A Theory Model of Digital Currency with Asymmetric Privacy

Katrin Tinn
McGill University
katrin.tinn@mcgill.ca

June 2024*

Abstract

This paper considers introducing asymmetric privacy in the design of central bank digital currencies (CBDC) and digital currencies more generally, to preserve the privacy of money spent while keeping the benefits of digital records for money received. It is shown that this feature would help minimize real distortions between consumers, firms, and financiers, while enabling tax optimization and better access to external financing. Protecting the privacy of consumers is always desirable from an aggregate standpoint as long as there exist some privacy concerns. Implementing asymmetric privacy is technologically feasible, using for instance Zero-Knowledge proofs or other privacy tools.

Keywords: Central bank digital currency design, data privacy, learning, real effects of privacy preferences, verification costs, technological tools for privacy.

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

*This paper stems from a paper co-authored with Christophe Dubach (at McGill University's Electrical Engineering and Computer Science Department) written for the Bank of Canada's Model X Challenge, which called for Finance and Computer Science collaboration to make proposals for a Canadian Central Bank Digital Currency design. The original policy and technical implementation-oriented version is available at SSRN (see Tinn and Dubach, 2021. We thank the Bank of Canada for their support and the many fruitful discussions. Note that the views expressed in this paper do not represent the official view of the Bank of Canada. Gratitude is extended to Claude Crépeau, Darrell Duffie, Gregory Dudek, Antonio Fatas, Rod Garratt, Nobuhiro Kiyotaki, Cyril Monnet as well as audiences at the Bank on Canada, McGill University, University of Toronto, Queen's University, University of Dauphine, Bocconi University, Morgan Stanley, Bank of Estonia, Federal Reserve Bank, Miami Herbert School of Business, and The Deutsche Bundesbank International Conference on Payments and Securities Settlement for valuable comments and suggestion.

1 Introduction

Central banks worldwide have been exploring the idea of issuing a Central Bank Digital Currency (CBDC), and some countries have already issued a version of it.[1] This paper focuses on retail CBDC[2], which is a central bank-issued digital equivalent of physical cash that can be used for making and receiving payments by individuals and firms. It centers on one key design choice which arises when implementing a retail Central Bank Digital Currency (CBDC): should the retail CBDC aim to replicate the level of privacy associated with physical cash, or should it harness the advantages of digital records to enhance efficiency? In other words, should it resemble "token money" or "account money" (see e.g., Brunnermeier et al., 2019 and Kahn et al., 2019 for definitions), or find a middle ground between these two extremes?

This paper emphasizes that a retail CBDC system will inevitably create a large database of transaction records. The chosen level of detail, or the lack thereof, in individual or firm identity-related information associated with transactions within this database has the potential to significantly impact fundamental economic interactions, either positively or negatively. Namely, the different privacy features of retail digital payment recording systems will influence the firm-consumer relationship.

Consequently, Central Banks or digital currency issuers face a trade-off. On the one hand, consumers seem to prefer a CBDC that is token money as privacy ranks highly in user preferences.[3] Token money comes with the advantage that consumers can make their purchasing decisions without needing to worry about outside parties collecting data on their purchasing choices, which may lead to unwanted advertising, adverse impact, and potential discrimination when it comes to their credit score or other unintended consequences.

On the other hand, the current view among many central banks seems to be that the future CBDC should not be as anonymous as physical cash for it not to become a tool for money laundering and other illicit activities, i.e., account money with some privacy protection is preferable.[4] Beyond

[1] Atlantic Council's CBDC Tracker (`https://www.atlanticcouncil.org/cbdctracker`) keeps up-to-date records on these developments. Bahamas and other smaller countries were the early movers, and larger central banks like the US Federal Reserve System, European Central Bank, Bank of England, Bank of Canada, Bank of Japan, People's Bank of China being at different stages of exploration, development and launch.

[2] An alternative is a wholesale CBDC that is focused on improving the transaction rails across central banks and financial institutions, while not directly used by individuals. See, e.g., Duffie et al., 2020, Auer et al., 2020, and Auer et al., 2022, for an overview of a broad set of crucial choices around CBDCs.

[3] For example, the results on public consultations by the European Central Bank indicated that "I want my payments to remain a private matter" ranked highest among the desired features of the digital euro (see European Central Bank, 2021. Similar preferences for privacy were highlighted in the UK's and Canada's public consultations (see the speech by Cunliffe, 2023 at the Federal Reserve Board, and the report by Bank of Canada, 2023)

[4] See e.g., the speech on digital euro by Panetta, 2022 and the report by UK Parlament Treasury Committee, 2023 on digital pound.

this preference, an account-money based CBDC system would add value by enabling programmable features, which can enhance the creation of innovative financing products, among other benefits of a digital database for efficiency.[5]

It has been argued that the trade-off between privacy and compliance may be unavoidable (Auer and Böhme, 2020), and that the main issue is to decide what limited degree of privacy to bestow on the new payment system. For instance, the digital euro project proposes to only ensure the privacy of small payments vs large sums of money. While the existence of trade-offs is undisputed, this paper argues that a better way to strike the right balance may not be to focus on amounts, but rather on which aspects of privacy need, and need not, be protected.

The main conclusion from the model presented in this paper is that a privacy asymmetric design where the privacy of consumers (senders of money) is maximally protected by design, while the privacy of firms (receivers of money) is less protected achieves the best of both the token and account-based systems. In short, the implementation of such a system could involve all retail CBDC users having unique ID-linked accounts, which allow consumers to generate anonymous tokens for payments. The consequence of this is that when an individual, say Alice, buys a good or service, say from Bob's firm, the database will have identity-associated information about the receipt of funds at Bob firm's ID-linked account, but not on Alice as the buyer (as it could have been Becca, John, Li Qiang, or anyone else).

Different technological possibilities exist for implementing these options: for instance, digital central bank token money could benefit from simple tools like prepaid cards up to sophisticated cryptography tools like Zero-Knowledge-proofs;[6] and the implementation of an account-based system could benefit from the existing and improved record keeping systems that already exist within the financial sector, and may further benefit from utilizing innovations based on digital ledger technologies, including blockchain.

The model starts by considering two extreme scenarios: 1) all payments between the firm and its consumers are made using token money and are anonymous; 2) all individual purchases are made using account money, and this information is visible to external parties beyond the firm, i.e., financiers, tax authorities, government, and other agents in the economy.

The analysis begins by highlighting the limitations of token money. While this system allows consumers to buy what they want without privacy concerns, traditional frictions become perpet-

[5] Programmable features, often called "smart contract" features in the context of CBDCs have been emphasized by Auer et al., 2022, Model X CBDC projects for the Bank of Canada Choi et al., 2021, Tinn and Dubach, 2021, and Veneris et al., 2021, and by Hyun Song Shin in his publicly available speeches as the Economic Advisor and Head of Research at the Bank of International Settlements (see also the literature section for further references), see further in the Literature section.

[6] See Section 6 for discussion on potential privacy tools.

uated. This part of the paper builds on the classical literature on costly verification and auditing and brings together a set of key insights from the finance and public economics literature.[7] While the key findings of this case are not surprising in the context of the classical literature on costly verification and auditing, its policy implications for a retail CBDC (or for other forms of digital currencies that may be widely used) are worth re-emphasizing. Incorporating both tax and financing considerations within the same framework reveals that a "token money" based system: 1) makes external financing less accessible and many projects that would be value-adding are not undertaken; 2) enables the most successful firms optimally evade taxes, because it is typically not optimal to audit their revenues beyond a threshold; 3) leads to optimal external financing contracts to favor debt contracts as these minimize auditing costs.

The analysis then points towards two benefits that a CBDC in the form of non-private "account money" would bring, both stemming from reduced auditing costs. First, lower auditing cost enable more efficient tax collection from firms, which in turn can lead to lower taxes, at least in a partial equilibrium context. Second, a richer set of financing contracts are possible and efficient in the context of verifiable records, including equity and equity-like contracts.[8]

Modeling the privacy concerns that would be associated with non-private account money deserves further clarification as there is no universal consensus in the academic literature on how different aspects of privacy concerns should be modeled. For the sake of clarity, the model in this paper does not consider the privacy concerns that affect the firm-consumer direct relationship, e.g., via special offers or price discrimination. The costs and benefits of this relationship can lead to mixed conclusions.

This paper focuses on the distortions between the firm and consumers, when both parties are aware that third parties may make unfavorable conclusions, should all purchasing decisions between the firms and consumers be known to these third parties.

The paper argues that worries about third parties, e.g., other firms and credit institutions making conclusions about consumer types, is of first-order importance when it comes to designing a large database of transactions (e.g., a retail CBDC). Recent empirical papers have established that even innocuous-seeming information about consumers' digital footprints (e.g., when and which device they use to access a website) can be used to assess creditworthiness[9], and the overall consumption choices consumers make may affect their access to credit and cost of credit.[10] Moreover,

[7] See references in the Literature Section 2

[8] See Modigliani and Miller, 1958 for the friction-less benchmark, and Holmstrom and Milgrom, 1987, Jovanovic and Szentes, 2013 and Tinn, 2018 for examples of settings where equity(-like) contracts help to sustain effort incentives.

[9] Berg et al., 2020

[10] see Vissing-Jorgensen, 2021, and also Parlour et al., 2022 and Ouyang, 2021 on information in payment flows.

the literature on differential privacy[11] highlights that the pseudonymization of identities does not guarantee the prevention of real identities from being revealed.

This paper also emphasizes another important feature of privacy that distinguishes it from settings with "good" and "bad" types, commonly analyzed in the adverse selection literature.[12] To make a clear case for privacy preferences, the model considers that there are no universally "good" and "bad" types, and all consumers dislike characteristic information about them, irrelevant for a particular buying decision, but possibly correlated with their type and affecting other interactions, being in any digital database that can be shared. Essentially, privacy preferences can be viewed as rational paranoia, worrying about the worst that can happen. One natural example is that someone buying medicine for their parent might worry that an external party (e.g., an insurance company) could conclude they are a high-risk customer for insurance. Another example to consider is a situation where someone, having paid for an expensive dinner, may worry that a travel agency, being aware of it, could conclude they are willing to accept higher prices for air travel and hotels.

Motivated by these examples, the paper models privacy concerns as a component in consumers' indirect utility. This enables to highlight the main trade-off between the extreme versions of token and account money: which one is better effectively boils down to the question how high privacy costs are compared to auditing costs. Both systems are inefficient in enabling all firms with positive net value projects to raise funds for their projects and may imply higher taxes in equilibrium. The reason why the extreme version of account money is inefficient is that the firm has to offer its product at a discount to compensate for the privacy loss or lose some demand from interested consumers who use mixed strategies.

The intuition derived from analyzing the extreme case further helps to understand why asymmetric privacy is welfare-enhancing in this setting. To benefit from improved financing contracting terms and lower taxes, firms do not require knowledge of the specific individuals who bought their products. Verifiable records of total purchases, however, play a crucial role in alleviating traditional contracting frictions. Simultaneously, asymmetric privacy makes it harder to learn consumer types, subsequently reducing consumers' rational "paranoia" and enhancing the firm's expected profits. This, in turn, facilitates the undertaking of more socially beneficial projects.

The rest of the paper is structured as follows. Section 2 discusses related literature, Sections 3 and 4 present the main model of the key trade-offs under pure strategies, Section 5 extends the model to allow mixed strategies to be played by consumers. It further presents a setting to rationalize the assumed functional form of consumers' privacy preferences. This section further highlights the

[11] see e.g., Dwork, Roth, et al., 2014 for an overview.

[12] See Mas-Colell et al., 1995 and Tirole, 2010 for classical arguments and literature on adverse selection.

connection between the assumed privacy tools with the concepts in differential privacy literature in computer science. Section 6 discusses some possible ways a CBDC with asymmetric privacy could be designed using available technologies. Finally, Section 7 concludes.

2 Related Literature

This paper relates to several streams of the literature, including considerations around the rationale and impact of CBDC issuance. There is a recent literature which emphasizes the macro-economic impact CBDC can have on banks and other mainstream financial institutions (see e.g., Li et al., 2023 Whited et al., 2022, Chiu and Davoodalhosseini, 2023, Keister and Sanches, 2023, Williamson, 2022 Andolfatto, 2021, Chiu et al., 2023, Fernández-Villaverde et al., 2021), on monetary policy and financial stability (see e.g., Barrdear and Kumhof, 2022, Brunnermeier and Niepelt, 2019, Niepelt, 2020, Davoodalhosseini, 2022, Schilling et al., 2020, Keister and Monnet, 2022), and international considerations around flows of capital adoption (see Minesso et al., 2022 Cong and Mayer, 2022). This literature focuses on the role of central bank-issued digital money as an alternative to deposits and savings held in banks, which in turn may crowd out bank deposits, affect financial stability, and become either locally, or internationally a compelling alternative to bank money and may affect the mainstream reserve currencies.

This paper's focus is closer to another stream of recent literature which focuses on the role of CBDCs as a compelling replacement of physical cash, especially when it comes to privacy considerations. For example, Ahnert et al., 2022 develop a model where privacy-preserving CBDC helps to resolve an otherwise rather complex problem that involves the seller's rational potential desire to hide information from their bank, and the merchants' choice to accept cash or digital payment. Agur et al., 2022 argues that households with heterogeneous preferences regarding anonymity and security benefit from the simultaneous existence of different means of payments, i.e., cash is anonymous and allows buying "sin goods" (e.g., alcohol), but can be lost. At the same time deposits are less private and more secure. Both are useful in the absence of a CBDC. The level of anonymity of CBDC is a choice variable and CBDC may pay interest. The paper argues that in the presence of network effects, it may be desirable to have an interest-paying CBDC that is somewhere in between (e.g. the ECB proposed a CBDC design which offers greater privacy of small payments). Garratt and Lee, 2021 highlights that a digital currency design that is not privacy-preserving tends to lead to data monopolies. Their setting allows prices to depend on the type of payment instrument used, and cash is costly to use but preserves anonymity. The main benefit of CBDC in this setting is that it offers low-cost means of digital payments and may be welfare-enhancing despite firms producing less valuable goods due to having less information. In Agur et al., 2024, lenders can exploit payment data to infer the creditworthiness of heterogeneous and privacy-concerned households. Their paper shows that data intermediaries with monopolistic power choose too little privacy compared to what is socially optimal, which provides a rationale for privacy-protecting regulations. Also Cheng and Izumi, 2024 explores the impact of the choice of means of payment in the context of bank lending and CBDC under anonymity preferences. These privacy considerations are complementary and different from the ones considered in this paper. Namely, these papers focus on comparing means of payment that are more or less (symmetrically) private rather than on which aspects of data should be kept private to harness the benefits of digital records without sacrificing the key benefits of privacy for consumers.

Furthermore, this paper considers privacy concerns that arise because consumers are worried about third parties drawing conclusions about their characteristics. These are aspects they would prefer not to share widely, even if these characteristics have no direct impact on the bilateral relationship between the producer/firm and the consumer. As highlighted in Jones and Tonetti, 2020 data is non-rival and firms that collect individuals' data (e.g., their medical records) may sell it to other parties. They show that giving individuals privacy tools and control over their data can be closer to socially optimal. In Jones and Tonetti, 2020 as well as in Garratt and Lee, 2021 and Cong et al., 2022 there are benefits of data sharing for producing goods that match consumer preferences better and for innovation. While this paper does not explicitly model this consideration to maintain the clarity of the main argument, it is worth noting that these benefits do not necessarily need to rely on an individual's purchase data. Similar benefits could arise when the data of total sales is publicly observable by relevant parties, a feature that the asymmetric privacy design would allow.

There are additional privacy considerations as the firms themselves may collect data about their customers to price-discriminate in the future, offer valuable complementary products, or share their data without consumers benefiting from this (see e.g., Garratt and Van Oordt, 2021, Acquisti and Varian, 2005, Calzolari and Pavan, 2006, Taylor, 2004, Varian, 1985). This literature highlights both the costs and benefits of this. While this model does not explicitly incorporate this consideration for the sake of clarity, note that consumer privacy-protecting means of payment would also alleviate the negative impact of these frictions. Furthermore, the consumers could still undo privacy (e.g., by using a loyalty card) if revealing information about past purchases from this firm is in their interest. Consumers may also prefer data privacy to conceal their behavioral vulnerabilities (see Liu et al., 2021), This paper further relates to the monetary economics literature[13] which considers how the privacy features of different means of payment and record-keeping systems affect the economic interactions between consumers and producers. In particular, it relates to Kahn et al., 2005 which considers money as a record-keeping device, and similarly to this paper highlights that cash has an advantage as it does not reveal the identity of the buyer, who may be more easily "robbed" or price-discriminated against otherwise. When rationalizing consumers' privacy preferences in Section 5, this paper proposes a setting that builds on Kahn et al., 2005. It additionally explicitly models the optimal arrangements for external credit and tax efficiency in an environment where verifying sales records requires costly auditing. This helps to highlight why asymmetric privacy is desirable to balance privacy concerns and the usefulness of data. It is also worth noting that unlike Kahn et al., 2005 and Kahn and Roberds, 2008 credit in this paper setting will not require repeated interactions, reputation costs, and "credit clubs", which may add realism as the firm's investors and consumers in the broad economy may often be very different entities. All this further enables us to derive some additional practical consequences for the type of external financing contracts that are optimal to use, as well as highlight the tax benefits of the asymmetric design while reconfirming some of the key insights of these papers in a different setting.

In addition to privacy, this paper considers that the data recorded within a CBDC system (e.g., on a blockchain or another shared digital ledger) can be used for more efficient financial contracting and tax optimization. The key features that these technologies enable are a reduction (if not elimination) of verification costs (see e.g., Catalini and Gans, 2020) and some programmable features, or smart contracts as these are referred to in the context of crypto-assets. While this paper considers applying smart contract functionalities for raising external funds and financial contracting as in Tinn, 2018 and Tinn, 2019, the recent literature has highlighted several other programmable features that a CBDC could include, including delayed payments (Kahn and Van Oordt, 2022), expiring cash for offline payments (Kahn et al., 2021), automated settlement (Lee et al., 2021). Different forms of smart contracts in the context of crypto-assets have also been analyzed in Bakos and Halaburda, 2019, Chiu and Koeppl, 2019, Cong and He, 2019, Cong et al., 2021b, and Gans, 2019, Gan et al., 2023 among others.

When modeling verification costs and the benefit of their reduction under digital ledger technologies, this paper builds on the large literature on costly auditing both in finance (see e.g., Tirole,

[13] For broader overviews of micro-founded models of money see e.g., Ostroy and Starr, 1990 and Williamson and Wright, 2010. For example, in Kiyotaki and Wright, 1993 money emerges as having value because it helps to overcome the lack of double coincidence of wants across goods, and in Kiyotaki and Moore, 2018 and Kiyotaki and Moore, 2002 money helps to transact across time in an environment with lack of commitment and lack resaleable of financial contracts. New digital ledger technologies, if widely adopted, have the potential to ease these frictions.

2010 for an overview, and seminal papers on costly state verification by Townsend, 1979, Gale and Hellwig, 1985, Diamond, 1984), and in public economics (see e.g., Slemrod and Yitzhaki, 2002 and Andreoni et al., 1998 for overview, and seminal papers by Graetz et al., 1986, Border and Sobel, 1987, Mookherjee and Png, 1989 Sanchez and Sobel, 1993). The modeling approach is closest to Sanchez and Sobel, 1993. In these papers, there is a need to audit because agents have an incentive to under-report their revenues and cash flows. A different auditing consideration features in Cao et al., 2019 where firms have incentives to over-report their revenues. While not explicitly modeled, these auditing considerations are likely to also benefit from the privacy asymmetric design of the payment system.

In addition to the finance and economics literature, this paper relates to the literature in computer science. Modeling privacy costs is at the center of the differential privacy literature, which analyzes how to make queries from a database without revealing too much information about individuals (see Dwork, Roth, et al., 2014 for an overview, and Abowd and Schmutte, 2019 and Schmutte and Yoder, 2022) in the context of econometric analysis. The paper will draw some parallels between modeling privacy costs in our setting and this literature in Section 5.

Implementing a privacy asymmetric design may benefit from the use of tools like Zero-knowledge proofs (see Goldwasser et al., 1985, Goldwasser et al., 2019, Fiege et al., 1987), and such approaches have also been proposed in a CBDC context, for example by Wüst et al., 2022 and Gross et al., 2021. These papers focus on the implementation of a privacy-preserving CBDC rather than developing an economic model. Zero-knowledge proofs are also used in the context of crypto-currencies like ZeroCoin and ZeroCash (Ben-sasson et al., 2014a and Ben-sasson et al., 2014b), and for getting real-world data on the blockchain, e.g., by ChainLink, and in the extensions of Ethereum system.[14] Also, Cao et al., 2019 analyzes Zero-Knowledge proofs in an economics context.

Finally, while this paper's narrative is built around the central bank issuing a retail digital currency, it is plausible that another entity, e.g., a technology firm or crypto-asset community, could issue such means of payment, especially as privacy considerations - as a driver of demand - feature prominently in this literature (see e.g., Pagnotta, 2018 and Zhou, 2021).[15]

[14] See the explanations of Zero-Knowledge proofs by Chainlink at https://chain.link/education/zero-knowledge-proof-zkp and by Ethereum at https://ethereum.org/en/zero-knowledge-proofs/ It will be also further discussed in Section 6

[15] The paper also relates to broader and expanding literature on different aspects of blockchain economics, see e.g., Yermack, 2017, Biais et al., 2019, Cong et al., 2021a, Saleh, 2021, Halaburda et al., 2022, John et al., 2022, Benhaim et al., 2023, Biais et al., 2023, Capponi et al., 2023a, Capponi et al., 2023b, Iyengar et al., 2023, Malinova and Park, 2023, among others including reviews and collections of articles, e.g., Fatás, 2019

3 The main model of key trade-offs: privacy concerns vs. efficient auditing

This section presents a stylized model to highlight the main trade-off between privacy concerns and efficient auditing. One key element which distinguishes various payment systems is the degree of information about consumers and firms which are recorded on digital ledgers. To underscore the consequences of this in the starkest way, assume that any information that is recorded on the payment ledger is available to any interested third parties.

After describing the setting, the paper proceeds to analyze two opposite systems: 1) pure token money, which could mean that payments are made in physical cash or by using a digital equivalent which maximally protects the identities of consumers and firms and the transactions between them - this system is called "symmetrically private"; 2) pure account money where details about all individual purchases and funds received are recorded on a publicly available digital ledger - this system is called "symmetrically non-private". After analyzing these two extreme cases, this section highlights the economic benefits of an "asymmetric privacy" design, where the firm's identity and funds it received (i.e., its total sales) is public, but who exactly bought their product remains private.[16]

3.1 Setting

Consider a firm which aims to produce an indivisible good or service and has a potential target market consisting of $N \geq 1$ individuals, indexed by n. Some of these consumers derive utility from consuming the good or service, and all consumers are aware that their purchasing decisions may reveal information about their privacy-sensitive traits, which they dislike having outside parties learn about. Production and trade take place on date 1, and when selling the product to the consumers, the firm has bargaining power and chooses the product price, p. Variable costs of production are normalized to zero for the sake of clarity.[17]

Before producing, the firm needs to make a fixed investment of $I > 0$ and needs to raise external financing from risk-neutral financiers to be able to produce. The financing contract is denoted by a function $\phi(.)$. If the firm is active, i.e., has invested, it is also liable for taxes, and the tax function $\tau(.)$ is set exogenously before the game starts. If the firm's total realized sales, S, are not costlessly

[16] There would be a theoretical fourth case where data on individual purchases is public, but the firm's total sales are not. This case is not explicitly analyzed because it would straightforwardly be sub-optimal to the other three cases in this setting.

[17] Should there be positive variable costs, the setting and related conclusions would straightforwardly extend from total sales to net sales, i.e., sales minus known variable costs per unit.

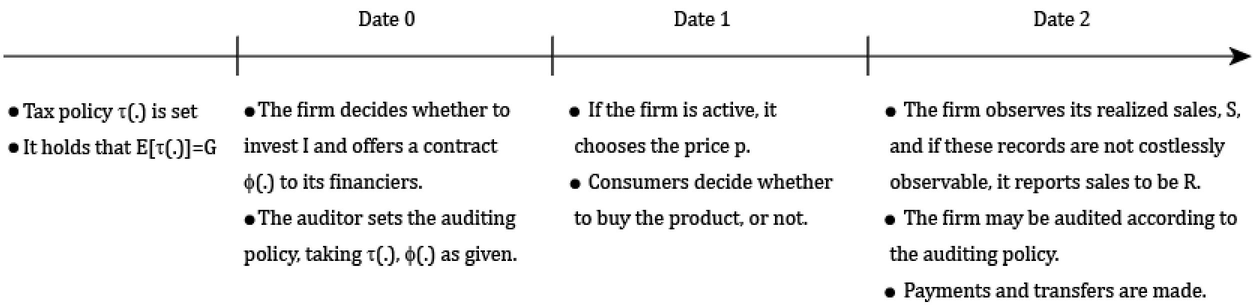
Date 0
Date 1
Date 2
• Tax policy τ(.) is set
• It holds that E[τ(.)]=G
• The firm decides whether to invest I and offers a contract ϕ(.) to its financiers.
• The auditor sets the auditing policy, taking τ(.), ϕ(.) as given.
• If the firm is active, it chooses the price p.
• Consumers decide whether to buy the product, or not.
• The firm observes its realized sales, S, and if these records are not costlessly observable, it reports sales to be R.
• The firm may be audited according to the auditing policy.
• Payments and transfers are made.

Figure 1: Sketch of the timing of events.

observable on the payment ledger, the firm self-reports R and may be audited. As in Sanchez and Sobel, 1993, auditing is delegated to a skilled, risk-neutral, and unbiased external party, i.e., an auditor, who can verify the value of S at the auditing cost $c \geq 0$, and is tasked with choosing the optimal auditing policy. This auditor acts on behalf of both the tax authority and financiers, and without loss of generality, it is assumed that the financiers are the ones who pay auditing costs, withhold taxes on behalf of the tax authority, and pass these on to the tax authority.[18] It is assumed that the government gives an active firm a lump-sum investment subsidy, denoted as G. While the firm takes G as given, the government aims for efficiency in expectations, i.e., G is set to be equal to the expected tax revenue. The discount rate is normalized to one.

While the full details of the setting are below, Figure 1 provides a sketch of the timing of events. The full game is solved using the concept of Perfect Bayesian Equilibrium and the date 1 contracting and auditing problem is solved using tools from Mechanism Design. The contracting setting is closest to Sanchez and Sobel, 1993. It also consolidates the financial contracting problem that builds on the costly state verification literature (see e.g., Chapter 3 in Tirole, 2010, Townsend, 1979, Gale and Hellwig, 1985, and Diamond, 1984).[19]

3.1.1 Target Consumers

The firm's target consumers are characterized by two parameters affecting their indirect utility. One parameter is their appreciation for consuming the firm's product, $v_n \in \{0, 1\}$, indicating whether consumer n values the product highly, or not. The second parameter reflects their privacy-sensitive

[18] As this setting does not aim to model the frictions between the financial sector and the government, it is implicitly assumed that these parties can arrange transfers among them optimally and that the tax authority and financiers have access to the information that the auditor has collected. There would be an extra benefit of reducing auditing costs should there be duplication of effort, and this would likely favor digital payment systems that reduce auditing costs.

[19] The main difference is to consider the optimal auditing problem separately and include the presence of taxes. It also allows probabilistic auditing.

traits, $T_n \in \{A, B\}$. Each consumer type is therefore characterized by a pair $\{v_n, T_n\}$. To emphasize pure privacy concerns, assume that consumers' privacy-sensitive traits are not directly relevant to the firm, which is solely concerned with its sales, promises to external financiers, and tax obligations.

Assume that the conditional probability that a consumer n values the firm's product highly is

$$\Pr(v_n = 1|T_n) = z_{T_n} \text{ for } T_n \in \{A, B\}, \tag{1}$$

where $0 < z_{T_n} < 1$, and consumer valuations for the product are conditionally independent, i.e., for any pair of consumers n and $k \neq n$, it holds that the conditional joint distribution $\Pr(v_n, v_k|T_n, T_k) = \Pr(v_n|T_n)\Pr(v_k|T_k)$. Without loss of generality, consider that $z_A > z_B$, such that the consumer types with $T_n = A$ are more likely to value the firm's product highly.

The prior probability $\Pr(T_n = A) \equiv q_A$, and $\Pr(T_n = B) = 1 - q_A$, where $0 < q_A < 1$, and the distribution of privacy-sensitive traits is independent across consumers. These distributional assumptions imply that the unconditional probability that a consumer values the product highly is

$$\Pr(v_n = 1) = q_A z_A + (1 - q_A) z_B \equiv \bar{z} \tag{2}$$

and the unconditional preferences for the good are also independent across consumers, i.e., $\Pr(v_n, v_k) = \Pr(v_n)\Pr(v_k)$.[20]

At the same time, it is crucial to note that this distributional setting implies that knowing a consumer's preference for the good would reveal information about their privacy-sensitive traits. For example, consider a consumer with $\{v_n, T_n\} = \{1, A\}$. By Bayes's rule, it holds that $\Pr(T_n = A|v_n = 1) = \frac{\Pr(v_n=1|T_n=A)\cdot\Pr(T_n=A)}{\Pr(v_n=1)} = \frac{z_A q_A}{\bar{z}} > q_A$, i.e., if an external party observes this consumer values the product, it would consider the consumer more likely to have privacy-sensitive traits, A.

The firm is aware that the willingness of its target consumers to buy the product may be influenced by their desire to keep their personal type, T_n, private. Denote consumer n's purchasing decision as

$$b_{\{v_n, T_n\}} = \begin{cases} 1 \text{ if consumer } n \text{ of type } \{v_n, T_n\} \text{ buys} \\ 0 \text{ if consumer } n \text{ of type } \{v_n, T_n\} \text{ does not buy} \end{cases}. \tag{3}$$

To develop the intuition, the results of the main setting in Section 4 consider that consumers follow pure strategies and $b_{\{v_n, T_n\}}$ also represents the consumer n's strategy there. Section 5.1 extends

[20] By the law of total probability $\Pr(v_n, v_k) = \sum_{T_n=\{A,B\}} \sum_{T_k=\{A,B\}} \Pr(T_n, T_k)\Pr(v_n, v_k|T_n, T_k) = \sum_{T_n=\{A,B\}} \sum_{T_k=\{A,B\}} \Pr(T_n)\Pr(T_k)\Pr(v_n|T_n)\Pr(v_k|T_k) = \sum_{T_n=\{A,B\}} \Pr(T_n)\Pr(v_n|T_n) \sum_{T_k=\{A,B\}} \Pr(T_k)\Pr(v_k|T_k) = \Pr(v_n)\Pr(v_k)$.

the analysis to mixed strategies by consumers, and mixed strategy $\sigma_{\{v_n,T_n\}}$ denotes the probability that the consumer uses pure strategy $b_{\{v_n,T_n\}} = 1$. The belief set $\mathcal{B}$ refers to equilibrium beliefs about the strategies used by consumers.

Let $V\left(v_n, T_n, b_{\{v_n,T_n\}}; \mathcal{B}\right)$ represent the indirect utility of a consumer n whose type is $\{v_n, T_n\}$ and makes a purchasing decision $b_{\{v_n,T_n\}}$, and assume that

$$V\left(v_n, T_n, b_{\{v_n,T_n\}}; \mathcal{B}\right) = \begin{cases} v_n - p - \epsilon \mathbb{E}\left[\mathcal{L}\left(\Omega, \mathcal{B}\right) | v_n, T_n, b_{\{v_n,T_n\}} = 1\right] \text{ if } b_{\{v_n,T_n\}} = 1 \\ -\epsilon \mathbb{E}\left[\mathcal{L}\left(\Omega, \mathcal{B}\right) | v_n, T_n, b_{\{v_n,T_n\}} = 0\right] \text{ if } b_{\{v_n,T_n\}} = 0 \end{cases}, \tag{4}$$

where p is the price charged by the firm, $\epsilon \geq 0$ is a parameter capturing the importance of privacy concern, and $\mathcal{L}\left(\Omega, \mathcal{B}\right)$ reflects learning about the consumer's privacy-sensitive traits. The latter in turn is determined by the information set Ω regarding the level of detail about consumer n's purchases recorded on the digital payment ledger and beliefs, $\mathcal{B}$. Learning about the consumer's privacy-sensitive traits is given by

$$\mathcal{L}\left(\Omega, \mathcal{B}\right) = \frac{\Pr\left(T_n | \Omega, \mathcal{B}\right)}{\Pr\left(T_n\right)} - \frac{\Pr\left(\overline{T_n} | \Omega, \mathcal{B}\right)}{\Pr\left(\overline{T_n}\right)}, \tag{5}$$

where $\overline{T_n}$ refers to privacy-sensitive traits that are different from consumer n's traits, i.e., if $T_n = A$ then $\overline{T_n} = B$, and vice versa.

While the main setting assumes the given functional form of indirect utility and privacy loss, Sections 5.2 and 5.3 will present a model in which this indirect utility emerges endogenously and draw parallels of this formulation to common privacy measures in the differential privacy literature. The essence of the assumed indirect utility, as defined in (4), is as follows: In the absence of privacy concerns ($\epsilon = 0$), the demand the firm faces from consumers is standard. Each consumer will purchase the product only if their consumption value for the product exceeds the price. The term $\mathcal{L}\left(\Omega, \mathcal{B}\right)$ highlights that any decision to buy or not may reveal information about their privacy-sensitive traits to external parties, and privacy-concerned consumers dislike this possibility of revealing information. By (5), the equilibrium effect of learning is evaluated using Bayes' rule and relative to the prior, and consumers would prefer external parties not to learn anything about their privacy-sensitive traits and if anything be mistaken about their true traits. This setting emphasizes that, unlike classical models of asymmetric information, there are no 'good' or 'bad' types; all privacy-concerned agents simply prefer not to reveal unneeded information.

Finally, the firm sets the same price, p, for all consumers.[21] Consequently, in its interactions

[21] As discussed in the literature section, price discrimination between the firm and its buyers can have both negative and positive consequences, and these considerations are not incorporated for the sake of clarity.

with consumers, the risk-neutral firm sets p to maximize its expected total revenues from sales

$$\mathbb{E}\left[Np\sum_{n=0}^{N}b_{\{v_n,T_n\}}|\Omega,\mathcal{B}\right] \tag{6}$$

taking the demand and payment recording system as given.

3.1.2 Firm's Interaction with Financiers and the Tax Authority

These interactions involve solving the Mechanism Design Problem for optimal auditing and financial contracting. Consistently with the Perfect Bayesian Equilibrium of the full game, consider that all relevant parties, i.e., the firm, financiers, the tax authority, and the auditor, have the same beliefs about the distribution of potential sales s at date 0. However, depending on the payment records system, the outsiders to the firm may not costlessly observe the realized sales, S, which will be observed by the firm at the time the contracts are settled on date 2. To simplify notation and allow for generality in this part of the problem, consider s as a discrete random variable with a probability mass function $\Pr(s=S)\equiv h_S$ and the cumulative distribution function $\Pr(s\leq S)\equiv H_S$. Naturally, this distribution is endogenous in equilibrium as it depends on the anticipated outcome of the game between the firm and its consumers on date 1. From (6), it must hold that S is proportional to the price and the number of consumers who buy the product, implying that the support of this distribution is $s\in\{s_0,s_1,\ldots,s_N\}$, where $s_0=0$, $s_n=np$, and $s_N=Np$.

It is useful to define some further simplifying notation and impose some realistic constraints. Recall that this setting assumes that tax policy, $\tau_s\equiv\tau(s)$, is a function of potential sales and has been set before the game starts. When raising external financing to cover its investment cost, I, the firm can aim to pledge its potential after-sales revenues, $s-\tau(s)$, and the contract it offers to risk-neutral financiers is denoted as a function $\phi(.)$. For analyzing efficient auditing, it is useful to define the total obligations of the firm to these external parties as $\varphi_s\equiv\varphi(s)=\tau(s)+\phi(s-\tau(s))$. As standard, it is assumed that a firm that sells nothing is not liable for taxes, i.e., $\tau(0)=0$, and there is limited liability, i.e., the firm's total obligations cannot exceed its revenues, $\varphi_s\leq s$. To keep the setting both general and realistic, assume that the tax function $\tau(.)$ is strictly increasing and conjecture that the financing contract with external financiers, $\phi(.)$, is weakly increasing. This implies that the firm's total obligations, $\varphi(.)$, are strictly increasing in s. Further assume that the total obligations function does not have too large upward jumps/is not too convex, i.e., $\frac{\varphi_{s+1}-\varphi_s}{\varphi_s-\varphi_{s-1}}\leq\frac{(s+1)(N-s+1)}{(N-s)s}$ for any s. Note that this condition holds for any total obligations that are linear or weakly concave in s, including proportional taxes and typical financing contracts like debt or equity.[22] As taxes could be convex (progressive), this formulation allows for this possibility and simply requires that the resulting $\varphi(.)$ is not too convex. This constraint on φ_s is sufficient, but not necessary, to guarantee the uniqueness of the solution to the optimal auditing problem.

The finance and taxation literature share insights on auditing efficiency and incentive-compatible reporting. One key insight is that, without an observable recording system for realized sales revenues, S, both implementable taxation and constrained efficient financing contracts cannot rely on self-reporting. Unless the payment technology makes the realized sales revenues observable to parties who need to know, there is a need for auditing. Namely, the firm chooses to make a report, denoted as R, after observing its sales revenues S, and this report may differ from actual sales. As in Sanchez and Sobel, 1993, there is an unbiased auditor, who takes the function $\varphi(s)$ as given and chooses the auditing policy $\gamma(R) \equiv \gamma_R$, i.e., the probability that report R is audited. There is a penalty for under-reporting, and the penalty rate, $\pi > 0$, is set outside the model (e.g., by law), and there is no reward for over-reporting. Namely, if the firm is audited, has sales S and has reported R it pays

$$\begin{cases} \varphi_R + (1+\pi)(\varphi_S - \varphi_R) \text{ if } S < R \\ \varphi_R \text{ if } S \geq R \end{cases} . \tag{7}$$

The auditor acts on behalf of both the tax authority and financiers, and the auditor chooses γ_R to maximize the expected payments to these external parties net of auditing cost, the optimal policy must solve

$$\max \mathbb{E}[\Phi(R)],$$

where expected payments conditional on the report, R, are

$$\Phi(R) \equiv \mathbb{E}[\varphi_R + (1+\pi)\gamma_R(\varphi_S - \varphi_R) - c\gamma_R | R], \tag{8}$$

and this optimization is subject to the firm's optimal reporting decision, i.e., for every S it must hold that reporting R is not more costly for the firm than any alternative report, R', i.e.,

$$\varphi_R + (1+\pi)\gamma_R(\varphi_S - \varphi_R) \leq \varphi_{R'} + (1+\pi)\gamma_{R'}(\varphi_S - \varphi_{R'}) \text{ for any } R' \neq R.$$

It is further assumed that the auditing cost $c \in \left(0, \frac{1+\pi}{(1-\bar{z})^N}\right]$, which guarantees that auditing is not prohibitively expensive. Assume without loss of generality that financiers are the ones who pay the auditing cost, collect all payments, and pass on the tax revenues to the government, and denote

[22] As $\varphi_{s+1} + \varphi_{s-1} - 2\varphi_s \leq 0 \Leftrightarrow \frac{\varphi_{s+1} - \varphi_s}{\varphi_s - \varphi_{s-1}} \leq 1$ and $\frac{(s+1)(N-s+1)}{(N-s)s} > 1$ for any s.

the expected tax revenues with

$$\boldsymbol{E}_\tau \equiv \mathbb{E}\left[\mathbb{E}\left[\tau_R + (1+\pi)\,\gamma_R\,(\tau_S - \tau_R)\,|R\right]\right].$$

The firm also negotiates the financing contract, $\phi\,(.)$, before it produces. It takes the auditing and tax policy as given and aims to maximize its expected profits

$$\mathbb{E}[S] - \mathbb{E}\left[\mathbb{E}\left[\phi\,(R-\tau_R) + (1+\pi)\,\gamma_R\,(\phi\,(S-\tau_S) - \phi\,(R-\tau_R))\,|R\right]\right] - \boldsymbol{E}_\tau,$$

subject to the auditing policy, γ_R, and the financier's participation constraint

$$\mathbb{E}\left[\mathbb{E}\left[\phi\,(R-\tau_R) + (1+\pi)\,\gamma_R\,(\phi\,(S-\tau_S) - \phi\,(R-\tau_R))\,|R\right]\right] + \boldsymbol{E}_\tau \geq I - G + \boldsymbol{E}_\tau + c\mathbb{E}[\gamma_R].$$

The last inequality simply requires that financiers accept the contract only if the expected payments they receive exceed the funds they lend, i.e., the funds needed in excess of the investment subsidy G, and pass on to the government and the auditor.

3.1.3 Equilibrium Conditions

The solution method uses backward induction and the concept of Perfect Bayesian Equilibrium (PBE), which requires both sequential rationality and consistency of beliefs. Throughout the analysis, it is assumed that the firm has full bargaining power. The results in Section 4 will only consider pure strategies for the game between consumers and the firm. The belief structure is rather simple: each consumer type $\{v_n, T_n\}$ is either believed to buy the product for sure, or not. All agents are assumed to share the same beliefs. The extension to mixed strategies is considered in Section 5.1.[23] It must also hold in the equilibrium that the realized sales at date 1 are equal to consumers' purchases

$$S = Np\sum_{n=0}^{N} b_{\{v_n,T_n\}},$$

and this will endogenously determine the distribution of sales, h_s, and H_s that will be correctly anticipated on the equilibrium path. The date 0 auditing and financial contracts solve the problem specified in Section 3.1.2.

Finally, the taxes are set to implement the efficiency condition $\boldsymbol{E}_\tau = G$ before the game starts. The welfare criterion used is the joint surplus of the firm and consumers as this setting considers

[23] See also Supplementary Appendix B.2 that considers all belief structures under pure and mixed strategies in the non-private payment system and derives belief structures that are consistent with PBE and the PBE that exist.

a partial equilibrium. Furthermore, as taxes paid by the firms are equal to the transfers to the firm, the tax policy does not affect investments and joint surplus directly. However, it may affect optimal auditing and financing contracts.

4 Results: The Trade-Off and Asymmetric Privacy

Before equilibrium analysis, note the frictionless benchmark: all projects with $I \leq N\bar{z}$ are worth pursuing. Provided this, the expected joint surplus between the firm and consumers is $JS^{benchmark} = N\bar{z} - I$, and this surplus is tax-neutral because expected taxes from the firm equal the government's subsidy. Auditing costs do not affect the first best, as auditing is only needed due to reporting frictions.

4.1 Symmetric Privacy of Payments Data, i.e., Traditional Private Cash-Based System

Suppose that consumers use physical cash (or an equivalent privacy-preserving digital token money) to buy the product from the firm. As these transactions do not reveal information about the consumer's character type to third parties, the following lemma characterizes the outcome of the interaction between the firm and consumers:

Lemma 4.1. *Provided that the firm is the one setting the price and will only operate when it makes non-zero profits, there exists a unique Perfect Bayesian Equilibrium (PBE) in pure strategies in the product market where all consumers with $v_n = 1$ buy the product, and none of the consumers with $v_n = 0$ buy the product. The firm sets the price $p = 1$.*
The firm's sales $S = \sum_{n=1}^{N} b_{\{v_n,T_n\}} = \sum_{n=1}^{N} v_n$ distribution is Binomial $(N, \bar{z})$, i.e.,

$$h_S = \frac{N!}{S!\,(N-S)!}\bar{z}^S\,(1-\bar{z})^{N-S}\,, \tag{9}$$

where $S \in [0, 1, \ldots, N]$, $\mathbb{E}[S] = N\bar{z}$, where $\bar{z}$ is defined in (2).

Proof. As there is no digital data that could be used to make inferences about consumers' character types, their indirect utility in (4) is the same for any character type T_n, i.e., $b_{\{v_n,T_n\}} = b_{\{v_n\}}$. It follows that $V\left(v_n, T_n, b_{\{v_n,T_n\}}; \mathcal{B}\right) = V\left(v_n, b_{\{v_n\}}\right)$, where $V\left(v_n, b_{\{v_n\}} = 1\right) = v_n - p$, and $V\left(v_n, b_{\{v_n\}} = 0\right) = 0$. As only consumers with $v_n = 1$ are willing to buy the product at any positive price, it is optimal for the firm to only sell to these consumers and to extract all consumer surplus by setting $p = 1$. The equilibrium sales distribution then follows from distributional assumptions in Section 3.1.1 as each target consumer buys with unconditional probability $\bar{z}$. □

The product market equilibrium in Lemma 4.1 is standard: consumers are willing to buy the product as long as the price is no higher than their consumption value from the product. It is optimal for the firm to only sell to consumers with high valuations and to extract all consumer surplus. The frictions in this environment arise due to taxation, fundraising, and auditing costs.

Building on Propositions 1-3 in Sanchez and Sobel, 1993, the optimal auditing policy takes the form[24]

$$\gamma_R = \begin{cases} \frac{1}{1+\pi} \text{ if } R < \bar{R} \\ 0 \text{ if } R \geq \bar{R} \end{cases}, \tag{10}$$

and it is optimal and incentive compatible for the firm to report

$$R = \begin{cases} S \text{ if } S \leq \bar{R} \\ \bar{R} \text{ if } S > \bar{R} \end{cases}. \tag{11}$$

That is, there exists an auditing threshold below which the firm's reports are optimally audited with a positive probability, while the highest reports are not necessarily audited. Consequently, it is optimal for the firm to report truthfully whenever its realized revenues are below this threshold and to report the minimum non-audited outcome when its revenues are higher.

While the full proofs do not need to be replicated, some key elements of the proof are worth re-emphasizing. First, it is never optimal for the firm to report $R > S$, as by (7) there is no reward from over-reporting and it leads to higher costs for the firm. Second, it is not optimal to set the auditing probability $\gamma_R > \frac{1}{1+\pi}$, as $\gamma_R = \frac{1}{1+\pi}$ already ensures that the firm prefers to report truthfully. Specifically, the difference between reporting $R = S$ and $R < S$ is $(\varphi_S - \varphi_R)(1 - (1+\pi)\gamma(R))$, which is non-negative whenever $\gamma(R) \leq \frac{1}{1+\pi}$ as $\varphi(.)$ is increasing. Thus the firm would already weakly prefer to report truthfully when audited with a probability of $\frac{1}{1+\pi}$, and auditing more frequently would only increase auditing costs. Third, it is easy to verify that given the above auditing policy, the firm cannot do better than to report according to the above-stated reporting policy (11). Finally, because there are greater gains from under-reporting a greater amount if the firm is not audited, it is intuitive that lower reports are the ones that need to be audited more frequently.

This implies that the optimal auditing policy boils down to the question of whether it is some-

[24] The only minor difference in this context is that Sanchez and Sobel, 1993 consider continuous distribution, while this paper considers a discrete distribution.

times optimal to tolerate some under-reporting at the highest sales outcomes instead of inducing truth-telling for all S?

Lemma 4.2. *The optimal auditing policy sets the auditing threshold*

$$\bar{R}^{*} = \max\left\{ \bar{R} \in \mathbb{Z}^{+} |\, \left(\varphi_{\bar{R}} - \varphi_{\bar{R}-1}\right) \frac{1 - H_{\bar{R}-1}}{h_{\bar{R}-1}} - \frac{c}{1+\pi} \geq 0 \right\}, \tag{12}$$

the threshold $\bar{R}^{*}$ is unique and weakly decreasing in c.

Proof. See Appendix A.1. □

Lemma 4.2 states that the optimal auditing policy is uniquely determined for any $\varphi\left(.\right)$ that satisfies minimal and realistic conditions (increasing and not too convex). The optimal auditing threshold, $\bar{R}^{*}$, is a positive integer equal to the highest value $\bar{R}$ for which the marginal extra revenues for the tax authority and financiers, multiplied with the inverse hazard rate, exceed the effective auditing costs. Lemma 4.2 also re-emphasizes insights from previous optimal taxation literature - when auditing costs are high enough, it is optimal to tolerate some non-compliance and tax evasion at higher revenues.

In what follows, consider that the auditing policy is always optimally set and denote the total funds raised under the optimal auditing policy with $\Phi\left(\bar{R}^{*}\right)$, where the function $\Phi\left(.\right)$ is defined in (8) and expressed in simpler form using (10) and (11) in (17) in Appendix A.1. It is also useful to define two additional thresholds:

Corollary 4.2.1. *There is an optimal auditing threshold $\bar{R}^{*}_{\tau}$ that is needed for covering the firm's tax obligations only, given by*

$$\bar{R}^{*}_{\tau} \equiv \left\{ \bar{R} \in \mathbb{Z}^{+} |\, \left(\tau_{S} - \tau_{S-1}\right) \frac{1 - H_{S-1}}{h_{S-1}} - \frac{c}{1+\pi} \geq 0 \right\}, \tag{13}$$

and an optimal auditing threshold $\bar{R}^{*}_{m}$ that corresponds to the firm pledging all its audited revenues to external parties, i.e., corresponding to the situation where $\varphi_{S} = S$ for any $S < \bar{R}^{*}_{m}$, given by

$$\bar{R}^{*}_{m} \equiv \left\{ \bar{R} \in \mathbb{Z}^{+} | \frac{1 - H_{S-1}}{h_{S-1}} - \frac{c}{1+\pi} \geq 0 \right\}. \tag{14}$$

The optimal auditing thresholds satisfy $\bar{R}^{*}_{\tau} \leq \bar{R}^{*} \leq \bar{R}^{*}_{m} \leq N$.

Proof. Follows from Lemma 4.2. See also Appendix A.1. □

Corollary 4.2.1 shows that there is a minimum auditing threshold, $\bar{R}\tau^{*}$, already needed for tax collection. As the tax policy is set exogenously, it may not have been designed to minimize auditing costs. Auditing required for tax purposes may also benefit external financiers. Consequently, the financing contract that the firm offers, $\phi(.)$, may or may not affect $\bar{R}^*$. Corollary 4.2.1 also highlights that when $\bar{R}\tau^* \leq \bar{R}\tau^* < N$, there may be some anticipated and optimally tolerated tax evasion, i.e., when $\tau(.)$ is increasing for revenues that exceed $\bar{R}^*$. Note that the firm must have at least enough expected revenues to cover its tax obligation. However, as there is a government investment subsidy, G, that an active firm invests in the project and which covers $\mathbf{E}\tau$ in equilibrium. The threshold $\bar{R}m^*$ is useful for identifying the maximum investment cost that could be covered in this environment.

The following lemma and the associated corollary explore the conditions and financing contracts (that satisfy limited liability) under which raising external financing to cover $I > 0$ is feasible and the the firm can operate.

Lemma 4.3. *The conditions for raising external financing and optimal contracts are the following:*
1) There exists a threshold for the auditing cost

$$c \leq c_{min} = (1+\pi)(\tau_N - \tau_{N-1}) \frac{\bar{z}}{N(1-\bar{z})}, \tag{15}$$

which determines whether the optimal contract is debt-like or there are many equivalent optimal financing contracts. When this condition holds, it is always optimal to audit all but the highest report, i.e., $\bar{R}^*_\tau = \bar{R}^* = \bar{R}^*_m = N$.
2) The firm can raise external financing as long as its investment cost,

$$I \leq \bar{I}^{symm\_privacy} = \Phi\left(\bar{R}^*_m\right) - H_{\bar{R}^*_m - 1} \frac{c}{1+\pi},$$

where $\bar{R}^*_m$ *is defined in (14). When* $I = \bar{I}^{symm\_privacy}$ *and* $c > c_{min}$*, the only optimal contract is a debt contract where* $\phi(S - \tau_S) = S - \tau_S$ *for any* $S \leq \bar{R}^*_m$*, and* $\phi(S - \tau_S) = \bar{R}^*_m - \tau_{\bar{R}^*_m}$ *when* $S > \bar{R}^*_m$.
3) When $0 < I < \bar{I}^{symm\_privacy}$ *and* $c > c_{min}$ *is high enough, it remains optimal for the firm to offer a debt-like contract*[25] *to reduce the equilibrium auditing threshold* $\bar{R}^*$.
4) When $0 < I < \bar{I}^{symm\_privacy}$ *and* $c \leq c_{min}$*, there are financing contracts (including simple debt and equity) that are feasible and optimal.*

[25] Because sales and thresholds are integers, there can be a few equivalent debt-like contracts that are equivalent. This is just a consequence of rounding. See Appendix A.2.

5) The firm's expected profit is

$$\Pi^{symm\_privacy} = N\bar{z} - I - H_{\bar{R}^*-1}\frac{c}{1+\pi},$$

where $\bar{R}^*$ is the auditing threshold under the best debt contract. When $c \leq c_{min}$, then

$$\Pi^{symm\_privacy}_{c\leq c_{min}} = N\bar{z} - I - \left(1-\bar{z}^N\right)\frac{c}{1+\pi}.$$

Proof. See Appendix A.2. □

Lemma 4.3 highlights the known insight that high enough auditing costs favor debt contracts, because under debt contracts auditing is only needed at low outcomes. At the same time, auditing costs also prevent all optimal investments from being undertaken. For later analysis, it is also worth noting that auditing costs just need to be small, but not necessarily zero to enable a broad set of financing contracts to be optimal. This is because tax policies, which are exogenous to the model, tend to have non-zero marginal tax rates. When auditing cost is small, it not only helps to reduce tolerance for tax evasion at the highest revenues, but also widens the set of optimal financing contracts which benefit from auditing that is needed for taxation purposes.

Corollary 4.3.1. *When $0 < I \leq \bar{I}^{symm\_privacy}$, and $c \leq c_{min}$, the optimal financing contract can be expressed as a simple equity contract that sets $\phi\left(S-\tau_S\right) = \bar{\phi}_E \cdot \left(S-\tau_S\right)$ for any S, and*

$$\bar{\phi}_E = \frac{I - G + \left(1-\bar{z}^N\right)\frac{c}{1+\pi}}{N\bar{z} - \boldsymbol{E}_\tau} \leq 1.$$

This corollary helps to highlight some key effects that are also present in Lemma 4.3. Higher auditing costs force the firm to pledge a greater share of their revenues. Naturally, the taxes also increase the equity that the firm has to offer to the investors.[26]

As lower auditing costs reduce optimal tax avoidance, it further enables the government to set lower taxes to collect the same expected revenues. To illustrate this, assume that the taxes are proportional, i.e., $\tau_s = \bar{\tau}\cdot s$, where $\bar{\tau}$ is a constant. When the auditing cost is low, i.e., $c \leq c_{\text{min}} = (1+\pi)\,\bar{\tau}\frac{\bar{z}}{N(1-\bar{z})}$, the tax rate that covers the cost G in expectations is $\bar{\tau} = \frac{G}{\mathbb{E}[S]} = \frac{G}{N\bar{z}}$, and no one firm has incentives to avoid taxes. In contrast, when c is high and the optimal auditing threshold $\bar{R}^* < N$, the tax rate that covers the cost G in expectation is $\bar{\tau} = \frac{G}{\Pr\left(S\leq\bar{R}^*\right)\mathbb{E}\left[S|S\leq\bar{R}^*\right]+\bar{R}^*\Pr\left(S>\bar{R}^*\right)} =$

[26] While the negative effect of taxes is offset by the corresponding investment subsidy in this setting, this is due to the assumption that taxes collected from the firm are distributed to the firm. Should the tax revenues be shared with other agents in the economy, the negative tax effect would dominate and there would be additional real and distributional effects.

$\frac{G}{\mathbb{E}[S]-\Pr\left(S<\bar{R}^*\right)\mathbb{E}\left[S-\bar{R}^*|S>\bar{R}^*\right]}$, which is clearly higher and comes with optimal tax underpayment by firms whose revenues are higher than $\bar{R}^*$.

For comparisons with the later setting, the following proposition holds.

Proposition 4.4. *Investments that can be undertaken must satisfy*

$$I \leq I^{symm\_privacy} = N\bar{z} - H_{\bar{R}^*-1}\frac{c}{1+\pi}$$

and the joint surplus between the firm and consumers is

$$JS^{symm\_privacy} = N\bar{z} - I - H_{\bar{R}^*-1}\frac{c}{1+\pi},$$

where the optimal $\bar{R}^*$ is defined above.

Proof. By Lemma 4.1 the consumer surplus is zero, which implies that the joint surplus equals the firm's profit given in Lemma 4.3. □

4.2 Symmetric lack of privacy of payments data, i.e., publicly shared records

Let us consider the other extreme, where all payments are recorded on an observable digital ledger, and there are digital records of every consumer's purchase of the good/service produced by the firm.

A common argument among the skeptics against a central bank digital currency is that it violates individual privacy. This section highlights that this consideration indeed has a real negative impact on the firm-consumer relationship even when the consumers' privacy sensitive traits are not of direct interest of the firm. Let us start again from the equilibrium in product market.

Lemma 4.5. *Provided that the firm is the one setting the price and will only operate when it makes non-zero profits, and $0 \leq \epsilon \leq \frac{1}{2} \cdot \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$, there exists a unique Perfect Bayesian Equilibrium (PBE) in pure strategies in the product market where all consumers with $v_n = 1$ buy the product, and none of the consumers with $v_n = 0$ buy the product. The firm sets the price*

$$p = 1 - \epsilon\frac{z_A - z_B}{\bar{z}\left(1-\bar{z}\right)},$$

and the firm's sales, S, distribution is Binomial $(N, \bar{z})$, i.e., $h_S = \frac{N!}{S!(N-S)!}\bar{z}^S\left(1-\bar{z}\right)^{N-S}$, where $S \in [0, p, ..., pN]$, $\mathbb{E}\left[S\right] = pN\bar{z}$, where $\bar{z}$ is defined in (2).
If $\epsilon > \frac{1}{2} \cdot \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$ there is no PBE in pure strategies where the firm operates.

Proof. See Appendix □

Proposition 4.6. *Investments that can be undertaken must satisfy*

$$I \leq I^{no\_privacy} = N\bar{z} - N\epsilon\frac{z_A - z_B}{(1-\bar{z})}$$

and the joint surplus between the firm and consumers is

$$JS^{no\_privacy} = N\bar{z} - I - 2N\epsilon\frac{(z_A - z_B)^2\, q_A\,(1-q_A)}{\bar{z}\,(1-\bar{z})}.$$

External financing can be obtained under a wide set of contracts including simple equity,

Proof. See Appendix A.3 and Supplementary Appendix B.2 □

Lemma 4.5 and Proposition 4.6 highlight that privacy considerations create distortions in the product market, necessitating firms to offer discounts when consumers play pure strategies to compensate them for the negative effects of third parties learning about their private traits. This, in turn, reduces the firms' profits and joint surplus.

At the same time, public records of individual purchases enable firms to have better access to external financing and multiple financing contracts, as there are no auditing frictions. Proposition 4.6 further indicates that the interaction between the firm and its consumers may break down when ϵ is high. In such cases, there are still PBE in mixed strategies (see Section 5.1). However, mixed strategies maintain a similar trade-off, as they imply a loss of revenue due to the firm losing some customers who strategically do not buy to maintain their privacy.

Proposition 4.6 also shows that there is under-investment compared to the first best. Additionally, as the firm in this setting cannot price-discriminate between different consumer types, consumer surplus is not zero, and privacy considerations affect different consumer types differently. Inefficiencies are higher for firms with a larger target market, N, and a more diverse consumer base, i.e., when the difference between z_A and z_B is larger.

4.3 The trade-off under symmetrically private versus symmetrically non-private payment systems

The comparison between $JS^{symm\_privacy}$ with $JS^{no\_privacy}$ implies a trade-off between the usefullness of the digital data in mitigating financing frictions and the firm's losses of revenues. Which environment is better, depends on the verification cost, c, and the importance of privacy concerns, which in turn depend on ϵ and the difference in consumers propensity to buy $(z_A - z_B)$, i.e., how diverse the firms target consumers are. It is also worth noting that $JS^{no\_privacy} - JS^{symm\_privacy} = H_{\bar{R}^*-1}\frac{c}{1+\pi} - 2N\epsilon\frac{(z_A - z_B)^2 q_A(1-q_A)}{\bar{z}(1-\bar{z})}$ is more likely to be negative (which implies that a payment system that does not preserve privacy is worse) if N is higher: the monitoring cost $H_{\bar{R}^*-1}\frac{c}{1+\pi}$ is bounded, while the importance of privacy concerns increases linearly with the size of the firm's target market.

4.4 Asymmetric privacy of payments data/P-hybrid CBDC

Suppose that there is a digital payment system such that the firm's sales data are digitally recorded, while the individual's purchasing decisions are not associated with their identity (the possible ways to create and implement such P-Hybrid CBDC is discussed in Section 6). This setting is not devoid of secondary effects on privacy, but these are greatly mitigated. Namely, there is still some learning about consumers' character types as aggregate sales data still enables learning about the prevalence of different consumer types in the population. Furthermore, when the target population of the firm's product is small, rational consumers are aware that their decision to buy or not to buy affects the aggregate sales, which affects their willingness to pay. However, as their individual decisions are no longer recorded, these effects are much smaller and become insignificant when N is large. At the same time, this system still enables frictionless contracting between the investors and the firm, and more efficient tax collection.

Proposition 4.7. *Whenever* $0 \leq \epsilon < \frac{N}{2} \cdot \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$, *there exists a unique PBE in pure strategies where all consumers with* $v_n = 1$ *buy the product, and none of the consumers with* $v_n = 0$ *buy the product. The firm sets the price*

$$p = 1 - \epsilon\frac{z_A - z_B}{N\bar{z}\,(1-\bar{z})}$$

Investments that can be undertaken satisfy

$$I \leq I^{assym\_privacy} = N\bar{z} - \epsilon\frac{z_A - z_B}{(1-\bar{z})}$$

and the aggregate surplus is

$$JS^{assym\_privacy} = N\bar{z} - I - 2\epsilon\frac{(z_A - z_B)^2\, q_A\,(1-q_A)}{\bar{z}\,(1-\bar{z})}$$

Whenever, $\epsilon > \frac{1}{2} \cdot \frac{\bar{z}(1-\bar{z})}{N(z_A - z_B)}$ *there is no PBE in pure strategies where the firm prefers is active.*

Proof See Appendix A.3.

The stark difference between Proposition 4.6 and 4.7 is that whenever $N > 1$, all distortions of privacy concern are mitigated, and therefore the outcomes under this setting always welfare dominate the outcomes under no privacy. Furthermore, as $N \to \infty$, the necessity of the firm to offer discounts disappears. This highlights that a privacy asymmetric CBDC, if widely adopted, enables frictionless financing contracts, while it limits the drawback of unnecessary information being recorded.

5 Extensions of the basic model

5.1 Consumers using mixed strategy for privacy

Sections 4.2 and 4.4 analyzed the real effects of privacy concerns by focusing on pure strategies. The Supplementary Appendix B.2 derives all PBE in pure and mixed strategies. It proves that in addition to the pure strategy equilibrium that only exists when privacy concerns are not too pressing, (i.e., when $\epsilon \leq \frac{1}{2} \cdot \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$), there exists PBE in mixed strategies for any value of ϵ. These equilibria differ from the one analyzed above in that some consumers who value the product do not always buy it. Namely, at least the consumer type $\{v_n, T_n\} = \{1, A\}$ buys the good with probability $\sigma < 1$ and is believed to do so on the equilibrium path.

Proposition B.5 the th Supplementary Appendix B proves that the welfare-dominant PBE in this class requires that consumer type $\{v_n, T_n\} = \{1, A\}$ buys the good with probability $\sigma = \frac{z_A}{z_B}$, while the type $\{v_n, T_n\} = \{1, B\}$ always buys the product, and the types $\{v_n, T_n\} = \{\{0, A\}, \{0, B\}\}$ never buy it. This randomization strategy allows the consumer type $1, A$ and all other consumers to perfectly hide their type, and the firm optimally sets the price at $p = 1$.

The Supplementary Appendix B.2 further shows that there exists a continuum of similar mixed strategy equilibria, where $p = 1$, the type $\{v_n, T_n\} = \{1, B\}$ buys the product with probability $\sigma_{1,B} \in (0, 1]$, and the type $\{v_n, T_n\} = \{1, A\}$ buys the product with probability $\sigma_{1,A} \in \sigma_{1,B} \frac{z_A}{z_B}$. The Supplementary Appendix also includes the proof that these mixed strategy equilibria and the PBE in pure strategies in Lemma 4.5 are the only PBE with $p > 0$ in this game.

While the firm does not need to offer discounts, there is still a welfare loss as the expected total demand is lower. Namely,the firm's profits and the joint surplus are at most $N z_B - I = N\bar{z} - N q_A (z_A - z_B) - I$. Consequently, there are still welfare losses and under-investment compared to the first best. Furthermore, it holds that $N z_B - I < JS^{\textit{no-privacy}}$ defined in Proposition 4.6 whenever the PBE in pure strategies exists, and thus the PBE with discounts welfare dominates the one with mixed strategies whenever it exists.[27]

[27] To see this, note $N z_B - I < JS^{no\text{-}privacy} \Leftrightarrow 2\epsilon \frac{(z_A - z_B)^2 q_A (1 - q_A)}{\bar{z}(1-\bar{z})} < \bar{z} - z_B = \epsilon < \frac{1}{2} \frac{1}{(1-q_A)} \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$, which holds for any $\epsilon \leq \frac{1}{2} \frac{\bar{z}(1-\bar{z})}{(z_A - z_B)}$.

A key message from this analysis is that interacting with privacy-concerned individuals whose purchase of a particular product is public information implies losses for a firm, either via offering noticeable discounts or via accepting to lose some consumers who may have an incentive to alter their consumption behavior.

5.2 Micro-founded rationale for assumed privacy preferences

As in the main model, consider that consumer n has privacy sensitive traits $T_n \in \{A, B\}$ and may derive utility from consuming an indivisible good.

Denote the purchasing decision of the consumer with $b_{n,T_n} \in \{0, 1\}$ and assume that the consumer n has quasilinear preferences, i.e.,

$$U_n(b_{n,T_n}) = u_n(b_{n,T_n}) + m_n,$$

where m_n is the monetary value of consumer n's other consumption and assets held, and $u_n(b_{n,T_n})$ is the consumption value of the product for this consumer. Assume that

$$u_n(b_{n,T_n}) = \begin{cases} u_n(1) > 1 & \text{if the consumer } n \text{ likes the product} \\ u_n(0) & \text{otherwise} \end{cases}.$$

There is a potential "thief" who aims to rob all agents, and the success of the theft depends on both the victim's type and the thief's type. The thief is specialized in targeting either types A or types B, and not both at the same time. Namely, the theft will fail when targeting the wrong type (e.g., suppose that the "thief" is an airline company, and the good is a dinner at an expensive restaurant, which may indicate an opportunity to charge a higher price from a consumer who seems willing to spend at restaurants and for traveling; but if the consumer is not interested in traveling, the attempt will not succeed). The probability that there is a thief is x, and the thief can be of type $T = \{A, B\}$, with both thief types being equally likely. Denote the event

$$\chi_{T_n,T} = \begin{cases} 1 & \text{if there is a thief that targets the consumer type } T_n = T \\ 0 & \text{if there is no thief, or } T \neq T_n \text{ as the thief does not have the matching type} \end{cases}.$$

The budget constraint of the consumer n is

$$w_n = p b_{n,T_n} + m_n - \chi_{T_n,T}\delta_{n,T},$$

where w_n is the wealth of consumer n and δ_T is the amount that the thief of type T will choose to steal from consumer n.

Assume that robbing is costly, and robbing more is increasingly costly, i.e., trying to rob an amount δ from one person costs $\kappa(\delta)$, where $\kappa(0) = 0$, $\kappa' > 0$, and $\kappa'' > 0$. For the sake of argument, assume that the cost of theft is quadratic $\kappa(\delta) = \kappa \frac{1}{2}\delta^2$.

First, suppose that the thief knows nothing apart from the prior. To decide the optimal amount to be stolen, a type T thief tries to steal the same amount from all consumers, i.e., $\delta_{n,T} = \delta_T$ for all n, and he/she chooses δ_T to maximize

$$\Pr(T_n = T)\delta_T - \frac{\kappa}{2}(\delta_T)^2,$$

which implies that the optimal amount stolen is $\delta_T = \frac{\Pr(T_n=T)}{\kappa}$, and the total amount stolen is $N\frac{\kappa}{2}(\Pr(T_n = T))^2$.

The consumer anticipates the optimal robbery decision, and his/her expected indirect utility is

$$V_n = \begin{cases} u_n(1) + w_n - p - \frac{x}{2\kappa}\Pr(T_n = T) & \text{if } b_{n,T_n} = 1 \\ u_n(0) + w_n - \frac{x}{2\kappa}\Pr(T_n = T) & \text{if } b_{n,T_n} = 0 \end{cases}.$$

Normalizing $u_n(0) = -w_n + \frac{x}{2\kappa}\Pr(T_n = T)$, and $u_n(1) = u_n(0) + v_n$, this corresponds to

$$V_n = \begin{cases} v_n - p & \text{if } b_{n,T_n} = 1 \\ 0 & \text{if } b_{n,T_n} = 0 \end{cases}.$$

Now consider that the thief observes whether the consumer bought the product, or not, and the purchasing decisions reveal information about T_n as in the main setting. The optimal amount stolen can now be conditioned on the purchasing decisions, and when choosing the amount to be stolen from consumer n, the thief solves

$$\Pr(T_n = T|b_{n,T_n})\delta_{n,T} - \frac{\kappa}{2}(\delta_{n,T})^2,$$

which implies that the optimal amount stolen is $\delta_{n,T} = \frac{\Pr(T_n=T|b_{n,T_n})}{\kappa}$.

The expected indirect utility of the consumer who anticipates this is now

$$V_n = \begin{cases} u_n(1) + w_n - p - \frac{x}{2\kappa}\Pr(T_n = T|b_{n,T_n} = 1) & \text{if } b_{n,T_n} = 1 \\ u_n(0) + w_n - \frac{x}{2\kappa}\Pr(T_n = T|b_{n,T_n} = 0) & \text{if } b_{n,T_n} = 0 \end{cases},$$

and using the same normalization for the model's parameters

$$V_n = \begin{cases} v_n - p - \frac{x}{2\kappa}\left(\Pr(T_n = T|b_{n,T_n} = 1) - \Pr(T_n = T)\right) & \text{if } b_{n,T_n} = 1 \\ -\frac{x}{2\kappa}\left(\Pr(T_n = T|b_{n,T_n} = 0) - \Pr(T_n = T)\right) & \text{if } b_{n,T_n} = 0 \end{cases}.$$

Furthermore, using that $\Pr(T_n = T) + \Pr(\overline{T_n} = T) = 1$, it holds that

$$\tfrac{\Pr(T_n=T|b_n)}{\Pr(T_n=T)} - \tfrac{\Pr(\overline{T_n}=T|b_n)}{\Pr(\overline{T_n}=T)} = \tfrac{\Pr(T_n=T|b_n)-\Pr(T_n=T)}{\Pr(T_n=T)\Pr(\overline{T_n}=T)}.$$

Replacing $\Pr(T_n = T|b_n) - \Pr(T_n = T)$ to the above, and defining the privacy preference parameter as

$$\epsilon \equiv \frac{x}{2\kappa}\Pr(T_n = T)\Pr(\overline{T_n} = T),$$

we obtain the assumed indirect utility in the main setting, see 4 for the case where the information set Ω includes observing every individual's purchases, and

$$V_n = \begin{cases} v_n - p - \epsilon\left(\frac{\Pr(T_n=T|b_{n,T_n}=1)}{\Pr(T_n=T)} - \frac{\Pr(\overline{T_n}=T|b_{n,T_n}=1)}{\Pr(\overline{T_n}=T)}\right) & \text{if } b_{n,T_n} = 1 \\ -\epsilon\left(\frac{\Pr(T_n=T|b_{n,T_n}=1)}{\Pr(T_n=T)} - \frac{\Pr(\overline{T_n}=T|b_{n,T_n}=1)}{\Pr(\overline{T_n}=T)}\right) & \text{if } b_{n,T_n} = 0 \end{cases}.$$

This extension provides possible interpretations for the magnitude of ϵ. Privacy concerns are higher if the success rate of theft by external parties, x, is higher and the cost of a robbery attempt, κ, is lower. Naturally, a similar derivation applies when the information external parties have about the individual is less precise, i.e., only based on total sales as under asymmetric privacy.

5.3 Parallel with Differential Privacy

The literature on differential privacy (see Dwork, Roth, et al., 2014 for a review)[28], adopts the following functional form of privacy concern where ϵ is a parameter that captures how important the privacy concern is:

$$\epsilon \ln\left(\frac{\Pr\left(\text{action}|\text{true\_type}\right)}{\Pr\left(\text{action}|\text{another\_type}\right)}\right). \tag{16}$$

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

Considering the case where all purchasing decisions are observed, this corresponds to:

$$\epsilon \ln \left(\frac{\Pr(b_n|\mathrm{T}_n)}{\Pr(b_n|\mathrm{T}_{-n})} \right) = \epsilon \ln \left(\frac{\Pr(\mathrm{T}_n|b_n)}{\Pr(\mathrm{T}_n)} \cdot \frac{\Pr(\mathrm{T}_{-n})}{\Pr(\mathrm{T}_{-n}|b_n)} \right)$$
$$= \quad \epsilon \left(\ln \left(\frac{\Pr(\mathrm{T}_n|b_n)}{\Pr(\mathrm{T}_n)} \right) - \ln \left(\frac{\Pr(\mathrm{T}_{-n}|b_n)}{\Pr(\mathrm{T}_{-n})} \right) \right),$$

and a linear approximation of this gives the privacy loss in the consumer's indirect utility (4).

An important common theme to note is that this model, as well as the differential privacy literature, considers that not only active actions (e.g., purchasing a good) but also passive actions (e.g., not purchasing a good) can enable external parties to make inferences about individuals who may dislike these inferences being made. The key difference is that differential privacy focuses on adding noise to queries from an existing database. As payment systems and related financial contracts need to involve precise data on total sales, it seems unlikely that these queries would be designed to allow adding randomness. At the same time, there is naturally a degree of randomness in this setting, as purchasing decisions are informative but imperfect signals about consumers' personal traits. Furthermore, consumers themselves can obscure information by playing mixed strategies.

6 Design of the asymmetrically private system

6.1 Transaction records

As shown in Sections 3-5 a desirable payment system features **intentional asymmetry regarding privacy**, between receiving and sending money. While this system could be implemented by the private sector, central banks may be in a better position and have incentives to implement such a system at a scale and simply engage technology firms working with digital ledger technologies and participants from decentralized blockchain communities to provide technology, and/or even engage as system actors to keep some parts of the records and validate transactions. In the context of CBDCs, we call it Privacy-Hybrid or P-Hybrid CBDC.[29] In what follows, we refer to CBDC, however, the same arguments would apply if a privacy asymmetric and commonly used payment system emerged from the private sector.

From an ideal technological point of view, a desirable P-Hybrid CBDC could aim to make

[29] Our term P-hybrid CBDC should not be confused with the term "hybrid CBDC" in *e.g.* Auer and Böhme, 2020, where it means that the central bank delegates retail payments to intermediaries while remaining the issuer of CBDC. In the case of P-hybrid CBDC the digital currency itself has a hybrid nature and treats money received and paid differently. Our proposed system enables both "direct" and "hybrid" CBDC as defined in Auer and Böhme, 2020, as the central bank can manage all functions the system architecture requires, or delegate some of these functions.

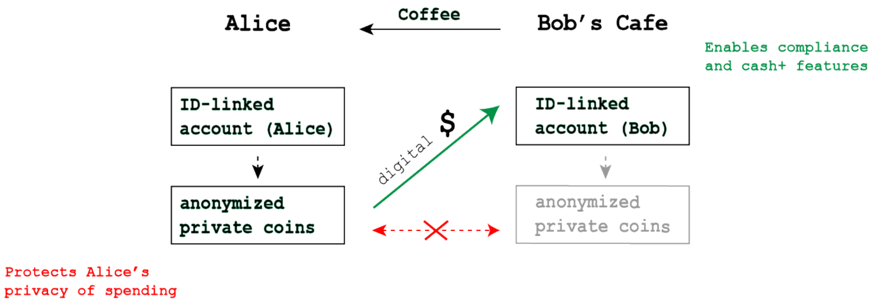

Figure 2: Illustration of a transaction using P-Hybrid CBDC

it technologically impossible (or more precisely, statistically close to impossible) for any party, including the keepers of the digital ledger used to associate an individual with their purchases of coffee, groceries, alcohol, entertainment, medicine, etc. At the same time, at least the parties who need to know should be able to observe the identity of the receivers of money, as it reduces auditing costs and enables more efficient taxation and external financing, and why not a more efficient implementation using "smart contracts" (i.e., incorporating programmable features).

To understand what makes the P-hybrid CBDC distinct from pure token or account money consider an example of a typical transaction as an illustration.

Figure 2 represents a transaction and associated records between Alice and Bob's Cafe for the purchase of a coffee. Alice and Bob have both gained access to the P-Hybrid CBDC system and its functionalities by creating beforehand an ID-linked account which is uniquely associated with their identity. However, before Alice buys her coffee, she has transferred some of her money to her "anonymized wallet" (discussion of technological possibilities on this follow in Section 6.2). No outside parties (i.e., not Bob, not her bank, no outside eavesdropper, and not even the system actors) are able to digitally observe and establish that it was Alice who bought the coffee. At the same time, all those who need to do so (e.g., system actors, tax authorities, anti-money laundering authorities, authorized firms who provide services to Bob at his consent) will observe that Bob received the digital payment from someone who used their private coins/"anonymous wallet". Furthermore, as private coins can only be sent to ID-linked accounts, this overcomes one weakness of physical cash and anonymous digital money as a tool for money laundering and terrorist financing. ID-linked accounts are also useful for making it harder to use the P-Hybrid CBDC as means of payment for other illegal goods and services.

6.2 Tools for anonymized wallets

One straightforward way to implement asymmetric privacy and anonymized wallets is to use a prepayment card that does not have identity data or a "name" on the card.[30] Such a card was implemented in the 1990s by the Bank of Finland, called Avant card[31] (see Grym, 2020). Such a system would easily implement the key features on Figure 2, where loading money on the nameless card protects the consumer's privacy, while the funds received via payment using such a card could be integrated with digital ledger technology to facilitate building on payments' data. Of course more contemporary technologies for cards, or QR-code-based payments would further enable inbuilt features that only allow authorized parties to accept transactions.

Creating anonymized wallets is possible also with purely digital forms of payment systems, which benefit from the use of blockchain (or another digital ledger). It is worth noting that many well known Blockchain systems, such as Bitcoin, are in fact pseudonymous. The identity of the users is not directly visible on the network, since users use anonymous addresses — the pseudonym — to send or receive transactions. However, despite these precautions, many techniques have been proposed to effectively de-anonymized Bitcoin users (Meiklejohn et al., 2016). Furthermore, pseudonymous accounts would not enable to harness the benefits of reduced auditing costs, and this can explain why much lending activity on decentralized blockchains is "over collateralized", i.e., includes particularly large amounts of collateral compared to typical loans as opposed to needing less collateral that technology enables efficiency gains allow. Another way to disassociate the digital coin from it's owner's identity from is to use "shuffling services" or "mixing services". These services connect groups of users wishing to mix their coins together, and enables users to coordinate to make a transaction that mixes their coins together and produces a set of new fresh coins.[32] The drawback of using shuffling services is that the level of privacy achieved is dependent on the number of users involved in the shuffling.

A technologically more sophisticated and possibly superior approach uses "Zero-Knowledge proofs." The key general idea behind Zero-Knowledge proofs is that it is a privacy tool that enables only necessary information to be communicated. For example, it could be used by someone to prove they know a secret code without directly communicating this code to anyone. In the context

[30] This idea is also discussed in Kahn and Roberds, 2008.

[31] Even though Avant Card project was ultimately not successful, it cannot be considered as an indication that asymmetric privacy is not valuable as data processing capabilities in 1990's were considerably lower, which likely made both privacy concerns less pressing, and did not come with a rich set of possibilities to create smart contracts for more efficient taxation or financing. Avant card's fee structure compared to bank cards has also been mentioned as a reason why this early attempt for a CBDC did not succeed.

[32] The very first generation of such services were not trustless since it was possible for the server to steal the users' funds. The second generation of approaches, such as CoinShuffle (Ruffing et al., 2014) are built in a trustless manner.

of payments, the main idea is that users *mint* their own coin by adding them to a list of minted coins. When they want to spend their coin, they simply provide a ZK (Zero-Knowledge) proof that the coin they want to spend has been minted, without having to reveal which one it is. The result is that when a coin is used, it could be anyone's coin. This contrasts with shuffling services where it is possible to infer that the coin must belong to the set of users involved in the shuffling at a given time.[33]

Zero-knowledge proofs are already quite widely used in the context of cryptocurrencies and cryptotokens. For example, ZeroCoin (Miers et al., 2013) and ZeroCash (Ben-sasson et al., 2014a) were among the first Zero-Knowledge based approaches applied to blockchain. They are also used in the context of prominent privately created blockchain and oracle systems like Ethereum and Chainlink. While the economic benefits of a CBDC design (or another payment system's) featuring cash-like privacy or asymmetric privacy do not require reliance on a specific method, technological tools enabling the preservation of crucial data *by design* already exist for partially or fully digital payment systems. Moreover, one could argue that *hardcoding* privacy directly into a widely adopted payment system (e.g., a successful retail CBDC) is preferable to merely trusting institutions to maintain consumer data anonymity. Even if an institution is genuinely trustworthy in this regard, cybersecurity vulnerabilities may still exist.

7 Conclusions

This paper argues that P-hybrid CBDC is a possible solution to the fundamental trade-off between protecting cash-like privacy and having digital records that can facilitate regulatory compliance and enable better financing contracts.

The paper highlights the costs of having a system which only considers one extreme. A symmetrically privacy-protecting payment system perpetuates existing frictions in taxation and financial contracting due to auditing costs. Simultaneously, a non-private payment system leads to real distortions in the consumer-producer relationship. Both are suboptimal, and determining which one is worse depends on the comparison of auditing costs and privacy losses.

Instead, a better system involves individuals making their CBDC spending wallet private, and such features are better built by design rather than left to an intermediary to protect. The desirable system should have the feature that neither private sector participants nor central authorities can track precisely how consumers spend their money. Such a design is technologically feasible in many

[33] More details on the potential use of Zero-knowledge proofs in the context of CBDC can be found in Tinn and Dubach, 2021.

circumstances.

Even though designing systems that are symmetric in privacy or lack of privacy is easier, the paper gave examples how an asymmetrically private CBCD (P-Hybrid CBDC) could be implemented using already existing tools from information and communication technology, and cryptography. The proposed system has the potential to ensure three important properties: compliance through the use of ID-linked registered accounts for money received and stored; reduced costs of verification that facilitated more efficient financial contracting and taxation; and privacy, which could be achieved through the use of Zero-Knowledge proofs, or other existing tools.

Finally, even though the paper focuses on CBDC, it is plausible that another private sector entity or community will design such a system. Nevertheless, central banks may be well-positioned to facilitate the design of a system that aims to maximize the joint utility of all parties, i.e., the firm, consumers, investors, and the tax authority.

References

Abowd, John M and Ian M Schmutte (2019). "An economic analysis of privacy protection and statistical accuracy as social choices". In: *American Economic Review* 109.1, pp. 171–202.

Acquisti, Alessandro and Hal R Varian (2005). "Conditioning prices on purchase history". In: *Marketing Science* 24.3, pp. 367–381.

Agur, Itai, Anil Ari, and Giovanni Dell'Ariccia (2024). *Bank competition and household privacy in a digital payment monopoly.* Tech. rep. International Monetary Fund.

Agur, Itai, Anil Ari, and Giovanni Dell'Ariccia (2022). "Designing central bank digital currencies". In: *Journal of Monetary Economics* 125, pp. 62–79.

Ahnert, Toni, Peter Hoffmann, and Cyril Monnet (2022). "The digital economy, privacy, and CBDC". In:

Andolfatto, David (2021). "Assessing the impact of central bank digital currency on private banks". In: *The Economic Journal* 131.634, pp. 525–540.

Andreoni, James, Brian Erard, and Jonathan Feinstein (1998). "Tax compliance". In: *Journal of economic literature* 36.2, pp. 818–860.

Auer, Raphael and Rainer Böhme (2020). "The technology of retail central bank digital currency". In: *BIS Quarterly Review, March.*

Auer, Raphael, Giulio Cornelli, and Jon Frost (2020). "Rise of the central bank digital currencies: drivers, approaches and technologies". In:

Auer, Raphael, Jon Frost, Leonardo Gambacorta, Cyril Monnet, Tara Rice, and Hyun Song Shin (2022). “Central bank digital currencies: motives, economic implications, and the research frontier”. In: *Annual review of economics* 14, pp. 697–721.

Bakos, Yannis and Hanna Halaburda (2019). “Smart Contracts, IoT Sensors and Efficiency: Automated Execution vs. Better Information”. In: *NYU Stern School of Business, available at https://ssrn. com/abstract* 3394546.

Bank of Canada (2023). *A Digital Canadian Dollar: What we heard 2020–23 and what comes next.* Avaliable at https://www.bankofcanada.ca/digitaldollar/a-digital-canadian-dollar-what-we-heard-2020-23-and-what-comes-next/.

Barrdear, John and Michael Kumhof (2022). “The macroeconomics of central bank digital currencies”. In: *Journal of Economic Dynamics and Control* 142, p. 104148.

Ben-sasson, Eli, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza (2014a). “Zerocash: Decentralized Anonymous Payments from Bitcoin”. In: *IEEE Symposium on Security and Privacy (SP).*

— (2014b). “Zerocash: Decentralized Anonymous Payments from Bitcoin (extended version)”. In: DOI: 10.1109/SP.2014.36.

Benhaim, Alon, Brett H Falk, and Gerry Tsoukalas (2023). “Scaling blockchains: Can committee-based consensus help?” In: *Management Science* 69.11, pp. 6525–6539.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri (2020). “On the rise of fintechs: Credit scoring using digital footprints”. In: *The Review of Financial Studies* 33.7, pp. 2845–2897.

Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta (2019). “The blockchain folk theorem”. In: *The Review of Financial Studies* 32.5, pp. 1662–1715.

Biais, Bruno, Agostino Capponi, Lin William Cong, Vishal Gaur, and Kay Giesecke (2023). “Advances in blockchain and crypto economics”. In: *Management Science* 69.11, pp. 6417–6426.

Border, Kim C and Joel Sobel (1987). “Samurai accountant: A theory of auditing and plunder”. In: *The Review of economic studies* 54.4, pp. 525–540.

Brunnermeier, Markus K, Harold James, and Jean-Pierre Landau (2019). *The digitalization of money.* National Bureau of Economic Research.

Brunnermeier, Markus K and Dirk Niepelt (2019). “On the equivalence of private and public money”. In: *Journal of Monetary Economics* 106, pp. 27–41.

Calzolari, Giacomo and Alessandro Pavan (2006). “On the optimality of privacy in sequential contracting”. In: *Journal of Economic theory* 130.1, pp. 168–204.

Cao, Sean, Lin William Cong, and Baozhong Yang (2019). "Financial reporting and blockchains: Audit pricing, misstatements, and regulation". In: *Misstatements, and Regulation (June 2019).*

Capponi, Agostino, Garud Iyengar, Jay Sethuraman, et al. (2023a). "Decentralized finance: Protocols, risks, and governance". In: *Foundations and Trends® in Privacy and Security* 5.3, pp. 144–188.

Capponi, Agostino, Sveinn Olafsson, and Humoud Alsabah (2023b). "Proof-of-work cryptocurrencies: Does mining technology undermine decentralization?" In: *Management Science* 69.11, pp. 6455–6481.

Catalini, Christian and Joshua S Gans (2020). "Some simple economics of the blockchain". In: *Communications of the ACM* 63.7, pp. 80–90.

Cheng, Yuteng and Ryuichiro Izumi (2024). "CBDC: Banking and anonymity". In: *Available at SSRN 4340866.*

Chiu, Jonathan and Seyed Mohammadreza Davoodalhosseini (2023). "Central bank digital currency and banking: Macroeconomic benefits of a cash-like design". In: *Management Science.*

Chiu, Jonathan, Seyed Mohammadreza Davoodalhosseini, Janet Jiang, and Yu Zhu (2023). "Bank market power and central bank digital currency: Theory and quantitative assessment". In: *Journal of Political Economy* 131.5, pp. 1213–1248.

Chiu, Jonathan and Thorsten V Koeppl (2019). "Blockchain-based settlement for asset trading". In: *The Review of Financial Studies* 32.5, pp. 1716–1753.

Choi, Kyoung Jin, Ryan Henry, Alfred Lehar, Joel Reardon, and Reihaneh Safavi-Naini (2021). "A Proposal for a Canadian CBDC". In: *Available at SSRN 3786426.*

Cong, Lin William and Zhiguo He (2019). "Blockchain disruption and smart contracts". In: *The Review of Financial Studies* 32.5, pp. 1754–1797.

Cong, Lin William, Zhiguo He, and Jiasun Li (2021a). "Decentralized mining in centralized pools". In: *The Review of Financial Studies* 34.3, pp. 1191–1235.

Cong, Lin William, Ye Li, and Neng Wang (2021b). "Tokenomics: Dynamic adoption and valuation". In: *The Review of Financial Studies* 34.3, pp. 1105–1155.

Cong, Lin William and Simon Mayer (2022). "The coming battle of digital currencies". In: *The SC Johnson College of Business Applied Economics and Policy Working Paper Series* 2022-04.

Cong, Lin William, Wenshi Wei, Danxia Xie, and Longtian Zhang (2022). "Endogenous growth under multiple uses of data". In: *Journal of Economic Dynamics and Control* 141, p. 104395.

Cunliffe, Jon (2023). *Money and payments: a 'black ships' moment? - speech by Sir Jon Cunliffe at at the Economics of Payments XII Conference at the Federal Reserve Board, Washington*

DC. Avaliable at https://www.bankofengland.co.uk/speech/2023/october/jon-cunliffe-speech-at-the-economics-of-payments-xii-conference.

Davoodalhosseini, Seyed Mohammadreza (2022). "Central bank digital currency and monetary policy". In: *Journal of Economic Dynamics and Control* 142, p. 104150.

Diamond, Douglas W (1984). "Financial intermediation and delegated monitoring". In: *The review of economic studies* 51.3, pp. 393–414.

Duffie, Darrell, Raghuram Rajan, Kenneth Rogoff, Hyun Song Shin, and G30 Working Group on Digital Currencies (2020). "Digital Currencies and Stablecoins: Risks, Opportunities and Challenges Ahead". In: URL: `https://group30.org/images/uploads/publications/auer_Digital_Currencies.pdf`.

Dwork, Cynthia, Aaron Roth, et al. (2014). "The algorithmic foundations of differential privacy." In: *Foundations and Trends in Theoretical Computer Science* 9.3-4, pp. 211–407.

European Central Bank (2021). *Eurosystem report on the public consultation on a digital euro*. Avaliable at https://www.ecb.europa.eu.

Fatás, Antonio (2019). *The Economics of fintech and digital currencies.* Centre for Economic Policy Research.

Fernández-Villaverde, Jesús, Daniel Sanches, Linda Schilling, and Harald Uhlig (2021). "Central bank digital currency: Central banking for all?" In: *Review of Economic Dynamics* 41, pp. 225–242.

Fiege, Uriel, Amos Fiat, and Adi Shamir (1987). "Zero knowledge proofs of identity". In: *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pp. 210–217.

Gale, Douglas and Martin Hellwig (1985). "Incentive-compatible debt contracts: The one-period problem". In: *The Review of Economic Studies* 52.4, pp. 647–663.

Gan, Jingxing, Gerry Tsoukalas, and Serguei Netessine (2023). "Decentralized platforms: Governance, tokenomics, and ICO design". In: *Management Science* 69.11, pp. 6667–6683.

Gans, Joshua S (2019). *The fine print in smart contracts.* Tech. rep. National Bureau of Economic Research.

Garratt, Rodney and Michael Lee (2021). "Monetizing privacy with central bank digital currencies". In: *Available at SSRN 3583949*.

Garratt, Rodney J and Maarten RC Van Oordt (2021). "Privacy as a public good: a case for electronic cash". In: *Journal of Political Economy* 129.7, pp. 2157–2180.

Goldwasser, Shafi, Silvio Micali, and Chales Rackoff (1985). "The knowledge complexity of interactive proof-systems". In: 18, pp. 186–208.

Goldwasser, Shafi, Silvio Micali, and Chales Rackoff (2019). "The knowledge complexity of interactive proof-systems". In: *Providing sound foundations for cryptography: On the work of Shafi Goldwasser and Silvio Micali*, pp. 203–225.

Graetz, Michael J, Jennifer F Reinganum, and Louis L Wilde (1986). "The tax compliance game: Toward an interactive theory of law enforcement". In: *The Journal of Law, Economics, and Organization* 2.1, pp. 1–32.

Gross, Jonas, Johannes Sedlmeir, Matthias Babel, Alexander Bechtel, and Benjamin Schellinger (2021). "Designing a Central Bank Digital Currency with Support for Cash-Like Privacy". In: *Available at SSRN 3891121.*

Grym, Aleksi (2020). *Lessons learned from the world's first CBDC.* Tech. rep. BoF Economics Review.

Halaburda, Hanna, Guillaume Haeringer, Joshua Gans, and Neil Gandal (2022). "The microeconomics of cryptocurrencies". In: *Journal of Economic Literature* 60.3, pp. 971–1013.

Holmstrom, Bengt and Paul Milgrom (1987). "Aggregation and linearity in the provision of intertemporal incentives". In: *Econometrica: Journal of the Econometric Society*, pp. 303–328.

Iyengar, Garud, Fahad Saleh, Jay Sethuraman, and Wenjun Wang (2023). "Economics of permissioned blockchain adoption". In: *Management Science* 69.6, pp. 3415–3436.

John, Kose, Maureen O'Hara, and Fahad Saleh (2022). "Bitcoin and beyond". In: *Annual Review of Financial Economics* 14, pp. 95–115.

Jones, Charles I and Christopher Tonetti (2020). "Nonrivalry and the Economics of Data". In: *American Economic Review* 110.9, pp. 2819–58.

Jovanovic, Boyan and Bálazs Szentes (2013). "On the market for venture capital". In: *Journal of Political Economy* 121.3, pp. 493–527.

Kahn, Charles M, James McAndrews, and William Roberds (2005). "Money is privacy". In: *International Economic Review* 46.2, pp. 377–399.

Kahn, Charles M, Francisco Rivadeneyra, and Tsz-Nga Wong (2019). "Should the central bank issue e-money?" In: *FRB St. Louis Working Paper* 2019-3.

Kahn, Charles M and William Roberds (2008). "Credit and identity theft". In: *Journal of Monetary Economics* 55.2, pp. 251–264.

Kahn, Charles M and Maarten RC Van Oordt (2022). "The demand for programmable payments". In: *Available at SSRN.*

Kahn, Charles M, Maarten RC Van Oordt, and Yu Zhu (2021). *Best before? Expiring central bank digital currency and loss recovery.* Tech. rep. Bank of Canada Staff Working Paper.

Keister, Todd and Cyril Monnet (2022). “Central bank digital currency: Stability and information”. In: *Journal of Economic Dynamics and Control* 142, p. 104501.

Keister, Todd and Daniel Sanches (2023). “Should central banks issue digital currency?” In: *The Review of Economic Studies* 90.1, pp. 404–431.

Kiyotaki, Nobuhiro and John Moore (2002). “Evil is the root of all money”. In: *American Economic Review* 92.2, pp. 62–66.

— (2018). “Inside money and liquidity”. In: *Manuscript, Princeton University.*

Kiyotaki, Nobuhiro and Randall Wright (1993). “A search-theoretic approach to monetary economics”. In: *The American Economic Review*, pp. 63–77.

Lee, Michael, Antoine Martin, and Robert M Townsend (2021). “Optimal design of tokenized markets”. In: *Available at SSRN 3820973.*

Li, Jiaqi, Andrew Usher, and Yu Zhu (2023). “Central Bank Digital Currency and Banking Choices”. In:

Liu, J, Michael Sockin, and Wei Xiong (2021). *Data privacy and consumer vulnerability.* Tech. rep. Working Paper.

Malinova, Katya and Andreas Park (2023). “Tokenomics: when tokens beat equity”. In: *Management Science* 69.11, pp. 6568–6583.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green, et al. (1995). *Microeconomic theory.* Vol. 1. Oxford University Press New York.

Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage (Mar. 2016). “A Fistful of Bitcoins: Characterizing Payments among Men with No Names”. In: *Commun. ACM* 59.4, pp. 86–93. ISSN: 0001-0782. DOI: 10.1145/2896384. URL: https://doi.org/10.1145/2896384.

Miers, Ian, Christina Garman, Matthew Green, and A.D. Rubin (May 2013). “Zerocoin: Anonymous Distributed E-Cash from Bitcoin”. In: pp. 397–411. ISBN: 978-1-4673-6166-8. DOI: 10.1109/SP.2013.34.

Minesso, Massimo Ferrari, Arnaud Mehl, and Livio Stracca (2022). “Central bank digital currency in an open economy”. In: *Journal of Monetary Economics* 127, pp. 54–68.

Modigliani, Franco and Merton H Miller (1958). “The cost of capital, corporation finance and the theory of investment”. In: *The American economic review* 48.3, pp. 261–297.

Mookherjee, Dilip and Ivan Png (1989). “Optimal auditing, insurance, and redistribution”. In: *The Quarterly Journal of Economics* 104.2, pp. 399–415.

Niepelt, Dirk (2020). “Monetary policy with reserves and CBDC: Optimality, equivalence, and politics”. In:

Ostroy, Joseph M and Ross M Starr (1990). “The transactions role of money”. In: *Handbook of monetary economics* 1, pp. 3–62.

Ouyang, Shumiao (2021). “Cashless payment and financial inclusion”. In: *Available at SSRN 3948925*.

Pagnotta, Emiliano (2018). “Bitcoin as decentralized money: prices, mining, and network security”. In: *SSRN (3264448)*.

Panetta, Fabio (2022). *A digital euro that serves the needs of the public: striking the right balance - Introductory statement at the Committee on Economic and Monetary Affairs of the European Parliament, Brussels*. Avaliable at https://www.ecb.europa.eu/.

Parlour, Christine A, Uday Rajan, and Haoxiang Zhu (2022). “When fintech competes for payment flows”. In: *The Review of Financial Studies* 35.11, pp. 4985–5024.

Ruffing, Tim, Pedro Moreno-Sanchez, and Aniket Kate (2014). “CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin”. In: *Computer Security - ESORICS*. Ed. by Mirosław Kutyłowski and Jaideep Vaidya. Cham: Springer International Publishing, pp. 345–364.

Saleh, Fahad (2021). “Blockchain without waste: Proof-of-stake”. In: *The Review of financial studies* 34.3, pp. 1156–1190.

Sanchez, Isabel and Joel Sobel (1993). “Hierarchical design and enforcement of income tax policies”. In: *Journal of public economics* 50.3, pp. 345–369.

Schilling, Linda, Jesús Fernández-Villaverde, and Harald Uhlig (2020). *Central bank digital currency: When price and bank stability collide*. Tech. rep. National Bureau of Economic Research.

Schmutte, Ian M and Nathan Yoder (2022). “Information Design for Differential Privacy”. In: *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 1142–1143.

Slemrod, Joel and Shlomo Yitzhaki (2002). “Tax avoidance, evasion, and administration”. In: *Handbook of public economics*. Vol. 3. Elsevier, pp. 1423–1470.

Taylor, Curtis R (2004). “Consumer privacy and the market for customer information”. In: *RAND Journal of Economics*, pp. 631–650.

Tinn, Katrin (2018). “Smart Contracts and External Financing”. In: *Available at SSRN 3072854*.

— (2019). “Distributed Ledger Technologies and Start-up Financing”. In: *In The economics of Fintech and digital currencies. Center for Economic Policy Research, London*, pp. 15–20.

Tinn, Katrin and Christophe Dubach (2021). “Central bank digital currency with asymmetric privacy”. In: *Available at SSRN 3787088*.

Tirole, Jean (2010). *The theory of corporate finance*. Princeton University Press.

Townsend, Robert M (1979). “Optimal contracts and competitive markets with costly state verification”. In: *Journal of Economic theory* 21.2, pp. 265–293.

UK Parlament Treasury Committee (2023). *The digital pound: still a solution in search of a problem?* Avaliable at https://publications.parliament.uk/pa/cm5804/cmselect/cmtreasy/215/report.html.

Varian, Hal R (1985). “Price discrimination and social welfare”. In: *The American Economic Review* 75.4, pp. 870–875.

Veneris, Andreas, Andreas Park, Fan Long, and Poonam Puri (2021). “Central bank digital loonie: Canadian cash for a new global economy”. In: *Osgoode Legal Studies Research Paper*.

Vissing-Jorgensen, Annette (2021). *Consumer credit: learning your customer’s default risk from what (s) he buys.*

Whited, Toni M, Yufeng Wu, and Kairong Xiao (2022). “Will Central Bank Digital Currency Disintermediate Banks?” In: *Available at SSRN 4112644*.

Williamson, Stephen (2022). “Central bank digital currency: Welfare and policy implications”. In: *Journal of Political Economy* 130.11, pp. 2829–2861.

Williamson, Stephen and Randall Wright (2010). “New monetarist economics: Models”. In: *Handbook of monetary economics*. Vol. 3. Elsevier, pp. 25–96.

Wüst, Karl, Kari Kostiainen, Noah Delius, and Srdjan Capkun (2022). “Platypus: A Central Bank Digital Currency with Unlinkable Transactions and Privacy-Preserving Regulation”. In: *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pp. 2947–2960.

Yermack, David (2017). “Corporate governance and blockchains”. In: *Review of finance* 21.1, pp. 7–31.

Zhou, Sophie L (2021). “Anonymity-Driven Demand for Cryptocurrency: Theory and Policy Implications”. In:

A Proofs

A.1 Proof of Lemma 4.2 and Corollary 4.2.1

Using (10) and (11) in (8), the optimal auditing policy simplifies to

$$\bar{R}^* = \arg\max_{\bar{R}} \Phi\left(\bar{R}\right) \equiv \mathbb{E}\left[\varphi_S - \frac{c}{1+\pi} \mid S < \bar{R}\right] \Pr\left(S < \bar{R}\right) + \varphi_{\bar{R}} \Pr\left(S \geq \bar{R}\right). \tag{17}$$

This problem has a unique maximum if the difference $\Phi\left(\bar{R}+1\right) - \Phi\left(\bar{R}\right)$ changes sign at most once from positive to negative. As $\Phi\left(\bar{R}\right)$ can be expressed as

$$\Phi\left(\bar{R}\right) = \sum_{S=0}^{\bar{R}-1} \left(\varphi_S - \frac{c}{1+\pi}\right) h_S + \varphi_{\bar{R}}\left(1 - H_{\bar{R}-1}\right),$$

it follows that

$$\Phi\left(\bar{R}+1\right) - \Phi\left(\bar{R}\right) = h_{\bar{R}}\left(\left(\varphi_{\bar{R}+1} - \varphi_{\bar{R}}\right)\frac{\left(1 - H_{\bar{R}}\right)}{h_{\bar{R}}} - \frac{c}{1+\pi}\right).$$

When $\bar{R} = 0$, $\Phi(1) - \Phi(0) = h_0\left(\frac{\varphi_1}{h_0} - \frac{c}{1+\pi}\right) = (1-\bar{z})^N\left(\frac{\varphi_1}{(1-\bar{z})^N} - \frac{c}{1+\pi}\right) > 0$ since $c \leq \frac{1+\pi}{(1-\bar{z})^N}$ by assumption and $\varphi_1 \leq 1$ due to limited liability.

Because $h_{\bar{R}} > 0$ for any $\bar{R}$, it is sufficient to show that $\left(\varphi_{\bar{R}+1} - \varphi_{\bar{R}}\right)\frac{\left(1-H_{\bar{R}}\right)}{h_{\bar{R}}}$ is strictly decreasing in $\bar{R}$, i.e., that

$$\Lambda\left(\bar{R}\right) \equiv \left(\varphi_{\bar{R}+1} - \varphi_{\bar{R}}\right)\frac{\left(1 - H_{\bar{R}}\right)}{h_{\bar{R}}} - \left(\varphi_{\bar{R}} - \varphi_{\bar{R}-1}\right)\frac{\left(1 - H_{\bar{R}-1}\right)}{h_{\bar{R}-1}} < 0$$

for any $\bar{R} > 0$.[34] As the Binomial distribution has an increasing Hazard function, it holds that $\frac{\left(1-H_{\bar{R}}\right)}{h_{\bar{R}}} < \frac{\left(1-H_{\bar{R}-1}\right)}{h_{\bar{R}-1}}$ (see Supplementary Appendix B.1 for proof). It follows that when $\varphi(.)$ is weakly concave, i.e., $\varphi_{\bar{R}+1} + \varphi_{\bar{R}-1} - 2\varphi_{\bar{R}} \leq 0$, $\Lambda\left(\bar{R}\right) < \frac{\left(1-H_{\bar{R}-1}\right)}{h_{\bar{R}-1}}\left(\varphi_{\bar{R}+1} + \varphi_{\bar{R}-1} - 2\varphi_{\bar{R}}\right) < 0$. Furthermore, a sufficient condition for a unique maximum is more general and allows for some convexity in the total obligations to external parties. As assumed in Section 3.1.2, $\varphi(.)$ is increasing

[34] It is immediate that if $\bar{R} = 0$, the firm is never audited, and no revenues are raised, which in turn makes raising any external financing impossible.

and $\frac{\varphi_{\bar{R}+1}-\varphi_{\bar{R}}}{\varphi_{\bar{R}}-\varphi_{\bar{R}-1}} \leq \frac{(\bar{R}+1)(N-\bar{R}+1)}{\bar{R}(N-\bar{R})}$, it follows that

$$\begin{aligned}
\Lambda\left(\bar{R}\right) &< \frac{\left(\varphi_{\bar{R}}-\varphi_{\bar{R}-1}\right)}{h_{\bar{R}}h_{\bar{R}-1}}\left(\frac{(\bar{R}+1)(N-\bar{R}+1)}{\bar{R}(N-\bar{R})}\left(1-H_{\bar{R}}\right)h_{\bar{R}-1}-\left(1-H_{\bar{R}-1}\right)h_{\bar{R}}\right) \\
&= \frac{\left(\varphi_{\bar{R}}-\varphi_{\bar{R}-1}\right)}{h_{\bar{R}}h_{\bar{R}-1}}\left(\frac{(\bar{R}+1)(N-\bar{R}+1)}{\bar{R}(N-\bar{R})}\sum_{S=\bar{R}+1}^{N}h_S h_{\bar{R}-1}-\sum_{S=\bar{R}}^{N}h_S h_{\bar{R}}\right) \\
&= \frac{\left(\varphi_{\bar{R}}-\varphi_{\bar{R}-1}\right)}{h_{\bar{R}}}\left(-h_N h_{\bar{R}}+\sum_{S=\bar{R}+1}^{N}\left(\frac{(\bar{R}+1)(N-\bar{R}+1)}{\bar{R}(N-\bar{R})}-\frac{S(N-\bar{R}+1)}{(N-S+1)\bar{R}}\right)h_S h_{\bar{R}-1}\right)<0,
\end{aligned}$$

where the derivation of the last equality used the functional form of the probability mass function of the Binomial distribution in (9), which implies that

$$\frac{h_{S-1}h_{\bar{R}}}{h_S h_{\bar{R}-1}}=\frac{S\left(N-\bar{R}+1\right)}{\left(N-S+1\right)\bar{R}}.$$

The last inequality follows from $\left(\frac{(\bar{R}+1)(N-\bar{R}+1)}{\bar{R}(N-\bar{R})}-\frac{S(N-\bar{R}+1)}{(N-S+1)\bar{R}}\right) \geq 0$ for any $S \geq \bar{R}+1$, with this inequality being strict for any $S > \bar{R}+1$.

As $\Lambda\left(\bar{R}\right) < 0$, it holds that $\Phi\left(\bar{R}\right)$ is quasi-concave and maximized at $\bar{R}^*$ defined in (12), which sets the auditing threshold to the highest value that $\bar{R}$ where $\Phi\left(\bar{R}+1\right)-\Phi\left(\bar{R}\right)>0$.

$\Lambda\left(\bar{R}\right) < 0$ also implies that $\left(\varphi_{\bar{R}+1}-\varphi_{\bar{R}}\right)\frac{(1-H_{\bar{R}})}{h_{\bar{R}}}-\frac{c}{1+\pi}$ is decreasing in $\bar{R}$, which implies that the optimal threshold is weakly decreasing in c. Furthermore, $\Lambda\left(\bar{R}\right)<0$ also implies that $\bar{R}^*_\tau \leq \bar{R}^*$ as $\varphi_{\bar{R}+1}-\varphi_{\bar{R}} \geq \tau_{\bar{R}+1}-\tau_{\bar{R}}$ for any financing contracts that are non-decreasing in $\bar{R}$, as it is a known result in the costly verification literature, and also shown to hold in Appendix A.2. Finally, pledging all reported sales revenues to external parties implies that $\varphi_{\bar{R}}=\bar{R}$. As $\varphi_{\bar{R}+1}-\varphi_{\bar{R}} \leq 1$, $\Lambda\left(\bar{R}\right)<0$ also implies that $\bar{R}^* \leq \bar{R}^*_m$.

A.2 Proof of Lemma 4.3 and Corollary 4.3.1

Given the optimal auditing and reporting policies (10) and (11), the firm's financing contract problem can be restated as choosing the function $\phi\left(.\right)$ to maximize

$$\mathbb{E}\left[S\right]-\Pr\left(S<\bar{R}^*\right)\mathbb{E}\left[\phi\left(S-\tau_S\right)|S<\bar{R}^*\right]-\Pr\left(S\geq\bar{R}^*\right)\mathbb{E}\left[\phi\left(\bar{R}^*-\tau_{\bar{R}^*}\right)|S\geq\bar{R}^*\right]-\boldsymbol{E}_\tau,$$

subject to the auditing threshold, (12), and the financier's break-even constraint

$$\begin{aligned}&\Pr\left(S<\bar{R}^{*}\right)\mathbb{E}\left[\phi\left(S-\tau_{S}\right)|S<\bar{R}^{*}\right]+\Pr\left(S\geq\bar{R}^{*}\right)\mathbb{E}\left[\phi\left(\bar{R}^{*}-\tau_{\bar{R}^{*}}\right)|S\geq\bar{R}^{*}\right]\geq\\&I-G+\Pr\left(S<\bar{R}^{*}\right)\frac{c}{1+\pi}\end{aligned}\qquad(18)$$

It is clearly optimal for the firm to offer a financing contract that sets the financier's breakeven constraint to hold with equality. Using this, we can rewrite the firm's objective function as

$$\mathbb{E}\left[S\right]-I-\boldsymbol{E}_{\tau}+G-\Pr\left(S<\bar{R}^{*}\right)\frac{c}{1+\pi},$$

which implies that the firm would benefit from offering a contract that sets $\bar{R}^{*}$ to be as low as possible to minimize the auditing costs, and all contracts that imply the same $\bar{R}^{*}$ are payoff equivalent for the firm.

Whenever there are positive investment costs, i.e., $I>0$, the financier's break-even constraint (18) implies that $\phi\left(S-\tau_{S}\right)$ must be positive at least for some values of S. Also note that the government subsidy is sufficient to cover the payments to the tax authority. As the tax policy is exogenous, and is not necessarily set to optimize auditing costs, it is possible that $\bar{R}^{*}=\bar{R}_{\tau}^{*}$, e.g., because of rounding. More importantly, this is always the case when auditing cost c is low enough such that $\bar{R}_{\tau}^{*}=N$, i.e., when it is optimal to audit all reports apart from the highest one. From $\bar{R}_{\tau}^{*}$ in Corollary 4.2.1 we find that the condition for this is

$$c\leq c_{min}\equiv\left(1+\pi\right)\left(\tau_{N}-\tau_{N-1}\right)\frac{1-H_{N-1}}{h_{N-1}}=\left(1+\pi\right)\left(\tau_{N}-\tau_{N-1}\right)\frac{\bar{z}}{N\left(1-\bar{z}\right)},$$

where we used the functional form of the binomial distribution to find that $\frac{1-H_{N-1}}{h_{N-1}}=\frac{h_{N}}{h_{N-1}}=\frac{\bar{z}^{N}}{N\bar{z}^{N-1}(1-\bar{z})}=\frac{\bar{z}}{N(1-\bar{z})}$.

When $\bar{R}_{\tau}^{*}=N$, then also $\bar{R}^{*}=N$ and and feasible financing contracts benefit from fully audited records. This in turn implies that any financing contract that sets the financier's breakeven constraint to hold with equality is optimal. For example, an equity contract that sets $\phi\left(S-\tau_{S}\right)=\bar{\phi}_{E}\cdot\left(S-\tau_{S}\right)$ for any S and $\bar{\phi}_{E}$ is a constant that solves

$$\begin{aligned}&\bar{\phi}_{E}\mathbb{E}\left[S-\tau_{S}\right]=I-G+\Pr\left(S<N\right)\frac{c}{1+\pi}\Rightarrow\\&\bar{\phi}_{E}=\frac{I-G+\left(1-h_{N}\right)\frac{c}{1+\pi}}{\mathbb{E}\left[S\right]-\mathbb{E}\left[\tau_{S}\right]}=\frac{I-G+\left(1-\bar{z}^{N}\right)\frac{c}{1+\pi}}{N\bar{z}-\boldsymbol{E}_{\tau}}\end{aligned}$$

is optimal. The firm's expected payoff under this contract is $N\bar{z}-I+G-\boldsymbol{E}_{\tau}-\left(1-\bar{z}^{N}\right)\frac{c}{1+\pi}=$

$N\bar{z} - I - \left(1 - \bar{z}^N\right)\frac{c}{1+\pi}$, which is very close to the first best when c is small, and equal to the first-best payoff if $c \to 0$.

Note that a debt contract that sets

$$\phi\left(S - \tau_S\right) = \begin{cases} S - \tau_S \text{ if } S < D \\ D \text{ if } S \geq D \end{cases},$$

where D solves $\Pr\left(S < D\right)\mathbb{E}\left[S - \tau_S | S < D\right] + D\Pr\left(S \geq D\right) = I + \left(1 - \bar{z}^N\right)\frac{c}{1+\pi}$ is optimal as well, because it does not affect the auditing threshold, and also sets the financier's breakeven constraint to hold with equality. There exist many financing contracts, $\phi\left(.\right)$, that take other shapes and imply the same payoffs.

Corollary 4.2.1 also defines another auditing threshold, $\bar{R}_m^*$, which corresponds to the maximum investment costs as the highest pledge per audited unit under limited liability is $\phi\left(S - \tau_S\right) = S - \tau_S$ for any $S \leq \bar{R}_m^*$. From the financiers' break-even constraint $\bar{I}^{symm\_privacy} \equiv \Phi\left(\bar{R}_m^*\right) - H_{\bar{R}_m^* - 1}\frac{c}{1+\pi}$. When $I = \bar{I}^{symm\_privacy}$, the only feasible contract is a simple debt contract, where the firm credibly pledges to pay the financier

$$\phi\left(S - \tau_S\right) = \begin{cases} S - \tau_S \text{ if } S < \bar{R}_m^* \\ \bar{R}_m^* - \tau_{\bar{R}_m^*} \text{ if } S \geq \bar{R}_m^* \end{cases}.$$

Suppose that $\bar{R}_\tau^* < N$ and $0 < I < \bar{I}^{symm\_privacy}$, such that auditing costs are high enough not to justify auditing all reports for taxation purposes only. As $\bar{R}^* \geq \bar{R}_\tau^*$, raising external financing weakly increases the auditing threshold. It is beneficial for the firm to pledge more funds for outcomes that are audited. To see this note that unless the threshold does not change due to sales distribution being discrete, offering more funds to financiers must increase the total funds raised at any fixed threshold $\bar{R}$, i.e., it must be the case that $\Phi\left(\bar{R}; \phi_H\left(.\right)\right) > \Phi\left(\bar{R}; \phi_L\left(.\right)\right)$ for some contracts $\phi_H\left(.\right)$ and $\phi_L\left(.\right)$ where the contract ϕ_H offers greater payments to financiers. Suppose that the corresponding optimal auditing thresholds are $\bar{R}_H^*$ and $\bar{R}_L^*$ and consider that the total funds raised under these contracts are the same, $\Phi\left(\bar{R}_H^*; \phi_H\left(.\right)\right) = \Phi\left(\bar{R}_L^*; \phi_L\left(.\right)\right) \iff \Phi\left(\bar{R}_H^*; \phi_H\left(.\right)\right) - \Phi\left(\bar{R}_H^*; \phi_L\left(.\right)\right) = \Phi\left(\bar{R}_L^*; \phi_L\left(.\right)\right) - \Phi\left(\bar{R}_H^*; \phi_L\left(.\right)\right)$. Because the left-hand side of the last equality is positive, the right-hand side must be positive as well. As Appendix A.1 proved the change in funds raised is proportional to a decreasing function (i.e., $\Lambda\left(\bar{R}\right) < 0$), it follows that $\bar{R}_L^* > \bar{R}_H^*$.

Therefore, to minimize the optimal auditing threshold (12), it is the best for the firm to have

$\varphi_S = \tau_S + \phi\left(S - \tau_S\right) = S$ for all $S < \bar{R}^*$, and the firm offers a debt-like contract

$$\phi\left(S - \tau_S\right) = \begin{cases} S - \tau_S \text{ if } S < \bar{R}^* \\ \bar{\phi}_D\left(\bar{R}^* - \tau_{\bar{R}^*}\right) \text{ if } S = \bar{R}^* \\ \bar{R}^* - \tau_{\bar{R}^*} \text{ if } S > \bar{R}^* \end{cases},$$

where $\bar{\phi}_D \in (0,1]$ is a constant that guarantees the financier's breakeven constraint holds with equality and is there because the sales distribution is discrete. The pair $\left(\bar{R}^*, \bar{\phi}_D\right)$ solves a system of equations, where $\bar{\phi}_D$ is set such that $\bar{R}^*$ is an integer and where for any tax system it holds that

$$\begin{cases} \left(1 - \left(1 - \bar{\phi}_D\right)\left(\bar{R}^* - \tau_{\bar{R}^*}\right)\right)\frac{1 - H_{\bar{R}^*-1}}{h_{\bar{R}^*-1}} - \frac{c}{1+\pi} = 0 \\ \bar{\phi}_D = \frac{I + H_{\bar{R}^*-1}\frac{c}{1+\pi} + \sum_{S=0}^{\bar{R}^*}(S - \tau_S)h_S}{h_{\bar{R}^*}\left(\bar{R}^* - \tau_{\bar{R}^*}\right)} \end{cases}. \tag{19}$$

It is straightforward to verify that the above contract is incentive compatible and consistent with the optimal reporting policy (11). While there exist other contracts that are payoff equivalent, these exist just because the payoffs are discrete in this setting. For example, instead of the above contract, the firm could offer

$$\phi\left(S - \tau_S\right) = \begin{cases} \bar{\phi}'_D\left(S - \tau_S\right) \text{ if } S \leq \bar{R}^* \\ \bar{R}^* - \tau_{\bar{R}^*} \text{ if } S > \bar{R}^* \end{cases},$$

where $0 < \bar{\phi}'_D < 1$ is a different constant than $\bar{\phi}_D$ and leads to the same payoffs as long as it implies the same $\bar{R}^*$.

Finally, from the above and $\boldsymbol{E}_\tau = G$, the firm's expected profit under the optimal financing contract is

$$\Pi^{symm\_privacy} = \mathbb{E}\left[S\right] - I - H_{\bar{R}^*-1}\frac{c}{1+\pi} = N\bar{z} - I - H_{\bar{R}^*-1}\frac{c}{1+\pi}.$$

When $c \leq c_{min}$ and thus $\bar{R}^* = N$, it holds that $\Pi_{c \leq c_{min}}^{symm\_privac} = \mathbb{E}\left[S\right] - I - \boldsymbol{E}_\tau - H_{N-1}\frac{c}{1+\pi} = N\bar{z} - \left(1 - \bar{z}^N\right)\frac{c}{1+\pi}$.

A.3 Proof of Lemma 4.5, Proposition 4.6 and Proposition 4.7

Supplementary Appendix B.2 derives all Perfect Bayesian Equilibria (PBE) in pure and mixed strategies in the product market for the setting where the payment system is such that all individual purchases are recorded and observable by third parties, i.e., $\Omega = \left\{b_{\{v_n, T_n\}}\right\}^N = \{0,1\}^N$. Namely, it considers all possible belief structures and identifies the ones that satisfy the requirements of PBE and have the feature that $p > 0$. The requirements for a PBE are that: 1) Consumers' strategies are optimal given their preferences (4) and equilibrium beliefs; 2) Beliefs are updated correctly according to Bayes' rule, at least on the equilibrium path; 3) The firm's pricing strategy maximizes its profit given the consumers' strategies and equilibrium beliefs. The additional feature that $p > 0$ is needed as before trading in the product market, the firm needs to raise external financing, and the firm also needs to raise taxes. For this, it needs to earn strictly positive profits in expectations.

Supplementary Appendix B.2 proves that the only belief structure that is consistent with PBE in pure strategies and $p > 0$ is the one where the types $\{v_n, T_n\} = \{\{1, A\}, \{1, B\}\}$ buy the product, i.e., $b_{\{1,A\}} = b_{\{1,B\}} = 1$, and the types $\{v_n, T_n\} = \{\{0, A\}, \{0, B\}\}$ do not buy the product, i.e., $b_{\{0,A\}} = b_{\{0,B\}} = 0$. The proof of Lemma 4.5 then directly follows from Proposition B.4 in Supplementary Appendix B.2.

To prove Proposition 4.6, note that under the PBE in pure strategies and on-path beliefs $\mathcal{B}^{NP\_P} = \left\{b_{\{0,A\}}, b_{\{0,B\}}, b_{\{1,A\}}, b_{\{1,B\}}\right\} = \{0, 0, 1, 1\}$, the information and belief structure implies that

$$\frac{\Pr(T_n|\Omega;\mathcal{B})}{\Pr(T_n)} = \frac{\Pr(T_n|b_n)}{\Pr(T_n)} = \frac{\Pr(b_n|T_n)}{\Pr(b_n)} = \begin{cases} \frac{z_{T_n}}{\bar{z}} \text{ if } b_n = 1 \\ \frac{(1-z_{T_n})}{1-\bar{z}} \text{ if } b_n = 0 \end{cases},$$

where b_n is consumer n's purchasing decision. Using this, (4), and (5) it follows that the equilibrium indirect utility $V^{NP\_P}_{\{v_n,T_n\}} \equiv V\left(v_n, T_n, b_{\{v_n,T_n\}}; \mathcal{B}^{NP\_P}\right)$ of different consumer types on equilibrium path is $V^{NP\_P}_{\{0,A\}} = -\epsilon\frac{z_A-z_B}{1-\bar{z}}$, $V^{NP\_P}_{\{0,B\}} = \epsilon\frac{z_A-z_B}{1-\bar{z}}$, $V^{NP\_P}_{\{1,A\}} = 1-p-\epsilon\frac{z_A-z_B}{1-\bar{z}}$ and $V^{NP\_P}_{\{1,B\}} = 1-p+\epsilon\frac{z_A-z_B}{\bar{z}}$, respectively. This implies that the total consumer surplus is $CS^{NP\_P} = Nq_A(1-z_A)V^{NP\_P}_{\{0,A\}} + N(1-q_A)(1-z_B)V^{NP\_P}_{\{0,B\}} + Nq_Az_AV^{NP\_P}_{\{1,A\}} + N(1-q_A)z_BV^{NP\_P}_{\{1,B\}}$.

Because auditing cost $c \to 0$, raising external funds and taxation are efficient, and the firm's expected revenues in the product market are $Np\bar{z}$, and its expected profit is $Np\bar{z} - I$.

The proof of Proposition 4.6 then follows by adding up the surpluses of consumers and the firm, i.e., $CS^{NP\_P} + Np\bar{z} - I$ to find the joint surplus, and from using $p = 1 - \epsilon\frac{z_A-z_B}{\bar{z}(1-\bar{z})}$ in the firm's profit.

To prove Proposition 4.7, note that the information set is now $\Omega = \{S\}$ as the third parties observe the total sales but do not observe individual purchases. From arguments akin to the ones in Supplementary Appendix B.2, it holds that the only plausible PBE in pure strategies with $p > 0$ are the ones where all consumers with $v_n = 1$ buy, and none of the consumers with $v_n = 0$ buy the product, i.e., the on-path beliefs are $\mathcal{B}^{AP\_P} = \left\{b_{\{0,A\}}, b_{\{0,B\}}, b_{\{1,A\}}, b_{\{1,B\}}\right\} = \{0, 0, 1, 1\}$. In such a case, the total sales of the firm are given by $S = \sum\limits_{n=1}^{N} v_n$, which is distributed according to $Binomial(N, \bar{z})$. It is also useful to define $S_{-n} = \sum\limits_{k=1,k\neq n}^{N} v_k$ as the demand by all other consumers except n. S_{-n} is distributed according to $Binomial(N-1, \bar{z})$, and all agents share beliefs about this distribution. We can then derive learning about the consumer types from the total sales:

$$\tfrac{\Pr(T_n=A|S)}{\Pr(T_n=A)}=\tfrac{\Pr(S|T_n=A)}{\Pr(S)}=\tfrac{\Pr(v_n=1|T_n=A)\Pr(S|T_n=A,v_n=1)+\Pr(v_n=0|T_n=A)\Pr(S|T_n=A,v_n=0)}{\Pr(S)}$$

As other consumers' preferences are independent of consumer n's preferences, it holds that

$$\begin{aligned}\Pr\left(S|T_n=A,v_n=1\right) &= \Pr\left(S_{-n}=S-1\right)=\tfrac{(N-1)!}{(S-1)!(N-S)!}\bar{z}^{S-1}\left(1-\bar{z}\right)^{N-S}\\ \Pr\left(S|T_n=A,v_n=0\right) &= \Pr\left(S_{-n}=S\right)=\tfrac{(N-1)!}{(S)!(N-1-S)!}\bar{z}^{S}\left(1-\bar{z}\right)^{N-1-S}\end{aligned}$$

Using these, distribution of S, and (2), we obtain that

$$\tfrac{\Pr(T_n=A|S)}{\Pr(T_n=A)}=\tfrac{\Pr(S|T_n=A)}{\Pr(S)}=\tfrac{Sz_A(1-\bar{z})+(N-S)(1-z_A)\bar{z}}{S\bar{z}(1-\bar{z})}=\tfrac{S(1-q_A)(z_A-z_B)+N\bar{z}(1-z_A)}{S\bar{z}(1-\bar{z})}.$$

Similarly, we find that

$$\tfrac{\Pr(T_n=B|S)}{\Pr(T_n=B)}=\tfrac{\Pr(S|T_n=B)}{\Pr(S)}=\tfrac{Sz_B(1-\bar{z})+(N-S)(1-z_B)\bar{z}}{N\bar{z}(1-\bar{z})}=\tfrac{-Sq_A(z_A-z_B)+N\bar{z}(1-z_B)}{N\bar{z}(1-\bar{z})}.$$

We can see that when S increases, the updated probability that a consumer has a character type A (B) increases (decreases). We can then derive the indirect utility $V^{AP\_P}_{\{v_n,T_n\}}\equiv V\left(v_n,T_n,b_{\{v_n,T_n\}};\mathcal{B}^{AP\_P}\right)$ of different consumer types on equilibrium path as $V^{AP\_P}_{\{0,A\}}=\epsilon\mathcal{L}_0$, $V^{AP\_P}_{\{0,B\}}=-\epsilon\mathcal{L}_0$, $V^{AP\_P}_{\{1,A\}}=1-p-\epsilon\mathcal{L}_1$ and $V^{AP\_P}_{\{1,B\}}=1-p+\epsilon\mathcal{L}_1$, respectively, where

$$\begin{aligned}\mathcal{L}_1 &\equiv \mathbb{E}\left[\tfrac{(S_{-n}+1)(1-q_A)(z_A-z_B)+N\bar{z}(1-z_A)}{N\bar{z}(1-\bar{z})}-\tfrac{-(S_{-n}+1)q_A(z_A-z_B)+N\bar{z}(1-z_B)}{N\bar{z}(1-\bar{z})}\right]\\ &= \tfrac{((N-1)\bar{z}+1)(1-q_A)(z_A-z_B)+N\bar{z}(1-z_A)}{N\bar{z}(1-\bar{z})}+\tfrac{((N-1)\bar{z}+1)q_A(z_A-z_B)-N\bar{z}(1-z_B)}{N\bar{z}(1-\bar{z})}=\frac{z_A-z_B}{N\bar{z}},\end{aligned}$$

$$\begin{aligned}\mathcal{L}_0 &\equiv \mathbb{E}\left[\tfrac{S_{-n}(1-q_A)(z_A-z_B)+N\bar{z}(1-z_A)}{N\bar{z}(1-\bar{z})}-\tfrac{-S_{-n}q_A(z_A-z_B)+N\bar{z}(1-z_B)}{N\bar{z}(1-\bar{z})}\right]\\ &= \tfrac{(N-1)\bar{z}(1-q_A)(z_A-z_B)+N\bar{z}(1-z_A)}{N\bar{z}(1-\bar{z})}+\tfrac{(N-1)\bar{z}q_A(z_A-z_B)-N\bar{z}(1-z_B)}{N\bar{z}(1-\bar{z})}=-\frac{z_A-z_B}{N\left(1-\bar{z}\right)}.\end{aligned}$$

Similarly to the derivations in Supplementary Appendix B.2, it must hold that $1-\epsilon\frac{z_A-z_B}{N\bar{z}(1-\bar{z})}\geq\epsilon\frac{z_A-z_B}{N\bar{z}(1-\bar{z})}\Leftrightarrow\epsilon\leq\frac{1}{2}\frac{N\bar{z}(1-\bar{z})}{z_A-z_B}$ for a PBE in pure strategies to exist, and that it is optimal for the firm to set the price $p=1-\epsilon\frac{z_A-z_B}{N\bar{z}(1-\bar{z})}$. It is also straightforward to verify that these strategies are optimal for all types of consumers. The firm's profit in this PBE is given by $N\bar{z}\cdot p-I=N\bar{z}\left(1-\epsilon\frac{z_A-z_B}{N\bar{z}(1-\bar{z})}\right)-I$, and the consumer surplus is $CS^{AP\_P}=Nq_A\left(1-z_A\right)V^{AP\_P}_{\{0,A\}}+N\left(1-q_A\right)\left(1-z_B\right)V^{AP\_P}_{\{0,B\}}+Nq_Az_AV^{AP\_P}_{\{1,A\}}+N\left(1-q_A\right)z_BV^{AP\_P}_{\{1,B\}}$. The proof of 4.7 then follows by adding up the surpluses of consumers and the firm, i.e., $CS^{NP\_P} + Np\bar{z} - I$ to find the joint surplus, and $p = 1$ to find the investment threshold.

B Supplementary Appendix

B.1 Proof of increasing Hazard function of Binomial Distribution

Claim B.1. *Suppose that $S \sim Binomial\,(N, z)$. It holds that*

$$\frac{h_S}{1 - H_S} > \frac{h_{S-1}}{1 - H_{S-1}}. \tag{20}$$

Proof. The inequality (20) can be rewritten as follows:

$$\frac{h_S}{1 - \sum\limits_{s=0}^{S} h_s} > \frac{h_{S-1}}{1 - \sum\limits_{s=0}^{S-1} h_s},$$

$$h_S \sum_{s=S}^{N} h_s > h_{S-1} \sum_{s=S+1}^{N} h_s = h_{S-1} \sum_{s=S}^{N-1} h_{s+1},$$

$$h_N h_S + \sum_{s=S}^{N-1} \left(h_s h_S - h_{s+1} h_{S-1}\right) > 0.$$

As $h_N h_S > 0$ for any S, we just need to show that $h_s h_S \geq h_{s+1} h_{S-1}$ for any $s \geq S$.

We find that

$$\begin{aligned} \frac{h_s h_S}{h_{s+1} h_{S-1}} &= \frac{\frac{N!}{s!(N-s)!} z^s \left(1-z\right)^{N-s} \frac{N!}{S!(N-S)!} z^S \left(1-z\right)^{N-S}}{\frac{N!}{(s+1)!(N-s-1)!} z^{s+1} \left(1-z\right)^{N-s-1} \frac{N!}{(S-1)!(N-S+1)!} z^{S-1} \left(1-z\right)^{S-V+1}} \\ &= \frac{\left(s+1\right)\left(N-S+1\right)}{\left(N-s\right)S}, \end{aligned}$$

which is clearly increasing in s, and bigger than zero when $s = S$ as

$$\frac{\left(S+1\right)\left(N-S+1\right)}{\left(N-S\right)S} > 1 \Longleftrightarrow S\left(N-S\right) + N + 1 > \left(N-S\right)S.$$

□

B.2 Perfect Bayesian Equilibrium in the product market

The goal of this appendix is to identify belief structures that are consistent with Perfect Bayesian Equilibrium (PBE) in the product market under both pure and mixed strategies that may be used by consumers and to derive the corresponding Perfect Bayesian Equilibria referred to in the main paper.

There are four consumer types $\{v_n, T_n\} = \{0,1\} \times \{A, B\}$, and each type follows and must be believed to follow the same (possibly mixed) strategy, $\sigma_{\{v_n,T_n\}}$, on the equilibrium path, where $\sigma_{\{v_n,T_n\}}$ is the probability that the consumer type $\{v_n, T_n\}$ buys the product (or equivalently plays a pure strategy $b_{\{v_n,T_n\}} = 1$). Denoting the belief that consumers of type $\{v_n, T_n\}$ buy the product with $\tilde{\sigma}_{\{v_n,T_n\}} = \Pr\left(b_{\{v_n,T_n\}} = 1 | T_n, v_n; \mathcal{B}\right)$, the belief set is $\mathcal{B} = \left\{\tilde{\sigma}_{\{0,A\}}, \tilde{\sigma}_{\{0,B\}}, \tilde{\sigma}_{\{1,A\}}, \tilde{\sigma}_{\{1,B\}}\right\}$. This belief set includes 16 ex-ante possibilities for pure strategies and a wider set of mixed strategies.

There is always a trivial PBE where $\sigma_{\{v_n,T_n\}} = \tilde{\sigma}_{\{v_n,T_n\}} = 0$ for all types, and the firm sets a very high price, guaranteeing that no consumer has an incentive to buy the product. However, because other parts of the model involve the firm needing to pay the investment cost and taxes if it is active, the only relevant and interesting PBE candidates are the ones where the firm makes a positive profit, i.e., it only considers PBE candidates where the price set by the firm, $p > 0$, and $\sigma_{\{v_n,T_n\}} = \tilde{\sigma}_{\{v_n,T_n\}} > 0$ at least for one type. This condition will be imposed from here onwards.

Defining

$$\Delta\mathcal{L}_{\{v_n,T_n\}} \equiv \mathbb{E}\left[\mathcal{L}\left(\Omega; \mathcal{B}\right) | T_n, v_n, b_{\{v_n,T_n\}} = 1\right] - \mathbb{E}\left[\mathcal{L}\left(\Omega; \mathcal{B}\right) | T_n, v_n, b_{\{v_n,T_n\}} = 0\right], \tag{21}$$

it follows from (4) that the consumer's optimal strategy is

$$\sigma_{\{v_n,T_n\}} = \begin{cases} 1 \text{ if } p < v_n - \epsilon\Delta\mathcal{L}_{\{v_n,T_n\}} \\ (0,1] \text{ if } p = v_n - \epsilon\Delta\mathcal{L}_{\{v_n,T_n\}} \\ 0 \text{ if } p > v_n - \epsilon\Delta\mathcal{L}_{\{v_n,T_n\}} \end{cases}. \tag{22}$$

For the sake of argument, let us focus on the setting where the payment system is such that all individual purchases are recorded and observable by third parties, i.e., $\Omega = \left\{b_{\{v_n,T_n\}}\right\}^N = \{0,1\}^N$. The setting where no information is available is akin to the special case where privacy concerns are moot, i.e., $\epsilon = 0$. Furthermore, while the PBE under asymmetric privacy is different, all the proofs identifying plausible belief structures are similar.

By Bayes's rule, the prior distribution of T_n being known, and the law of total expectations,

$$\frac{\Pr\left(T_n|\Omega; \mathcal{B}\right)}{\Pr\left(T_n\right)} = \frac{\Pr\left(\Omega|T_n; \mathcal{B}\right)}{\Pr\left(\Omega; \mathcal{B}\right)} = \frac{\left(1 - z_{T_n}\right)\Pr\left(\Omega|v_n = 0, T_n; \mathcal{B}\right) + z_{T_n}\Pr\left(\Omega|v_n = 1, T_n; \mathcal{B}\right)}{\sum_{v_n=\{0,1),T_n=\{A,B\}}\Pr\left(v_n|T_n\right)\Pr\left(T_n\right)\Pr\left(\Omega|v_n, T_n; \mathcal{B}\right)},$$

where $z_{T_n} = \Pr\left(v_n = 1|T_n\right)$ has been defined in (1).

Denote $\Omega_{-n} = \{0,1\}^{N-1}$ the information that is observable on the payment ledger about all other consumers except n, such that we can express $\Omega = \left\{\Omega_{-n}, b_{\{v_n,T_n\}}\right\}$. Furthermore, as individual preferences and purchasing decisions are independent conditional on the type (and must be believed to be independent along any PBE), it holds that $\Pr\left(\Omega|T_n, v_n; \mathcal{B}\right) = \Pr\left(\Omega_{-n}|\mathcal{B}\right) \Pr\left(b_{\{v_n,T_n\}}|T_n, v_n; \mathcal{B}\right)$. As the consumer of type $\{T_n, v_n\}$ is believed to buy the product with probability $\tilde{\sigma}_{\{T_n,v_n\}}$, it follows that

$$\frac{\Pr\left(T_n|\Omega; \mathcal{B}\right)}{\Pr\left(T_n\right)} = \frac{\Pr\left(b_{\{v_n,T_n\}}|T_n; \mathcal{B}\right)}{\Pr\left(b_{\{v_n,T_n\}}; \mathcal{B}\right)} = \begin{cases} \frac{(1-z_{T_n})\tilde{\sigma}_{\{0,T_n\}}+z_{T_n}\tilde{\sigma}_{\{1,T_n\}}}{\tilde{Z}} \text{ if } b_{\{v_n,T_n\}} = 1 \\ \frac{(1-z_{T_n})\left(1-\tilde{\sigma}_{\{0,T_n\}}\right)+z_{T_n}\left(1-\tilde{\sigma}_{\{1,T_n\}}\right)}{1-\tilde{Z}} \text{ if } b_{\{v_n,T_n\}} = 0 \end{cases},$$

where

$$\tilde{Z} \equiv \left(1 - z_A\right) q_A \tilde{\sigma}_{\{0,A\}} + \left(1 - z_B\right)\left(1 - q_A\right) \tilde{\sigma}_{\{0,B\}} + z_A q_A \tilde{\sigma}_{\{1,A\}} + z_B \left(1 - q_A\right) \tilde{\sigma}_{\{1,B\}}. \tag{23}$$

Using this and (5) in (21) and simplifying

$$\Delta\mathcal{L}_{\{v_n,A\}} = \frac{z_A\left(\tilde{\sigma}_{\{1,A\}} - \tilde{\sigma}_{\{0,A\}}\right) - z_B\left(\tilde{\sigma}_{\{1,B\}} - \tilde{\sigma}_{\{0,B\}}\right) + \left(\tilde{\sigma}_{\{0,A\}} - \tilde{\sigma}_{\{0,B\}}\right)}{\tilde{Z}\left(1 - \tilde{Z}\right)} = -\Delta\mathcal{L}_{\{v_n,B\}}. \tag{24}$$

We can see from (24) that $\Delta\mathcal{L}_{\{1,T_n\}} = \Delta\mathcal{L}_{\{0,T_n\}}$ as the privacy depends on the core parameters of the model and beliefs, and not on v_n directly. This enables us to shorten the notation and define $\Delta\mathcal{L} \equiv \Delta\mathcal{L}_{\{1,A\}} = \Delta\mathcal{L}_{\{0,A\}}$, such that $\Delta\mathcal{L}_{\{1,B\}} = \Delta\mathcal{L}_{\{0,B\}} = -\Delta\mathcal{L}$. Using this, we can express (22) for different consumer types as

$$\sigma_{\{v_n,T_n\}} = \begin{cases} 1 \text{ if } \left\{p < v_n - \epsilon\Delta\mathcal{L} \wedge T_n = A\right\} \vee \left\{p < v_n + \epsilon\Delta\mathcal{L} \wedge T_n = B\right\} \\ (0, 1] \text{ if } \left\{p = v_n - \epsilon\Delta\mathcal{L} \wedge T_n = A\right\} \vee \left\{p = v_n + \epsilon\Delta\mathcal{L} \wedge T_n = B\right\} \\ 0 \text{ if } \left\{p > v_n - \epsilon\Delta\mathcal{L} \wedge T_n = A\right\} \vee \left\{p > v_n + \epsilon\Delta\mathcal{L} \wedge T_n = B\right\} \end{cases} \tag{25}$$

Using these results together with the PBE requirement of consistency between strategies and beliefs, we can narrow the set of belief structures that are consistent with PBE, i.e., satisfy $\sigma_{\{v_n,T_n\}} = \tilde{\sigma}_{\{v_n,T_n\}}$.

Lemma B.2. *There are no Perfect Bayesian Equilibria where* $\Delta\mathcal{L} < 0$.

Proof. Suppose that there is a PBE where $\Delta\mathcal{L} < 0$. As we are seeking PBE where $p > 0$, by (25) it must hold that the type $\{0, B\}$ never buys, i.e., $\sigma_{\{0,B\}} = 0$. As beliefs must be consistent with optimal strategies, it must also hold that $\tilde{\sigma}_{\{0,B\}} = 0$. From (25) it also follows that the type $\{1, A\}$ must be most willing to buy, and whether the type $\{1, B\}$ or $\{0, A\}$ is more willing to buy depends on $|\Delta\mathcal{L}|$. We have the following possibilities:

1) Consider that the firm has set $p \leq \min\left[-\epsilon\Delta\mathcal{L}, 1 + \epsilon\Delta\mathcal{L}\right]$ to ensure that both of these types buy, or $p \leq \max\left[-\epsilon\Delta\mathcal{L}, 1 + \epsilon\Delta\mathcal{L}\right]$ so that only one of these types buys. In either case, it must hold that the

type $\{v_n, T_n\} = \{1, A\}$ buys for sure and must be believed to do so on-path, i.e., $\sigma_{\{1,A\}} = \tilde{\sigma}_{\{1,A\}} = 1$. Using $\tilde{\sigma}_{\{0,B\}} = 0$ and $\tilde{\sigma}_{\{1,A\}} = 1$ in (24) we obtain that $\Delta\mathcal{L} \propto z_A - z_B + z_B\left(1 - \tilde{\sigma}_{\{1,B\}}\right) + (1 - z_A)\,\tilde{\sigma}_{\{0,A\}}$, which is strictly positive as $z_A > z_B$, and $\tilde{\sigma}_{\{1,B\}}, \tilde{\sigma}_{\{0,A\}} \in [0, 1]$. This contradicts $\Delta\mathcal{L} < 0$ and the belief structures that take the form $\mathcal{B} = \left\{\tilde{\sigma}_{\{0,A\}}, \tilde{\sigma}_{\{0,B\}}, \tilde{\sigma}_{\{1,A\}}, \tilde{\sigma}_{\{1,B\}}\right\} = \left\{\tilde{\sigma}_{\{0,A\}}, 0, 1, \tilde{\sigma}_{\{1,B\}}\right\}$ cannot be consistent.

2) Consider that the firm has set $\max\left[-\epsilon\Delta\mathcal{L}, 1 + \epsilon\Delta\mathcal{L}\right] < p \leq 1 + \epsilon\Delta\mathcal{L}$. From (25) it is then clear that $\sigma_{\{v_n,T_n\}} = 0$ for types $\{v_n, T_n\} = \{0, A\}, \{0, B\}, \{1, B\}$, and by consistency of beliefs and strategies, it must also hold that $\tilde{\sigma}_{\{0,A\}}, \tilde{\sigma}_{\{0,B\}}, \tilde{\sigma}_{\{1,B\}} = 0$. Using this in (24) we obtain that $\Delta\mathcal{L} \propto z_A\tilde{\sigma}_{\{1,A\}} > 0$, which again contradicts $\Delta\mathcal{L} < 0$ and a belief structure where only the type $\{1, A\}$ buys the product with positive probability cannot be consistent.

3) Finally, if the firm has set $p > 1 + \epsilon\Delta\mathcal{L}$, no consumer buys and $\Delta\mathcal{L} = 0$, which again leads to a contradiction. □

Lemma B.2 shows that belief structures that may be consistent with PBE cannot have the property that types with $T_n = A$ suffer a lesser privacy loss if they buy. This is intuitive as these consumers are more likely to value the product highly, i.e., $z_A > z_B$. Hence it must be the case in any PBE that $\Delta\mathcal{L} \geq 0$, and this is taken as given in what follows.

Lemma B.3. *There are no Perfect Bayesian Equilibria where* $\Delta\mathcal{L} \geq \frac{1}{2\epsilon}$.

Proof. Suppose that there is a PBE where $\Delta\mathcal{L} \geq \frac{1}{2\epsilon}$ and $p > 0$. As we are seeking PBE where $p > 0$, by (25) it must hold that the type $\{0, A\}$ never buys and must be believed to never buy, i.e., $\sigma_{\{0,A\}} = \tilde{\sigma}_{\{0,A\}} = 0$. Furthermore, when $\Delta\mathcal{L} > \frac{1}{2\epsilon}$ it holds that $\min\left[\epsilon\Delta\mathcal{L}, 1 - \epsilon\Delta\mathcal{L}\right] = 1 - \epsilon\Delta\mathcal{L}$. We have the following possibilities to consider:

1) Consider that the firm has set $p \leq 1 - \epsilon\Delta\mathcal{L}$. From (25) it then follows the type $\{1, A\}$ is willing to buy the product with a positive probability, while the types $\{0, B\}$ and $\{1, B\}$ buy the product with probability 1. As beliefs must be consistent with strategies, it must hold that $\tilde{\sigma}_{\{0,B\}} = \tilde{\sigma}_{\{1,B\}} = 1$. However, by (24), it must then also be the case that $\Delta\mathcal{L} \propto z_A\tilde{\sigma}_{\{1,A\}} - 1 < 0$, which contradicts $\Delta\mathcal{L} > \frac{1}{2\epsilon}$.

2) Consider that the firm has set $1 - \epsilon\Delta\mathcal{L} < p \leq \epsilon\Delta\mathcal{L}$. From (25) it then follows the type $\{v_n, T_n\} = \{1, A\}$ never buys the product, i.e., $\sigma_{\{1,A\}} = \tilde{\sigma}_{\{1,A\}} = 0$, the type $\{0, B\}$ is willing to buy it at a positive probability, and the type $\{1, B\}$ must buy the product with probability 1, i.e., $\sigma_{\{1,B\}} = \tilde{\sigma}_{\{1,B\}} = 1$. However by (24) this implies that $\Delta\mathcal{L} \propto -z_B - (1 - z_B)\,\tilde{\sigma}_{\{0,B\}} < 0$, which again contradicts $\Delta\mathcal{L} > \frac{1}{2\epsilon}$.

3) Setting $p > \epsilon\Delta\mathcal{L}$ cannot also be consistent. The only type who could buy at $p > \epsilon\Delta\mathcal{L}$ is the

type $\{v_n, T_n\} = \{1, B\}$, i.e., it must be the case that $\tilde{\sigma}_{\{0,A\}}, \tilde{\sigma}_{\{0,B\}}, \tilde{\sigma}_{\{1,A\}} = 0$, which implies that $\Delta\mathcal{L} \propto -z_B \tilde{\sigma}_{\{1,B\}} \leq 0$ and contradicts $\Delta\mathcal{L} \geq \frac{1}{2\epsilon}$. □

Lemmas B.2 and B.3 eliminate a wide set of belief structures as implausible and highlight that any plausible PBE must have $0 \leq \Delta\mathcal{L} \leq \frac{1}{2\epsilon}$. By (25), this clearly establishes constraints that a plausible belief structure must follow. Namely, the type $\{v_n, T_n\} = \{0, A\}$ never buys and must be believed to follow this strategy, i.e., $\tilde{\sigma}_{\{0,A\}} = 0$, the type $\{0, B\}$ cannot be more likely to buy than the type $\{1, A\}$, and the corresponding beliefs must have $0 \leq \tilde{\sigma}_{\{0,B\}} \leq \tilde{\sigma}_{\{1,A\}}$. Finally, it also holds that the type $\{1, B\}$ must be the type who is most likely to buy the product in any plausible PBE.

We can then identify a set of plausible PBE as follows.

Proposition B.4. *Provided that $\epsilon \leq \frac{1}{2}\frac{\bar{z}(1-\bar{z})}{z_A - z_B}$, there exists a PBE in pure strategies where $\sigma_{\{1,A\}} = \sigma_{\{1,B\}} = 1$, $\sigma_{\{0,A\}} = \sigma_{\{0,B\}} = 0$, $p = 1 - \epsilon\frac{z_A - z_B}{\bar{z}(1-\bar{z})}$, where $\bar{z}$ is defined in (2) and equals the expected demand by one consumer in equilibrium. Furthermore, this equilibrium is a unique equilibrium in pure strategies, and there is no PBE in pure strategies when $\epsilon > \frac{1}{2}\frac{\bar{z}(1-\bar{z})}{z_A - z_B}$.*

Proof. As shown, a PBE could only exist when $0 \leq \Delta\mathcal{L} \leq \frac{1}{2\epsilon}$. By (25) it follows that there are only three potentially plausible belief structures remaining to be considered: these are $\mathcal{B} = \left\{\tilde{\sigma}_{\{0,A\}}, \tilde{\sigma}_{\{0,B\}}, \tilde{\sigma}_{\{1,A\}}, \tilde{\sigma}_{\{1,B\}}\right\} = \{\{0,0,0,1\}, \{0,0,1,1\}, \{0,1,1,1\}\}$. Using these in (24), it follows that the last belief structure cannot be consistent with $0 \leq \Delta\mathcal{L} \leq \epsilon\Delta\mathcal{L}$ as $\Delta\mathcal{L}_{\{0,0,0,1\}} \propto -z_B < 0$ and $\Delta\mathcal{L}_{\{0,1,1,1\}} \propto -(z_A - 1) < 0$. Hence, the only plausible belief structure is $\mathcal{B} = \{0,0,1,1\}$, which by (24) and (23) gives $\Delta\mathcal{L}_{\{0,0,1,1\}} = \frac{z_A - z_B}{\bar{z}(1-\bar{z})}$. Clearly, $\frac{z_A - z_B}{\bar{z}(1-\bar{z})} > 0$. However, we also need $\Delta\mathcal{L} \leq \frac{1}{2\epsilon}$, which implies that this equilibrium only exists if $\epsilon \leq \frac{1}{2}\frac{\bar{z}(1-\bar{z})}{z_A - z_B}$. Furthermore, $0 \leq \Delta\mathcal{L} \leq \frac{1}{2\epsilon}$ also implies that the firm must set the price $p \in \left(\epsilon\frac{z_A - z_B}{\bar{z}(1-\bar{z})}, 1 - \epsilon\frac{z_A - z_B}{\bar{z}(1-\bar{z})}\right]$ to induce only the types $\{1, A\}$ and $\{1, B\}$ to buy. As replacing beliefs with actual demand in (23) gives the expected demand by a random consumer, the firm's expected profit under these strategies is $Np\bar{z}$, setting $p = 1 - \epsilon\frac{z_A - z_B}{\bar{z}(1-\bar{z})}$. Notice that as $\epsilon \leq \frac{1}{2}\frac{\bar{z}(1-\bar{z})}{z_A - z_B}$, $p > 0$. □

Proposition B.4 shows that if the privacy concerns are not too pressing, there is an equilibrium where the firm offers enough of a discount to induce all consumers who value the product at $v_n = 1$ to buy the product. However, this comes at a cost of having to offer the product at a lower price.

While there is no PBE in pure strategies when ϵ is high, there always exist mixed strategy equilibria. In particular, there are always mixed strategy equilibria that fully conceal the consumer types.

Proposition B.5. *For any ϵ, there exist PBE in mixed strategies where the type $\{1, B\}$ buys the product with probability $\sigma_{\{1,B\}} = \breve{\sigma} \in (0, 1]$, the type $\{1, A\}$ buys the product with probability*

$\sigma_{\{1,A\}} = \frac{z_B}{z_A}\breve{\sigma}$*, and the types* $\{0, A\}$ *and* $\{0, B\}$ *do not buy the product. The Pareto dominant equilibrium is the one where* $\{1, B\}$ *plays a pure strategy* $\sigma_{\{1,B\}} = 1$ *and* $\{1, A\}$ *plays a mixed strategy* $\sigma_{\{1,A\}} = \frac{z_B}{z_A}$*. The expected demand by one consumer in this Pareto dominant equilibrium is* $z_B < \bar{z}$*.*

Proof. Consumer types are fully concealing only when $\Delta\mathcal{L} = 0$. From (25), it then follows that the types $\{0, A\}$ and $\{0, B\}$ would not buy the product at $p > 0$, while the types $\{1, A\}$ and $\{1, B\}$ are willing to buy it as long as $p \leq 1$. Given this and the requirement of consistency of beliefs, it must be the case that $\Delta\mathcal{L} = \frac{z_A\tilde{\sigma}_{\{1,A\}} - z_B\tilde{\sigma}_{\{1,B\}}}{\tilde{Z}(1-\tilde{Z})} = 0$, where $\tilde{Z} = q_A z_A \tilde{\sigma}_{\{1,A\}} + (1 - q_A)\, z_B \tilde{\sigma}_{\{1,B\}}$. Clearly there are multiple solutions to $z_A\tilde{\sigma}_{\{1,A\}} = z_B\tilde{\sigma}_{\{1,B\}}$ and as $\tilde{\sigma}_{\{1,B\}} \in (0, 1]$ and $z_A > z_B$, it must be the case that $\{1, A\}$ uses a mixed strategy that is proportional to the strategy used by $\tilde{\sigma}_c = \breve{\sigma}$. It then follows that the expected demand by all consumers is $N\left(q_A z_A \left(\frac{z_B}{z_A}\breve{\sigma}\right) + (1 - q_A)\, z_B\breve{\sigma}\right) = N z_B \breve{\sigma}$ and the firm's profit is $Npz_B\breve{\sigma}$. It is clear that it is optimal for the firm to set $p = 1$ and it is easy to confirm that both types $\{1, A\}$ and $\{1, B\}$ consumers are indifferent between buying and not buying and are willing to buy at any probability. As consumer surplus is zero, the joint surplus equals the firm's profit, which is maximized when $\breve{\sigma} = 1$. □

Proposition B.5 shows that there are always equilibria where some consumers forego the opportunity to buy a product they like to conceal their type. While there is no need for discounts in these PBE, the firm faces lower demand.

One may wonder if there are other mixed strategy equilibria. Given the above analysis, there are only two possible candidate belief structures to consider:

1. One where the types $\{0, A\}$ and $\{0, B\}$ are believed to follow pure strategies not to buy, i.e., $\tilde{\sigma}_{\{0,A\}} = \tilde{\sigma}_{\{0,B\}} = 0$, the type $\{1, B\}$ is believed to follow a pure strategy to buy, i.e., $\tilde{\sigma}_{\{1,B\}} = 1$, and the type $\{1, A\}$ is indifferent and is believed to buy with probability $\tilde{\sigma}_{\{1,A\}} = \tilde{\sigma}'_A$, where $\tilde{\sigma}'_A \in (0, 1)$ and $\tilde{\sigma}'_A \neq \frac{z_B}{z_A}$.

2. One where the type $\{0, A\}$ is believed to follow a pure strategy not to buy, i.e., $\tilde{\sigma}_{\{0,A\}} = 0$, the types $\{1, A\}$ and $\{1, B\}$ are believed to follow pure strategies to buy, i.e., $\tilde{\sigma}_{\{1,A\}} = \tilde{\sigma}_{\{1,B\}} = 1$, and the type $\{0, B\}$ is indifferent and is believed to buy with probability $\tilde{\sigma}_{\{0,B\}} = \tilde{\sigma}'_B$, where $\tilde{\sigma}'_B \in (0, 1)$.

We exclude both of these as being inconsistent with optimal strategies by consumers and firms.

Consider the first case. From (24) it follows that $\Delta\mathcal{L} = \frac{z_A\tilde{\sigma}'_A - z_B}{\tilde{z}'_A(1-\tilde{z}'_A)}$, where $\tilde{z}'_A = q_A z_A \tilde{\sigma}'_A + (1 - q_A)\, z_B$. As we have shown that it must hold that $\Delta\mathcal{L} > 0$ in any PBE equilibrium that does not involve consumers perfectly concealing their type, it must be the case that $\tilde{\sigma}'_A < \frac{z_A}{z_B}$. Note that this further implies that $\tilde{z}'_A \in (z_B, \bar{z})$. Furthermore, as the type $\{1, A\}$ must be indifferent, it must be the case that $p = 1 - \epsilon \frac{z_A \tilde{\sigma}'_A - z_B}{\tilde{z}'_A \left(1 - \tilde{z}'_A\right)} = 1 - \epsilon \frac{\tilde{z}'_A - z_B}{q_A \tilde{z}'_A \left(1 - \tilde{z}'_A\right)}$. While we can verify using (25) that all consumers' optimal strategies are consistent with the conjectured beliefs, these beliefs are not consistent with the firm's optimal strategy. To see this, notice that when setting the price on-path, the firm must believe that the total demand is $N\tilde{z}'_A$. Hence its profit is $Np\tilde{z}'_A = N\left(1 - \epsilon \frac{\tilde{z}'_A - z_B}{q_A \tilde{z}'_A \left(1 - \tilde{z}'_A\right)}\right)\tilde{z}'_A$. Differentiating this, we obtain that $\frac{\partial\left(Np\tilde{z}'_A\right)}{\partial \tilde{z}'_A} = N - N\epsilon \frac{1 - 2\tilde{z}'_A + z_B}{q_A \left(1 - \tilde{z}'_A\right)^2}$ and $\frac{\partial^2\left(Np\tilde{z}'_A\right)}{\partial \tilde{z}'_A \partial \tilde{z}'_A} = 2N\epsilon \frac{\tilde{z}'_A - z_B}{q_A \left(1 - \tilde{z}'_A\right)^2} > 0$ as $\tilde{z}'_A > z_B$. Because the firm's profit is convex in $\tilde{z}'_A$, it follows that it would set the price such that $\tilde{z}'_A$ is outside the range $(z_B, \bar{z})$, i.e., it would either have an incentive to set $p = 1$, in which case $\tilde{z}'_A = z_B$ and this choice corresponds to the PBE in Lemma B.5, or to set $p = 1 - \epsilon \frac{\bar{z} - z_B}{q_A \bar{z}(1 - \bar{z})} = 1 - \epsilon \frac{z_A - z_B}{\bar{z}(1 - \bar{z})}$, and this choice corresponds to the PBE in Lemma B.4.

Consider then the second case. From (24) it follows that $\Delta\mathcal{L} = \frac{z_A - z_B - (1 - z_B)\tilde{\sigma}'_B}{\tilde{z}'_B \left(1 - \tilde{z}'_B\right)}$, where $\tilde{z}'_B = (1 - q_A)(1 - z_B)\tilde{\sigma}'_B + \bar{z}$. As it must hold that $\Delta\mathcal{L} > 0$, it must further hold that $\tilde{\sigma}'_B < \frac{z_A - z_B}{(1 - z_B)}$. And as the type $\{0, B\}$ must be indifferent, it must hold that $p = \epsilon\Delta\mathcal{L} = \epsilon \frac{z_A - z_B - (1 - z_B)\tilde{\sigma}'_B}{\tilde{z}'_B \left(1 - \tilde{z}'_B\right)}$. Using that firm must expect the relationship between p and $\tilde{\sigma}'_B$ on path, its expected profit is $Np\tilde{z}'_B = N\epsilon \frac{z_A - z_B - (1 - z_B)\tilde{\sigma}'_B}{\left(1 - \tilde{z}'_B\right)} = N\epsilon \frac{z_A - z_B - (1 - z_B)\tilde{\sigma}'_B}{\left(1 - \bar{z} - (1 - q_A)(1 - z_B)\tilde{\sigma}'_B\right)}$. Differentiating this gives $\frac{\partial Np\tilde{z}'_B}{\partial \tilde{\sigma}'_B} = -N\epsilon \frac{(1 - z_B)(1 - z_A)}{\left(1 - \tilde{z}'_B\right)}$, which implies that the firm would set the price such that $\tilde{\sigma}'_B$ and $\tilde{z}'_B$ are as small as possible. In particular, the firm would set $p = \epsilon \frac{z_A - z_B}{\bar{z}(1 - \bar{z})}$ such that $\tilde{\sigma}'_B = 0$. However, this is outside the range. Furthermore, when $\tilde{\sigma}'_B = 0$ the firm would have an incentive to deviate and set $p = 1 - \epsilon \frac{z_A - z_B}{\bar{z}(1 - \bar{z})}$, which is higher and leads to the same demand.