

Staff Working Paper/Document de travail du personnel 2019-24

Privacy as a Public Good: A Case for Electronic Cash



by Rodney J. Garratt and Maarten R. C. van Oordt

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.

Bank of Canada Staff Working Paper 2019-24

July 2019

Privacy as a Public Good: A Case for Electronic Cash

by

Rodney J. Garratt¹ and Maarten R. C. van Oordt²

¹ Department of Economics
University of California Santa Barbara
Santa Barbara, California, USA 91306
garratt@ucsb.edu

² Currency Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
mvanoordt@bankofcanada.ca

Acknowledgements

We are grateful to Charles M. Kahn, Yu Zhu and participants in presentations at the Bank of Canada (Ottawa, 2019) for helpful comments and suggestions. This paper was written while Garratt was a visiting scholar at the Bank of Canada. The views in this paper do not necessarily reflect those of the Bank of Canada.

Abstract

Privacy is a feature inherent to the use of cash for payments. With steadily increasing market shares of commercial digital payments platforms, privacy in payments may no longer be attainable in the future. In this paper, we explore the potential welfare impact of reductions in privacy in payments in a dynamic framework. In our framework, firms may use data collected through payments to price discriminate among future customers. A public good aspect of privacy in payments arises because individual customers do not bear the full cost of failing to protect their privacy. As a consequence, they may suboptimally choose not to preserve their privacy in payments. When left to market forces alone, the use of privacy-preserving means of payments, such as cash, may decline faster than is optimal.

Bank topics: Digital currencies and fintech; Bank notes; Payment clearing and settlement systems; Central bank research

JEL codes: E42, G28

“This [central bank digital] currency could satisfy public policy goals, such as (i) financial inclusion, (ii) security and consumer protection; and to provide what the private sector cannot: (iii) privacy in payments.”

– Christine [Lagarde](#), Speech, 14 November 2018.

1 Introduction

Privacy in payments is a feature inherent to the use of cash, but transactional usage of cash is in decline. Cash cannot be used for the increasing share of online transactions, and the share of cash payments at the point of sale is also gradually decreasing.¹ Going forward, transaction privacy may no longer be attainable due to steadily growing market shares of commercial payments platforms that generate datasets with payment histories at the user level.

Commercial payments platforms can monetize user data in multiple ways. They can share data with third parties, who may combine it with data from other sources and use the resulting dataset for marketing purposes or to generate predictions on creditworthiness.^{2,3} Alternatively, without sharing user-generated data, commercial payments providers can employ payments data to cluster users into profiles that, in turn, can be shared with third parties for marketing purposes or could be used for marketing purposes on behalf of those third parties.⁴ The effectiveness and profitability of such methods is expected to grow with advances in prediction tools (e.g., machine learning) and increases in computational power and the scale and scope of data gathering.

¹[Khiaonarong and Humphrey \(2019\)](#) report this trend in cash shares for various countries based on aggregate statistics. For national estimates of downward trends based on survey data, see, e.g., [Kumar et al. \(2018, United States\)](#), [Deutsche Bundesbank \(2017, Germany\)](#) and [Henry et al. \(2018, Canada\)](#).

²See, e.g., the [Bloomberg \(2018\)](#) news report, which reads “For the past year, select Google advertisers have had access to a potent new tool to track whether the ads they ran online led to a sale at a physical store in the U.S. That insight came thanks in part to a stockpile of Mastercard transactions that Google paid for.”

³Zhima (aka Sesame) credit, an affiliate of the Chinese Alibaba Group, reportedly uses data from Alipay to determine credit scores ([Financial Times, 2018](#)).

⁴Patent applications provide evidence of ample research and development investments of major commercial payment providers in this area; see the excerpts provided in Appendix A.

Lack of privacy in payments can also lead to consequences beyond marketing and credit provision. Credit scores generated in part from payments data are used by governments in “Orwellian” ways. In China, Zhima credit scores determine eligibility for discounts and deposit waivers, while low scores can reportedly lead to punishments such as career restrictions and penalties on asset holdings ([CNBC, 2017](#)). Some have even argued that payments data might someday play a role in China’s proposed new “Social Credit System” ([WSJ, 2019](#); [Wired, 2017](#)). Finally, concern that some countries may deny entry to individuals if they know they have purchased cannabis in Canada – even though cannabis is legal in Canada – led the [Office of the Privacy Commissioner of Canada \(2018\)](#) to advise people to use cash for such transactions.

This paper shows that the failure of individuals to preserve their privacy in payments by using privacy-enhancing techniques, such as cash, can lead to socially suboptimal outcomes. In the environments we consider, firms use information extracted during payments of one consumer without privacy-enhancing techniques to cluster potential future customers into groups with different reservation prices. Hence, failure on the part of one individual to preserve his or her information imposes a negative externality on others. In that sense, our approach to modelling privacy in payments is complementary to, but distinct from, the seminal work of [Kahn et al. \(2005\)](#) on the private cost of information disclosure. In the Kahn et al. model, individuals incur a private cost when information is disclosed during payment (modeled as the possibility of the purchase being stolen). Consumers are not given a choice to pay privately in the [Kahn et al.](#) framework; however, if they were, they would choose to use a privacy-enhancing technique if the private benefit exceeds the cost. In our set-up, information disclosure during the payment does not affect the terms and conditions of the current transaction, and the customer does not face any negative consequences from disclosing information during the payment. Our model emphasizes the possibility that individuals may suboptimally choose not to use privacy-enhancing techniques because the cost of foregoing privacy in payments is primarily borne by the rest of society.

Our paper connects to the so-called privacy paradox, which refers to the mismatch between stated preferences for privacy and actual behaviour to preserve privacy (Norberg et al., 2007; Athey et al., 2017). The literature offers several explanations such as unawareness of the extent of data collection and behavioural targeting and unfamiliarity with privacy-enhancing measures (Acquisti et al., 2016). Our paper adds an additional explanation by providing a mechanism that leads to differences between the private and social cost of information disclosure: Information disclosed by others allows inference about you. Hence, inference about you occurs regardless of your own choices to protect your information. Moreover, you do not bear the full or any cost of the information you inadvertently reveal about others. Both aspects induce a low willingness to take costly measures to protect privacy.

The notion that there may be a public good aspect to privacy has been noticed by legal scholars (MacCarthy, 2010; Fairfield and Engel, 2015). Choi et al. (2019) formulate an economic model where a ‘nuisance cost’ of foregoing privacy appears in the utility function. This cost is assumed to be an increasing function in the number of agents who reveal information, which imposes a negative externality that may lead to too little privacy in equilibrium. Similarly, Kahn et al. (2000) consider the case where a third party derives a fixed utility benefit from learning about a transaction. In both cases, the privacy externality itself is imposed as a model assumption rather than an economic outcome. Our model provides an explicit mechanism that generates a negative externality in the context of privacy in payments. Data revealed by one customer can be used to price discriminate among future customers. Hence, information disclosed during the payment is associated with a negative externality on future customers.

We seek to make these points in a model in which money is essential (Wallace, 2001). However, the analysis of privacy in payments in a model where money is essential is complicated by the fact that such a model requires a dynamic structure, so that intrinsically worthless money can have a positive price. The existence of monetary equilibria requires balanced modelling that ensures that agents form rational expectations regarding endoge-

nous transaction probabilities.⁵ Achieving this feature makes it harder to see the tension between private and social benefits and costs to maintaining privacy in payments. As such, in Section 2, we describe a one-period, three-cohort model that can be used to derive simplified versions of the main results. In this model, agents only live one period, and we give money value by including it in the utility function. The money in this model could be interpreted as a commodity money. The results for the one-period, three-cohort model are described in Section 3. Section 4 describes an infinite-horizon, overlapping-generations model, which builds off of the one-period model and justifies a number of assumptions made in the one-period model. Most notably, an assumption in the one-period model that ensures the entire market is served in the absence of price discrimination emerges endogenously in the full monetary equilibrium. In this alternative, dynamic model, money no longer appears directly in the utility function. Agents live multiple periods, and the incentive to acquire money when they are young is to be able to purchase items when they are middle-aged or old (Wallace, 1980; Balasko and Shell, 1981). Section 5 provides a discussion of policy issues and the potential role of government regulation. Section 6 discusses implications of some alternative modelling assumptions. Concluding remarks appear in Section 7.

2 One-Period, Three-Cohort Model

The economy consists of equal numbers of three cohorts of agents: young merchants (y), middle-aged consumers (m) and old consumers (o). Young merchants are endowed with two units of an indivisible good that they do not wish to consume themselves but would like to sell for money to the middle-aged and old consumers. Middle-aged and old consumers are each endowed with at least $r_H > 0$ units of money, where r_H denotes the highest possible reservation price consumers have for the good. The aggregate endowment of money is denoted by M .

⁵See Kovenock and De Vries (2002, pp. 147–149) for an overview of different modelling approaches.

First, each merchant is randomly matched with a middle-aged consumer and makes a take-it-or-leave-it offer that the consumer can either accept or refuse. Then, each merchant is randomly matched with an old consumer and makes a take-it-or-leave-it offer that the consumer can either accept or refuse.

Middle-aged and old consumers can either have a H [igh] or L [ow] value for the good. A consumer's type $s \in \{H, L\}$ determines her reservation price r_s for the indivisible good.⁶ The known quantities of types H and L are the same in each age cohort and are given by zn and $(1 - z)n$, respectively, where n is the population size of each of the three cohorts of agents and $0 < z < 1$. The utility of consumer i in cohort $a \in \{m, o\}$ depends on her consumption $c_{ia} \in \{0, 1\}$ of the indivisible good and her terminal money holdings $m_{ia} \geq 0$.⁷ A consumer of type $s_{ia} \in \{H, L\}$ weighs consumption of the indivisible good by the reservation price r_s , where $zr_H < r_L$. Consumer ia 's utility is augmented if she chooses to take an action $e_{ia} \in \{0, 1\}$ to protect her privacy during payments. The action involves a private benefit $\beta > 0$ and a cost of effort $\delta > 0$. The utility function of consumer i in cohort $a \in \{m, o\}$ of type $s_{ia} \in \{H, L\}$ is given by

$$u_a(c_{ia}, m_{ia}, e_{ia}; s) = c_{ia}r_s + m_{ia} + e_{ia}(\beta - \delta).$$

Merchants only care about the amount of money they end up with from selling goods. The utility function of young merchant i is given by $u_y(m_{iy}) = m_{iy}$. To maximize this amount, merchants seek to price discriminate. Ideally, a merchant would like to charge a price of r_H to consumers of type H and a price r_L to consumers of type L . However, consumer type is not directly observable. Merchants can, however, profile middle-aged and old customers based on an observable characteristic $h_{ia} \in \{X, Y\}$, for $a = m, o$, respectively, that is related to consumer type. Observable characteristics refer to consumer traits that

⁶In the dynamic model, reservation prices are determined by consumers' endogenously determined money holdings.

⁷We use a (short for age) to denote cohorts in the one-period model. In the dynamic model, cohorts will be identified by the generation t (for time) in which they were born.

are observable prior to the transaction. These can be things that are physically observed in the case of in-person transactions, either because they relate to the consumer's appearance or because the consumer uses coupons or vouchers that were sent to targeted areas. For an online store, examples of observable characteristics include online profiles or other customer information that can be revealed when customers access websites.

Consider the interaction between the merchant and the middle-aged consumer. The observable characteristic X is associated with all middle-aged consumers of a particular type, and the observable characteristic Y is associated with all middle-aged consumers of the other type. The merchant's problem is that, before making an offer to the middle-aged consumer, the merchant does not know which characteristic is associated with each type, and there is no credible way for consumers to signal their type. The merchant only knows that either of the two associations is equally likely.

The merchant cannot price discriminate against middle-aged consumers based on observable types. However, her interaction with the middle-aged consumers may help her to price discriminate against old consumers. This can happen in two ways. First, by charging the high reservation price to middle-aged consumers, she can learn which observable characteristic is possessed by old customers that are likely to accept a high offer. Second, the merchant can observe the consumer's type, regardless of the amount they pay, if consumers do not act to preserve their privacy in payments when they pay.⁸ Either of these methods allows the merchant to discover the relationship between the observed characteristic of the consumer and her type.

Knowledge of the correlation between the observable characteristic and type of the middle-age consumer allows the merchant to imperfectly price discriminate against old consumers. Price discrimination against old consumers is imperfect because we assume that only a fraction $1 - \varepsilon$, where $0 < \varepsilon < 1/2$, of old consumers has the same correlation between the

⁸Privacy in payments is not protected in reality when people use payment cards or mobile payment services that are linked to their broader payment history and other financial metrics, and this information is fed back to the retailer in some form.

observable characteristic and type as the middle-aged consumers. The remaining fraction ε has the opposite correlation.⁹ Profiling of old consumers becomes increasingly uninformative as ε approaches 1/2 and closer to perfect as ε approaches 0.

We can now specify social welfare in terms of the total utility of the three generations of agents. Utility is assumed to be linear in money for both merchants and consumers so that the “consumption” of money is irrelevant for social welfare: the money in the utility function of consumers balances out the presence of money in the utility function of the merchants. As a consequence, we can ignore the distribution of money holdings and write social welfare as

$$W = M + znr_H \Pr(c_{im} = 1|H) + (1 - z)nr_L \Pr(c_{im} = 1|L) + \\ + znr_H \Pr(c_{io} = 1|H) + (1 - z)nr_L \Pr(c_{io} = 1|L) + 2n\mathbb{E}e_{ia}(\beta - \delta), \quad (1)$$

where $\Pr(c_{im} = 1|H)$ and $\Pr(c_{im} = 1|L)$ denote the probabilities that a middle-aged consumer of, respectively, type H and L accepts the offer from the merchant and where $\Pr(c_{io} = 1|H)$ and $\Pr(c_{io} = 1|L)$ reflect these probabilities for the merchant’s encounter with the old consumer.

3 Results

Middle-aged consumers will not act to preserve their privacy if $\beta < \delta$ – that is, whenever the private benefit to preserving privacy is less than the private cost. However, this outcome may not be socially optimal. Actions taken by middle-aged consumers to preserve privacy benefit old consumers by reducing the ability of merchants to price discriminate against them. So the social benefit to the middle-aged consumers of acting to preserve privacy is

⁹The perfect correlation between the observable characteristic and type of the middle-aged consumers is not essential to our results. The same results hold true when we assume the same imperfect correlation for middle-aged and old consumers, although the expressions for θ^U and θ^W will be more complex.

greater than the private benefit. Because middle-aged consumers do not fully internalize the social benefit, they may make socially suboptimal decisions.

Suboptimal outcomes occur if the equilibrium that arises when middle-aged consumers take action to preserve their privacy has higher welfare than the equilibrium that occurs when they do not. Below we characterize equilibrium for the two relevant cases, $\beta > \delta$ and $\beta < \delta$. The difference in social welfare across these two cases can be used to define the social benefit of a subsidy designed to induce the socially optimal outcome.

3.1 Equilibrium with $\beta > \delta$

First, we consider the strategy where merchants quote the same price to all consumers in both encounters. It is optimal for merchants to quote customers a price just below or at their reservation price. So merchants will either quote a price r_H or r_L . When merchants quote the higher price r_H to all consumers, then only type H consumers will accept the offer, and the expected revenue of each merchant equals $2zr_H$. When quoting the lower price r_L to all consumers, then all the offers will be accepted. This will generate an expected revenue of $2r_L$ for each merchant, which is strictly higher since we assumed $zr_H < r_L$.

The expected utility of the young merchant when quoting the same price r_L to all consumers equals

$$\mathbb{E}u_y^* = 2r_L.$$

Moreover, in the absence of price discrimination, social welfare equals

$$W^* = M + 2znr_H + 2(1 - z)nr_L + 2en(\beta - \delta), \quad (2)$$

where $e = 1$ if $\beta > \delta$ and $e = 0$ otherwise. Note that this strategy maximizes social welfare, because it is equivalent to setting all probabilities in Eq. (1) equal to one.

Second, we consider the strategy where merchants try to differentiate their price quotes in the second encounter after learning the relationship between the observable characteristic and consumer type based on observable behaviour in the first encounter.

If the merchant quotes the lower price r_L in the first encounter, then the merchant cannot infer the relationship between the observable characteristics and consumer type from consumer behaviour, because both consumer types react in the same way. So, the merchant can only learn the relationship between the observable characteristics and consumer type when quoting a price larger than r_L in the first encounter. If the merchant quotes a price r_H in the first encounter, then only type H customers will take the offer. If the customer accepts the offer, then the merchant knows that consumers with the same observable characteristic are of type H . If the customer declines the offer, then the merchant knows that consumers with the same observable characteristic are type L customers.

In the second encounter, the merchant can use this information to quote the higher price r_H to consumers with the observable characteristic associated with type H customers, while quoting price r_L to consumers with the characteristic associated with type L . The only consumers who will decline the offer in the second encounter are the fraction ε of type L consumers who randomly adopted the characteristic associated with type H customers. Moreover, a fraction ε of type H consumers, who randomly adopted the characteristic associated with type L customers, will be paying the low price r_L . Hence, the merchant's expected utility from price discriminating (D) while consumers act to preserve universal privacy (U) in payments equals

$$\mathbb{E}u_y^{D,U} = z(2 - \varepsilon)r_H + [z\varepsilon + (1 - z)(1 - \varepsilon)]r_L.$$

Comparing $\mathbb{E}u_y^*$ and $\mathbb{E}u_y^{D,U}$ gives that, in a situation of privacy in payments, the merchant will price discriminate only if the profiling of consumers is sufficiently precise, i.e., only if

$\varepsilon < \theta^U$, where

$$\theta^U = \frac{zr_H - zr_L - (r_L - zr_H)}{r_L + zr_H}. \quad (3)$$

If merchants choose to price discriminate in a situation with privacy in payments, i.e., if $\varepsilon < \theta^U$, then social welfare will be lower and equals

$$W^{D,U} = W^* - n(1 + \varepsilon)(1 - z)r_L. \quad (4)$$

A social welfare loss occurs as a consequence of price discrimination with privacy in payments for two reasons. First, to learn about the link between consumer profiles and their behaviour, the merchant needs to quote prices that some customers will not accept in the first meeting. Second, the merchant may quote high prices that some customers will not accept in the second meeting due to imprecise profiling ($\varepsilon > 0$).

Finally, note that the threshold θ^U will be negative when the reservation price r_H is sufficiently low such that $r_H < r_L(1 + z)/2z$. If so, then price discrimination will never occur when transaction privacy is given, regardless of the precision of the profiling technique as indicated by ε . This implies that the minimum level of r_H at which price discrimination occurs is lower for higher values of z . In other words, price discrimination may occur with a lower dispersion in reservation prices if a larger fraction of consumers is willing to pay a high reservation price.

3.2 Equilibrium with $\beta < \delta$

Without universal transactions privacy, the merchant receives information on the consumer type in the first encounter once the payment for the consumption good is settled when the consumer makes no effort to protect her privacy in payments. Otherwise, the merchant does not receive this information. The merchant can use the information to infer the relationship between the observable characteristic and consumer type. This information may be used by the merchant to price discriminate in the second encounter.

Would the consumers in the first encounter make an effort to protect their privacy in payments? The information is only disclosed to the merchant once the payment for the consumption good is settled. Hence, the disclosure of information to the merchant cannot adversely affect the outcome for the consumer who makes the payment, although subsequent consumers could be worse off. Moreover, there is a net utility cost of maintaining privacy in payments when $\beta < \delta$. In other words, individual consumers in the first encounter have no incentives to maintain privacy in payments whenever $\beta < \delta$, because doing so strictly reduces their utility.


Merchants can follow a different strategy that leads to price discrimination when consumers do not maintain privacy in payments. First, they can quote the low price r_L to all customers in the first encounter. They collect information on the consumer type during this first transaction since there is no privacy in payments. In the second encounter, the merchant uses this information to quote the higher price r_H to consumers with the observable characteristic associated with type H customers, while quoting price r_L to consumers with the characteristic associated with type L . The merchant's expected utility from following this strategy in an environment without universal privacy (W) is

$$\mathbb{E}u_y^{D,W} = r_L + z(1 - \varepsilon)r_H + [z\varepsilon + (1 - z)(1 - \varepsilon)]r_L.$$

The merchant's expected utility from relying on this method to price discriminate, $\mathbb{E}u_y^{D,W}$, strictly exceeds $\mathbb{E}u_y^{D,U}$, since $r_L > zr_H$. Moreover, comparing $\mathbb{E}u_y^{D,W}$ to $\mathbb{E}u_y^*$ gives that, without privacy in payments, the merchant will price discriminate only if the profiling technique is sufficiently precise such that $\varepsilon < \theta^W$, where

$$\theta^W = \frac{zr_H - zr_L}{r_L + zr_H}. \quad (5)$$

Figure 1: Probability of Profiling Errors ε and Price Discrimination



If the merchant chooses to price discriminate without privacy in payments, i.e., if the profiling of consumers is sufficiently precise such that $\varepsilon < \theta^W$, then social welfare equals

$$W^{D,W} = W^* - n\varepsilon(1-z)r_L. \quad (6)$$

Price discrimination in the absence of universal transaction privacy leads to a welfare loss whenever the profiling of consumers is imperfect ($\varepsilon > 0$).

3.3 Price discrimination and profiling errors

Whether price discrimination occurs in equilibrium depends on the accuracy of consumer profiling and whether consumers choose to protect their privacy in payments. Figure 1 summarizes the relationship between price discrimination and the probability of profiling errors ε .

Starting at the right of Figure 1 is the nonempty region $\theta^W \leq \varepsilon < 1/2$, where the social optimum is attained regardless of whether consumers protect their privacy in payments.¹⁰ In this region, price discrimination does not occur in equilibrium because the profiling of consumers is too imprecise. Price discrimination is not profitable because the expected

¹⁰The region is nonempty because $r_L > zr_H$ implies $\theta^W < 1/2$.

loss in sales from erroneously quoting high prices to consumers with low reservation prices outweighs the potential benefit of earning more on some transactions.

In the middle region, where $\theta^U \leq \varepsilon < \theta^W$, price discrimination only occurs when consumers do not protect their privacy in payments. The equilibrium achieves the social optimum with privacy in payments, but it leads to a suboptimal outcome in the absence of privacy in payments because not all potential win-win situations will lead to transactions between merchants and consumers.

The reason the social optimum may not be achieved in the middle region is that consumers do not internalize the benefit to society of protecting their personal privacy during the transaction. It would be socially optimal if all consumers would maintain their privacy in payments in the first encounter when the net private cost is smaller than the social benefit of each consumer doing so, that is, when

$$\delta - \beta < \varepsilon(1 - z)r_L.$$

Nevertheless, the consumers choose not to do so during the first encounter because the private net benefit is not positive. Society would benefit from the universal provision of privacy in payments in this region.

Finally, there *may* be a third region to the left where $0 \leq \varepsilon < \theta^U$. This region only exists when the reservation price r_H is sufficiently high such that $r_L(1/2 + 1/(2z)) < r_H < r_L(1/z)$. In this region, the social optimum will not be attained regardless of whether there is universal transaction privacy.

4 Overlapping Generations Model

Time is discrete and continues without end. A new generation of n identical three-period lived agents starts at each date t . “Generation t ” is young at t , middle-aged at $t + 1$, and old at $t + 2$ (the three cohorts from the one-period model). Agents maximize their life-time

utility, which depends negatively on the amount of labour to produce consumable goods when they are young and positively on their consumption pattern when they are middle-aged and old. Moreover, as before, agents face a the net private cost (or benefit) of preserving privacy in payments. Agent i from generation t can produce Q_{it} perishable indivisible consumption goods at a utility cost $Q_{it}f$, but derives no utility from consuming the good.¹¹ Middle-aged and old agents cannot produce goods but derive utility from consuming them. Each agent has three opportunities to consume a good: one opportunity when middle-aged (referred to as meeting A), and an early and late opportunity to consume when old (referred to as meetings B and C , respectively). Their utility function reads

$$u(c_{it}^A, c_{it}^B, c_{it}^C, Q_{it}, e_{it}) = c_{it}^A + c_{it}^B b + c_{it}^C c - Q_{it}f - e_{it}(\beta - \delta). \quad (7)$$

Preferences of the agents are such that they enjoy consumption more at certain points in their lives. This is indicated by the parameters in the utility function: $\frac{1}{3} < c < 1 < b$. The utility cost of production is assumed to be low but positive, and satisfies $0 < f < \min\{\frac{1}{3}, (b - c)/2\}$. The assumption $f < \frac{1}{3}$ implies that the benefit of consumption in meeting C exceeds the utility cost of producing the consumption good. The assumption $f < (b - c)/2$ places a restriction on the minimum level of impatience agents have when old.

There is no technology for record-keeping. However, there is a durable, indivisible, non-consumable good and a transaction technology that allows exchanging the durable asset for the consumable good whenever an offer is made and accepted. We refer to the durable asset as “money.” The total money stock is fixed at $4n$. At all times, individual agents cannot carry more than three units of money, and the amount agents carry is not observable.¹²

¹¹Throughout this section we use the subscript position to identify a particular individual (i) born at a particular date (t).

¹²Lagos and Wright (2005, p. 464) set out the complexities of solving monetary equilibria without constraints on money holdings. Their solution is to make assumptions that lead to a degenerate distribution of money holdings, so all agents hold the same amount of money. An alternative is to resort to finding numerical solutions for the distribution of money holdings (e.g., Molico, 2006). We opt for an approach where agents can hold up to three units of indivisible money, which allows us to derive analytical solutions in an environment where agents can hold different amounts of money in equilibrium.

Figure 2: Timeline from the Perspective of Young Agent i Born in Period t




Figure 2 shows a timeline of the only interactions between agents in the economy. Meeting A involves a random matching between a young agent and a middle-aged agent. Meetings B and C involve random matchings between young agents and old agents (each old agent has both an early and a late encounter with two different young agents). During each of these meetings, the young agent may make a take-it-or-leave-it offer regarding the number of units of money at which the young agent is willing to produce a unit of the consumable good for the older agent.

The price offers that are *received* by agent it during the meetings are denoted by p_{it}^A , p_{it}^B and p_{it}^C . Hence, the offer made by young agent it to middle-aged agent $j, t-1$ is denoted as $p_{j,t-1}^A$, etc. Similarly, the amount of money that agent it carries into the meetings when middle-aged and old is denoted as m_{it}^A , m_{it}^B and m_{it}^C , respectively. If an agent accepts an offer, then the young agent will produce the good, a transaction occurs and the older agent consumes the good.

After the meetings, consumable goods and old agents perish. Young agents may inherit one or more units of money from old agents. Money left by old agents is allocated in rounds.

Young agents who own less than three units of money may receive a unit of money in each round. Whenever there are insufficient funds to allocate a unit of money to every eligible young agent, the agents who do not inherit a unit of money are randomly selected. Another allocation round occurs whenever there are funds left. The inheritance that each young agent it receives is denoted as m_{it}^I .

The amount of money each agent of generation t accumulates when young determines their lifetime status $s_{it} \in \{H, L\}$. All agents who accumulate three units of money attain the lifetime status $s_{it} = H$, and agents who accumulate less than three units of money have a lifetime status $s_{it} = L$. An agent's lifetime status is not directly observable to other agents. Instead, at the start of each period t , generations born in periods $\tau \in \{t-1, t-2\}$ take on an observable characteristic $h_{i\tau}^t \in \{X, Y\}$, which is assigned to all middle-aged and old agents depending on their lifetime status. The correlation between $s_{i\tau}$ and $h_{i\tau}^t$ is perfect for middle-aged agents ($\tau = t-1$) but imperfect for old agents ($\tau = t-2$) due to ε . The young generation does not know the sign of the correlation, which may be different in every period t .

From the model it follows that a few equalities must hold true in the aggregate. First, the total amount of money that agent it earns and inherits should equal the amount that the agent brings to meeting A when middle-aged. As a direct consequence, it must hold true for any young agent it who meets a random middle-aged agent $j, t-1$ and two different random old agents indexed $k, t-2$ and $l, t-2$, that

$$\Pr [c_{j,t-1}^A p_{j,t-1}^A + c_{k,t-2}^B p_{k,t-2}^B + c_{l,t-2}^C p_{l,t-2}^C + m_{it}^I = w] = \Pr [m_{it}^A = w] =: g_t(w) \quad (8)$$

for $w = 0, 1, 2, 3$. Function $g_t(w)$ is the density of money holdings for young agents (the wealth distribution). Moreover, since the aggregate money holdings must equal the total quantity of money, and given that young agents are born without any money, it holds true

at the start of every period t that

$$\sum_i m_{i,t-1}^A + \sum_i m_{i,t-2}^B = 4n. \quad (9)$$

Agents need to decide on their strategies regarding which offers to make during the meetings when young. These strategies may depend on whether offers in previous meetings were accepted as well as on the observable characteristics of the agents they encounter, and which offers they expect and accept when middle-aged and old.

Definition 1 *A steady-state Nash equilibrium is defined as the state of the economy where none of the agents can increase their expected utility in Eq. (7) by deviating from their current strategy given the strategies of all other agents and rational expectations regarding the steady-state wealth distribution in Eq. (8), and where the restriction on aggregate money holdings in Eq. (9) holds true. A monetary equilibrium is defined as a steady-state Nash equilibrium in which meetings occur where the consumable good is exchanged for the durable asset.*

4.1 Non-monetary equilibrium

There exists a steady-state Nash equilibrium where no exchanges of consumable goods for money occur. Young agents do not make any offer to exchange consumable goods for money given that they rationally don't expect any offers when being middle-aged and old. Middle-aged and old agents have no opportunities to spend their money holdings and keep their money until they die. Young agents inherit those money holdings. The steady-state distribution of money holdings depends on the initial wealth distribution.

4.2 Optimal consumption decisions

In a monetary equilibrium, agents need to decide which offers to accept and which to decline when middle-aged and old. The decision for old agents is simple. Since old agents

have no further use for money after their last meeting, they will accept any offer that they can afford in meeting C . So, agent it accepts any offer where $p_{it}^C \leq m_{it}^C$ when old. Moreover, since agents prefer early consumption when old ($c < b$), they will also accept any offer that they can afford in meeting B . Hence, agent it accepts any offer where $p_{it}^B \leq m_{it}^B$ when old. Therefore, the reservation price of an old agent depends on the agent's money holdings, where larger money holdings are associated with a higher reservation price.

The decision to consume is more complicated for middle-aged agents. They prefer early consumption when old above consumption when middle-aged ($b > 1$), but accepting an offer when middle-aged may affect the ability to pay for consumption when old. The expected utility from consumption when old depends on which offers an old agent with a given amount of money and given status can afford (early, late, or both). For an old agent it with money m_{it}^B and status s_{it} , this can be written as

$$\begin{aligned} V^o(m_{it}^B; s_{it}) = & c \Pr [p_{it}^C \leq m_{it}^B < p_{it}^B | s_{it}] + b \Pr [p_{it}^B \leq m_{it}^B < p_{it}^B + p_{it}^C | s_{it}] \\ & + (b + c) \Pr [p_{it}^B + p_{it}^C \leq m_{it}^B | s_{it}]. \end{aligned} \quad (10)$$

Accepting offers when middle-aged may reduce the probabilities in Eq. (10) by lowering money holdings m_{it}^B . Agent it will accept an offer of p_{it}^A when middle-aged if and only if the utility of consuming when middle-aged outweighs the expected loss in consumption when old, i.e., if and only if

$$1 \geq V^o(m_{it}^A; s_{it}) - V^o(m_{it}^A - p_{it}^A; s_{it}). \quad (11)$$

Whether this holds true depends on the agent's money holdings m_{it}^A , the offer under consideration p_{it}^A and the anticipated distribution of future offers p_{it}^B and p_{it}^C given the agent's status s_{it} .

4.3 Best social outcome without free goods

The ability to benefit from offers received when middle-aged and old increases in the amount of money collected when young. Young agents face a positive cost of producing the consumable good. Since there is no record of whether an agent has lead a productive life while young, except for the agent's money holdings, young agents will only propose offers where the potential benefits from additional monetary earnings exceed the cost of producing the consumable good. As a consequence, free goods cannot exist in a monetary equilibrium. Or conversely, any transfer of the consumable good to the middle-aged and old consumers will require at least one unit of money in a monetary equilibrium.

One can ask what the best feasible economic outcome could look like given the restriction that each transfer of the consumable good involves at least one unit of money, while further ignoring the choices made by optimizing agents. Obviously, more goods can be consumed when the price offers received by agents are as low as possible, i.e., if $p_{it}^A = p_{it}^B = p_{it}^C = 1$ for all it . Moreover, more goods can be consumed when they are consumed early in life: When all agents earn one unit of money when young and directly spend that unit when middle-aged, then this requires n units of money in total. By contrast, when all agents earn one unit of money when young, save that unit when middle-aged, and spend that unit on a consumable good when old, then this requires $2n$ units of money in total. Finally, exchanges in meeting C are the least attractive, because they require a relatively large amount of money and yield the lowest utility.

Altogether, it requires $3n$ units of money in total to allow all young agents to earn two units of money and to spend that money in meetings A and B . Given that there are $4n$ units of money, there is still room to facilitate exchanges in meeting C for half of the agents. The following lemma describes this economic situation:

Lemma 1 *The best achievable social outcome in the absence of free goods is that where all young agents offer the consumable good at a price of 1 in each meeting. Half of the agents earn 2 units of money when young and accept offers in meetings A and B when middle-*

aged and old. Half of the agents earn 3 units of money when young and accept offers in all meetings. The steady-state wealth distribution is $g(2) = g(3) = 1/2$. Expected utility of new agents is

$$\mathbb{E}u^* = 1 + b - 2f + \frac{1}{2}(c - f).$$

Social welfare, in terms of the aggregate per-period utility, equals $W^* = n\mathbb{E}u^*$.

4.4 Monetary equilibrium

In what follows, we show that the best achievable social outcome in the absence of free goods in Lemma 1 is a monetary equilibrium conditional on ε being sufficiently large. As in the single-period model, price discrimination may become optimal for a low probability of profiling errors ε . Moreover, the threshold for ε at which price discrimination becomes optimal for young agents will depend on whether agents take action to preserve their privacy in payments in a similar manner. That is, the decision will depend on the private rather than the social costs and benefits of privacy in payments.

4.4.1 Offers without considering price discrimination

We start by showing that it is optimal for a young agent it , who meets a random middle-aged agent $j, t-1$ and two random old agents indexed $k, t-2$ and $l, t-2$, to make offers $p_{j,t-1}^A = p_{k,t-2}^B = p_{l,t-2}^C = 1$ when young agents do not consider price discrimination based on observable consumer characteristics. For convenience, we will drop the subscript of the prices when discussing the optimal offers made by the young agent it (i.e., p^A refers to $p_{j,t-1}^A$, etc.) since the index of the matched agent is irrelevant to the pricing decision.

Start with the optimal choice for the last offer (meeting C). Could it be optimal to offer a price $p_{l,t-2}^C > 1$? Given the strategies of the other agents, half of the old agents will carry no money and half of the old agents will carry one unit of money at the start of meeting C . So, none of the old agents can accept an offer $p_{l,t-2}^C > 1$. Would it be optimal for the

young agents to make offers when their money holdings are smaller than three? Producing the consumable good costs f . Given the strategies of other agents, one unit of money will facilitate buying a consumable good when middle-aged or old, which yields at least utility $c > f$ (there is no chance of receiving an inheritance because all agents spend all their money during their lifetime). Hence, it is optimal to offer a price $p_{l,t-2}^C = 1$ in meeting C whenever cumulative earnings from meetings A and B are less than three.

Now consider the optimal choice for the second offer (in meeting B). Could it be optimal to offer a price $p_{k,t-2}^B > 1$? Again, since all agents have spent one unit of money in meeting A , no agent is expected to have sufficient funds to accept an offer $p_{k,t-2}^B = 3$. Moreover, agents who started with three units of money would accept an offer of $p_{k,t-2}^B = 2$ because $b > c$.

Suppose that the young agent earned one unit of money in meeting A . If the offer $p_{k,t-2}^B = 2$ is accepted, which occurs with 50 percent probability, then the agent ends with three units of money by producing only two goods. If the offer is not accepted, then it would be optimal to offer a price $p_{l,t-2}^C = 1$ in the next meeting, which would lead to a 25 percent probability of earning two units of money when young (facilitating consumption in meetings A and B) and a 25 percent probability of ending up with one unit of money (facilitating consumption in meeting B only). Altogether, the expected utility derived from this strategy is

$$\frac{1}{2} [b + 1 + c - 2f] + \frac{1}{2} \left[\frac{1}{2}(b + 1 - 2f) + \frac{1}{2}(b - f) \right] = \mathbb{E}u^* - \frac{1}{4} + \frac{3}{4}f, \quad (12)$$

which is less than $\mathbb{E}u^*$ since $f < 1/3$.

Moreover, suppose that the agent earned no unit of money in meeting A . In this situation, there is a 50 percent probability that the offer $p_{k,t-2}^B = 2$ will be accepted and a 50 percent probability that the offer $p_{l,t-2}^C = 1$ will be accepted. The expected utility of this strategy is

$$\frac{1}{2} \left[\frac{1}{2}(b + 1 + c - 2f) + \frac{1}{2}(b + 1 - f) \right] + \frac{1}{2} \left[\frac{1}{2}(b - f) \right] = \mathbb{E}u^* - \frac{1}{4}(b + c) - \frac{1}{2} + \frac{3}{2}f,$$

which is also less than $\mathbb{E}u^*$ since $f < 1/3$. Since $f < \frac{1}{2}(b - c)$, this is also strictly less than the expected utility of earning no money in meeting A and subsequently offering $p_{k,t-2}^B = 1$ and $p_{l,t-2}^C = 1$, which gives $(b - f) + \frac{1}{2}(1 - f) < \mathbb{E}u^*$. In summary, regardless of how much the agent earned in meeting A , it is optimal to offer $p_{k,t-2}^B = 1$ in the second meeting.

Finally, we consider the optimal choice for the first offer (in meeting A). Could it be optimal to offer a price $p_{j,t-1}^A > 1$? No middle-aged agent would accept an offer $p_{j,t-1}^A = 3$ because agents prefer to consume when old in meeting B . A middle-aged agent with two units of money would not accept an offer $p_{j,t-1}^A = 2$ for the same reason, but an agent with three units of money would accept an offer $p_{j,t-1}^A = 2$ because they prefer to consume in meetings A and B rather than in meetings B and C since $1 + b > b + c$. If the offer $p_{j,t-1}^A = 2$ is not accepted, then it would be optimal for the young agent to offer prices $p_{k,t-2}^B = 1$ and $p_{l,t-2}^C = 1$, which leads to a 50 percent probability of earning one unit of money and a 50 percent probability of earning two units of money as discussed above. If the offer $p_{j,t-1}^A = 2$ is accepted, then it is optimal to offer $p_{k,t-2}^B = 1$, which ensures earning three units of money while producing only two goods. Altogether, the expected utility of offering $p_{j,t-1}^A = 2$ is the same as that in expression (12), which is less than $\mathbb{E}u^*$. Hence, an offer of $p_{j,t-1}^A = 2$ is not optimal. Could it be optimal to make no offer in meeting A ? As we have seen, it is then subsequently optimal to have $p_{k,t-2}^B = 1$ and $p_{l,t-2}^C = 1$, which yields a lower utility than $\mathbb{E}u^*$. In summary, it is also optimal to offer $p_{j,t-1}^A = 1$ in the first meeting.

This concludes the argument that it is optimal for young agents to set $p_{j,t-1}^A = p_{k,t-2}^B = p_{l,t-2}^C = 1$ if young agents do not consider price discrimination based on the observable characteristic h_{it}^t .

4.4.2 Price discrimination with privacy in payments

All agents act to preserve their privacy when $\beta > \delta$. The status of an agent it , s_{it} , is therefore not revealed when making a payment. However, it is still possible for young agents to learn the status of customers based on their purchasing behaviour.

Young agent it observes the characteristic of the middle-aged agent, $h_{j,t-1}^t$, in meeting A , but they do not know how it is correlated with that agent's status, $s_{j,t-1}$, at date t . As derived above, the middle-aged agent will accept an offer $p_{j,t-1}^A = 2$ when $m_{j,t-1}^A = 3$, and, hence, $s_{j,t-1} = H$, but not when $m_{j,t-1}^A < 3$, and, hence, $s_{j,t-1} = L$. So, young agents could try to learn the relationship between the observable characteristic $h_{j,t-1}^t$ and status $s_{j,t-1}$ by making an offer $p_{j,t-1}^A = 2$. Is it optimal for young agents to do so given that they may benefit from price discrimination in subsequent encounters?

We established that none of the agents has funds to accept an offer $p_{k,t-2}^C > 1$. So, a young agent may only benefit from price discrimination in meeting B in cases where those with initial balances $m_{k,t-2}^A = 3$ (i.e., $s_{k,t-2} = H$) still have two units of money and those with initial balances $m_{k,t-2}^A = 2$ (i.e., $s_{k,t-2} = L$) still have one unit of money. Young agents could aim at offering a high price $p_{k,t-2}^B = 2$ to old agents profiled as bearing status $s_{k,t-2} = H$.

If the young agent decides to offer $p_{j,t-1}^A = 2$, two potential outcomes may occur, each with probability $1/2$. First, the offer of $p_{j,t-1}^A = 2$ may be accepted when the young agent meets a middle-aged agent with status $s_{j,t-1} = H$ in meeting A . In this situation, it is optimal for the young agent to make an offer of $p_{k,t-2}^B = 1$ in meeting B – which will certainly be accepted – and no offer in the third encounter.

Second, the offer of $p_{j,t-1}^A = 2$ may be rejected because the young agent meets a middle-aged agent with status $s_{j,t-1} = L$. The young agent can then infer which observable characteristic $h_{k,t-2}^t$ tends to be associated with status $s_{k,t-2} = L$. If the old agent encountered in meeting B bears that same observable characteristic, then the young agent makes an offer $p_{k,t-2}^B = 1$. This offer is accepted not only by old agents with status $s_{k,t-2} = L$, but also by the fraction ε of old agents with status $s_{k,t-2} = H$ who are incorrectly profiled as poor customers based on the observable characteristic. If the old agent encountered in meeting B bears the other observable characteristic, then the young agent infers that the old agent is likely to have status $s_{k,t-2} = H$ and makes an offer $p_{k,t-2}^B = 2$. This offer will be accepted only if the old agent's status is indeed $s_{k,t-2} = H$, which, given the wealth distribution,

occurs with probability $1 - \varepsilon$. Regardless of the outcome in meeting B it will be optimal for the young agent to make an offer $p_{l,t-2}^C = 1$ in meeting C .

Altogether, the expected utility of deviating to the price discrimination strategy with privacy of payments is

$$\mathbb{E}u^{D,U} = \frac{1}{2} [1 + b + c - 2f] + \frac{1}{4} [b - f + \frac{1}{2}(1 - f)] + \frac{1-\varepsilon}{4} [1 + b - f + \frac{1}{2}(c - f)] + \frac{\varepsilon}{4} [\frac{1}{2}(b - f)].$$

Comparing the levels of $\mathbb{E}u^{D,U}$ and $\mathbb{E}u^*$ gives that price discriminating in a situation with privacy of payments is only optimal whenever the profiling error ε is sufficiently small; i.e., $\varepsilon < \theta^U$, where

$$\theta^U = \frac{6f - (1 - c)}{2 - 2f + b + c}.$$

This result is summarized in the following proposition:

Proposition 1 *The economic situation described in Lemma 1 is a monetary equilibrium whenever it is privately optimal for agents to preserve privacy in payments (i.e., $\beta > \delta$) and profiling errors are sufficiently large (i.e., $\varepsilon \geq \theta^U$).*

4.4.3 Price discrimination without privacy in payments

When $\beta < \delta$, none of the agents acts to preserve their privacy. As a consequence, it is possible for young agents to learn the status of the agents in meeting A whenever a payment occurs.

This allows young agents to deviate to a new strategy where they offer $p_{j,t-1}^A = 1$, learn the period t relationship between $s_{i\tau}$ and $h_{i\tau}^t$ from the payment of the middle-aged agent in meeting A , and then use this information to price discriminate against old agents in meeting B at the risk of making a profiling error because of the imperfect correlation due to ε . As before, practicing price discrimination can only result in additional money for a young agent in meeting B because no agent will accept an offer $p_{l,t-2}^C > 1$.

The offer of $p_{k,t-2}^B = 1$ to agents with the observable characteristic associated with $s_{k,t-2} = L$ will be accepted by old agents who have status L and the fraction ε of old agents who have status H and are incorrectly profiled as having status $s_{k,t-2} = L$. The offer of $p_{k,t-2}^B = 2$ to an agent who is profiled as having status $s_{k,t-2} = H$ is only accepted by old agents who truly have status $s_{k,t-2} = H$. The young agent benefits from collecting a total of three units of money with effort $2f$ if an offer of $p_{k,t-2}^B = 2$ is accepted. If the offer is declined, then there might still be a possibility to sell a second good to an old agent in the third encounter. Hence, the expected utility of this strategy is

$$\mathbb{E}u^{D,W} = \left[\frac{1}{2}(1 - \varepsilon)\right] (b + 1 + c - 2f) + \left[\frac{1}{2}\right] (b + 1 - 2f + \frac{1}{2}(c - f)) + \left[\frac{1}{2}\varepsilon\right] (b - f + \frac{1}{2}(1 - f)).$$

Comparing the levels of $\mathbb{E}u^{D,W}$ and $\mathbb{E}u^*$ reveals that price discriminating without privacy of payments is only optimal when the profiling error is sufficiently small; i.e., $\varepsilon < \theta^W$, where

$$\theta^W = \frac{c + f}{2c + 1 - f}.$$

This result is summarized in our second proposition:

Proposition 2 *The economic situation described in Lemma 1 is a monetary equilibrium whenever it is privately optimal for agents not to act to preserve privacy in payments (i.e., $\beta < \delta$) and profiling errors are sufficiently large (i.e., $\varepsilon \geq \theta^W$).*

4.5 Welfare and price discrimination

Lemma 1 shows that any monetary equilibrium that does not involve young agents offering $p_{j,t-1}^A = p_{k,t-2}^B = p_{l,t-2}^C = 1$ must be inferior from a welfare perspective. Moreover, any form of price discrimination must involve potential transactions at a price of more than one unit of money. Hence, any equilibrium that involves price discrimination is inferior from a welfare perspective relative to the economic situation described in Lemma 1. Whether a

socially inferior outcome occurs depends on whether agents act to preserve privacy in payments and on whether the probability of profiling errors ε is sufficiently small. In particular, price discrimination will occur when consumer profiling is less precise if there is no privacy in payments. This result is expressed in our third proposition:

Proposition 3 *The range of ε for which the best achievable social outcome corresponds to a monetary equilibrium is smaller when agents do not act to preserve their privacy.*

Proof. Proposition 1 shows that the economic outcome described in Lemma 1 is a monetary equilibrium when $\varepsilon \geq \theta^U$ and agents act to preserve their privacy in payments. Proposition 1 shows that the economic outcome described by Lemma 1 is a monetary equilibrium when $\varepsilon \geq \theta^W$ and agents do not act to preserve their privacy in payments. Appendix B shows that $\theta^U < \theta^W$. Hence, the best achievable social outcome in the absence of free goods is a monetary equilibrium for a smaller range of ε when agents do not act to preserve their privacy. ■


5 Policy

Figure 3 summarizes the results regarding the feasibility of the economic outcome described in Lemma 1 for two key parameters in the model: the probability of profiling errors based on observable data, ε , on the horizontal axis and the net private benefit of privacy in payments, $\beta - \delta$, on the vertical axis.¹³ Moving along the line from past to present illustrates two trends over the past decades that are likely to continue going forward.

The first trend is the increasing cost of maintaining privacy in payments that has resulted from the expansion of ecommerce (in general, one cannot order goods online using cash), improvements in digital payment methods, such as contactless payments, which reduce processing time at point of sale, and reduced acceptance of cash. These factors have led to a

¹³As $\beta - \delta$ becomes an increasingly small negative number there eventually reaches a point where it is no longer socially optimal to protect privacy in payments because the benefits in terms of avoiding price discrimination are capped by the consumption loss. We do not show this region in Figure 3.

Figure 3: Privacy in Payments and Big Data: Past, Present and Future?



reduction in the private net benefit of privacy in payments from pre-e-commerce revolution levels to the current day.

The second trend is the enhanced ability to profile consumers. This enhanced ability results from the progress in prediction techniques and computational power, the increase in the scale and scope of data gathering and the enhanced ability to manage and combine big datasets. These developments reduce the amount of errors in attempts to profile consumers, leading to reductions in ε and a movement to the left of the past situation.

Continuation of these trends may lead the economy into a situation where the optimum level of consumption may no longer be feasible, as shown in Figure 3. Such an outcome would involve price discrimination where firms use payments data to predict consumers' willingness to pay from observable consumer data, such as address information and online profiles.

There are several policies that could promote privacy in payments among consumers and hence avoid undesirable outcomes. One option that could, at least temporally, reduce

the cost of privacy in payments is the promotion or subsidization of cash use. Naturally, such policies would require a careful balancing of all social costs and benefits of cash usage (Rogoff, 2016; McAndrews, 2017). However, such a policy leaves the net private benefit or cost of maintaining privacy unchanged for the increasing share of online purchases.¹⁴

Another option is the promotion of *electronic cash*. Transactions using electronic cash would differ from payments using methods facilitated by today’s mainstream commercial payments providers in the sense that it would be impossible for any commercial parties to track payments of individuals. Electronic cash could rely on privacy-enhancing techniques where information such as the identity and characteristics of the payer remain shielded from commercial parties such as the payee, the providers of the payments infrastructure and financial institutions. Arguably, electronic cash could be issued by a commercial party (e.g., the Digicash initiative of Chaum, 1985; 1992), in a decentralized collective (e.g., cryptocurrencies such as Bitcoin of Nakamoto, 2008) or by a central bank. Each of these variations has its own costs and benefits in terms of providing privacy in payments. Of course, a prerequisite for electronic cash to have an impact is that it has the ability to preserve privacy in payments while also allowing consumers to trust the technique (Kahn, 2018). Moreover, for electronic cash to reduce the costs of privacy in payments, δ , it needs to be widely adopted and accepted.

In discussions on electronic cash, it is often conjectured that the government will need some possibility to access payments data in order to protect society and enforce criminal law (e.g., after the court issues a warrant). Others focus predominantly on the risks of foregoing privacy in payments with respect to the government, as has been raised in relation to China’s Social Credit System. It is possible that both these concerns can be addressed without government or the private sector having full information about payments. If commercial parties were to provide the front-end to government-issued electronic cash, then the processing of

¹⁴Sometimes, one has to forego anonymity in online transactions. For example, providing an address may be necessary in order to have goods delivered. However, foregoing anonymity is different from foregoing privacy in payments: foregoing anonymity means that the counterparty can identify you but does not allow insight into information related to a person’s past payment history with third parties.

payments and storage of payments data could be designed such that the reconstruction of someone’s payments history would require data disclosure by both the government entity and the commercial entity.

A third policy option is to counter the declining trend in ε . A government could limit the availability of data on observable characteristics by carefully weighing privacy concerns when releasing populational statistics ([Abowd and Schmutte, 2019](#)) and by imposing privacy regulations (e.g., the General Data Protection Directive in the European Union). However, the path to effective privacy regulation is filled with obstacles, and it is not clear how long the journey will take. For example, since privacy is a public good, one cannot simply assume that a consent-based approach will lead to socially optimal choices of privacy. Even with good privacy regulations there can be security breaches that undermine their effectiveness. Ultimately, slowing the trend in ε may prove difficult, because success requires countering advances made in the areas of computing and prediction techniques that are constantly evolving.

The mechanism in the present paper illustrates a positive externality of privacy in payments. More data-sharing may in certain situations also lead to more efficient outcomes. One could think, for example, about situations where information allows producers to help consumers make better choices ([Ichihashi, 2019](#)) or where data is used to construct credit ratings that could potentially improve the allocation of credit ([Pagano and Jappelli, 1993](#)).

6 Discussion

6.1 Bargaining power for consumers

Market power is essential for price discrimination. Privacy in payments, and privacy more generally, prevents young merchants in the model from exploiting situations in which they have market power. Since the young agents have all the bargaining power in the model,

one may consider the economic framework in the present paper as a relatively extreme case where all retailers in the economy have monopoly pricing power.

The main results are unchanged in a more balanced specification of the model where bargaining power is distributed more evenly between merchants and customers. In particular, we explored a setting where a coin toss decides whether a new agent i has all bargaining power when young, or whether it will be the middle-aged and old customers who will make a take-it-or-leave-it offer to the young agent i . This alternative specification does not change the equilibrium conditions for the economic situation in described Lemma 1. In particular $p_{j,t-1}^A = p_{k,t-2}^B = p_{l,t-2}^C = 1$ will still be an equilibrium whenever ε is sufficiently large with exactly the same thresholds θ^U and θ^W . The bargaining power does not allow middle-aged and old consumers to improve upon the unit price, so they will offer to pay a unit price for the consumable good in each meeting. The sole difference is that only the proportion of young merchants who received bargaining power based on the coin toss will be inclined to deviate from the equilibrium strategy by price discriminating if the probability of profiling errors is sufficiently small.

6.2 Increasing the quantity of money

One may be tempted to believe that the reason exchanges do not occur in every meeting in the monetary equilibrium described in Lemma 1 is a lack of money. This is not the case. Rather, some scarcity of money is needed to achieve the consumption gains that are realized in the monetary equilibrium without price discrimination. Rational behaviour requires the young merchant to be sufficiently uncertain about the money holdings of old agents in meeting B in order for any consumption to occur as a consequence of encounters in meeting C . If all middle-aged agents were able purchase a good for one unit of money and were to enter meeting B with two units of money when old (this would be a feasible outcome with an aggregate quantity of money of $5n$), then each of the young merchants could simply make a take-it-or-leave-it offer of two in meeting B . All old agents would optimally choose to accept

this offer. As a consequence, old agents would run out of money after accepting the offer in meeting B , and no consumption would occur in meeting C . Hence, increasing the quantity of money does not necessarily lead to an outcome with higher levels of consumption.

6.3 Credit arrangements

Credit arrangements cannot improve the allocation in the model. Middle-aged and old agents have nothing to offer to young agents besides money, since agents in the model can produce and sell goods only when young. As a consequence, there is no way that they could return a favour that was granted on credit unless they already had money to pay for the favour in the first place.

7 Concluding Remarks

Privacy is a public good in our model because it protects consumers from price discrimination. Consumers do not have to contribute to *provide* this public good. It is not like national defence or a public park. Rather, they have to take costly actions to *preserve* it. In that sense privacy is like clean air. This analogy is not perfect in our model because the actions that agents in our model take to preserve the protection against price discrimination only help others. That is, if an agent in our model fails to protect her privacy, she herself is not price discriminated against. Rather, price discrimination occurs to other customers based on observable characteristics. This is not an essential feature of our model, however. We could have allowed individuals to make repeated transactions with the same merchant. It simply was not necessary to make our general point, which is that information collected from one consumer may reveal something about other consumers with similar observable characteristics that can be exploited by retailers. The private costs of individuals failing to protect their privacy will be less than the social costs as long as some of the cost associated with each individual's failure to preserve their privacy is borne by others.

We contend that the increased speed and convenience of commercially provided digital payment options, which reduce the private incentives to protect privacy in payments, together with advancements in big data processing, not only promote price discrimination, but also lead to a reduction in social welfare. The notion that price discrimination decreases social welfare may seem surprising. This result arises in our model for two reasons. First, due to unavoidable errors in customer profiling, retailers will mistakenly quote high prices to some customers who have low reservation prices. As a consequence, some transactions that would increase social welfare do not occur. This first effect particularly hurts the agents in the model who have less money. Second, price discrimination reduces the purchasing power of agents who earn more money, which results in a lower overall level of consumption and also weakens incentives for young agents to earn money in the first place.

One remedy to the current trend in declining privacy in payments would be the widespread adoption of a digital cash substitute that offers users a similar level of privacy in payments as physical cash. Decentralized cryptocurrencies like Bitcoin meet some of the privacy requirements but have other features, such as costly proof-of-work transaction validation or scalability issues, that may limit their usefulness as a universal cash substitute. Other solutions include the addition of privacy features to conventional private-sector payment platforms. Tokenization (e.g., Apple Pay) is a step in this direction, but it only provides privacy from the retailer and the provider of the payment infrastructure. It does not provide privacy from the credit card company. As such, the risk that payments data might be exploited remains. The epigraph to this article includes an assertion by IMF Managing Director Christine [Lagarde](#) that privacy in payments cannot be provided by the private sector. A compelling reason for a public solution is that the government has no profit incentive to exploit payments data. This unique feature of payment instruments provided by central banks may become increasingly important as private-sector alternatives, such as Facebook's initiative to launch an electronic payment instrument called Libra ([Bloomberg, 2019](#)), begin to gain momentum.

Appendix A Patents and Patent Applications

[Mastercard \(2011\)](#): “Systems and methods for analyzing and segregating payment card account profiles into clusters and targeting offers to cardholders. (...) Customers who have no transaction history with a merchant may be selected for offers based on similarities with respect to other customers of the merchant.”

[Mastercard \(2018\)](#): “Disclosed herein are systems and methods of individual level learning that include receiving purchase event data from a merchant device that indicates that a purchase event occurred by a user on a user device, and transmitting the purchase event data to an analytics server. The methods may also include processing the purchase event data. (...) When the purchase hazard probability is above a threshold, the system may push a message to the user device.”

[American Express \(2014\)](#): “End consumer-facing business entities tend to have limited data relating to their customers. While a business entity may have a customer list containing demographic information about a customer and/or a set of prior transactions conducted by the business entity and the customer, many business entities may not have additional information about their customers. (...) It would thus be useful for a business entity to enhance the value of this limited dataset to gain additional insights into its consumer base through, for example, a cooperative data exchange.”

“In various embodiments, a node comprises a private data store and/or a public data store that comprises internal data. Internal data may be gathered from a transaction system (...) Consumer transactional data may include any data pertaining to the particular transactions in which a consumer engages during any given time period.”

“Node 204 may supplement received data (...) with third party data sources. For example, a third party data source may provide customer credit scores, social network histories (which include any information a social network may gather regarding a consumer, for example, posted messages, approximate age and gender of spouse, children and other members of household, pictures, past consumer geographic locations, patterns of past consumer geographic locations, propensity to engage in risky behaviors and the frequency of engaging in the same, marital status, substance use history, dating history, education level, present and past health status including disease status), public records, consumer transactions conducted using alternate payment systems, consumer health status, and any other data relating to consumers who may appear in the internal data.”

Visa (2018): “Systems, apparatus, and methods for determining incentives based on consumer history. When, how, and to whom incentives are sent can be determined. For example, an incentive can be sent to a consumer to encourage a transaction at a time when the particular consumer is predisposed to initiate the transaction. Also, an incentive for a transaction can be sent to a consumer when that transaction has a high likelihood of leading to other transactions.”

Appendix B Proofs

Proof that $\theta^U < \theta^W$ in the dynamic model

From $f < (b - c)/2$, which implies $b > 2f + c$, we have that

$$\theta^U < \frac{6f - (1 - c)}{2 + 2c}.$$

Proving that this quantity is smaller than θ^W is sufficient to prove $\theta^U < \theta^W$. Using the expression for θ^W , this is equivalent to proving

$$(6f + c - 1)(1 + 2c - f) < (c + f)(2 + 2c).$$

This expression holds true, because rewriting gives

$$\begin{aligned} 0 &< 3c - 9cf + 1 - 5f + 6f^2, \\ 0 &< \underbrace{3c(1 - 3f)}_{>0 \text{ for } f < 1/3} + \underbrace{(1 - 2f)(1 - 3f)}_{>0 \text{ for } f < 1/3}. \blacksquare \end{aligned} \tag{13}$$

References

- J.M. Abowd and I.M. Schmutte. An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices. *American Economic Review*, 109(1):171–202, 2019.
- A. Acquisti, C. Taylor, and L. Wagman. The Economics of Privacy. *Journal of Economic Literature*, 54(2):442–492, 2016.
- American Express. Systems and Methods for Cooperative Data Exchange. Inventors: Bonalle et al. *U.S. Patent*, 8,676,992 B2, 2014.
- S. Athey, C. Catalini, and C. Tucker. The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. *NBER Working Paper*, 23488, 2017.
- Y. Balasko and K. Shell. The Overlapping-Generations Model. II. The Case of Pure Exchange with Money. *Journal of Economic Theory*, 24(1):112–142, 1981.
- Bloomberg. Google and Mastercard Cut a Secret Ad Deal to Track Retail Sales. *News Report (30 August)*, 2018. URL <https://www.bloomberg.com/news/articles/2018-08-30/google-and-mastercard-cut-a-secret-ad-deal-to-track-retail-sales>.

- Bloomberg. Facebook’s Libra Gambit Forces Washington’s Hand on Crypto Policy. *News Report (20 June)*, 2019. URL <https://www.bloomberg.com/news/articles/2019-06-20/facebook-s-bold-move-spurs-d-c-to-confront-its-crypto-dithering>.
- D. Chaum. Security without Identification: Card Computers to Make Big Brother Obsolete. *Communications of the ACM*, 28(10):96–101, 1985.
- D. Chaum. Achieving Electronic Privacy. *Scientific American (August)*, pages 96–101, 1992.
- J.P. Choi, D.-S. Jeon, and B.-C. Kim. Privacy and Personal Data Collection with Information Externalities. *Journal of Public Economics*, 173:113–124, 2019.
- CNBC. FICO with Chinese Characteristics: Nice Rewards, but Punishing Penalties. *News Report (March 16)*, 2017. URL <https://www.cnbc.com/2017/03/16/china-social-credit-system-ant-financials-sesame-credit-and-others-give-scores-that-go-beyond-fico.html>.
- Deutsche Bundesbank. Payment Behaviour in Germany in 2017. *Statistical Report*, 2017.
- J.A.T. Fairfield and C. Engel. Privacy as a Public Good. *Duke Law Journal*, 65:385, 2015.
- Financial Times. Alibaba’s Social Credit Rating Is a Risky Game. *News Report (February 21)*, 2018. URL <https://www.ft.com/content/99165d7a-1646-11e8-9376-4a6390addb44>.
- C.S. Henry, K.P.H. Huynh, and A. Welte. 2017 Methods-of-Payment Survey Report. *Bank of Canada Staff Discussion Paper*, 2018-17, 2018.
- S. Ichihashi. Online Privacy and Information Disclosure by Consumers. *Bank of Canada Staff Working Paper*, 2019-22, 2019.
- C.M. Kahn. The Threat of Privacy. *Journal of Financial Market Infrastructures*, 6(2/3): 1–10, 2018.

- C.M. Kahn, J. McAndrews, and W. Roberds. A Theory of Transactions Privacy. *Federal Reserve Bank of Atlanta Working Paper*, 2000-22, 2000.
- C.M. Kahn, J. McAndrews, and W. Roberds. Money Is Privacy. *International Economic Review*, 46(2):377–399, 2005.
- T. Khiaonarong and D. Humphrey. Cash Use Across Countries and the Demand for Central Bank Digital Currency. *IMF Working Paper*, 2019-46, 2019.
- D. Kovenock and C.G. De Vries. Fiat Exchange in Finite Economies. *Economic Inquiry*, 40(2):147–157, 2002.
- R. Kumar, T. Maktabi, and S. O’Brien. 2018 Findings from the Diary of Consumer Payment Choice. *Statistical Report*, 2018.
- C. Lagarde. Winds of Change: The Case for New Digital Currency. *Speech*, 2018. URL <https://www.imf.org/en/News/Articles/2018/11/13/sp111418-winds-of-change-the-case-for-new-digital-currency>.
- R. Lagos and R. Wright. A Unified Framework for Monetary Theory and Policy Analysis. *Journal of Political Economy*, 113(3):463–484, 2005.
- M. MacCarthy. New Directions in Privacy: Disclosure, Unfairness and Externalities. *I/S Journal of Law and Policy for the Information Society*, 6:425–512, 2010.
- Mastercard. Techniques for Targeted Offers. Inventors: Alfonso et al. *U.S. Patent Application*, 2011/0106607 A1, 2011.
- Mastercard. Individual Level Learning Mechanism. Inventors: J.P. Gerard, S.X. Meng and P. Hu. *U.S. Patent Application*, 2018/0211267 A1, 2018.
- J. McAndrews. The Case for Cash. *ADB Working Paper*, 679, 2017.

- M. Molico. The Distribution of Money and Prices in Search Equilibrium. *International Economic Review*, 47(3):701–722, 2006.
- S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. *White Paper*, 2008.
- P.A. Norberg, D.R. Horne, and D.A. Horne. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.
- Office of the Privacy Commissioner of Canada. Protecting Personal Information: Cannabis Transactions. *Guidance Document (December 17)*, 2018.
- M. Pagano and T. Jappelli. Information Sharing in Credit Markets. *Journal of Finance*, 48(5):1693–1718, 1993.
- K.S. Rogoff. *The Curse of Cash*. Princeton University Press, 2016.
- Visa. Determining Targeted Incentives Based on Consumer Transaction History. Inventors: P. Faith and K.P. Siegel. *U.S. Patent*, 9,984,379 B2, 2018.
- N. Wallace. The Overlapping Generations Model of Fiat Money. In J.H. Kareken and N. Wallace, editors, *Models of Monetary Economies*, pages 49–82. Federal Reserve Bank of Minneapolis, 1980.
- N. Wallace. Whither Monetary Economics? *International Economic Review*, 42(4):847–869, 2001.
- Wired. Inside China’s Vast New Experiment in Social Ranking. *News Report (December 14)*, 2017. URL <https://www.wired.com/story/age-of-social-credit/>.
- WSJ. China’s New Tool for Social Control: A Credit Rating for Everything. *News Report (17 February)*, 2019. URL <https://www.wsj.com/articles/chinas-new-tool-for-social-control-a-credit-rating-for-everything-1480351590>.