

Best Before? Expiring Central Bank Digital Currency and Loss Recovery

by Charles M. Kahn,¹ Maarten R. C. van Oordt² and Yu Zhu³

¹ University of Illinois

² Tinbergen Institute
Vrije Universiteit Amsterdam

³ Banking and Payments Department
Bank of Canada

cmkahn@illinois.edu, m.oordt@vu.nl, yuzhu@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

We are grateful to Zamid Aligishiev, David Andolfatto, Rakesh Arora, Jonathan Chiu, Carl Andreas Claussen, Martin Diehl, Jean Flemming, Maximilian Guennewig, Janet Jiang, Anneke Kosse, Jiaqi Li, Cyrus Minwalla, Francisco Rivadeneyra and Dinesh Shah, Xinyue Wen, participants in the 96th Annual Conference of the Western Economic Association International (2021), the Annual Central Bank Research Association Meeting (2021), the Federal Reserve Bank of Philadelphia's Workshop on Digital Currencies (2021), the 52nd Konstanz Seminar on Monetary Theory and Monetary Policy (2021), the 7th Joint Bank of Canada and Payments Canada Symposium on Shaping the Future of Payments (2021), the Economics of Payments X Conference (2021), the 3rd Crypto Asset Lab Conference (2021), the 3rd Tokenomics Conference (2021), the 4th UWA Blockchain and Cryptocurrency Conference (2021), and seminar participants at the Bank of Canada (2021), De Nederlandsche Bank (2021), the Hong Kong Baptist University (2021) and the University of Illinois (2021) for helpful comments and suggestions. With thanks to Ramin Shahabadi for excellent research assistance. Views expressed do not necessarily represent official positions of the Bank of Canada.

Abstract

An important feature of physical cash payments is resilience, which is due to their indifference to power outages or network coverage. Many central banks are exploring issuing digital cash substitutes with similar online payment functionality. Such substitutes could incorporate novel features, making them more desirable than physical cash. This paper considers introducing an expiry date for online digital currency balances to automate personal loss recovery. We show that this functionality could substantially increase consumer demand for digital cash, with the time to expiration playing an important role. Having more information available to the central bank improves accuracy of loss recovery but may decrease welfare.

Topics: Digital currencies and fintech

JEL codes: E41, E42

1 Introduction

As the transactions demand for physical cash wanes, many central banks are considering issuing digital cash substitutes (central bank digital currencies, or retail CBDCs) (Boar et al., 2020). An important property of a digital cash substitute is its operational resilience: Physical cash enables economic exchange even in conditions without electrical power or network access. Policy makers are seeking to build similar offline payments functionality into digital cash substitutes, which could be used in remote communities and serve as a backstop system when any disruption occurs (Bank of Canada, 2020; Bank of England, 2020; Group of Seven Central Banks, 2020).¹ At the same time, a digital cash substitute with offline functionality could incorporate features that would make it more desirable to use than physical cash.

This paper considers one possible feature: the introduction of an expiry date on offline digital currency balances. Although such functionality might seem inconvenient at first glance, it has one major advantage, previously unconsidered: It would facilitate personal recovery of funds accidentally lost.² An inconvenience of a bearer instrument such as cash is that it is easily lost, with little possibility of recovery by the owner.³ One reason for the lack of opportunity for recovery is that it is usually difficult for the owner to prove that cash is truly lost and will never be used for payments in the future.⁴ This is different for a digital currency that is allowed to expire over time. Since money balances that remain unspent after their expiry date cannot be spent in the future, it would be safe for the central bank to reimburse the (most likely) owner in terms of online balances. The reimbursement of expired funds could be implemented in a fully automatic fashion without the need for the owner to file a loss claim. To further simplify end-user experience, the expiry date could automatically be refreshed before the funds expire whenever users' devices connect to the network or are used to pay at a point-of-sale terminal with a network connection.

¹Offline payments functionality may be even more important in developing countries where substantial shares of the population have unreliable access to electricity or no access at all (data collected by the World Bank suggest that about 10 per cent of the world population had no access to electricity in 2019).

²Protecting individuals against accidental loss of cash is a new rationale for an expiry date. Others have investigated alternative rationales, such as stimulating spending at the macro level; see the literature review for details.

³One important role of modern banks is to keep deposits safe from loss or theft. In the electronic world, the loss problem can be even more dramatic. According to the *New York Times*, “Of the existing 18.5 million Bitcoin, around 20 percent – currently worth around \$140 billion – appear to be in lost or otherwise stranded wallets, according to the cryptocurrency data firm Chainalysis.” (Popper, Jan. 12, 2021).

⁴This statement does not consider mutilated and contaminated bank notes (in which case, often, the loss can be proven). Some central banks spend considerable effort to reimburse owners of damaged bank notes. An infamous anecdote is the story about a Dutch cow who ate a wallet without realizing that it would be her most expensive, and alas, also her last meal. The cow was slaughtered and her tripe was delivered to the Dutch Central Bank where currency recovery experts retrieved the remains of seven bank notes of a thousand guilders each (Dutch News Report, 1974).

The features that are necessary to rule out double-spending in a payment system for offline payments also cause digital currency balances to be subject to loss events. To completely rule out double-spending in offline environments, it is necessary to uniquely store offline digital currency balances in a single local device.⁵ To see this, consider the situation where copies of the same offline digital currency balances could be stored in multiple local devices. In this situation, the payer could simply double-spend the same funds by using different devices to pay different offline payees.⁶ By definition, offline payments do not allow verification of a payment based on real-time information in a central ledger, so it would be impossible for offline payees to be informed in a timely manner as to whether the payer is attempting to double-spend.⁷ Similarly, it is necessary to separate (or “ earmark ” / “ lock ”) digital currency balances available for offline spending with a local device from balances that can be spent without that device. Otherwise, after storing funds in a local device for offline payments, a payer could continue to spend the same funds to pay an online payee while using the offline device to pay a different offline payee. There would be no way for either payee to be aware of the attempt to double-spend at the time of the payment. The separation of offline digital currency balances while storing them uniquely in a single local device makes those funds subject to loss due to, e.g., malfunctioning, physical theft, or loss of the device.

The present paper starts from the position that the central bank would like to rule out double-spending, so that offline money balances will be separated and uniquely stored in a single local device, and, hence, will be subject to loss. We do so for two reasons. First, financial inclusiveness and universal access are regularly raised as core public policy goals for issuing digital cash substitutes (Miedema et al., 2020). Limiting fraud in a system that does not rule out double-spending may be prohibitively expensive if there are no possibilities to exclude bad actors who abuse possibilities to double-spend.⁸ Second, this is the combination of features that is most like the existing

⁵In practice, the design of a payment instrument for a generally accepted money is subject to a trilemma. As we will see, it can only have two out of three properties: “offline payment functionality,” “no double-spending” and “loss of funds not implied by device loss.”

⁶Note that uniquely storing funds in a single local device also rules out the possibility for the owner to create a backup. The absence of the possibility to restore a backup is necessary to prevent double-spending. Otherwise, an agent could pretend to have lost a local device and restore the funds on a second local device, which would allow the agent to use the two devices to double-spend at two different offline locations.

⁷A separate but related point is that the device to store funds needs to be tamper-resistant so that it is prohibitively expensive to restore a previous level of balances after making a payment or to copy the device’s contents to another device. Physical cash achieves this through its physical nature and security features that make it difficult to copy (and difficult to counterfeit). An offline CBDC system may require additional security measures to mitigate the impact of a breach of tamper-resistance (Minwalla, 2020; Armelius et al., 2021a), but the technological challenges do not seem insurmountable, given past experience (Grym, 2020).

⁸Some traditional systems for offline payments, such as cheques and store-and-forward payments with credit cards, do not completely rule out double-spending. In these settings, individuals can be excluded after abuse by denying them as client for a chequing account or credit card. Fraud may also be mitigated by introducing penalties for bad actors through law enforcement. Even with exclusion and law enforcement, the costs of fraud to payees can be considerable. Fraud with paper cheques continues to increase even as they are used less and less for payments (American Bankers Association, 2020). Moreover, Adyen (2020) indicates that authorization for offline store-and-forward payments with credit cards may be as low as 95 per cent.

arrangements for physical cash: Offline balances (cash) are separated from online balances when the bank debits the withdrawer’s account for the amount withdrawn and are then held uniquely in the form of physical notes or coins.

Given this design choice, the present paper studies the economics of introducing an expiry date to facilitate loss recovery for offline money balances.⁹ Consumers in our model need to choose the optimal distribution of their money between offline and online balances. Both types of balances can be used to pay in environments where consumers have network connectivity (“centralized meetings”). During an outage, consumers can exclusively trade in an environment without network connectivity (“decentralized meetings”), in which only offline money balances can be used. A difference with typical monetary models (e.g., the centralized/decentralized markets models such as [Lagos and Wright, 2005](#)) is that both the occurrence and the length of decentralized periods are stochastic. In our model, an outage occurs occasionally and lasts for a number of periods. A disadvantage of offline money balances is that they can be lost. Users may be reimbursed automatically for lost offline balances after they expire.¹⁰

Without an expiry date, balances are irrecoverable, as is usually the case with cash. With an expiry date, users can be reimbursed for losses, but merchants may not always be willing to accept offline money.

Our calibration results suggest that the introduction of an expiry date to facilitate loss recovery can have a substantial positive impact on consumer demand for offline digital currency balances. The reason is straightforward: The ability to reimburse individuals for personal losses once the lost digital currency expires reduces the cost of losses from the full amount to the inconvenience of temporarily not having access to the funds (until after they expire). This addresses a potentially significant personal cost to the users of offline money balances: A small survey suggests the per annum probability of personal losses from offline devices to be in a range of 8 per cent (for funds

⁹Physical means of payment can also be issued with expiry dates, and sometimes are. Grocery store coupons typically have expiry dates. Cheques issued by governments or financial institutions also include expiry dates, arguably to facilitate redemption if the recipient loses the cheque. Countries have sometimes imposed an expiry date on certain denominations or vintages of their cash or coin, often as an anti-crime measure. Sudden aggressive banknote changeovers with inconvenient redemption rules and old notes quickly becoming invalid have been blamed for contributing to the unwillingness of the public to continue holding physical cash in both Sweden and Norway ([Armelius et al., 2021b](#)). Such changeovers could be less painful for individuals who hold cash for emergency situations if serial numbers on bank notes were tied to the individuals who withdrew them – potentially on a voluntary basis – and if those individuals would be automatically reimbursed when the notes expire without being deposited by someone else. In principle, expiry dates for physical cash could also be used to facilitate loss recovery; however, such an arrangement is much more attractive in a digital environment, with automatic refreshment.

¹⁰For our mechanism to be effective, it requires that spending digital cash stored in a local device will require some form of user authentication (e.g., unlocking a phone when digital cash is stored in a smart phone or entering a pin code when digital cash is stored in a card), so that it is unlikely that someone finding or stealing a local device will result in the stored digital cash being spent.

stored locally on a smart phone) to 16 per cent (for funds stored locally in a payment or stored value card).¹¹

Starting from the optimal expiry date, we find that the cost of small deviations is strongly asymmetric. There is a high cost associated with setting an expiry date that is shorter than optimal, while the cost of setting an expiry date somewhat longer than optimal is limited. The high cost associated with setting an expiry date that is somewhat too short is that it may prevent the ability to conduct any transactions when payees expect to remain offline for a period that exceeds the expiry date. Payees will refuse to accept offline transactions, and in the limit an extremely short expiry date is equivalent to a situation without offline cash. The only inconvenience to the users from setting an expiry date that is longer than optimal is the additional delay in recovering lost offline digital currency balances.

Determining the likely owner of expired offline digital currency balances requires exchange of information between the central bank and the devices of consumers when the devices are online. We consider both a low- and a high-information model for how the central bank infers whether an agent's expired digital currency balances remained unspent. The low-information model places the onus fully on the payee to deposit offline digital currency balances before the expiry date. This model could support a higher level of privacy in that it does not require the payer's device to reveal whether and where offline digital currency balances have been spent. However, loss recovery in this model is less precise in that the payer may be reimbursed for expired funds that a payee failed to deposit in a timely manner.¹² The high-information model requires the payer's device to reveal whether and where funds have been spent when connecting to the central bank. This model is likely to be conceived as more invasive in terms of privacy. Under this model, the payee is still required to deposit received offline digital currency balances before the expiry date. However, whenever a payee fails to do so, then disclosure by the payer's device may allow the central bank to reimburse the payee rather than the payer.

Our results suggest that enabling more information sharing on spending has an ambiguous impact on social welfare. Whether it improves social welfare depends on whether payers choose to reconnect their devices after making offline payments. The choice will depend on comparing the forgone interest on unspent offline balances to the opportunity of a windfall gain if the payee fails to deposit the spent balances. If payers were to choose to reconnect their devices, then loss recovery is more precise and social welfare may be higher in the high information model (it could also be lower). However, in extreme cases it can be socially optimal for payers not to reconnect, in which

¹¹The details of the survey are reported in Appendix A.

¹²A comparable windfall profit for the payee occurs in a traditional payment setting when a payee fails to deposit a cheque written by the payer.

case the high information model can only lead to lower social welfare. Payers will forgo interest over their offline balances for a longer period of time and are therefore inclined to carry lower offline balances, resulting in lower spending during offline periods.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 illustrates the major trade-offs involving cash related to outages and loss events in the context of simple finite-time model. Section 4 introduces a more complex infinite-time model with a stochastic length for offline periods to obtain a better understanding of the quantitative impact of introducing an expiry date to facilitate loss recovery. Section 5 discusses the results of calibrating the more complex model. Section 6 concludes.

2 Related Literature

Our paper fits into the rapidly expanding economic literature on CBDC.¹³ Early economic research on this topic focused primarily on whether it would be beneficial if central banks were to issue CBDCs (Barrdear and Kumhof, forthcoming; Brunnermeier and Niepelt, 2019; Keister and Sanches, 2019; Andolfatto, 2021; Chiu et al., 2019; Fernández-Villaverde et al., 2021; Schilling et al., 2020; Kahn et al., 2020a; Zhu and Hendry, 2019). More recently, the focus has increasingly shifted towards design aspects of CBDC. These include the security features of a CBDC (Kahn et al., 2020b), the privacy it provides to its users (Garratt and Van Oordt, 2021; Lee and Garratt, 2021; Chaum et al., 2021), whether CBDC should generally be more deposit-like or more cash-like (Agur et al., forthcoming), the programmability of payments (Kahn and Van Oordt, 2021), and whether CBDC balances should pay interest (Keister and Sanches, 2019; Jiang and Zhu, 2021; Garratt and Zhu, 2021). Li (2021), Bijlsma et al. (2021) and Huynh et al. (2020) estimate how some of these features could affect the demand for CBDC based on survey data. Auer et al. (2020) study the technological approaches and policy stances of central banks on the issuance of CBDC. Our paper contributes to this literature by studying whether a digital cash substitute should be designed with a potential expiry date where users are automatically reimbursed for expired balances in order to enable the recovery of lost balances.

The traditional Baumol-Tobin model of cash demand suggests that cash holdings should explode as the interest rate approaches zero (Baumol, 1952; Tobin, 1956). More recent literature has recognized that the cost of carrying cash consists not only of the forgone interest but also of the risk of losing cash balances. Alvarez and Lippi (2009) approximate the costs of carrying cash as

¹³See Kiff et al. (2020) and Carapella and Flemming (2020) for early surveys.

the sum of the nominal interest rate and the probability of cash theft – as a source of cash losses – when estimating a dynamic cash inventory model.¹⁴ He et al. (2005, 2008) study the implication of banks’ safe-keeping role on monetary equilibrium and policy if cash can be lost. Sanches and Williamson (2010) explain theoretically how credit can co-exist with cash in an environment with limited commitment when cash can be subject to theft. Moreover, the model of Williamson (2019) considers CBDC while explicitly assuming that CBDC comes with the advantage over cash that it cannot be stolen. Arguably, in the current low-interest-rate environment, one could reasonably take the position that the risk of losing cash has become one of the major reasons why individuals don’t carry around substantial amounts of cash. Our paper analyzes how a cost that may have become one of the major costs of carrying funds for offline payments can be limited for a digital cash substitute.

The motivation for an expiry date on a digital cash substitute to enable recovery of lost balances is distinct from that of putting an expiry date on stimulus money in order to encourage consumer spending in recessions (Andolfatto, 2020).¹⁵ Imposing an expiry date on a digital cash substitute without reimbursing the owner for expired balances effectively increases the cost of carrying those balances: It imposes the threat of a tax on its owner if the funds are not spent before the expiry date. This has quite the opposite effect of the introduction of an expiry date to enable recovery of lost balances, which reduces the cost of holding balances. Imposing an expiry date without reimbursing owners for expired balances leads to worse outcomes in our model, as our model includes no rationale for the government to stimulate spending.

3 A Model of Offline Money and Outages

The major trade-offs involving offline money related to outages and loss events can be illustrated in a simple discrete finite-time model. In this setup, we analyze how design factors affect the incentives of a consumer to hold offline money issued by a central bank. The essence of offline money in our model is that it is a bearer instrument that can be used for making offline payments during outages. Offline money may be physical cash in the form of coins and bank notes, or it may be stored-value in a payment card or smartphone. For convenience, in the model we will refer to offline money as “cash.” Carrying cash provides the ability to purchase consumption goods

¹⁴Kosse (2013) and Kahn and Liñares-Zegarra (2016) document some empirical evidence on the relationship between perceived safety and cash use.

¹⁵Some stimulus programs provided shopping vouchers with expiry dates (Kan et al., 2017).

during an outage, but entails the cost of forgone interest as well as the risk of losing the cash. The properties of the cash issued by the central bank determine whether cash lost once is lost forever.

All agents are assumed to have quasi linear preferences; in the simple model this assumption will reduce to risk neutrality. Agents discount time with factor $\beta < 1$. There are two categories of agents—producers (she) and consumers (he)—and two types of consumers, denoted by $s \in \{1, 2\}$. A consumer of type s will have a demand for at most s units of the good manufactured by the producer. A consumer starts out with m units of money balances. He will divide them between cash and online money holdings. The online money holdings could be interpreted as balances held in an account at the central bank or a commercial bank. Cash pays no interest; online holdings pay interest at the rate $i = \beta^{-1} - 1$.¹⁶


The timeline of the model is shown in Figure 1. At the initial date $t = 0$, the consumer decides what part of his money balances to hold as cash, denoted by z . (For simplicity, we treat cash holdings in excess of m as borrowings of online money balances at interest rate i .) At the end of $t = 0$, the consumer discovers his own type (we assume the two types are equally probable). Producers cannot observe the consumer’s type.

At $t = 1$, there is a possibility for the consumer to purchase units of the good from producers and to consume them. Each producer can supply at most one unit of the good at a cost of β . The value of each unit to the consumer is v , up to the capacity determined by the consumer’s type (1 or 2). Every consumer has the possibility to meet multiple producers at $t = 1$. Trading during a meeting works as follows: The consumer makes a take-it-or-leave-it offer that consists of a price per unit p_j and a method of payment $j \in \{c, d\}$, where c stands for cash and d stands for online money. Then the producer makes an acceptance decision $a \in \{0, 1\}$, where $a = 1$ stands for accepting the offer and producing the good. The transaction technology allows for an exchange of the good and money based on the agreed-upon price and method of payment.

Two types of adverse events may occur before the consumer gets to make an offer at $t = 1$. First, with probability δ the consumer may lose his cash. The occurrence of the loss is denoted by $\tilde{\delta} \in \{0, 1\}$. Second, with probability λ an outage may occur. The occurrence of the outage is denoted by $\tilde{\lambda} \in \{0, 1\}$. If there is no outage, then online money balances and cash are both feasible choices. Payments with online money take one period to settle, so that an online payment made by the payee at $t = 1$ starts earning interest for the payee from $t = 2$ onward. If there is an outage, then the consumer can only pay with cash. Thus no consumption is possible in the unhappy state

¹⁶Our framework is not directly suitable to study optimal monetary policy, which involves broader objectives than the ones studied in this paper. For an approach relating monetary policy to cash loss, see [He et al. \(2005, 2008\)](#).

Figure 1: Model of Lost Cash and Outages: Timeline



where there is an outage and the consumer has lost his cash. If the consumer pays with cash, then there is a probability η that the producer will lose the cash.¹⁷ The occurrence of the loss is denoted as $\tilde{\eta} \in \{0, 1\}$. All chance events are drawn independently. We assume the following parameter restriction:

$$v > 1/(1 - \eta). \quad (1)$$

Otherwise, the utility from consumption would be so small that the consumer would have no incentive to make an acceptable cash offer during an outage.

Formally, the payoff of the consumer as a function of consumption at $t = 1$ and the consumer's terminal money holdings w_c at $t = 3$ is given as

$$u(q, w_c; s) = \beta \min\{q, s\}v + \beta^3 w_c, \quad (2)$$

where q denotes the number of accepted offers, which is the sum of the number of accepted cash offers, q_c , and the number of accepted offers involving online payments, q_d .

The consumer chooses a combination of (z, p_j, j) to maximize the payoff function in (2) subject to a cash-in-advance constraint for offers involving cash payments

$$q_c p_c \leq z(1 - \tilde{\delta}) \quad (3)$$

and a no-outage constraint for offers involving online payments

$$\tilde{\lambda} q_d p_d \leq 0. \quad (4)$$

¹⁷Part of η may also be thought of as the cost to the producer of handling cash; this reinterpretation leaves the results in conditions (6) and (7) unaffected.

A producer simply needs to decide whether to accept the offer she receives; the payoff function of a producer is given by

$$u_p(a, w_p) = -\beta^2 a + \beta^3 w_p. \quad (5)$$

The income of the central bank consists of the seigniorage from cash losses, forgone interest on cash holdings and forgone interest on online money in transit. All the cost from money and payments for the consumer and producer are income for the central bank. Social welfare is defined as the sum of the payoffs of the consumers, the producers and the central bank.

3.1 Cash lost is lost forever

We first consider the equilibrium in the case where, as with physical cash, there is no possibility for the consumer to recover lost balances. In this situation, the ultimate money holdings of the consumer are given by

$$w_c = (1+i)^3 [m - z] + (1+i)^2 \left[(1 - \tilde{\lambda})((1 - \tilde{\delta})z - qp_j) \right] + (1+i) \left[\tilde{\lambda}(1 - \tilde{\delta})(z - qp_j) \right].$$

The elements in the first pair of brackets reflect the consumer's online money holdings at $t = 0$, which earn interest over the entire horizon. The elements in the second pair of brackets reflect the cash the consumer can deposit at $t = 1$ in the absence of an outage – if he didn't lose the cash – minus the consumer's expenditures. The elements in the third pair of brackets reflect unspent cash that could only be deposited by the consumer after an outage at $t = 2$. Similarly, the ultimate money holdings of the producer are given by the sales revenue – unless sales were paid with cash and lost – i.e.,

$$w_p = ap_j(1 - \tilde{\eta}\mathbf{1}_{j=c})(1+i).$$

Online money will be the default payment instrument of choice for the consumer, since the producer needs to be reimbursed for the risk of losing cash: If the consumer were to offer an online payment, then the producer accepts any offer with a price greater than or equal to $p_d = 1$. If the consumer were to offer a cash payment, then the producer rejects the offer unless the price is greater than or equal to

$$p_c = \frac{1}{1 - \eta}.$$

It is optimal for the consumer to make a take-it-or-leave-it-offer of an infinitesimal amount above these prices. So, it is cheaper from the consumer's point of view to pay with online money whenever

possible, that is, whenever there is no outage. Moreover, restriction (1) implies that, if an outage were to occur so that the consumer cannot pay with online money, he would always be willing to buy the quantity of goods he wants to consume at the higher price when carrying enough cash.

The only remaining question is whether the consumer is willing to carry enough cash to be able to pay during an outage. Carrying cash is costly because of the risk of losing it and because of the interest forgone. So, the consumer will either choose not to carry cash at all, or to carry exactly enough to pay p_c for one unit during an outage, or to carry the amount to pay $2p_c$ for two units during an outage. Which amount the consumer chooses depends on whether the cost of carrying cash is less than the benefit of the insurance for consumption during outages. The consumer is willing to carry enough cash to purchase at least one unit of the good during an outage if and only if

$$i + \delta \leq \lambda(1 - \delta) \times [\nu(1 - \eta) - 1]. \quad (6)$$

In other words, for the consumer to hold the first unit of cash, the cost of carrying cash – the sum of forgone interest and expected costs of losing cash, i.e., $i + \delta$ – on the left-hand side needs to be less than the probability that it can be used in an outage multiplied with the marginal benefit of spending with cash during an outage on the right-hand side.

For the consumer to be willing to carry more cash, a stronger condition is necessary. In particular, the consumer would be willing to carry enough cash to purchase two units of the goods during an outage if and only if

$$i + \delta \leq \frac{\lambda}{2}(1 - \delta) \times \left[\nu(1 - \eta) - 1 - \frac{i}{1 + i} \right]. \quad (7)$$

This condition is stronger for two reasons. First, the probability that the consumer wants to consume the second unit during an outage is half the probability that he will want to purchase the first unit. So carrying additional cash insures against a smaller probability event. Second, and more subtly, if the consumer turns out not to need the larger amount during the outage, the incremental cash will sit idle and forgo interest until after the outage.

3.2 Social welfare

Social welfare – the sum of the payoffs of the consumers, the producers and the central bank – depends on the expected level of consumption. Each unit consumed in equilibrium provides a net social welfare benefit of $v - \beta$ at $t = 1$. The impact of forgone interest and cash losses on the payoffs

of consumers and producers cancel out in the aggregate due to their positive impact on the payoff of the central bank.

The number of goods purchased per consumer depends on whether an outage occurs ($\tilde{\gamma} = 1$) and on the number of goods the consumer could purchase with his cash balances (z/p_c). Formally, the expected number of goods purchased per consumer is

$$\mathbb{E}(q|z/p_c) = (1 - \gamma)\mathbb{E}(q|z/p_c, \tilde{\gamma} = 0) + \gamma\mathbb{E}(q|z/p_c, \tilde{\gamma} = 1). \quad (8)$$

During normal periods, the level of consumption is unaffected by cash balances, and consumers choose to purchase the goods they would like to consume, and, hence, $\mathbb{E}(q|z/p_c, \tilde{\gamma} = 0) = \frac{3}{2}$. During outages, consumption may be limited by the available cash to purchase goods. As mentioned before, consumers choose to hold no more cash in equilibrium than they need to purchase either none, one or two units, so $z/p_c \in \{0, 1, 2\}$. Since transactions are not possible without cash when offline, we have $\mathbb{E}(q|0, \tilde{\gamma} = 1) = 0$. If a consumer decides to carry cash, then there is a probability δ that the cash will be lost and no purchases can be made during an outage. The consumer would be able to purchase at most one unit if he were to hold enough cash for one unit, so $\mathbb{E}(q|1, \tilde{\gamma} = 1) = 1 - \delta$. Similarly, if a consumer decides to carry enough cash to purchase two units, then the consumer may buy all the goods he would like to consume during an outage, unless he loses the cash, so we have $\mathbb{E}(q|2, \tilde{\gamma} = 1) = \frac{3}{2}(1 - \delta)$.

Thus social welfare in the model depends entirely on whether consumers carry sufficient cash balances to facilitate purchases during outages. Social welfare is lowest in the scenario where consumers carry no cash and no transactions between producers and consumers occur during outages. Then the expected number of units sold per consumer equals $\frac{3}{2}(1 - \lambda)$. Social welfare is intermediate in the scenario where consumers carry sufficient cash to purchase at most one unit when there is an outage. Then the expected number of units sold per consumer equals $\frac{3}{2}(1 - \frac{2}{3}\lambda\delta - \frac{1}{3}\lambda)$. Social welfare is highest in the scenario where consumers carry sufficient cash to purchase two units when there is an outage. Then the expected number of units sold per consumer equals $\frac{3}{2}(1 - \lambda\delta)$.

Which scenario occurs without an expiry date depend on conditions (6) and (7). Welfare is highest if both conditions hold. Welfare is intermediate if only the weaker condition (6) holds. Finally, if neither condition holds true, i.e., if the cost of carrying cash without an expiry date, $i + \delta$, is too great, then social welfare is lowest. Whether alternative cash designs with an expiry date can improve social welfare compared to a situation without expiry date will depend on whether those designs relax the conditions governing consumers' cash-holding decisions.

3.3 Cash with loss recovery and no information exchange

The probability of cash losses may induce consumers to reduce the cash holdings that could otherwise insure them against outages, as shown in the left-hand side of conditions (6) and (7). Next, we consider schemes involving an expiry date on cash that could help to weaken those constraints by alleviating the consequences of cash losses. The idea of introducing an expiry date on digital cash is that all cash that expires will be treated as lost and can be automatically reimbursed by the central bank in terms of online money balances to the agent who (most likely) lost it. Deriving which agent most likely lost the cash depends on the information exchange that occurs between the devices of consumers and the central bank when the devices are online.

We first consider the environment without any information sharing on spending between the consumer's device and the central bank, so that connecting the consumer's device does not reveal to the central bank whether and where the consumer spent offline balances. As a consequence, the onus will be on the producer to deposit the cash received from the consumer before it expires (note that the producer may fail to do so if she lost the cash). If no one deposits the cash before the expiry date, then the central bank will infer that the consumer still owns the cash when it expires.

A very short (one-period) shelf life would not be useful for facilitating transactions during outages. If cash withdrawn at $t = 0$ were to expire during an outage at $t = 1$, then the producer would always be too late when she tries to deposit the cash after the outage. Therefore the producer would not accept it, and, hence, the consumer would not be willing to hold it. Social welfare in an environment with cash with a one-period shelf life would be comparable to social welfare in an environment without cash.

On the other hand, a two-period shelf life allows the consumer to pay the merchant during an outage at $t = 1$ with money withdrawn at $t = 0$ while leaving enough time for the merchant to deposit the money in her online account at $t = 2$ before it expires. Since no information on spending is provided to the central bank, the consumer's decision to deposit unspent balances after an outage at $t = 2$ will not be affected by the expiry date. If the money is not deposited at $t = 2$ by the merchant or the consumer, then the central bank reimburses the consumer at $t = 3$ without any risk of losses to the central bank.

With an expiry date and no information sharing on spending, the consumer would be willing to hold enough cash to purchase a single unit during outages if and only if

$$i + \delta(1 - \beta^2) < \lambda(1 - \delta) \times [\nu(1 - \eta) - (1 - \eta\beta^2)]. \quad (9)$$

This condition based on cash with an expiry date is weaker than condition (6) for two reasons. The main reason is the lower cost of cash losses to the consumer, which reduces the cost of holding cash on the left-hand side of the condition. Before, cash lost would be lost forever. With loss recovery, there is only a cost of a delay during which the consumer cannot access the cash, as he has to wait until after the expiry date before the central bank can reimburse him. The second reason is the more subtle impact of putting the onus on the producer to deposit the cash before it expires. The producer will fail to do so if the cash is lost by the producer, in which case the consumer will have the luck of being reimbursed for cash that was lost by the producer. Note from before that the potential of cash losses lead the producer to require higher prices for cash payments. With the expiry date, the consumer has a small chance of being reimbursed without having lost cash, which reduces the wedge between the costs of cash and electronic payments to the consumer. The reduction in the wedge increases the marginal benefit of spending with cash during an outage, as shown on the right-hand side.

With an expiry date and no information sharing on spending, the consumer would be willing to hold enough cash to purchase two units during outages if and only if

$$i + \delta (1 - \beta^2) < \frac{\lambda}{2}(1 - \delta) \times \left[\nu(1 - \eta) - (1 - \eta\beta^2) - \frac{i}{1 + i} \right]. \quad (10)$$

This condition corresponds to condition (7) from the case with no expiration date.

The introduction of an expiry date with the objective of loss recovery has the potential to improve social welfare in an environment where there is no information exchange on spending between the consumer's devices and the central bank. The design weakens the conditions that must be satisfied for consumers to hold cash compared to the design without expiry date (conditions (9) and (10) compared to conditions (6) and (7)). The following proposition summarizes the previous discussion on welfare.

Proposition 1 *If there is no information exchange on spending between consumers' offline devices and the central bank, then personal loss recovery through a sufficiently long expiry date can only improve social welfare compared to the case without expiry date. It does so whenever it increases consumers' holdings of offline balances.*

3.4 Cash with loss recovery and information exchange

A potential disadvantage of the design with an expiry date but no information exchange on spending between the consumer's device and the central bank is the imprecise nature of loss recovery in that consumers may receive windfall reimbursements after spending cash if producers lose the cash. As a potential solution to the imprecise nature of loss recovery without information exchange, we explore an alternative scheme where the consumer's device reveals whether and where cash balances have been spent. The information released by the consumer's device can then be used by the central bank to reimburse the producer for cash she received and lost, rather than causing a windfall profit for the consumer. Whether this improves social welfare depends on the incentives for consumers to connect their devices to the central bank, since they cannot be required to do so.¹⁸

Suppose the environment is such that the consumer holds only enough cash to purchase a single unit. In this situation, the consumer has no incentive to reconnect after an outage. (The consumer would have spend all his money, so reconnecting only eliminates the probability to get the windfall profit of having the spent money returned when lost by merchants.) So, enabling information exchange on spending between the consumer's device and the central bank does not change condition (9), which determines whether consumers carry any cash at all.

Things are different if the environment is such that consumers hold enough cash for two units. In this situation, consumers who spent all their cash still have no incentives to reconnect. However, a consumer who does not spend all his cash may have incentives to reconnect, since depositing unspent cash allows him to earn interest. This also comes at a cost to him as he forgoes the windfall from being reimbursed by the central bank for cash previously spent during an outage that was lost by the producer. So, the consumer decides, based on trading off forgone interest on unspent cash balances, against the chance of getting back spent cash balances. In equilibrium, it is optimal for a consumer who bought a single unit and held enough cash to purchase two units to reconnect if the cost of forgone interest exceeds the potential windfall profit when the consumer is incorrectly reimbursed for the producer's cash losses, that is, if

$$i \geq \eta. \tag{11}$$

If no consumers reconnect because the reconnect condition (11) is violated, then loss allocation is not improved and the cash prices will be unchanged from before, i.e., $1/(1 - \eta)$. In this scenario,

¹⁸Requiring consumers to reconnect their device before receiving a reimbursement for cash losses would defeat the purpose of loss recovery, as lost devices are unavailable for reconnecting.

the average cost of carrying cash balances to consumers increases, because they have an incentive to delay depositing unspent cash balances after an outage. Not reconnecting leaves them with a chance of obtaining cash lost by retailers. Under information exchange on spending, if condition (11) is violated, consumers will hold enough cash to purchase two units of consumption only if

$$i + \delta(1 - \beta^2) < \frac{\lambda}{2}(1 - \delta) \times \left[\nu(1 - \eta) - (1 - \eta\beta^2) - \frac{i}{1+i} \frac{2+i}{1+i} \right]. \quad (12)$$

The higher cost of carrying cash compared to the case without information exchange is captured in the last element of the inequality, which reflects the delay in depositing unspent cash balances, which makes this a stronger condition than the one without information exchange in (10). Hence, if no consumers were to reconnect, i.e., if the reconnect condition in (11) is violated, then a system with more information exchange on spending can only reduce social welfare.

If consumers who did not spend all their cash balances were to reconnect to the central bank because the reconnect condition (11) holds true, then loss allocation can be improved, as the central bank will be informed when and where some consumers spend their cash. When the consumers who purchased a single unit at $t = 1$ reconnect at $t = 2$, then the lowest cash price that risk-neutral producers would accept for any unit sold equals

$$p_c^I = \frac{1}{1 - \eta(1 - \frac{1}{3}\beta)}.$$

The reduction in the cash price reflects the reduced cost of accepting cash due to the possibility of the central bank reimbursing the producer for lost cash. The one-third in the denominator reflects the fact that this only happens for one-third of the cash lost by the producer: There is a probability of two-thirds that the producer sells a unit to a consumer who spends all his cash on two units, and, hence, has no incentive to reconnect, while one-third of the goods are sold to users who purchase a single unit and reconnect to deposit unspent cash.¹⁹ If the producer loses the cash received from consumers who reconnect, then the producer will be reimbursed by the central bank, albeit with a delay – the cash must expire first – which explains the β in the denominator.

Given the lower cash price p_c^I with information exchange on spending and the reconnect condition (11) holding true, consumers would be willing to hold enough cash to purchase two units during an outage if

$$i + \delta(1 - \beta^2) < \frac{\lambda}{2}(1 - \delta) \times \left[\nu(1 - \eta(1 - \frac{1}{3}\beta)) - 1 - \frac{i}{1+i} \right]. \quad (13)$$

¹⁹A producer cannot differentiate between the two types of consumers, because a consumer with preferences to consume two units purchases the goods from two different producers.

Whether social welfare improves with information exchange on spending if the reconnect condition (11) holds true depends on whether condition (13) is weaker than condition (10). Which condition will be the weaker one depends on how two effects that reduce the cost of cash balance out. With information exchange and reconnecting consumers, producers can accept more favorable prices thanks to the better targeted loss recovery, which reduces the cost to consumers of carrying cash because they need to carry less of it to purchase the same amount of goods. Without information exchange on spending, the cost to consumers of using cash is reduced because consumers may benefit in the form of a windfall profit when producers lose the spent cash. If $v/3 > \beta$, then the first effect is stronger, and condition (13) is weaker than condition (10). Individuals will be more inclined to hold sufficient cash balances to pay for two goods with information exchange than without information exchange. Hence, if the reconnect condition (11) holds true and $v/3 > \beta$, then enabling better targeted loss recovery through information exchange on spending improves social welfare. If, however, $v/3 < \beta$, then the second effect is stronger, and condition (10) will be weaker than condition (13). Hence, if the reconnect condition (11) holds true and $v/3 < \beta$, then enabling better targeted loss recovery through information exchange on spending reduces social welfare. If some consumers reconnect, then information exchange may improve social welfare, but this is not necessarily the case, and welfare could be reduced.

Enabling information exchange on spending can make matters both better and worse compared to the case with no information exchange. This comes from the fact that consumers may choose strategically to avoid reconnecting devices for offline payments if information exchange is enabled. Doing so increases the effective cost of carrying cash, which may induce consumers to carry less cash compared to the case where there is no information exchange and, as a consequence, reduce social welfare. If some consumers do reconnect, then information exchange may improve social welfare due to more precise loss recovery, but it does not necessarily do so. But regardless of whether information exchange occurs, loss recovery is never fully precise, as some consumers have no incentives to reconnect (in the model, those who spent all their cash). The following proposition summarizes the discussion on the social welfare impact of enabling information exchange on spending.

Proposition 2 *If enabling information exchange on spending deters all consumers from reconnecting, then it can only reduce social welfare compared to the case with no information exchange. If some consumers reconnect, then enabling information exchange on spending may increase or reduce social welfare compared to the case with no information exchange.*

4 An Infinite Horizon Model

To better study the impact of an expiry date on the demand for cash and the optimal length of the period before cash expires, it is convenient to consider a model with no limit on the number of periods. Time is discrete as before, and periods are numbered $0, 1, 2, \dots$. An outage begins at a random date $t \geq 1$. Once the outage is over, no further outages occur. (This assumption allows us to avoid some technical complications arising when successive surprise outages occur with no gaps. Given that outages occur infrequently, this assumption is a reasonable approximation.) Conditional on the outage not having occurred by period $t - 1$, the probability that the outage begins in period t is a constant λ . The outage is of stochastic length; g_τ is the probability it lasts τ periods. The length of the outage is revealed at the first period of the outage and known to all agents.

Following [Lucas \(1982\)](#) and [King et al. \(1992\)](#), we assume that all agents are household pairs consisting of a producer, who is in charge of selling, and a shopper, who is in charge of buying. Every shopper can use only the offline cash that his partner accumulated in the last normal period. Households cannot consume their own production. In a normal period, households derive utility $u(x)$ from consuming x units of the numeraire good, which can be produced at a constant marginal cost which is normalized to one. We assume that $u(x)$ is concave and strictly increasing in x and $u(0) = 0$. Agents decide the amount of offline money brought into the next period. To focus on the essence of the payment design problem, we assume that post-outage, the real value of offline money is constant over time, i.e., inflation is zero. As in the normal periods, households want to consume the good in the outage, but the utility of the consumption of x units depends on the length of the outage. At the beginning of the outage, buyers buy all goods needed from sellers for consumption during the outage. We assume that in each period in the outage, households consume the same amount.²⁰ Hence, their indirect utility from consuming x during the outage is $U(x, \tau)$, where

$$U(x, \tau) = u(x/\tau) \frac{1 - \beta^\tau}{1 - \beta}.$$

During the outage, any transaction must be facilitated by offline cash due to the lack of double coincidence. We focus on the regime without information exchange: Offline cash expires after T periods. If it has not been deposited by another agent, expired offline cash will be reimbursed to the account of the buyer. As in the simple model, reimbursement of expired cash occurs in the first period of the outage. If a seller receives offline cash, then she can deposit it by connecting to the central bank before the cash expires. But if the seller has no opportunity to connect to the central bank before the money expires, then she will not be willing to accept cash. Therefore, no

²⁰This assumption is not necessary but makes computation simpler.

producer is willing to accept the offline money if $\tau \geq T$ in an outage (Figure 2, panel a). On the other hand, if $\tau < T$, there can be trade (Figure 2, panel b). Each period, buyers lose their devices with probability δ and sellers lose their devices with probability η . As in the simple model, the event of loss occurs at the beginning of each period before any other actions.

We solve the model backwards. First, we consider the first period after the outage. Let a be the value in the buyer's online account, plus the expected value of reimbursements still due to be received for cash losses up to the present, discounted by the wait time until those reimbursements will be made, and let z be the current offline cash holding. Then the buyer's value function is

$$Q(a, z) = (1 - \delta)\bar{Q}(a, z) + \delta\bar{Q}(\bar{a}, 0),$$

where $\bar{Q}(a, z)$ and $\bar{Q}(\bar{a}, 0)$ are respectively the value functions when the buyer keeps her device or loses her device at the beginning of this period. Suppose the buyer loses the remaining cash balances in the first period after the outage ends (that is, just before reconnection is established), then the buyer will be reimbursed for those losses by the central bank after $\max(T - \tau, 0)$ periods, where τ is the realized length of the outage period. So, the discounted value of reimbursements increases by $z\beta^{\max(T-\tau, 0)}$ when he loses the cash in the first period after an outage, hence, $\bar{a} = a + z\beta^{\max(T-\tau, 0)}$. After the outage, there are no further incentives for the buyer to hold cash because no further outages occur and cash is costly to hold due to discounting and the possibility of a loss.

In a more general formulation it might be necessary to keep track of the entire history of cash losses by the buyer, since reimbursements of different losses could occur in different future periods. Fortunately, given the quasi-linear utility, only the expected present value of the reimbursement is relevant: actual value functions are equal to those in a scenario where agents get offline money reimbursed immediately after the loss, as long as the reimbursed value equals the expected discounted value of the future reimbursement. This allows us to write

$$\begin{aligned} \bar{Q}(a, z) &= \max_{x, \ell} u(x) - \ell + \beta\bar{Q}(0, 0), \\ \text{s.t. } x &= \ell + z + a. \end{aligned}$$

In this formulation, agents are reimbursed the expected discounted value a right away and have no future reimbursement; thus a enters the budget equation but not the continuation value function \bar{Q} . Also, notice that agents do not hold offline money after the outage. Therefore, the continuation value function is \bar{Q} . An immediate consequence is that $\bar{Q}_a(a, z) = \bar{Q}_z(a, z) = 1$. Therefore, $Q_a(a, z) = 1$ and $Q_z(a, z) = (1 - \delta) + \delta\beta^{\max(T-\tau, 0)}$.

Now we consider the problems before and during the outage. Let $W(a, z)$ be the value function at the beginning of a period. It equals the sum of the value function of a normal period and the value function of an outage weighted by the probability of an outage:

$$W(a, z) = (1 - \lambda)W^n(a, z) + \lambda W^o(a, z),$$

where W^n and W^o are the value functions in a normal period and an outage period, respectively.

The value function of a buyer in a normal period can be written as

$$W^n(a, z) = (1 - \delta)\bar{W}^n(a, z) + \delta\bar{W}^n(\bar{a}, 0),$$

where a buyer's value function is $\bar{W}^n(\bar{a}, 0)$ if he loses his device and $\bar{W}^n(a, z)$ otherwise. And $\bar{a} = a + \mathbb{E}\beta^{\phi(T)}z$, where $\phi(T)$ is the random time required to get reimbursement of lost cash. It depends on T and is random because it can be impacted by a potential outage, of which the arrival time and length are random. The calculation of the value of $\mathbb{E}\beta^{\phi(T)}$ is given in Appendix B. One can write

$$\begin{aligned} \bar{W}^n(a, z) &= \max_{\hat{z}, \ell, x} u(x) - \ell + \beta W(0, \hat{z}), \\ \text{s.t. } x &= \ell + a + z. \end{aligned} \tag{14}$$

This implies that $W_a^n(a, z) = 1$ and $W_z^n(a, z) = (1 - \delta) + \delta\mathbb{E}\beta^{\phi(T)}$.


Now move to the value function at the beginning of an outage. One implication of (14) is that we can assume that buyers carry no online value a into the next period if the outage does not occur, i.e., they use the discounted value of cash waiting for future reimbursement to consume now. Therefore, without loss of generality, we only need to consider the case where $a = 0$ at the beginning of the outage:

$$W^o(0, z) = \sum_{\tau} g_{\tau} J^{\tau}(z), \tag{15}$$


where J_{τ} is the value function at the beginning of an outage conditional on an outage length of τ periods. If $\tau \geq T$, offline cash expires in the middle of the outage and is reimbursed to the buyer after the outage (Figure 2, panel a). Therefore, no transaction occurs because a seller anticipates that she could not deposit the money before the expiration date. This implies that $J^{\tau}(z) = \beta^{\tau}Q(z, 0)$. If $\tau < T$, a seller can deposit the offline balances she receives if she does not lose her device (Figure 2, panel b). Therefore, a buyer can buy goods if he does not lose his device. This implies that if

Figure 2: Shelf life and outages

Panel (a): Shelf life does not exceed length of outage



Panel (b): Shelf life exceeds length of outage



$\tau < T$,

$$J^\tau(z) = (1 - \delta) \max_{x,p} [U(x, \tau) + \beta^\tau (1 - \delta_\tau) \bar{Q}(0, z - p) + \beta^T \delta_\tau \bar{Q}(z - p, 0)] + \delta \beta^T z, \text{ s.t. } p \leq z. \quad (16)$$

Here $\delta_\tau = 1 - (1 - \delta)^\tau$ is the probability that a buyer loses his device after trading in the first period of the outage but before he can deposit the cash. This formula follows because a buyer can lose his device with δ probability in each of the remaining $\tau - 1$ outage periods and also in the first period after the outage. The buyer decides his consumption x and payment p given that the terms of trade are determined by buyers making take-it-or-leave-it offers.

Sellers lose their devices with probability $\eta_\tau = 1 - (1 - \eta)^\tau$ after the trade and before they can deposit the cash. If a seller loses her device, she cannot claim the money, which will instead be reimbursed to the buyer. Taking this into account, the buyer's offer makes the seller indifferent between a trade and no trade, which implies $\beta^\tau (1 - \eta_\tau) p = x$. The β^τ captures the fact that sellers use the received funds to consume only after the outage is over. Then the envelope condition of J^τ is

$$\frac{\partial}{\partial z} J^\tau(z) = (1 - \delta) (\Lambda(z, \tau) + \beta^\tau (1 - \delta_\tau) + \delta_\tau \beta^T) + \delta \beta^T, \quad (17)$$

where

$$\Lambda(z, \tau) = \max\{\beta^\tau (1 - \eta_\tau) U_x(\beta^\tau (1 - \eta_\tau) z, \tau) - (\beta^\tau (1 - \delta_\tau) + \delta_\tau \beta^T), 0\}. \quad (18)$$

Thus in a normal period before the outage, a buyer brings z^* into the next period, where z^* solves the following equation in z :

$$1 = \beta \lambda \sum_{\tau=1}^{T-1} g_{\tau} [(1 - \delta)(\Lambda(z, \tau) + \beta^{\tau}(1 - \delta_{\tau}) + \delta_{\tau}\beta^T) + \delta\beta^T] + \lambda \sum_{\tau=T}^{\infty} g_{\tau}\beta^{\tau+1} + (1 - \lambda)\beta[(1 - \delta) + \delta\mathbb{E}\beta^{\phi(T)}]. \quad (19)$$

Because Γ is decreasing in z , the right-hand side of (19) is decreasing in z . Therefore, the solution to this equation is positive and unique if $T > 1$. The equation highlights the trade-offs involved in the choice of T . On the one hand, a higher T allows agents to trade in longer outages, which is reflected by the first summation on the right-hand side of (19). This increases the right-hand side of (19) and encourages the use of cash. On the other hand, it delays reimbursement of the lost offline cash, which is reflected by the terms β^T and $\beta^{\phi(T)}$. This decreases the right-hand side of (19) and discourages the use of cash. As a result, both the optimal cash holdings (z^*) and consumption in the outage are non-monotonic in T . Given z^* , the welfare of the buyer at period 0 is

$$W^n(0, 0) = u(x^*) - x^* - z^* + \beta W(0, z^*), \quad (20)$$

where $u'(x^*) = 1$.

5 Quantitative Analysis

Theoretically, offline cash with a longer T enables consumers to consume in more outage states, but consumers also need to wait longer to retrieve it if they lose their device. It remains an empirical question as to how to optimally set T . In this section, we calibrate the infinite horizon model to data to provide some initial insights on this question.

5.1 Calibration

We first parameterize the model. We assume the utility function has the familiar form $u(x) = x^{1-\sigma}/(1 - \sigma)$ and the duration of the outage τ follows a Poisson distribution with a parameter γ . Then the model is calibrated by specifying values for the unknown parameters of the model $(\beta, \lambda, \sigma, \delta, \eta, \gamma)$.

Table 1: Calibration

Parameter	Daily value	Annualized level
Discount factor (β)	0.99990	0.96
Risk aversion (σ)	0.70	
Loss probability consumer (δ)	0.0004845	0.162
Loss probability producer (η)	0.0004845	0.162
Outage probability (λ)	0.00061	0.200
– Length: Poisson distribution (γ)	9.555	0.0262

We consider a daily calibration. We set β to match an annual discount factor of 0.96. The risk aversion parameter σ determines the demand for offline cash. It is analogous to the curvature of the utility function in the decentralized market in the literature on money search, e.g., [Lagos and Wright \(2005\)](#). Many papers in this literature try to calibrate the curvature of this utility function; the resulting risk aversion parameter ranges from around 0.2 to more than 0.9. We set σ to be 0.7 in our benchmark calibration, but we also experimented with other values as discussed below.

Next, we calibrate the loss probability for consumers, δ . To pin down this parameter, we commissioned an online survey to estimate the probability of a consumer losing offline digital currency balances that would be locally stored in a payment card (see [Appendix A](#) for more details). The results suggest that the annual probability of a consumer losing digital currency balances stored locally in a payment card is around 16 per cent.²¹ We calibrate δ to match this probability. We assume that the probability that a seller loses cash stored in her device (η) is the same as that for the buyer. In a sensitivity analysis, we calibrate the model such that the loss probability corresponds to that when digital currency balances were locally stored in a phone and obtain qualitatively similar results (see [Appendix C](#)).

Lastly, we calibrate the parameters related to the likelihood and the length of an outage, λ and γ . For the likelihood, we choose λ such that someone is expected to enter into an extended offline period about once every five years. For the length of the outage, we choose γ such that the offline money balances are well over twice the level of daily spending as $T \rightarrow \infty$. The idea is that if $T = \infty$ – that is, if it takes infinitely long before lost cash is returned to the owners – then the

²¹This is higher than the number used by [Alvarez and Lippi \(2009\)](#) to approximate the probability of losing physical cash. They calibrate their model based on the annual probability that someone loses physical cash as a consequence of crime, which was around 2 per cent in Italy in 2002. The difference can be explained by a substantial probability of losing cash as a consequence of chance or carelessness.

cost of carrying offline digital currency balances to consumers in the model should be equivalent to the cost of carrying physical cash.²² Survey evidence of [Greene and Stavins \(2020\)](#) suggests that, on average, the level of precautionary cash holdings is well over three times the amount of daily spending on purchases.²³ [Table 1](#) summarizes the calibration.

5.2 Effects of the expiry date


We now use the calibrated model to analyze how different expiry dates affect the demand for offline money balances and the utility of buyers during outages. [Figure 3](#), panel (a) shows the demand for offline balances as T increases. Recall that a higher T enables consumers to consume in outages that last longer (“trade effect”) but also makes consumers wait longer to be reimbursed if they lose their device (“delay effect”). The former effect increases the demand for offline money, while the latter effect reduces the demand. If T is small, then the trade effect dominates because it is relatively likely that the outage is longer than T . Therefore, the demand for offline money increases. But if T is large, the delay effect dominates, and the demand decreases. Notice that when T is small, the demand rises sharply with T , while the demand decreases slowly if T is large. This is because consumption in outage is very valuable, and delay in reimbursement is not very costly because agents are patient. The demand for offline money peaks at $T = 24$ days, but this is sensitive to the shape of the distribution function for the length of the offline periods. Importantly, the demand for offline balances is substantially higher with an expiry date and loss recovery. For our specific calibration, the maximum demand is more than twice as high as the demand for offline money when there is no expiry date, i.e., when $T = \infty$. Panel (b) shows the expected daily utility of consumers during an outage. We also normalize the utility without an expiry date to 1. The utility-maximizing expiry date T is 27 days. It increases the expected daily utility during outages by about 20 per cent compared to the case of physical cash without an expiry date. Similar to the impact on demand, the cost of setting a longer expiry date is small, while setting an expiry date that is too short has a large negative impact. Lastly, panel (c) presents the expected per-period consumption in an outage. Its pattern is similar to that of offline money balances and utility during outages. At the utility-maximizing T , the expected per-period consumption during outages is on

²²A subtle difference between physical cash and the offline digital cash in the model is that spending physical cash does not require user authentication. This prevents a fortunate person who finds lost cash from spending it. In spite of this difference, the cost of carrying physical cash to the consumer will still be the same as the cost of carrying offline digital cash if the expiry date is infinitely far in the future. The reason is that, in the limit, the consumer would have to wait infinitely long for reimbursement, so the consumer is unaffected by someone else spending the balance once it is lost.


²³Survey evidence suggests that consumers in the United States spend on average \$50.32 per day on purchases (\$1559.9 divided by 31), while the level of precautionary cash holdings measured as cash on person and cash held elsewhere is estimated at on average \$180.30 ([Greene and Stavins, 2020](#), Tables 3a, 6 and 7). This suggests a ratio of about $180.30/50.32 \approx 3.6$.

Figure 3: Cash holdings with expiry date and privacy


Panel (a): Optimal cash holdings as a function of the expiry date



Panel (b): Daily utility during outages as a function of the expiry date



Panel (c): Daily consumption during outages as a function of the expiry date




Note: Panel (a) shows the optimal cash holdings as a multiple of daily consumption for different lengths of the expiry dates. Panel (b) measures the expected utility in the outage minus the cost of bringing offline money. Panel (c) measures the expected daily consumption in an outage. The calibration of the model is reported in Table 1.

Figure 4: Expiry date with a low discount factor

Panel (a): Optimal cash holdings

Panel (b): Consumption during outages



Note: Panel (a) shows the optimal cash holdings as a multiple of daily consumption for different lengths of the expiry dates. Panel (b) measures the expected daily consumption in an outage. The discount factor is set at an annualized level of 0.76. The other parameter values are set as reported in Table 1.

average around 80 per cent of the consumption during normal periods. It is less than the level of consumption during a normal period, because some of the costs associated with offline money balances remain after implementing the expiry date. However, consumption during outages almost doubles when comparing it to the case where $T = \infty$.

The main insights are not very sensitive to the choice of the risk aversion parameter σ , but quantitative implications can be. Suppose, for example, $\sigma = 0.5$. Then, the γ needs to be set at 14.169 to match the diary data. The values of T that maximize offline money demand and the daily utility during outages are 30 and 33, respectively. Compared to the case without an expiry date, the utility-maximizing T increases offline money holdings and expected per-period consumption during an outage by almost a factor three and raises expected utility during outages by more than half. Again, setting T larger than the utility-maximizing value leads to a small decrease in utility during outages, while setting T smaller can be very costly.

5.2.1 Cash-constrained households or high inflation

The relative flatness of the optimal cash holdings for longer expiry dates is partially driven by the choice for the discount factor. Our baseline calibration with an annualized level of 0.96 implies

that the waiting cost of the consumer for a refund after two years equals about 8 per cent of the lost balance. One may argue that this calibration does not reflect scenarios where consumers are financially constrained or the situation in some developing countries with high inflation rates. As an alternative, we set the discount factor at the more extreme annualized level of 0.76 – which implies that the consumer’s waiting cost for a refund after two years equals about 42 per cent of the lost balance – while keeping all other parameter values identical to those in Table 1. Figure 4 summarizes the results for this alternative calibration. The main difference with the baseline results is that the levels of the optimal cash holdings and consumption during outages with expiry date (blue lines) converge more quickly to their levels without expiry date (dashed lines). In other words, it becomes more costly to set an expiration date that is longer than optimal, so the discount factor is an aspect that the policy maker would need to take into account. That said, the asymmetry in deviations from the optimal expiry date remains because optimal cash holdings and consumption during outages converge to zero as the expiry date tends to zero on the left side of the chart.

6 Concluding Remarks

The robustness of physical cash as a means of payment comes at a cost: it is essentially impossible for a user of cash to convincingly demonstrate to the issuer that it has been lost and should be replaced. In this paper we argue that central bank digital currency can be designed to improve on physical cash – combining offline robustness with loss recovery – by including an expiry date, automatically renewable whenever the holder is online.

We have provided a simple model of the process and used it to examine the incentive issues entailed by such an arrangement. While a facility for recovering lost cash would be welfare improving, the details of its design matter. Increasing the information shared between consumers and the central bank in the loss recovery process could discourage consumers to carry cash.

We have also provided a more complicated dynamic model of outages and cash loss, one amenable to calibration. Our results show that including provision for loss recovery through expiry dates can have a significant welfare effect during outages, although it is clear that these calculations are only a first step in such an analysis. We have also examined the question of the optimal expiry date and shown that the benefits are asymmetric: Small upward deviations from the optimal duration have only minor welfare effects, while small deviations downward can entail substantial welfare losses. The preliminary conclusion then is that while a facility for limiting the life of offline CBDC is a

desirable part of the design, given the inherent uncertainties it will be safest to make the offline CBDC relatively long-lasting.

In principle, it would also be possible to use an expiry date to implement personal loss recovery for physical cash. However, relative to digital cash, such a scheme for physical cash would provide less protection against theft and accidental losses, since a different person can make payments with bank notes simply passing them around, while payments with a card or smart phone may be protected from spending by someone else by requiring some form of user authentication. Moreover, refreshing an expiry date on physical cash is not as simple as refreshing digital cash, so that a user may need to wait for reimbursement until the central bank issues the next bank note series.

Our paper has not touched upon the question of how the presence of offline balances with various expiry dates in a user's device could complicate negotiations between the payer and payee. While such analysis would be important in the case of physical cash with a variety of expiration dates, we foresee less relevance of these considerations in an electronic setting where expiry dates of offline digital currency balances would typically be refreshed regularly and in an essentially costless fashion. As noted earlier, refreshment of the expiry date of all offline digital currency balances stored in a user's device could be implemented with minimum end-user impact by implementing a regular refresh cycle (e.g., daily) if the device appears online, or, alternatively, whenever the device is used to make a payment at a point-of-sale terminal with a connection. This feature would limit the typical variation in the expiry dates both within and across the wallets of consumers as long as the expiry date is set at some point sufficiently far in the future.

As noted in the literature review, the introduction of CBDC will entail a host of decisions on multiple dimensions of design, of which expiry date is just one. Some of these issues will interact in significant ways with expiry date; in particular, it will be important to engage in further investigation of the relation between expiry date and interest on CBDC, as well as their joint impact on monetary policy.

A second dimension which is of central importance to the question of design is anonymity. However, the extent to which CBDC would provide anonymity is not immediately affected by whether an expiry date is implemented to achieve loss recovery. Although we considered two schemes that differed in the amounts of information on spending shared between the user devices and the central bank, neither scheme requires that the offline CBDC balances or the online account be associated with specific natural or legal persons. In principle, either scheme could be implemented in a CBDC system where the offline CBDC balances are not associated with specific natural or legal persons and where loss recovery occurs to a pseudonymous online account. Another possibility is

a hybrid system where owners of offline CBDC balances only benefit from loss recovery if their devices are linked to an online account of which the ownership is registered. In other words, whether loss recovery for offline balances through an expiry date would be a useful feature is a separate question from the question of whether a fully anonymous CBDC would be desirable, and both design dimensions will need to be considered by any CBDC issuer.²⁴

References

- Adyen. Offline Payments. *Online documentation (retrieved on 21 October 2020)*, 2020. URL <https://docs.adyen.com/point-of-sale/offline-payments>.
- I. Agur, A. Ari, and G. Dell’Ariccia. Designing Central Bank Digital Currencies. *Journal of Monetary Economics*, forthcoming.
- F. Alvarez and F. Lippi. Financial Innovation and the Transactions Demand for Cash. *Econometrica*, 77(2):363–402, 2009.
- American Bankers Association. Deposit Account Fraud Survey: 2019 Highlights. *Online report (retrieved on 16 April 2021)*, 2020. URL <https://www.aba.com/news-research/research-analysis/deposit-account-fraud-survey-report>.
- D. Andolfatto. Hot Money for a Cold Economy. *Federal Reserve Bank of St. Louis Working Paper*, 2020-019, 2020.
- D. Andolfatto. Assessing the Impact of Central Bank Digital Currency on Private Banks. *Economic Journal*, 131(634):525–540, 2021.
- H. Armelius, C.A. Claussen, and I. Hull. On the Possibility of a Cash-Like CBDC. *Sveriges Riksbank Staff Memo*, 2021a.
- H. Armelius, C.A. Claussen, and A. Reslow. Withering Cash: Is Sweden Ahead of the Curve or Just Special? *Working Paper*, 2021b.
- R.A. Auer, G. Cornelli, and J. Frost. Rise of the Central Bank Digital Currencies: Drivers, Approaches and Technologies. *BIS Working Paper*, 880, 2020.

²⁴As an example, the Central Bank of the Bahamas implemented a hybrid system for the Sand dollar (also known as the digital Bahamian dollar) where government-issued identification is not an enrolment requirement, but fewer restrictions are in place for registered accounts ([Central Bank of the Bahamas, 2021](#)).

- Bank of Canada. Contingency Planning for a Central Bank Digital Currency. *Technical Document*, 2020. URL <https://www.bankofcanada.ca/2020/02/contingency-planning-central-bank-digital-currency>.
- Bank of England. Central Bank Digital Currency: Opportunities, Challenges and Design. *Discussion Paper*, 2020.
- J. Barrdear and M. Kumhof. The Macroeconomics of Central Bank Issued Digital Currencies. *Journal of Economic Dynamics & Control*, forthcoming.
- W.J. Baumol. The Transactions Demand for Cash: An Inventory Theoretic Approach. *Quarterly Journal of Economics*, 66(4):545–556, 1952.
- M. Bijlsma, C. van der Cruijsen, N. Jonker, and J. Reijerink. What Triggers Consumer Adoption of CBDC? *De Nederlandsche Bank Working Paper*, 709, 2021.
- C. Boar, H. Holden, and A. Wadsworth. Impending Arrival: A Sequel to the Survey on Central Bank Digital Currency. *BIS Working Paper*, 118, 2020.
- L.D. Brown, T.T. Cai, and A. DasGupta. Interval Estimation for a Binomial Proportion. *Statistical Science*, 16(2):101–117, 2001.
- M.K. Brunnermeier and D. Niepelt. On the Equivalence of Private and Public Money. *Journal of Monetary Economics*, 106:27–41, 2019.
- F. Carapella and J. Flemming. Central Bank Digital Currency: A Literature Review. *FEDS Notes (9 November)*, 2020. URL <https://doi.org/10.17016/2380-7172.2790>.
- Central Bank of the Bahamas. How Can I Get a Sand Dollar Wallet? *Website (retrieved on 15 September 2021)*, 2021. URL <https://www.sanddollar.bs/individual>.
- D. Chaum, C. Grothoff, and T. Moser. How to Issue a Central Bank Digital Currency. *Swiss National Bank Working Paper*, 2021-03, 2021.
- J. Chiu, M. Davoodalhosseini, J.H. Jiang, and Y. Zhu. Central Bank Digital Currency and Banking. *Bank of Canada Staff Working Paper*, 2019-20, 2019.
- Dutch News Report. Wij Verbranden Elke Dag Vijftien Miljoen Gulden (translation: We Burn Fifteen Million Guilders Every Day). *Reformatorsch Dagblad (20 September)*, page 17, 1974.
- J. Fernández-Villaverde, D. Sanches, L. Schilling, and H. Uhlig. Central Bank Digital Currency: Central Banking For All? *Review of Economic Dynamics*, 41:225–242, 2021.

- R.J. Garratt and M.R.C. Van Oordt. Privacy as a Public Good: A Case for Electronic Cash. *Journal of Political Economy*, 129(7):2157–2180, 2021.
- R.J. Garratt and H. Zhu. On Interest-Bearing Central Bank Digital Currency with Heterogeneous Banks. *Working Paper*, 2021.
- C. Greene and J. Stavins. The 2019 Diary of Consumer Payment Choice. *Federal Reserve Bank of Atlanta Research Data Reports*, 20-4, 2020. URL https://www.atlantafed.org/-/media/documents/banking/consumer-payments/diary-of-consumer-payment-choice/2019/depc2019_tables-pdf.pdf.
- Group of Seven Central Banks. *Central Bank Digital Currencies: Foundational Principles and Core Features*. Bank of International Settlements, 2020. ISBN 978-92-9259-427-5. URL <https://www.bis.org/publ/othp33.pdf>.
- A. Grym. Lessons Learned from the World’s First CBDC. *Bank of Finland Economics Review*, 8, 2020.
- P. He, L. Huang, and R. Wright. Money and Banking in Search Equilibrium. *International Economic Review*, 46(2):637–670, 2005.
- P. He, L. Huang, and R. Wright. Money, Banking, and Monetary Policy. *Journal of Monetary Economics*, 55(6):1013–1024, 2008.
- K. Huynh, J. Molnar, O. Shcherbakov, and Q. Yu. Demand for Payment Services and Consumer Welfare: The Introduction of a Central Bank Digital Currency. *Bank of Canada Staff Working Paper*, 2020-7, 2020.
- J. Jiang and Y. Zhu. Monetary Policy Pass-Through with Central Bank Digital Currency. *Bank of Canada Staff Working Paper*, 2021-10, 2021.
- C.M. Kahn and J.M. Liñares-Zegarra. Identity Theft and Consumer Payment Choice: Does Security Really Matter? *Journal of Financial Services Research*, 50(1):121–159, 2016.
- C.M. Kahn and M.R.C. Van Oordt. On the Demand for Programmable Payments. *Working Paper*, 2021.
- C.M. Kahn, F. Rivadeneyra, and T.-N. Wong. Should the Central Bank Issue E-money? *Journal of Financial Market Infrastructures*, 8(4), 2020a.
- C.M. Kahn, F. Rivadeneyra, and T.-N. Wong. Eggs in One Basket: Security and Convenience of Digital Currencies. *Federal Reserve Bank of St. Louis Working Paper*, 2020-32, 2020b.

- K. Kan, S.-K. Peng, and P. Wang. Understanding Consumption Behavior: Evidence from Consumers' Reaction to Shopping Vouchers. *American Economic Journal: Economic Policy*, 9(1): 137–53, 2017.
- T. Keister and D.R. Sanches. Should Central Banks Issue Digital Currency? *Federal Reserve Bank of Philadelphia Working Paper*, 19-26, 2019.
- J. Kiff, J. Alwazir, S. Davidovic, A. Farias, A. Khan, T. Khiaonarong, M. Malaika, H.K. Monroe, N. Sugimoto, H. Tourpe, and P. Zhou. A Survey of Research on Retail Central Bank Digital Currency. *IMF Working Paper*, 2020-104, 2020.
- R.G. King, N. Wallace, and W.E. Weber. Nonfundamental Uncertainty and Exchange Rates. *Journal of International Economics*, 32(1):83–108, 1992.
- A. Kosse. The Safety of Cash and Debit Cards: A Study on the Perception and Behavior of Dutch Consumers. *International Journal of Central Banking*, 9(4):77–98, 2013.
- R. Lagos and R. Wright. A Unified Framework for Monetary Theory and Policy Analysis. *Journal of Political Economy*, 113(3):463–484, 2005.
- M. Lee and R.J. Garratt. Monetizing Privacy. *Federal Reserve Bank of New York Staff Report*, 958, 2021.
- J. Li. Predicting the Demand for Central Bank Digital Currency: A Structural Analysis with Survey Data. *Working Paper*, 2021.
- R.E. Lucas. Interest Rates and Currency Prices in a Two-Country World. *Journal of Monetary Economics*, 10(3):335–359, 1982.
- J. Miedema, C. Minwalla, M. Warren, and D. Shah. Designing a CBDC for Universal Access. *Bank of Canada Staff Analytical Note*, 2020-10, 2020.
- C. Minwalla. Security of a CBDC. *Bank of Canada Staff Analytical Note*, 2020-11, 2020.
- N. Popper. Lost Passwords Lock Millionaires out of Their Bitcoin Fortunes. *New York Times*, Jan. 12, 2021. URL <https://www.nytimes.com/2021/01/13/business/tens-of-billions-worth-of-bitcoin-have-been-locked-by-people-who-forgot-their-key.html>.
- D.R. Sanches and S. Williamson. Money and Credit with Limited Commitment and Theft. *Journal of Economic Theory*, 145(4):1525–1549, 2010.
- L.P. Santoso, R. Stein, and R. Stevenson. Survey Experiments with Google Consumer Surveys: Promise and Pitfalls for Academic Research in Social Science. *Political Analysis*, 24(3):356–373, 2016.

- L. Schilling, J. Fernández-Villaverde, and H. Uhlig. Central Bank Digital Currency: When Price and Bank Stability Collide. *NBER Working Paper*, 28237, 2020.
- K. Sostek and B.S. Slatkin. How Google Surveys Works. *White Paper (June 2018)*, 2018. URL https://services.google.com/fh/files/misc/white_paper_how_google_surveys_works.pdf.
- J. Tobin. The Interest-Elasticity of Transactions Demand for Cash. *Review of Economics and Statistics*, 38(3):241–247, 1956.
- S. Williamson. Central Bank Digital Currency: Welfare and Policy Implications. *Working Paper*, 2019.
- World Bank. Access to Electricity. *Data (retrieved on 3 June 2021)*, 2021. URL <https://data.worldbank.org/indicator/EG.ELC.ACCS.ZS>.
- Y. Zhu and S. Hendry. A Framework for Analyzing Monetary Policy in an Economy with E-money. *Bank of Canada Staff Working Paper*, 2019-1, 2019.

Appendix

A Survey

We obtain a rough estimate of the probability of individuals losing offline digital currency balances based on two single-question online surveys. We do so for two potential modes to store offline digital currency balances: when offline digital currency balances would be stored in a secure element in a payment card and when offline digital currency balances would be stored in a secure element in a smart phone. In either case, we will presume that the devices require some form of user authentication (e.g., a pin code or unlocking the device) so that the balances cannot be spent by others when the device is stolen or lost. The survey questions and the responses are reported in Table 2.

Our service provider is Google Surveys. The responses are provided by users on websites in the Google Surveys publisher network who are asked to fill out a survey before they can continue reading the content they would like to view (a so-called “survey wall”). Methodological details are provided by [Sostek and Slatkin \(2018\)](#). Generally speaking, the service provider implements several mitigation strategies to deal with response biases and provides weights to weight responses by age, gender and region. [Santoso et al. \(2016\)](#) provide a relatively positive assessment of the service for academic research in social science, albeit with some cautions. One concern they identify is the potentially less substantive engagement of respondents when facing a survey wall. We include both a “Don’t know” option and a “Don’t want to answer” option as potential responses to our survey questions in order to mitigate the risk of this affecting our outcomes. We find that the percentage of respondents who choose one of these options in our surveys is comparable to the percentage of responses for the “Don’t know” option observed in the assessment of Google Surveys by [Santoso et al. \(2016, p. 364\)](#).

The responses in Canada and the United States are generally quite comparable, as are the unweighted and weighted responses. Based on the weighted responses in the United States, the fraction of respondents who would not have lost stored-value in a payment card based on our survey question is estimated to be about $0.584/(0.584 + 0.086 + 0.027) \approx 0.837$ on an annual basis, which corresponds to an annual loss probability of 16.3 per cent. For stored-value in a phone, the corresponding estimate is about $0.716/(0.716 + 0.057 + 0.009) \approx 0.916$ on an annual basis, which corresponds to an annual loss probability of about 8.4 per cent. These are the loss probabilities that are used for the baseline calibration and the calibration of the robustness check. For Canadians,

Table 2: Survey questions and response rates

Panel (a): Over the past 12 months, did you replace or cancel a payment card (for example, a debit or credit card) because it was damaged, physically stolen or lost?

Answer:	Canada:		United States:	
	Weighted	Unweighted	Weighted	Unweighted
“No”	63.7%	60.6%	58.4%	59.2%
	(-2.6%, +2.5%)	(-2.2%, +2.1%)	(-2.9%, +2.8%)	(-2.2%, +2.1%)
“Once”	8.0%	7.1%	8.6%	8.2%
	(-1.3%, +1.6%)	(-1.1%, +1.2%)	(-1.5%, +1.8%)	(-1.1%, +1.3%)
“Twice, or more”	1.8%	1.6%	2.7%	2.2%
	(-0.6%, +0.8%)	(-0.5%, +0.7%)	(-0.8%, +1.1%)	(-0.6%, +0.8%)
“Don’t know”	4.2%	4.3%	5.1%	5.1%
	(-0.9%, +1.2%)	(-0.8%, +1.0%)	(-1.1%, +1.4%)	(-0.9%, +1.1%)
“Don’t want to answer”	22.3%	26.3%	25.2%	25.1%
	(-2.1%, +2.3%)	(-1.9%, +2.0%)	(-2.4%, +2.6%)	(-1.9%, +1.9%)
Respondents	1,376	2,001	1,146	2,001

Panel (b): Over the past 12 months, was your smart phone stolen, permanently lost, or broken so that you could no longer start it?

Answer:	Canada:		United States:	
	Weighted	Unweighted	Weighted	Unweighted
“No”	70.2%	67.1%	71.6%	72.7%
	(-2.4%, +2.3%)	(-2.1%, +2.0%)	(-2.7%, +2.6%)	(-2.0%, +1.9%)
“Once”	4.4%	4.5%	5.7%	4.0%
	(-1.0%, +1.2%)	(-0.8%, +1.0%)	(-1.2%, +1.5%)	(-0.8%, +1.0%)
“Twice, or more”	2.4%	2.7%	0.9%	1.2%
	(-0.7%, +0.9%)	(-0.6%, +0.8%)	(-0.4%, +0.7%)	(-0.4%, +0.6%)
“Don’t know”	2.9%	3.5%	2.4%	3.2%
	(-0.8%, +1.0%)	(-0.7%, +0.9%)	(-0.8%, +1.1%)	(-0.7%, +0.9%)
“Don’t want to answer”	20.1%	22.2%	19.4%	18.8%
	(-2.0%, +2.2%)	(-1.8%, +1.9%)	(-2.2%, +2.4%)	(-1.7%, +1.8%)
Respondents	1,419	2,001	1,118	2,001

Note: The table reports the responses for two single-question surveys held in both Canada and the United States. Responses are provided by users of websites included in the Google Surveys publisher network over the period from 5 to 29 May 2021. The weighted responses weigh respondents by age, gender and region and assign a zero weight to respondents for which this information is not fully available (hence, the higher count of respondents for the unweighted responses). The table reports the 95 per cent confidence intervals using the modified Wilson method (Brown et al., 2001) in parentheses.

the estimated annual loss probabilities are respectively 13.3 per cent for stored-value in a card and 8.8 per cent for stored-value in a phone. The pattern in both jurisdictions is that correspondents are less likely to lose stored-value in phones, potentially due to features that allow them to locate their devices when lost.

B Value of recovery with stochastic outage length

This appendix derives the value of $\mathbb{E}\beta^{\phi(T)}$, which is the expected value of loss recovery for a single dollar to the consumer as a function of the expiry date in an environment where the occurrence and the length of the outage are stochastic. The value of loss recovery is not straightforward because outages with stochastic length introduce uncertainty around the moment when the consumer can access recovered funds. Three different scenarios may materialize that need to be accounted for: (i) no outage may occur until the moment of recovery, (ii) an outage may occur that ends before the moment of loss recovery, and (iii) an outage may start before the moment of recovery but may continue until after the moment of recovery. The probabilities in the equation related to these scenarios are indicated below the equation

$$\mathbb{E}\beta^{\phi(T)} = \beta^{T+1} \left(\underbrace{(1-\lambda)^{T+1}}_{\substack{\text{No outage} \\ \text{until } T+1 \\ \text{(inclusive)}}} + \underbrace{\sum_{s=1}^T \sum_{t=1}^{T+1-s} \lambda(1-\lambda)^{s-1} g(t)}_{\substack{\text{Probability of outage} \\ \text{between now and } T+1 \\ \text{that ends before } T+1}} \right) + \sum_{i=1}^{\infty} \beta^{T+1+i} \left(\underbrace{\sum_{s=1}^{T+1} \lambda(1-\lambda)^{s-1} g(T+1-s+i)}_{\substack{\text{Probability of outage} \\ \text{between now and } T+1 \text{ (inclusive)} \\ \text{that ends at } T+i}} \right).$$


In scenarios (i) and (ii), the value of a recovered dollar simply equals β^{T+1} . In scenario (iii), the value of a recovered dollar depends on when it can be accessed by the consumer (i.e., β^{T+1+i} for access at $t = T + 1 + i$). The function calculates the expected value of loss recovery by summing the product of β^{T+1+i} and the probability that an outage starts before $T + 1$ (inclusive) and ends at $t = T + 1 + i$ for each $i = 1, 2, \dots, \infty$.

C Sensitivity analysis


According to the survey in Appendix A, the probability of losing value stored in a secure element in a smart phone is around 8.44 per cent, which is lower than the probability of losing value stored in a payment card. In this appendix, we analyze the sensitivity of our results to a lower probability of losing cash. If the offline money is stored on a secure element in a smart phone, a buyer loses

Figure 5: Cash holdings with expiry date and privacy: Sensitivity to lower loss probability


Panel (a): Optimal cash holdings as a function of the expiry date



Panel (b): Daily utility during outages as a function of the expiry date



Panel (c): Daily consumption during outages as a function of the expiry date



Note: Panel (a) shows the optimal cash holdings as a multiple of daily consumption for different lengths of the expiry dates. Panel (b) measures the expected utility in the outage minus the cost of bringing offline money. Panel (c) measures the expected daily consumption in an outage.

the offline balance with a daily probability of approximately $\delta = 2.42 \times 10^{-4}$. We again assume that $\delta = \eta$ and recalibrate the average length of the outages to match the holdings of physical cash as $T \rightarrow \infty$. We then obtain $\gamma = 6.82$ days for $\sigma = 0.7$. All other parameters in this calibration are the same as in those reported in Table 1.

Figure 5 reports the results for this alternative calibration with a lower loss probability. The qualitative results are similar to those obtained using the baseline calibration. As one may expect, the quantitative impact of implementing an expiry date is smaller when the loss probability is set at a lower value, but the impact is still substantial. The value of T that maximizes offline money holdings is 20 days, while the one that maximizes utility during outages is 23 days. Compared to the case without an expiry date, the utility-maximizing T increases offline money holding by more than one-half, daily consumption during an outage by more than one-third and daily utility during outages by slightly more than 10 per cent.