated with the national ID-card system. It is slightly more complex in countries like Canada (or USA, UK), where national identity cards and passports are not compulsory and specific to individuals' needs. However, if the policy choice is to maintain a decentralized and voluntary system of identity documents, it would still be the case that most individuals already have documents that can certify their identity (for example, passports and driving licences), and that there are institutions like authorized financial intermediaries, postal offices and institutions like public notaries, who are authorized to confirm and certify an individual's identity. Furthermore, a number of Provinces of Canada (e.g., Alberta, British Columbia, and Ontario) are already in the process of introducing digital identity certificates, which would make identity verification easier.

When it comes to certifying new residents and returning citizens (i.e., new consumers with some medium or long-term engagement to Canada), the identity certification can easily be done at the border (and facilitated by the Border Authorities of Canada). Namely, when entering Canada they must already provide some form of identification, and could choose to have their ID and rights to access the P-hybrid CBDC verified at the border. Tourists and foreign investors with a repeated interest in Canada could have the option to be certified similarly, while one-time visitors may rather stick to various international means of payments.

It is therefore natural to enable the system of Identity Verifiers for the P-hybrid CBDC to be built on existing structures, and it is natural that this task be delegated to institutions who are already managing standard forms of identity verification.

One remaining important concern regarding universal access is that some citizens and residents of Canada may not have a formal identification document that has strong enough security features to uniquely identify them. As a national identity document is not an obligation, they may not have needed one (e.g., do not

[29] This is one important concern against a introducing a retail CBDC as highlighted in [DSR+20]

drive a car, and do not travel internationally) or do not wish to have one. At the aggregate level, this may correlate with individuals being less well off. There are three approaches that can be taken to facilitate universal access to the functionalities of the P-hybrid CBDC for these individuals:

1. The system's design enables differential rights to the benefits of P-hybrid CBDC. As seen in section 3.4.2, the system could support limits in the value and frequency of incoming payments. As a result, the certification of individuals that may spend small amounts may be subject to less scrutiny (*e.g.* a photo and mobile number), provided limits on receiving digital cash.

2. There could be a new form of identification document enabling individuals without other formal ID-document to obtain a "P-hybrid CBDC specific ID-card or digital identity" by providing their photo (and any other biometric data deemed necessary for standard scrutiny). To achieve the policy goal of universal access, this service should be free for individuals least well off (subsidized via public funds, seigniorage income or cross-subsidies) and could perhaps be managed by postal offices or other institutions with a natural country-wide presence.

3. Outside the direct P-hybrid CBDC system, an another intermediary could provide them with this service, as see in section 3.8.1.

While both approaches are feasible within our system's architecture, we would argue for the benefits of the second approach, as it better preserves the integrity of the system. In particular, it is better for preventing money laundering tactics such as "smurfing" (see section 2.2.1). If access to the P-hybrid CBDC were to only require a mobile number and some photos, a rich enough criminal could acquire enough mobile numbers and photos to game the system.

We further note that the adoption of the P-hybrid CBDC is ultimately voluntary, and individuals who do not trust the privacy preserving features of our proposed system, and are unwilling to verify their identity, can continue to use physical cash and other existing means of payments.

5.1.2 ID-linked Database Keeper(s) and Transaction Verifiers

Instead of the central bank being solely responsible for managing the verification systems, it could delegate these functions to the respective consortiums formed to act on behalf of the central bank. In addition to the central bank, which could be a partner in both consortiums (if it wishes), natural private sector members are likely to be technology firms and financial institutions (be it established institutions or newcomers), and public sector supported agents may be needed to manage externalities.

We highlight three principles that we consider important for the choice of partners when the central bank delegates these roles:

1. Separation of ID-linked Database Keeper(s) and Transaction Verifiers

2. Multiple institutions engaged as ID-linked Database Keeper(s) and Transaction Verifiers

3. Avoiding potential conflicts of interest - in particular avoiding a situation where a key partner's main business model would benefit from acquiring and withholding superior information about individuals' (or firms') ID-linked accounts.

The privacy benefits of separating the roles of Database Keeper(s) and Transaction Verifiers were already highlighted in section 3.3.

Point 2 is about the resilience of the system and the very nature of the benefits of a distributed ledger system - should one member of the consortium (which here only means a manager and holder of a database) fail or face an exogenous shock, the system would continue to function. Assuming point 1 separation, recall that our system's design is such that Transaction Verifiers do not observe who made transactions. This function can be delegated to private sector participants. To maintain quality incentives, there is a case for giving them some oligopolistic power (i.e., limiting their number) and the threat of exclusion, such that their incentives to participate and maintain their quality of service is determined by a trade-off between maintaining their quality of service and earning oligopolistic fees, or being removed from the consortium if their service is sub-standard.

The management of ID-linked Database Keeper(s) involves a public goods dimension, as these records come with many positive aggregate externalities (see Section 2.2). For this reason, it would be best managed by either the public sector or by private-sector-financed pure technology firms. Furthermore, point 3 above implies that ID-linked Database Keeper(s) should not have conflicting interests. Therefore, both financial institutions (whose business model relies on information) and the subset of technology firms (whose business model focuses on acquiring individuals' data for advertising), are not ideal private sector partners for being ID-linked Database Keeper(s).

5.1.3 CBDC issuers

As highlighted above, we consider that the central bank will maintain control over the high-powered money supply to conduct monetary policy. As the manifested goal of the Bank of Canada (as well as of the European Central Bank) is to maintain physical cash in circulation in parallel, and to ensure convertibility, there are a few options.

First, the central bank could be the only agent who has the right to convert CBDC to physical currency and vice versa. Second, it could delegate the rights and obligations of conversion to authorized financial institutions, who would be able to temporarily "issue" and "withdraw" CBDC while reporting these transactions to the Central Bank. Third, the "issuance" and "withdrawals" could be conducted by the Central Bank (close to real time) at the request of authorized financial institutions. [30]

The second and third option seem more natural and practical, but also imply from the system's architecture perspective that authorized financial institutions either have the rights to temporarily "issue" and "withdraw" CBDC, or that they can communicate their individual customers' requests of conversions frequently. While the first option is straightforward, the second and third options warrant further explanations. Note that under the current system of cash distribution and record keeping, financial institutions need to predict cash demand, keep track of notes that exist but are not in circulation, and send the end of day balance to the Central Bank's ledger. A similar system for physical cash can easily continue when there is also a CBDC. However, unlike physical cash which is a form of token money and not linked to individuals, our proposed P-hybrid CBDC is issued to (or withdrawn from) a specific individual's or a firm's ID-linked account. Should the individual who converts cash to CBDC need to wait until this conversion is confirmed with the central bank at the end of a day, there would be a time delay that limits the CBDC's attractiveness.

To explain the second and third option, consider the following example: Alice (with her ID-linked account) goes to her bank and asks CAD $100 of physical cash to be put on her digital CAD (P-Hybrid CBDC) account. Neutrality with respect to the central bank's accounts implies that the aggregate supply of digital CAD has increased by CAD $100, and that the physical cash in circulation has decreased by CAD $100. From the system's perspective, there has been CAD $100 worth of new digital currency issued. Later on the day, there is a similar opposite transaction when an individual Bob (with his ID-linked account) goes to the same bank and asks CAD $50 digital CBDC to be converted to physical cash. From the system's perspective, there is now CAD $50 worth of new digital currency withdrawn.

Under the second option, the authorized financial institution (Alice's and Bob's bank) has a right to "issue" $100 worth of digital CADs to Alice's account, and to "withdraw" $50 worth of digital CAD from Bob's account. At the end of the day the bank needs to have $50 more physical CAD notes in its vault of notes that are not in circulation. Of course, the central bank sees both of these transactions as they are recorded in the ID-linked blockchain-based database, and must still receive reports on the cash balance at the end of the day to ensure that these transactions are central bank balance sheet neutral, and that financial institutions do not influence the overall high-powered money balance. Both Alice and Bob can manage and spend their new digital CAD and physical CAD privately.

Under the third option, an authorized financial institution/the bank would not have the right to "issue" and "withdraw" digital CAD, but would still manage conversions. The bank would then have to send two digital requests to the central bank to issue $100 worth of CAD to Alice's account and withdraw $50 from Bob's account. The central bank would still have a reporting system to ensure that the bank now has $50 more physical CAD notes in its vault of non-valid notes.

[30] Our system architecture of P-hybrid CBDC is compatible with either one or multiple institutions issuing (or withdrawing) digital cash.

We would argue that although the reporting requirements and the ability of the central bank to control the money supply is similar under the second and thirds option, the third option is more cumbersome and can create administrative time lags. This is because individuals need to wait until their requests have been cleared at the central bank before they can be recorded in the system.

Compared to option one where the central bank manages all conversions, options two and three still ensure that such transactions remain central bank balance sheet neutral. Options two and three simply require some off-line agreements, reporting requirements and auditing that are not fundamentally different from the current system of cash management. Similar trade-offs are also discussed in [AB20], [Car21].

5.2 Viewers of ID-linked information

As discussed in Section 2.2, ID-linked accounts bring a number of benefits, as long as regulators, innovative platform developers, the government and the central bank, can make "queries", build new interfaces based on this information, and make targeted transfers based in the information that is recorded in the ID-linked Database (off the blockchain).

These economic agents are the ones who can bring value to the users (individuals and firms), but they are not system actors, so that we call them the "viewers". Their minimum rights to observe private information on ID-linked accounts to maintain "cash+" features can be summarized as follows:

1. To ensure tax compliance and the detection of money laundering, the tax and money laundering authorities should be able to analyze patterns of incoming flows to ID-linked accounts, as well as make more detailed queries on individuals and firms acting fraudulently.

2. Developers of platforms that facilitate basic user interfaces and tax calculations should be able to link their code to the ID-linked accounts.

3. Authorized mainstream financial institutions and FinTechs should be able to make queries and link their code to the relevant subset of ID-linked accounts relevant to their proposed financing contracts, *e.g.* wages of former students for student financing providers, a firm's (possibly product specific) sales records for providers of external financing.

4. Financial institutions should be able to send returns on investments to ID-linked accounts. They only need to be able to access a subset of ID-linked accounts that are associated with their clients.

5. Central bank and the government should be able to identify target consumers or firms to whom it wishes to send funds.

As different institutions have different needs for information, there is a case for differentiated rights, where some institutions can only view a part of the full database. For instance, financial intermediaries focused on firms do not need to have access to an individual's ID-linked data.

6 Adoption of the CBDC and initial investment

There are three reasons why a P-hybrid CBDC (or any other form of CBDC) would be adopted at a large enough scale: privacy considerations discussed in section 2.1, convenience of the user interface, network effects.

Privacy considerations in section 2.1 indicate that richer individuals have a greater incentive to use and adopt a P-hybrid CBDC, and when the average income increases, more individuals have an incentive to adopt the P-hybrid CBDC. Because the P-hybrid CBDC competes with other means of payments, there is a necessity to provide a good initial user interface and basic service that enable the basic payment functionality.

From a technical and economic perspective there is a strong case for the basic user software development to be publicly funded, available as open source software, and procured via a price competition or auction mechanism. The basic user software should enable easy — and possibly automated — transfer from an individual's interest bearing account to his/her P-hybrid CBDC account. The analysis of auction mechanisms and precise details of privacy properties developer's access rights to the system is an interesting topic for future research.

Another feature of our proposed system is that users would be charged per transaction, but these fees are minimal (sub-cent) and borne by the senders (see section 3.2.1). This makes cost-wise the P-hybrid CBDC a very competitive means of payment compared to alternatives (be it VISA or MasterCard supported bank payments, physical cash, private sector generated digital currencies).

The initial system's design will require substantial investment; it is worth highlighting some sources of project specific income and the benefits of externalities.

First, as highlighted in section 3.2.1 each transaction involves a sub-cent transaction cost, which is needed for a well functioning system. These small costs are negligible for individuals and firms, but can easily sum up to a substantial amount for the Transaction Verifiers. Second, all providers of novel services built on ID-linked accounts (e.g., financial institutions, technology firms and FinTechs) can be charged a license fee to have access to a subset of ID-linked accounts. Third, providers of all other services that are better or more elaborate than free basic services can be charged a license fee. For example, these initiatives could include offering tourists (or local individuals who do not want to have an id-linked account) a prepaid card that is cheaper to use within Canada than an international bank card, and that can be sold back to the issuer. Some other services that have a public goods dimension (such as better tax collection, the prevention of money laundering, better control over monetary and fiscal policy tools, and the reduction of bureaucratic costs) can further justify public investment in the CBDC's development.

7 Conclusions

This paper has proposed a P-hybrid CBDC as a possible solution for the fundamental trade-off between cash-like privacy for consumer spending and regulatory compliance, while enabling "cash+" features. It has made an economic case for the value of such currency, and proposed a system architecture that can achieve this vision. As seen, the proposed system aims to ensure three important properties: *compliance* through the use of ID-linked registered accounts; *transparency* through the use of an open block-chain where any users can observe that the rules of the system are enforced; and *privacy* which is achieved through the use of Zero-Knowledge proofs.

In our proposed system, individuals can make their CBDC spending wallet private and neither the private sector participants nor central authorities can track how exactly they spend their money. However, transfers between anonymous accounts are not allowed, and the instances of money received are not fully private, and are observable at least by authorities who need to track money laundering incidences, tax evasion etc. Such viewing rights can be extended to other regulated parties and can facilitate innovation. We argue some lack of privacy of money received does not harm individuals and firms who aim to act within legal boundaries and innovative systems built on these records can give them further benefits. We also discussed numerous practical and policy choices regarding system actors and the roles that the central bank could delegate.

8 Acknowledgments

This paper has been completed as part of the Bank of Canada's Model X Challenge, which called for Finance and Computer Science collaboration for proposals for Canadian Central Bank Digital Currency design. We thank the Bank of Canada for their support and the many fruitful discussions. Note that the views expressed in this paper do not represent the official view from the Bank of Canada. Finally, we wish to thank our colleagues Prof. Gregory Dudek and Prof. Claude Crépeau from McGill University for their valuable insights and suggestions.

References

[AB20] Raphael Auer and Rainer Böhme. The technology of retail central bank digital currency. *BIS Quarterly Review, March*, 2020.

[ACF+20] Raphael Auer, Giulio Cornelli, Jon Frost, et al. *Rise of the central bank digital currencies: drivers, approaches and technologies.* Centre for Economic Policy Research, 2020.

[AV05] Alessandro Acquisti and Hal R Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, 2005.

[Ban20] Bank of Canada. Contingency planning for a central bank digital currency. `https://www.bankofcanada.ca/2020/02/contingency-planning-central-bank-digital-currency/`, 2020.

[Bau52] William J Baumol. The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, pages 545–556, 1952.

[BBGP20] Tobias Berg, Valentin Burg, Ana Gombović, and Manju Puri. On the rise of fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, 33(7):2845–2897, 2020.

[BF89] OJ Blanchard and S Fisher. Lectures on macroeconomics". the mit press, cambridge, mass. 1989.

[BJL19] Markus K Brunnermeier, Harold James, and Jean-Pierre Landau. The digitalization of money. National Bureau of Economic Research, 2019.

[BLLV98] Ahto Buldas, Peeter Laud, Helger Lipmaa, and Jan Villemson. Time-stamping with binary linking schemes. In Hugo Krawczyk, editor, *Advances in Cryptology — CRYPTO '98*, pages 486–501, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[BMSW19] Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace. Consumer lending discrimination in the fintech era. *Manuscript,University of California, Berkeley*, 2019.

[BN19] Markus K Brunnermeier and Dirk Niepelt. On the equivalence of private and public money. *Journal of Monetary Economics*, 106:27–41, 2019.

[BsCG+14a] Eli Ben-sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *IEEE Symposium on Security and Privacy (SP)*, 2014.

[BsCG+14b] Eli Ben-sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin (extended version). 2014.

[But14] Vitalik Buterin. A next-generation smart contract and decentralized application platform. Technical report, 2014.

[Car21] Agustin Carstens. Digital currencies and the future of the monetrary system. Remarks by Mr Agustín Carstens, General Manager of the BIS, at the Hoover Institution policy seminar, Basel, 27 January 2021. Available at `https://www.bis.org/speeches/sp210127.htm`, 2021.

[CG16] Christian Catalini and Joshua S Gans. Some simple economics of the blockchain. Technical report, National Bureau of Economic Research, 2016.

[CH19] Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797, 2019.

[DMS04] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *13th USENIX Security Symposium*, 2004.

[DR+14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[DSR+20] Darrell Duffie, Hyun Song Shin, Raghuram Rajan, Arminio Fraga, Kenneth Rogoff, Jacob Frenkel, Agustín Carstens, Jaime Caruana, Tharman Shanmugaratnam, Masaaki Shirakawa, and Zhou Xiaochuan. Digital currencies and stablecoins: Risks, opportunities, and challenges ahead. Group of Thirty, Washington, D.C., 2020.

[Eth18] Ethereum. Design rationale – accounts and not utxos. https://eth.wiki/en/fundamentals/design-rationale, 2018.

[Eur20] European Central Bank (ECB). Report on a digital euro. https://www.ecb.europa.eu/euro/html/digitaleuro-report.en.html, 2020.

[FGPRW20] Andreas Fuster, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. Predictably unequal? the effects of machine learning on credit markets. *Journal of Finance, forthcoming*, 2020.

[Fin19] Financial Transactions and Reports Analysis Centre of Canada (FINTRAC). Guideline 1: Backgrounder. https://www.fintrac-canafe.gc.ca/guidance-directives/overview-apercu/Guide1/1-eng, 2019.

[GL20] Rodney Garratt and Michael Lee. Monetizing privacy. *Available at SSRN*, 2020.

[GV19] Rod Garratt and Maarten Van Oordt. Privacy as a public good: A case for electronic cash. *Staff Working Paper 2019-24*, 2019.

[KM02] Nobuhiro Kiyotaki and John Moore. Evil is the root of all money. *American Economic Review*, 92(2):62–66, 2002.

[KM18] Nobuhiro Kiyotaki and John Moore. Inside money and liquidity. *Manuscript, Princeton University*, 2018.

[Kom14] Matt Komorowski. A history of storage cost. https://mkomo.com/cost-per-gigabyte-update, accessed 7 Feb 2021, 2014.

[KRW18] Charles M. Kahn, , Francisco Rivadeneyra, and Tsz-Nga Wong. Should the central bank issue e-money? *Staff Working Paper 2018-58*, 2018.

[KW93] Nobuhiro Kiyotaki and Randall Wright. A search-theoretic approach to monetary economics. *The American Economic Review*, pages 63–77, 1993.

[LLW15] Eric Lombrozo, Johnson Lau, and Pieter Wuille. Bip-141: Segregated witness. Bitcoin Improvement Proposals, 2015.

[MGGR13] Ian Miers, Christina Garman, Matthew Green, and A.D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. pages 397–411, 05 2013.

[MPJ+16] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. *Commun. ACM*, 59(4):86–93, March 2016.

[Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Technical report, 2008.

[OAL+19] A. Pinar Ozisik, Gavin Andresen, Brian N. Levine, Darren Tapp, George Bissias, and Sunny Katkuri. Graphene: Efficient interactive set reconciliation applied to blockchain propagation. In *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19, page 303–317, New York, NY, USA, 2019. Association for Computing Machinery.

[PVD+13] Lori Ann Post, Federico E Vaca, Kelly M Doran, Cali Luco, Matthew Naftilan, James Dziura, Cynthia Brandt, Steven Bernstein, Liudvikas Jagminas, and Gail D'Onofrio. New media use by patients who are homeless: The potential of mhealth to build connectivity. *J Med Internet Res*, 15(9):e195, Sep 2013.

[RMSK14] Tim Ruffing, Pedro Moreno-Sanchez, and Aniket Kate. Coinshuffle: Practical decentralized coin mixing for bitcoin. In Mirosław Kutyłowski and Jaideep Vaidya, editors, *Computer Security - ESORICS*, pages 345–364, Cham, 2014. Springer International Publishing.

[Rom86] David Romer. A simple general equilibrium version of the baumol-tobin model. *The quarterly journal of economics*, 101(4):663–685, 1986.

[RtC18] Canadian Radio-television and Telecommunications Commission. Communications monitoring report, 2018.

[SCS+17] Andrea Suisani, Andrew Clifford, Andrew Stone, Erik Beijnoff, Peter Rizun, Peter Tschipper, Alexandra Fedorova, Chen Feng, Victoria Lemieux, and Stefan Matthews. Measuring maximum sustained transaction throughput on a global network of bitcoin nodes. Scaling Bitcoin conference presentation `https://scalingbitcoin.org/stanford2017/Day1/Stanford_2017.pptx.pdf`, 2017.

[Tin18] Katrin Tinn. Smart contracts and external financing. *Available at SSRN 3072854*, 2018.

[Tin19] Katrin Tinn. Distributed ledger technologies and start-up financing. *In The economics of Fintech and digital currencies. Center for Economic Policy Research, London*, pages 15–20, 2019.

[Tir10] Jean Tirole. *The theory of corporate finance*. Princeton University Press, 2010.

[Tob56] James Tobin. The interest-elasticity of transactions demand for cash. *The review of Economics and Statistics*, pages 241–247, 1956.

[Wic13] Knut Wicksell. *Lectures on Political Economy (Routledge Revivals): Two Volumes*. Routledge, 2013.

A Technical details for Section 2.1

A.1 Real effects of privacy preferences on the consumer-producer relationship

Assume a stylized setting with one good, and N consumers, and one producer with monopolistic power. For the sake of stylized argument, assume that the monopolistic producer has zero marginal cost. Each consumer j has two characteristics: their valuation of the product $v_j = \{v_H, v_L\}$, where $v_H > v_L \geq 0$, and their "character" type $T_j = \{A, B\}$, which is not directly relevant for the producer, who only cares about v_j. However, according to (1), it is publicly known that character type "A" (resp. "B") gets dis-utility if their decision to buy the good reveals information about their character type. Namely, consumer of type "A" buys the product if

$$\begin{aligned} U_j^A(buy) = v_j - p - \epsilon \ln\left(\frac{\Pr(buy|A,\Omega)}{\Pr(buy|B,\Omega)}\right) \geq \\ (-\epsilon)\ln\left(\frac{\Pr(dont\_buy|A,\Omega)}{\Pr(dont\_buy|B,\Omega)}\right) = U_j^A(dont\_buy) \end{aligned} \tag{4}$$

and consumer of type "B" buys the product if

$$\begin{aligned} U_j^B(buy) = v_j - p - \epsilon \ln\left(\frac{\Pr(buy|B,\Omega)}{\Pr(buy|A,\Omega)}\right) \geq \\ (-\epsilon)\ln\left(\frac{\Pr(dont\_buy|B,\Omega)}{\Pr(dont\_buy|A,\Omega)}\right) = U_j^B(dont\_buy), \end{aligned} \tag{5}$$

where p is the price of the good and Ω reflects the on-path beliefs about consumer purchasing behavior.

Assume a general form of joint distribution of consumption and character types

$$\begin{array}{c|cc} & v_j = v_H & v_j = v_L \\ \hline T_j = A & x_A q_A & (1-x_A)\,q_A \\ T_j = B & x_B\,(1-q_A) & (1-x_B)\,(1-q_A) \end{array}, \tag{6}$$

where $q_A \equiv \Pr(T_j = A) = 1 - \Pr(T_j = B)$, and $x_A \equiv \Pr(v_j = v_H | T_j = A)$; $x_B \equiv \Pr(v_j = v_H | T_j = A)$.

We seek to find the Perfect Bayesian Equilibrium (PBE) of this game, where

- All consumes are rational, make their purchasing decisions according to (4) and (5).
- The firm set the price, p, to maximize $\Pi \equiv Np \Pr(j\_buys|\Omega)$.
- The firm and consumers know (6), and their beliefs about Ω are consistent.
- Equilibrium beliefs, Ω, follow Bayes' rule.

When the privacy concern is mute, i.e., $\epsilon = 0$, the outcome is the standard textbook case. The firm either sets $p = v_L$, and sells to everyone: profit Nv_L, or sets the price $p = v_H$ and sells to all consumers who value the product highly: profit $Nv_H\,(x_A q_A + x_B\,(1-q_A))$. Selling to high value customers is strictly preferred whenever

$$\frac{v_L}{v_H} < (x_A q_A + x_B\,(1-q_A)), \tag{7}$$

which always holds when $v_L \to 0$.

When $\epsilon > 0$, there exist multiple equilibria, all of which make the firm worse off, unless inequality (7) does not hold. Namely, if the firms sets $p = v_L$, all consumers buy and the privacy relevant terms in (4) and (5) do not have a bite.

In all other equilibria, the firm sets the price, $p > v_L$, and PBE outcome depends on beliefs, Ω. One possible equilibrium is the "pure discount equilibrium", where beliefs are such that that all consumers who have high valuation for the product buy, i.e., $\Pr(j\_buys|\Omega) = \Pr(v_j = v_H|\Omega) = (x_A q_A + x_B\,(1-q_A))$. However, the incentive compatible price the firm can charge is strictly lower that v_H. Namely under these

beliefs, $\frac{\Pr(buy|A,\Omega)}{\Pr(buy|B,\Omega)} = \frac{\Pr(v_j=v_H|A)}{\Pr(v_j=v_H|B)} = \frac{x_A}{x_B}$ and $\frac{\Pr(dont\_buy|A,\Omega)}{\Pr(dont\_buy|B,\Omega)} = \frac{\Pr(v_j=v_L|A)}{\Pr(v_j=v_L|B)} = \frac{(1-x_A)}{(1-x_B)}$, which by 4) and (5) implies that the price the firm optimally sets is

$$p = v_H - \epsilon \left| \ln \left(\frac{x_A}{x_B} \frac{(1-x_B)}{(1-x_A)} \right) \right|, \tag{8}$$

where the last term is always strictly negative unless T_j and v_j are statistically independent ($x_A = x_B$).

Another set of equilibria involve mixed strategies, where at least some consumers who have a high valuation of the product, $v_j = v_H$ do not buy the product. Without loss of generality, assume that $x_A > x_B$, such that $\ln \left(\frac{x_A}{x_B} \frac{(1-x_B)}{(1-x_A)} \right) > 0$, and under the "pure discount equilibrium" described above, type A consumers are less willing to pay. Let us consider PBE, where only type A consumers use a mixed strategy, and buy the product with probability $z_A \equiv \Pr(buy|A, v_j = v_H)$. On the equilibrium path, all agents share the beliefs and consumers with $v_j = v_L$ never buy the product, and conditional on $v_j = v_H$, those with $T_j = A$ (resp. $T_j = B$), buy the product with probability z_A (resp. 1). Given these beliefs, we can find that on the equilibrium path

$$\frac{\Pr(buy|A,\Omega)}{\Pr(buy|B,\Omega)} = \frac{z_A x_A}{x_B}$$
$$\frac{\Pr(dont\_buy|A,\Omega)}{\Pr(dont\_buy|B,\Omega)} = \frac{(1-x_A) + (1-z_A)\, x_A}{(1-x_B)}$$

For the consumer with $\{T_j = A, v_j = v_H\}$ to be willing to use a mixed strategy, it must be the case that his/her incentive compatibility constraint (4) holds with equality. This implies that the price the firm charges in such PBE is

$$p = v_H - \epsilon \ln \left(\frac{z_A x_A}{x_B} \cdot \frac{(1-x_B)}{(1-x_A) + (1-z_A)\, x_A} \right) \tag{9}$$

First, notice that $z_A = \frac{x_B}{x_A}$ implies that there is no learning about character types in equilibrium, and the firm optimally sets the price $p = v_H$. It is straightforward to verify that all other incentive compatibility constraints in (4) and (5) are satisfied as well. However, as some customers who like the product do not buy it, the firm's sales and profits are always lower compared to an environment where $\epsilon = 0$. Namely, the firm's profit in this case is $N v_H (x_A z_A q_A + x_B (1 - q_A)) = N v_H x_B < N v_H (x_A q_A + x_B (1 - q_A))$ for any $x_A > x_B$.

Second, consider that $z_A \in \left[\frac{x_B}{x_A}, 1 \right]$. There exists a set of PBE where A uses the described mixed strategy, and the equilibrium price is given by (9),which is a function of beliefs, z_A (which in turn need to be consistent with equilibrium beliefs). In all these cases, it is also straightforward to verify that the incentive compatibility constraint for consumer type $\{T_j = B, v_j = v_H\}$, who adopts a pure strategy, is also satisfied. Furthermore, as long as the difference between v_H and v_L is high enough, all consumers with $v_j = v_L$ do not have an incentive to deviate, and buy the product. Also notice that this case nests the "pure discount equilibrium" and the equilibrium where $z_A = \frac{x_B}{x_A}$, where the firm does not sell at a discount, but loses consumers.

We summarize the described outcomes in the following proposition

Proposition *When $\epsilon > 0$, $x_A > x_B$, and $\frac{v_L}{v_H} < x_B$, there exists a PBE for any $z_A \in \left[\frac{x_B}{x_A}, 1 \right]$, where consumer types $\{T_j, v_j = v_L\}$ do not buy the product; those with $\{T_j = A, v_j = v_H\}$ buy the product with probability z_A; and those with $\{T_j = B, v_j = v_H\}$ buy the product with probability one.*
The firm sets the price

$$p = v_H - \epsilon \ln \left(\frac{z_A x_A}{x_B} \cdot \frac{(1-x_B)}{(1-x_A) + (1-z_A)\, x_A} \right) \leq v_H$$

and its profit is

$$N p (z_A x_A q_A + x_B (1 - q_A)) < N v_H (x_A q_A + x_B (1 - q_A)).$$

A similar proposition holds for the case where $x_A < x_B$. To complete this part of the analysis, notice that PBE allows arbitrary beliefs as long as there is a consistency of beliefs in the equilibrium. It can be

shown that all such equilibria are Pareto dominated by the described one and we do not include the complete analysis of these PBE, are they are less interesting.[31]

Finally, it is interesting to highlight that when ϵ is high, setting the price at v_H and losing some customers is better than setting the price low enough to preserve all customers. Namely, from Proposition 1, the firm's profits under $z_A = \frac{x_B}{x_A}$ and $p = v_H$ do not depend on ϵ, while they are decreasing in ϵ under $z_A = 1$ and $p = v_H - \epsilon \ln\left(\frac{x_A}{x_B} \cdot \frac{(1-x_B)}{(1-x_A)}\right)$, which implies that with high enough ϵ, the firm would rather lose consumers than set a lower price. For some specific values of ϵ, the best general outcome may be that the firm offer some discount and lose some customers.

A.2 Micro-foundation for our modeling of the representative consumer

We derive here the instantaneous utility of the representative consumer with privacy concerns, which is motivated by our model in Appendix A.1. .

Consider two baskets of goods, and privacy-concerned individuals. As above, we consider two character types (A and B), and two possible consumption baskets (a and b). To emphasize privacy concerns as separate from other dimensions of heterogeneity, assume that the price of basket a and b is the same, and both character types will spend the same amount, c, to buy either basket a or b, during some period t. Assume that consumer type A likes basket a more, and consumer type B likes basket b more, but both of them are willing to substitute their basket for the one they favor less, at a utility cost parameterized by δ.

We assume that consumers make a choice between consumption baskets denoted with $\mathbf{1}_A = \{1, 0\}$ and $\mathbf{1}_B = \{1, 0\}$, where $\mathbf{1}_A = 1$ (resp. $\mathbf{1}_B = 1$) denotes consumer A's (resp. B's) decision to buy their preferred basket a (resp. basket b). For the sake of comparison with seminal models of transaction demand of money (and for simpler analytical expressions), we assume logarithmic preferences.

We further denote $c^{T,L}$ as the demand for basket $L \in \{a, b\}$ by consumer type $T = \{A, B\}$. We also assume that parameter $\delta \in (0, 1)$ measures the relative utility loss of substituting one's favorite basket for one's less favorite basket. To complete the setting, $\Pr(L|T)$ reflects the information revealed by consumer type T choosing basket L in equilibrium.

Given the above, a type A consumer chooses $\mathbf{1}_A$ and his/her demand for goods in basket a and b ($c^{A,a}$ and $c^{A,b}$), taking into account the aggregate privacy impact of their choice, to maximize

$$\begin{aligned} U^A = {} & \mathbf{1}_A \ln\left(c^{A,a}\right) + (1 - \mathbf{1}_A)\,\delta \ln\left(c^{A,b}\right) \\ & - \mathbf{1}_A \epsilon \ln\left(\frac{\Pr(a|A)}{\Pr(a|B)}\right) - (1 - \mathbf{1}_A)\,\epsilon \ln\left(\frac{\Pr(b|A)}{\Pr(b|B)}\right) \end{aligned}$$

subject to

$$\mathbf{1}_A c^{A,a} + (1 - \mathbf{1}_A)\, c^{A,b} = c$$

Similarly, a type B consumer chooses $\mathbf{1}_B$ and his/her demand for goods in basket a and b ($c^{B,a}$ and $c^{B,b}$), taking into account the aggregate privacy impact of their choice, to maximize

$$\begin{aligned} U^B = {} & (1 - \mathbf{1}_B)\,\delta \ln\left(c^{B,a}\right) + \mathbf{1}_B \ln\left(c^{B,b}\right) \\ & - \mathbf{1}_B \epsilon \ln\left(\frac{\Pr(b|B)}{\Pr(b|A)}\right) - (1 - \mathbf{1}_A)\,\epsilon \ln\left(\frac{\Pr(a|B)}{\Pr(a|A)}\right) \end{aligned}$$

subject to

$$\mathbf{1}_B c^{A,a} + (1 - \mathbf{1}_B)\, c^{A,b} = c$$

In the benchmark case where $\epsilon = 0$, it is straightforward that it is optimal to choose $\mathbf{1}_A, \mathbf{1}_B = 1$, i.e., both types of consumers only buy their favorite basket. That is $c^{A,a} = c$, and $c^{B,b} = c$.

[31] For example, when $z_A < \frac{x_B}{x_A}$, consumer type $\{T_j = B, v_j = v_H\}$ will no longer have an incentive to buy the product with probability one, and them using a pure strategy is no longer consistent PBE. There also exist PBE where both $\{T_i = B, v_i = v_H\}$ and $\{T_i = A, v_i = v_H\}$ use mixed strategies and the price is $p = v_H$. However, it can be shown that such mixed strategy equilibria are Pareto dominated by the one where all type B consumers always buy the product with probability one. We suppress the full analysis of possible equibria, as it adds less value to the overall message that $\epsilon > 0$ limits the firm's profits and the joint surplus.

However when $\epsilon > 0$, buying one's favorite basket is no longer optimal as $\frac{\Pr(a|A)}{\Pr(a|B)} = \frac{1}{0} \to \infty$ and $\frac{\Pr(b|B)}{\Pr(b|A)} = \frac{1}{0} \to \infty$. Note that this is true even when ϵ is arbitrarily small.

As in Appendix A.1, consider that consumers use a mixed strategy to hide their type, and sometimes opt for buying their less favored basket. Consider that A buys his/her favored basket with probability z_A and B buys their favorite basket with probability z_B.

It then follows that

$$\frac{\Pr(a|A)}{\Pr(a|B)} = \frac{z_A}{(1-z_B)}; \frac{\Pr(b|A)}{\Pr(b|B)} = \frac{(1-z_A)}{z_B}$$

The expected utility (internalizing the budget constraint) of type A from following this strategy is

$$\mathbb{E}\left[U^A\right] = (z_A + \delta(1-z_A))\ln(c) - z_A\epsilon\ln\left(\frac{z_A}{1-z_B}\right) - (1-z_A)\epsilon\ln\left(\frac{1-z_A}{z_B}\right)$$

and similarly, the expected utility of type B from following this strategy is

$$\mathbb{E}\left[U^B\right] = (\delta(1-z_B) + z_B)\ln(c) - (1-z_B)\epsilon\ln\left(\frac{(1-z_B)}{z_A}\right) - z_B\epsilon\ln\left(\frac{z_B}{1-z_A}\right)$$

In a static Bayes-Nash equilibrium, type A chooses z_A taking z_B as given, and type B chooses z_B taking z_A as given. After some derivations both optimization problems give the same first order condition/best-response

$$c^{(1-\delta)}\left(\frac{z_A z_B}{(1-z_B)(1-z_A)}\right)^{-\epsilon} = 1$$

Focusing on the symmetric equilibrium where $z_A = z_B = z$, we obtain that

$$\begin{aligned} z &= \frac{c^{\frac{(1-\delta)}{2\epsilon}}}{\left(1 + c^{\frac{(1-\delta)}{2\epsilon}}\right)} \\ 1 - z &= \frac{1}{1 + c^{\frac{(1-\delta)}{2\epsilon}}} \end{aligned} \tag{10}$$

Plugging these into $\mathbb{E}\left[U^A\right]$ and $\mathbb{E}\left[U^B\right]$, and simplifying, we obtain

$$\mathbb{E}\left[U^A\right] = \mathbb{E}\left[U^B\right] = \frac{1}{2}(\delta+1)\ln(c) = (1-\varepsilon)\ln(c),$$

where $\varepsilon \equiv \frac{1-\delta}{2}$, which highlights that the instantaneous impact is determined by δ, the relative dis-utility consumers obtain from switching to their less favoured basket. Furthermore, as the utility of type A and B is the same in this stylized example, we can model both of them as a representative consumer with a familiar instantaneous utility function from consumption, which takes the form

$$u(c_t) = (1-\varepsilon)\ln(c_t).$$

A.3 Transaction demand of money

We build on Romer (1986). Assume that a consumer has real income Y, which he/she spends during interval $[0, T]$. When the consumer keeps his/her money in the form of financial assets, he/she earns a nominal interest rate $i = r + \pi$, where the real interest rate (r) and inflation (π) are exogenous and constant during interval $[0, T]$. The consumer only uses cash (physical or digital) to buy goods, chooses N (the number of times to convert financial assets to cash during interval T), and chooses his/her cash demand during each interval between conversions. We consider a per transfer fixed cost b, which may converge to 0 when the cost of converting digital assets to a digital means of payment (e.g., CBDC) is negligible. As novel twist (motivated in section 2.1), we assume that the privacy cost depends on the number of withdrawals, and is $\xi(N)$. The

function $\xi(N)$ has the following realistic properties: $\xi(N=0)=0$ and $\lim_{N\rightarrow\infty}\xi(N)=\varepsilon$, and $\frac{\partial\xi(N)}{\partial N}>0$. Following this, the consumer utility with N conversions is

$$U=\int_0^T(1-\xi(N))\ln(c_t)\,dt-Nb$$

As in Romer (1986), we first solve the optimal consumption choice in an arbitrary interval $[s,s+\Delta t_k]$ where the consumer chooses c_t to maximize

$$u^c_{\Delta t_k}=\int_s^{s+\Delta t_k}(1-\xi(N))\ln(c_t)\,dt$$

subject to $\int_s^{s+\Delta t_k}P_tc_t=P_sm_k$.

$$u^c_{\Delta t_k}=(1-\xi(N))\left(\Delta t_k\ln\left(\frac{m_k}{\Delta t_k}\right)-\frac{\pi}{2}(\Delta t_k)\right).$$

As in Romer (1986), it then follows that it is optimal for the consumer to make conversions of financial asset to money in equal intervals, i.e., $\Delta t_k=\frac{T}{N+1}$.

The optimal demand for money and N is then the solution of choosing m_k and N to maximize

$$U=\sum_{k=1}^{N+1}\frac{T}{N+1}(1-\xi(N))\ln\left(\frac{m_k(N+1)}{T}\right)-(1-\xi(N))\frac{\pi}{2}\left(\frac{T}{N+1}\right)^2-Nb$$

subject to

$$Y=m_1+...+m_{N+1}\exp\left(-r\frac{NT}{N+1}\right)=\sum_{k=1}^{N+1}m_k\exp\left(-r\frac{(k-1)T}{N+1}\right).$$

We obtain

$$m_i=\frac{Y}{N+1}\exp\left(-r\frac{T(i-1)}{N+1}\right).$$

Using this in the utility function, we obtain

$$T(1-\xi(N))\left(\ln\left(\frac{Y}{T}\right)-\frac{i}{2}\frac{T^2}{N+1}+\frac{rT^2}{2}\right)-Nb.$$

Finally, maximizing with respect to the number of transfers, we obtain

$$\frac{\partial U}{\partial N}=\left(1-\xi(N)+\frac{\partial\xi(N)}{\partial N}(N+1)\right)\frac{i}{2}\left(\frac{T}{N+1}\right)^2-b-\frac{\partial\xi(N)}{\partial N}\left(\ln\left(\frac{Y}{T}\right)+\frac{rT^2}{2}\right).$$

If $1-\xi(N)+\frac{\partial\xi(N)}{\partial N}(N+1)\leq 0$, we have a corner solution as $\frac{\partial U}{\partial N}<0$. It is then optimal to set $N=0$, to keep/convert all assets into cash at date 0. If $1-\xi(N)+\frac{\partial\xi(N)}{\partial N}(N+1)>0$, we have a new privacy modified optimal time between conversions given by our model

$$\mu\equiv\frac{T}{N+1}=\sqrt{\frac{2b+2\frac{\partial\xi(N)}{\partial N}\left(\ln\left(\frac{Y}{T}\right)+\frac{rT^2}{2}\right)}{i\cdot\left(1-\xi(N)+\frac{\partial\xi(N)}{\partial N}(N+1)\right)}}. \tag{11}$$

A.4 Choice between means of payment

In this section we analyze the choice between means of payment, and further include a consumption basket that is not privacy sensitive. The latter is needed for re-introducing a possibility of real effects from the production side in the spirit of Appendix A.1. We focus on analyzing three distinct options for the means of payment: 1) using financial assets as means of all payments, which comes at a privacy cost, but maximizes interest income; 2) using a form of cash to buy all goods; 3) a mixed case where consumers use money to buy goods that are privacy sensitive and financial assets to buy non-privacy sensitive goods.

We do not consider here the nuances of timing when using a form of cash analyzed in section A.3, and assume that all consumers decide at date 0 which one of the above three options they will follow. Without loss of generality, and for the sake of comparison across models, we assume that consumers are initially endowed with Y units of a privacy sensitive good, which has initial price P_0^P. Consumers also take the initial price of the non-privacy sensitive good, P_0^N as given. Both prices follow the same exogenous inflation process in the interval $[0, T]$, such that $P_t^P = P_0^P \exp(\pi t)$ and $P_0^N = P_0^N \exp(\pi t)$, which consumers take as given. Consumers also take the real interest rate, r, and the nominal interest rate $i = \pi + r$, as given and constant.

A.4.1 Using financial assets as means of payment

We assume here that consumers are affected by the instantaneous privacy cost rationalized in Section 4, and solve the following problem

$$\max U^{FA} = \int_0^T \left\{\beta(1-\varepsilon)\ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^{NP}\right)\right\} dt,$$

subject to a law of motion

$$P_t^P c_t^P + P_t^N c_t^N + \dot{b}_t = i b_t,$$

initial condition $b_0 = P_0^P Y$, and end condition $b_T = 0$.

c_t^P and c_t^N are the consumer's demand for private and non-private baskets of goods at date t, respectively, b_t is their holdings of financial asset at date t, $\beta \in (0,1)$ is the preference parameter for the privacy-sensitive goods basket. $\varepsilon \in (0,1)$ captures the instantaneous dis-utility due to privacy concerns (see section 4).

The Hamiltonian of this problem is

$$\mathcal{H} = \beta(1-\varepsilon)\ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right) + \eta_t\left(i b_t - P_t^P c_t^P - P_t^N c_t^N\right),$$

and the implied optimality conditions are $\frac{\beta(1-\varepsilon)}{c_t^P} = \eta_t P_t^P$; $\frac{(1-\beta)}{c_t^N} = \eta_t P_t^N$; $-\dot{\eta}_t = \eta_t i$. From here, it follows that $\eta_t = \eta_0 \exp(-it)$, where η_0 is to be found from the differential equation for financial asset holdings (stemming from the budget constraint) which is

$$-\frac{(1-\beta\varepsilon)}{\eta_0 \exp(-it)} = \dot{b}_t - i b_t.$$

We find that $\eta_0 = \frac{(1-\beta\varepsilon)}{P_0^P Y} T$, which implies that

$$c_t^P = \left(\frac{Y}{T}\right)\frac{(\beta-\beta\varepsilon)}{(1-\beta\varepsilon)}\exp(rt),$$
$$c_t^N = \left(\frac{Y}{T}\right)\frac{P_0^P}{P_0^N}\frac{(1-\beta)}{(1-\beta\varepsilon)}\exp(rt).$$

It follows that the demand for the privacy sensitive good is decreasing in ε, and the demand for the privacy non-sensitive basket is increasing in ε.

Consumer utility under using financial assets as means of payment is

$$U^{FA} = T\ln\left(\left(\frac{Y}{T}\right)^{1-\beta\varepsilon}\left(\frac{P_0^P}{P_0^N}\right)^{(1-\beta)}\frac{(\beta-\beta\varepsilon)^{(\beta-\beta\varepsilon)}(1-\beta)^{(1-\beta)}}{(1-\beta\varepsilon)^{(1-\beta\varepsilon)}}\right) + (1-\beta\varepsilon)\frac{rT^2}{2}$$

A.4.2 Using cash as the only means of payment

We assume here that the consumer only uses cash to buy both privacy-sensitive and privacy-non-sensitive baskets of goods. The immediate benefit of this is that privacy concerns, ε, no longer affect consumer preferences (see Section 5), and the consumer solves the following problem:

$$\max U^M = \int_0^T \left\{\beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^P\right)\right\} dt,$$

subject to the constraint of sufficient resources to cover spending

$$\int_0^T \left(P_t^P c_t^P + P_t^N c_t^N\right) dt = P_0 Y.$$

The Lagrangian of this problem is

$$\mathcal{L} = \int_0^T \left\{\beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right)\right\} dt - \lambda\left[\int_0^T \left(P_t^P c_t^P + P_t^N c_t^N\right) dt\right] - P_0 Y$$

We obtain the following first order conditions: $\frac{\beta}{c_t^P} = \lambda P_t^P$ and $\frac{(1-\beta)}{c_t^N} = \lambda P_t^N$. Using these in the budget constraint implies that $\lambda = \frac{T}{P_0^P Y}$, and

$$c_t^P = \left(\frac{Y}{T}\right)\beta \exp\left(-\pi t\right)$$
$$c_t^N = \left(\frac{Y}{T}\right)\left(\frac{P_0^P}{P_0^N}\right)(1-\beta)\exp\left(-\pi t\right)$$

Unlike the case in section 6.1 where consumption of both baskets of goods is increasing over time (provided that the real interest rate is positive), the consumption in this case is decreasing over time (provided that the inflation rate is positive).

Consumer utility under using cash as means of payment

$$U^M = T\ln\left(\left(\frac{Y}{T}\right)\left(\frac{P_0^P}{P_0^N}\right)^{(1-\beta)}\beta^\beta\left(1-\beta\right)^{(1-\beta)}\right) - \frac{\pi T^2}{2}.$$

A.4.3 Using cash to buy the privacy sensitive basket, and financial assets to buy the non-privacy sensitive basket

As above, privacy concerns, ε, do not affect consumer preferences. The consumer now solves the following problem

$$\max U^{M/FA} = \int_0^T \left\{\beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^P\right)\right\} dt,$$

subject to the following set of constraints

$$P_t^N c_t^N + \dot{b}_t = i b_t$$
$$P_t^P c_t^P = m_t$$
$$m_0 = \int_0^T m_t dt,$$

with the end condition that $b_T = 0$, and the initial condition $m_0 + b_0 = P_0 Y$. While this problem is slightly more complex than the ones in Sections 6.1 and 6.2, it is solved as follows.

The Hamiltonian of this problem is

$$\mathcal{H} = \beta \ln\left(c_t^P\right) + (1-\beta)\ln\left(c_t^N\right) + \eta_t\left(ib_t - P_t^N c_t^N\right) + \lambda\left(m_t - P_t^P c_t^P\right).$$

We obtain the following optimality conditions, $\frac{\beta}{c_t^P} = \lambda P_t^P$; $\frac{(1-\beta)}{c_t^N} = \eta_t P_t^N$; $-\dot{\eta}_t = \eta_t i$. The differential equation from the budget constraint for financial assets is

$$\frac{(1-\beta)}{\eta_0 \exp(-it)} + \dot{b}_t = ib_t.$$

From here, $\frac{(1-\beta)}{\eta_0}T = P_0^P Y - m_0$, while from the budget constraint for money, $m_0 = \int_0^T m_t dt = \int_0^T \frac{\beta}{\lambda} dt = T\frac{\beta}{\lambda}$.

Using these relationships to maximize consumer utility with respect to η_0 and λ gives

$$\eta_0 = \lambda = \frac{T}{P_0^P Y}.$$

Using this we obtain

$$c_t^P = \left(\frac{Y}{T}\right)\beta \exp(-\pi t)$$
$$c_t^N = \left(\frac{Y}{T}\right)(1-\beta)\left(\frac{P_0^P}{P_0^N}\right)\exp(rt)$$

and the consumer utility under using cash as means of payment for the privacy-sensitive basket and financial assets for the privacy-non-sensitive is

$$U^{M/FA} = T\ln\left(\left(\frac{Y}{T}\right)\left(\frac{P_0^P}{P_0^N}\right)^{(1-\beta)}\beta^\beta(1-\beta)^{(1-\beta)}\right) - \beta\frac{\pi T^2}{2} + (1-\beta)\frac{rT^2}{2}$$

A.4.4 Comparison of means of payments

First, as arguably expected, using cash-like money only for the privacy-sensitive basket is strictly better than using cash-like money to buy all goods, as

$$U^{M/FA} - U^M = (1-\beta)\frac{(\pi+r)T^2}{2} = (1-\beta)\frac{iT^2}{2}.$$

The utility difference is proportional to the nominal interest rate and to consumers' relative preference for the privacy-non-sensitive basket. There are some important outside-the-model considerations here: this utility gain is rather small when the nominal interest rate is low, and we have not incorporated behavioral aspects such as the inconvenience of using multiple means of payments.

Second, a more interesting comparison involves the use of money compared to the use of financial asset as means of payment. We find that

$$U^{M/FA} - U^{FA} = T\ln\left(\frac{\left(\frac{Y\beta}{T}\right)^{\beta\varepsilon}(1-\beta\varepsilon)^{(1-\beta\varepsilon)}}{(1-\varepsilon)^{\beta(1-\varepsilon)}}\right) - \beta\frac{\pi T^2}{2} - \beta(1-\varepsilon)\frac{rT^2}{2}$$
$$U^{M} - U^{FA} = T\ln\left(\frac{\left(\frac{Y\beta}{T}\right)^{\beta\varepsilon}(1-\beta\varepsilon)^{(1-\beta\varepsilon)}}{(1-\varepsilon)^{\beta(1-\varepsilon)}}\right) - \frac{iT^2}{2} + \beta\varepsilon\frac{rT^2}{2}$$

We see that greater privacy concerns (higher ε) lead to a greater benefit, and thus greater demand for privacy-preserving cash. This observation explains why the demand for cash is still there and perhaps increasing (higher ε), even though it is difficult to argue that the Clower constraint holds.

Another immediate effect from the above is that a lower inflation increases cash demand. Interestingly, the impact of the real interest rate on cash demand depends on whether consumers use cash to buy all goods or privacy sensitive consumption baskets only. From the above, we also see that consumers with greater wealth (i.e., higher Y) are more likely to use cash-like assets as means of payment.