



HAL
open science

Who Pays Whom? Anonymous EMV-Compliant Contactless Payments

Charles Olivier-Anclin, Ioana Boureanu, Liqun Chen, Christopher Newton,
Tom Chothia, Anna Clee, Andreas Kokkinis, Pascal Lafourcade

► **To cite this version:**

Charles Olivier-Anclin, Ioana Boureanu, Liqun Chen, Christopher Newton, Tom Chothia, et al.. Who Pays Whom? Anonymous EMV-Compliant Contactless Payments. 34th USENIX Security Symposium 2025, Aug 2025, Seattle, United States. hal-04917364

HAL Id: hal-04917364

<https://hal.science/hal-04917364v1>

Submitted on 29 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Who Pays Whom? Anonymous EMV-Compliant Contactless Payments

Charles Olivier-Anclin

*Universite de Clermont Auvergne, LIMOS
INSA CVL, LIFO, Université d’Orléans, Inria
be ys Pay*

Ioana Boureanu, Liqun Chen, Chris Newton

Surrey Centre for Cyber Security, University of Surrey

Tom Chothia, Anna Clee, Andreas Kokkinis

University of Birmingham

Pascal Lafourcade

Universite de Clermont Auvergne, LIMOS

Abstract

EMV is the de-facto worldwide payment system used by Mastercard, Visa, American Express, and such. *In-shop* EMV contactless payments are not anonymous or private: the payers’ long-term identification data leaks to Merchants or even to observers. *Anti-Money Laundering (AML)*, *Know Your Customer (KYC)* and *Strong Customer Authentication (SCA)* are payment regulations protecting us from illegal activities, but –in so doing– contribute chiefly to this lack of privacy in EMV payments. Threading the tightrope of AML, KYC and SCA regulations, we provide two *privacy-enhancing*, *EMV-compatible*, *law-abiding* and practicable contactless-payments protocols: *PrivBank* and *PrivProxy*.

We do not use privacy-enhancing technology, like homomorphic encryption, that would break backwards-compatibility with current EMV, but rather we do privacy by engineering design, *adhering to the existing EMV infrastructure, as is*. So, *PrivBank* and *PrivProxy* *provably* achieve strong notions of payers and merchant privacy, anonymity and unlinkability as seen in e-cash or shopping vouchers, whilst being *implementable in EMV as it stands*.

1 Introduction

No Privacy in Contactless EMV by Default. EMVCo [5] is the largest consortium of payment providers, including Visa, Mastercard, American Express and UnionPay. *In-shop* EMV transactions form 94% of the whole payments market [47], with 12.8 billion bankcards in circulation. Despite its modern features, EMV payments do not provision privacy guarantees, for *in-shop*, contactless EMV payments, be it by “plastic”/physical card or mobile device: in these type of transactions, payees and payers identifiable data can be tracked. For starters, payment providers (e.g., Visa, Mastercard) link together mobile and plastic-card payment data, in their standard form, for instance, in the context of loyalty schemes and statistics [19]. Also, card-issuing banks as well as the EMV payment networks always know where we shop, with

which merchants, at which location and time, and could be using this to profile us. Worse, despite the fact that all modern, EMV-compliant mobile-apps generate and use a new “account number” with every payment one makes, merchants can link our mobile purchases together, even if we make some with GooglePay and others with ApplePay or SamsungPay. And, there is no complexity to this: just observing the in-shop transaction between the payment device (*i.e.*, card or phone) and the merchant’s payment-terminal/PoS (Point of Sale) suffices.

Existent Privacy-aware EMV: Online Only & Potential for Usability Improvements. Certain banks (such as Revolut [17]) have taken steps towards better EMV-payment security and give some (weak) form of pseudonymity by generating cards for one-time use. However, these solutions are only for online shopping and lack a certain degree of end-to-end usability: a one-time card is generated on an app, then the payer has to go online separately and pay with it. Meanwhile, anonymous payments have been considered since the introduction of electronic payments in the 1980 in the form of e-cash [36, 59], and large-scale projects such as GNU TALER [34, 41] aim to revive that; their drawback (with respect to our goals) is that they are not EMV-compliant, and in fact they would need to replace in-shop EMV transactions with a new *online* payment system. Deploying that at a large scale would be costly.

Adding Practicable Privacy to EMV, In-Shop Contactless Payments. There is no in-shop/by-PoS contactless payment solution providing privacy preserving properties, *e.g.*, pseudonymity (payments traceable to longterm identifier) and unlinkability (preventing the linking of two payments) for payers and for merchants, whilst being EMV-compliant, or as practicable as the user-friendly, in-shop EMV contactless payments. Achieving this is not trivial: due to anti-money laundering and fraud-protection regulations, legal requirements inherently do not align with privacy-preservation in EMV contactless in-shop payments, making it hard to attain. Pseudonymity in EMV is hindered by the need for EMV payments to be auditable by the decision-making entities. There are several *Anti-Money Laundering (AML)* regulations [51],

including *Know Your Customer (KYC)* and further there are *Strong Customer Authentication (SCA)* [49] regulations to protect customers from fraud. There are also further fraud-protection mechanisms which are not mandatory but lack thereof increases the risk of financial losses for banks and financial bodies, as they have to reimburse customers for unauthorised use of their cards. So, there is one added complexity to proposing a scheme for privacy-preserving EMV contactless, in-shop payments: making that scheme abide by the laws and regulations governing EMV.

Our Research Questions (RQ) & Contributions.

We state below our contributions as answers to research questions (RQ), providing intuitions for these answers.

RQ1: How can one create in-shop, contactless payments that are privacy preserving, and EMV-compliant from the viewpoint of system requirements engineering, and transaction flow of current in-shop, contactless EMV?

When answering this question, the solution would not add heavy privacy-enhancing cryptographic machinery (such as homomorphic encryption, *etc.*) on top of EMV payments, as that would immediately lead to a system that is not compliant with today’s EMV back-ends. This would also likely increase the duration of a payment, infringing on the timing constraints that exist in place today (*i.e.*, a contactless payment end-to-end takes only 12 seconds [20]). So, a payment scheme answering RQ1 will likely follow closely existing EMV payments, adding to it not privacy-enhancing cryptography but rather privacy-enhancing parties such as anonymising proxies (also called instant escrows). Indeed, privacy proxies are a well-known privacy-enhancing technology (PET), accepted and catered for even by internet standardisation bodies [54]. In our solutions, we pursue this idea; see Section 6.

RQ2: If EMV-compliant privacy preserving payments rely on privacy-enhancing proxies being added to the infrastructure, how do these proxies fit in with pseudonymity-hindering AML, SCA laws and fraud prevention mechanisms?

At present time, these pseudonymity-hindering AML and SCA laws and fraud prevention mechanisms are primarily observed and implemented by the card-issuing banks. With this in mind, there are two avenues forward for payment-anonymising proxies in EMV: (a) the issuing banks provision these proxies themselves; (b) third-party proxies are used, but then there will be a trade-off of liability and risks between them and the issuing banks.

In this work, we provide two privacy enhancing designs, *PrivBank* (see Section 6.1) and *PrivProxy* (see Section 6.3), which can be plugged directly into the banking system. They modulate liability and risks differently, as per (a) and (b) above. We evaluate and compare their alignment with existing contactless EMV-based in-shop payment systems and prevailing regulations; see Section 6.2, 6.4 and 6.5.

RQ3: What type of privacy protection that complies with EMV and the law would be desirable and possible to achieve?

We formulate this notion of privacy with adequate

pseudonymity and unlinkability requirements of payment schemes (as well as a way to reason about them) in a way that is clear and easy to understand by laypersons, since they will need to choose the levels of privacy (formalised as pseudonymity and unlinkability of various entities) and products suited to them; see Section 3 (pseudonymity notions) and Section 7 (analysis for *PrivBank* and *PrivProxy*). Moreover, these privacy properties need to be general enough to fit not just EMV but also other payment systems, enabling direct comparison; see Section 4. That said, our privacy properties can be made generic in some ways (*e.g.*, akin to those applying to large classes of protocols [33]), and are limited in others (*e.g.*, not covering metadata attacks [39], or counter-based attacks [33]). The limitations of our privacy notions and comparisons with other models are discussed in Section 7.5.

2 Related Work

Research into EMV is vast, ranging from applied works such as [52, 61, 64] to formal treatments [25, 28–31, 40, 56, 58].

Formal Analysis of EMV. Most formalisms for EMV analysis are based on the symbolic/Dolev-Yao model [42], very few are computational (*e.g.*, [32]) like the one we give in Appendix D. Moreover, most formalisms focus on security, with few addressing privacy. We are the first to give a model based on mathematical relations (see Section 7) to encode privacy in EMV. The closest idea to this, not in EMV, speaks of traceability relations [33]; but these are complex links made between protocol layers. Next, we will cover directly related works, on privacy and traditional payment systems.

EMV Payments & Privacy. The card-to-PoS channel is insecure as per current EMV specifications, and [35, 53] work on next-generation EMV, where this channel will be secure. In this setting, considering a corruption model and linkability attack stronger than one against *Unlnk*, [53] find next-generation EMV payments to be linkable. The authors of [35] build on [53] and extend their model to dishonest terminals, achieving unlinkability and pseudonymity for smart card-based payments. Both proposals are non-EMV compliant.

If we go back to standard-EMV systems and their insecure card-to-PoS channels, then we should mention [29]; they showed that long-term *Primary Account Numbers* (PANs, identifier written on the payment card) can be used to track people. This gave rise to mobile-devices having “tokenised PANs”, which would go towards payers’ anonymity, were it not for the introduction of an element called “PAR (Payment Account Reference)” to mobile EMV payments. Such payments were studied, in particular w.r.t. tokenization, in [25, 38], but from security perspectives, not privacy ones.

We fill in these gaps in comparisons and analyses of EMV-payments’ privacy. Concretely, in Section 4 and Appendix B, we compare *our solutions* with the most common and contemporary payment schemes — ranging from cash and PayPal to

mobile payments via platforms like Curve and shop vouchers — evaluating them based on the privacy they offer and the regulations they are subject to.

3 Payments’ Privacy Notions

We propose privacy notions from the perspective of different entities in the payment systems: a *payer* who pays for goods/services, a *merchant* selling them, an *issuer* who gives the payer a means of payments (bank account, cards, banknotes, *etc.*), and a *proxy* who mediate the purchase from the payer to the merchant.

3.1 Entities Identification in EMV

Law-Enforced Payer Identifications. In EMV payments, the payer has to be identified at the on-boarding phase with their banks. To obtain a bank card, customers must provide a piece of identification such as a passport and proof of address to the card-issuer. These measures fall under the *Know Your Customer (KYC)* regulations. Secondly, in EMV payments, the payers have to be identified during transactions: knowledge (*e.g.*, PIN, Personal Identification Number), possession (*e.g.*, of the card) or physiological characteristics (biometrics) *etc.* must be checked as they pay; this is known as *Strong Customer Authentication (SCA)* and it is governed by Payments Security Directive (PSD2) regulations [50]. Thirdly, payers must be identified (even beyond SCA) for all transfers or payments amounting to certain values over a certain period (*e.g.*, EUR/GBP1000 per month in the EU/UK). Together, it is all driven by fraud protection (*e.g.*, in the case of SCA) and/or AML which includes the Money Laundering and Terrorist Financing Regulations (AML) regulations [51] [24] [23]. The Financial Conduct Authority (FCA)’s Handbook FCG 3.2.5 [27] requires regulated firms such as issuers to perform (real-time) *monitoring* of transactions and submit “Suspicious Activity Reports” (SAR) (containing the payer identity with the payment details) to the FCA, if concerns arise. Such reports lead to the full identification of the payers and their actions, rendering our privacy notions inapplicable under SAR.

Merchant Identification. In EMV payment systems, the issuers usually get the following merchant information with each of their customers’ transactions: *Merchant Category Code (MCC)*, *merchant’s name*, *Merchant Risk Index (MRI)*, *Merchant Location (ML)*.

Issuer Identification. The identity of the issuer is generally disclosed to the entities participating in a payment, and is of no interest to third parties. If there is a proxy in the system, this proxy may learn who someone banks with, as is the case in many existing settings, such as those offered by Curve, Monzo, and similar services.

3.2 Our Payment-Privacy Notions

Now, we are in a position to put forward a set of desirable privacy notions, general enough that they can be meaningful in the context of most payment systems. These properties can be from the perspective of the different relevant entities involved: merchants, issuers, proxies and observers.

Payer Pseudonymity (P_{PS}). An instance of P_{PS} holds if a given entity cannot recover the payer’s identity ID or a long-term pseudonym based on actions its sees.

Also, P_{PS} comes in various flavours w.r.t. what is the identifiable “object”. For instance, in EMV-like payment systems:

P_{PS}^{ID} : imposes that the merchant does not learn the payer’s long-term identity ID ;

P_{PS}^{CID} : imposes that the merchant does not learn the payer’s long-term card-number or bank account.

However, in EMV payments, even those made via mobile phones, the merchant always learns a long-term pseudonym of the payer. While there is generally no direct relationship between P_{PS}^{ID} and P_{PS}^{CID} , as shown below, our constructions achieve both simultaneously.

$P_{PS}^{CID} \not\Rightarrow P_{PS}^{ID}$. For instance, paying for goods with the long-term card currently reveals the PAN/PAR of the card CID to the (PoS of the) merchant, but the identity of the payer ID is not disclosed in the payment.

$P_{PS}^{ID} \not\Rightarrow P_{PS}^{CID}$. A payer utilising a payment-proxy system (*e.g.*, Revolut, see Appendix B) may undergo a KYC procedure with the proxy thus leaking their ID , but said proxy may never know a long-term card CID of the payer as they may always top up their proxy account from one-time cards.

Payments’ Unlinkability (Unlnk). An instance of Unlnk holds if a relevant entity in payment systems will stay unable to link payments made by the same payer.

So, we ask a merchant not distinguish if two payments are made by the same payer, as with cash transactions, say.

Merchant/Seller Pseudonymity (M_{PS}). An instance of M_{PS} holds if a payment-system entity (*e.g.*, an issuer) cannot infer the identity of the merchant involved in a payment.

Discussions. P_{PS} and Unlnk are considered in front of the merchant or an observer, since the issuer may always know the payer’s identity due to KYC. Similarly, M_{PS} sits most naturally in front of the issuer.

Additional relations between these properties apply: payer pseudonymity is required for unlinkability, *i.e.*, $Unlnk \Rightarrow P_{PS}$. Indeed, an entity deducing the identity of the payer or their information persistent across transactions can link said transactions. However, the reverse does not hold, *i.e.*, $P_{PS} \not\Rightarrow Unlnk$. For P_{PS}^{CID} specifically, an example of this is provided by EMV payment tokenisation [46]: there, each payment yields an ephemeral identifier for the card/payer called *tokenised Primary Account Number (PAN)*: *i.e.*, P_{PS}^{CID} holds. But payments by the same payer are linkable together (by the merchant and any observer) via an EMV element called *Payment*

Payment method	SCA		KYC		Pseudonymity		
	Issuer	Proxy	Issuer	Proxy	Issuer	Merchant	Proxy
1. Cash	no	×	×	×	M_{PS}	P_{PS}, Unlnk	×
2. Cheque	yes	×	yes	×	$\neg M_{PS}$	$\neg P_{PS}, \neg \text{Unlnk}$	×
3. E-cash	yes	×	yes	×	M_{PS}	P_{PS}, Unlnk	×
4. Physical cards	yes	×	yes	×	$\neg M_{PS}$	$\neg P_{PS}, \neg \text{Unlnk}$	×
5. Google, Apple Pay, etc.	yes	×	yes	×	$\neg M_{PS}$	$P_{PS}, \neg \text{Unlnk}$	×
6a. Top-up cards	yes	(yes)	yes	(yes)	$\neg M_{PS}$	$P_{PS}/\neg P_{PS}, \neg \text{Unlnk}$	$(\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS})$
6b. Pre-Paid/gifts cards	no	(no)	no	(no)	$M_{PS}/\neg M_{PS}$	$P_{PS}, \text{Unlnk}/\neg \text{Unlnk}$	$(\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS})$
7. Virtual cards	yes	(yes)	yes	(yes)	$M_{PS}/\neg M_{PS}$	P_{PS}, Unlnk	$(\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS})$
8a. PayPal	yes	yes/no	yes	yes/no	M_{PS}	$P_{PS}, \neg \text{Unlnk}$	$\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS}$
8b. Curve	yes	yes	yes	yes	$\neg M_{PS}$	$\neg P_{PS}, \neg \text{Unlnk}$	$\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS}$
9. Online Marketplaces	yes	no	yes	yes/no	M_{PS}	$P_{PS}, \neg \text{Unlnk}$	$\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS}$
10. PrivBank	yes	no	yes	no	M_{PS}	P_{PS}, Unlnk	$P_{PS}, \neg \text{Unlnk}, \neg M_{PS}$
11. PrivProxy	yes	yes	no	yes	M_{PS}	P_{PS}, Unlnk	$\neg P_{PS}, \neg \text{Unlnk}, \neg M_{PS}$

Table 1: SCA, KYC and pseudonymity properties of payment methods from the point of view of the Issuer, Merchant and the Proxy when it exists. $\neg P_{PS}$ and $\neg \text{Unlnk}$ holds in all systems for the Issuer and $\neg M_{PS}$ holds in all systems for the Merchant. Detailed explanations are provided in the Appendix B. (Notation: \times stands for “not applicable”, P_{PS} for P_{PS}^{ID} and P_{PS}^{CID} , brackets are used when the proxy may not necessarily exist in all systems and / when deployment can lead to different properties.)

Account Reference (PAR): i.e., Unlnk does not hold.

3.3 Threat Model

For the privacy-preserving properties of PrivBank and PrivProxy, we assume that parties, payers, issuers, merchants, proxies, can be corrupted, at any point in the execution, and they can be made to behave arbitrarily. Yet, due to AML auditability requirement, the issuer and proxy involved in a given executions cannot be simultaneously corrupted, otherwise, in compliance with the law, trivially breaking any pseudonymity property. Thus, excluding all scenarios where identity disclosure is deemed necessary, such as through *suspicious activity reports* mechanisms or other mandatory auditing requirements. Our corruption and adversaries vary with the properties. In general, any party can be corrupted, except for specific ones designated for each property.

For P_{PS} and Unlnk:(1) the payer against whom the property is considered and at least one other banking with the same issuer; (2) an adversary can eavesdrop payments made by the payers to merchants’ PoS¹;

For M_{PS} :(1) the merchant against whom M_{PS} is considered and at least one other; (2) an adversary can eavesdrop all payments submitted by the proxies to the issuers².

We also assume two realistic aspects. One, the (application implementing our solution on the) payer’s phone will not be corrupted, unless the Payer is also corrupted. Second, other

¹This is pertinent in the case where the channel between the PoS and payer is public/un-encrypted, which is the case today; in this setting, the eavesdropper adversary is weaker than a corrupted merchants/payers. However, there are proposals to make this channel secure in EMV 2nd Gen [44].

²This is a very strong adversary in practice, as it would mean that there is breach in the backend of the payment networks.

data (values, dates, etc.) in funds’ transfers are independent of the Payer’s long-term identity; it means that the Payer does not encode their identity in the time or value of the payment. Moreover, we assume that spending caps, applicable to all payment methods and potentially enforced in PrivBank and PrivProxy due to AML legal requirements, are never exceeded by payers when assessing the Unlnk property. Similarly, we assume that the two honest payers requesting a payment always share the same acceptance profile, leading to indistinguishable acceptance of their transactions.

4 Scrutinising KYC & Privacy in Payments

We now empirically evaluate *traditional payment systems* and some *modern* ones, and discuss where they sit w.r.t. KYC and SCA regulations and how they fare against our privacy requirements P_{PS} , Unlnk, and M_{PS} . We exclude crypto currencies [57] and QR-code-based payments [12]. The reason for this is that their infrastructure is totally different from the rest of the long-established payments, especially the EMV-based systems, which we aim to augment here. Thus, the systems of interest in our analysis here are: 1. cash, 2. cheque, 3. e-cash [36], 4. physical/classical bankcards, 5. mobile-phone apps [11, 15], 6. top-up and pre-paid cards [7, 13], 7. virtual/one-time cards [17], 8. payment service providers such as PayPal [6] and Curve [8], 9. online marketplaces [9, 16]. The readers are likely familiar with most of these systems and can judge if they have SCA, KYC, P_{PS} , Unlnk, M_{PS} , so we left most of their full descriptions in Appendix B. We give below just the description of Curve, as it is less well-known.

Curve. Curve [8] is a payment proxy providing a payer with a card and a payment application. Curve users must

satisfy the KYC rules. The payer registers multiple bankcards issued by one or several banks in the Curve app. When paying with a Curve card, authorisation by the payment network (*e.g.*, Visa, Mastercard) and the bank is carried out on the basis of Curve card data and merchant information, as it is the Curve card interacting with the merchant’s PoS. Curve then pays the merchant on behalf of the payer on the spot and one of the payer’s Curve-registered bankcards is charged. So, there is an intricate payment authorisation process between Curve, traditional bankcard issuers and the payment network.

Analysis of KYC, P_{PS} , Unlnk, M_{PS} in Payments. Table 1 surveys the way payment systems fare against SCA and KYC (described in Section 3.1), and our notions P_{PS} , Unlnk, M_{PS} (Section 3.2). We got to these results via an empirical analysis, brief justifications for our claims are provided in Appendix B. For instance, e-cash, physical cards, payment apps (Google Pay, Apple Pay, *etc.*), top-up and virtual cards are payment methods for which the SCA/KYC apply to the issuer unlike pre-paid cards which fall under exemption cases. PayPal applies strict limits unless the customer is identified, rules application therefore depends on the usage scenario.

Takeaway Message. Cash (rows 1 in Table 1) offers the best pseudonymity and unlinkability, with no SCA/KYC, but is arguably not the modern commodity. In turn, modern payment methods (*e.g.*, rows 4, 5, 7 in Table 1) need to adhere to SCA and/or KYC, and in doing so lose most privacy. Also, when a payment proxy is used (*e.g.*, rows 6 to 9 in Table 1), merchant pseudonymity is typically lost, in part due to AML requirements. Given the variety of results in Table 1, this section suggests, that meeting in the middle between proxied solutions, cash and modern payments methods (like EMV cards, mobile wallets) in terms of P_{PS} , M_{PS} , Unlnk properties, KYC/SCA regulations, while maintaining the EMV infrastructure and complying to AML as well, is likely non-trivial.

5 Our Main EMV Ingredients

We recall the notions³ related strictly to EMV payments which are relevant to us.

5.1 From Card Issuing to Payment Processing

We divide EMV card-based payments in 4 main stages.

1. Card-issuing, KYC & AML. A future Payer opens a bank-account with a *card issuer* (*i.e.*, bank). We discuss the case where they receive a credit/debit card, associated with the account. The card is supplied by one of the current *card providers* (*e.g.*, Mastercard, Visa, American Express, *etc.*). We will refer to the collection of Issuers, card providers and the proxies linking them loosely as *payment networks*. To obtain such a card, the customers must provide a piece of

³Capital letters are used to refer formally to entities in the system: *e.g.*, “payer” – a personal paying, in Sections 1-4, vs “Payer” – a formal algorithmic party, from Section 5 onwards.

identification such a passport and proof of address to the card issuer, in line with KYC regulations. KYC falls under the AML regulations [51] (see Section 3.1).

2. Making a Card Payment & Relevant Payment Data: PAN, MCC, ML. A *cardholder* goes to pay with their card to a *Merchant*, using their card readers, known also as *points of sale*: the cardholder inserts the card into the PoS or taps the PoS in the case of a contactless transaction. The basic operations of the protocol between the card and the PoS are defined in the set of standards from EMVCo. A partial example of the data exchanged between the card and the reader, called the *payment transcript*, is available in Appendix F.

Due to the transmission of the PAN, plastic-card EMV does not have P_{PS} from the viewpoint of the Merchant, or an observer between the card and the Merchant’s PoS.

At the end of the protocol, certain transaction data is signed by the card and returned to the PoS for its checks. Equally, the card sends a MAC of certain transaction data to the PoS to forward it to the card-issuers for checks therein. They check this data and based on it, they approve/decline the transaction.

Alongside with the card-centric data sent on the back-end from the Merchant’s PoS to the Issuer, payment networks and others, the Merchant’s PoS also adds all or some of the following merchant-identifying details, relevant to us: *Merchant category code (MCC)*, *Merchant’s name (MN)*, *Merchant risk index (MRI)*, *Merchant location (ML)*. Thus, due to their transmission of the MN, plastic-card EMV does not have M_{PS} from the viewpoint of the Issuer.

3. Customer Identification During Payments. When the card is presented to the Merchant’s PoS for payment, the SCA/PSD2 [49, 50] regulation require two factor authentication of the Payer (*e.g.*, possession of card and associated PIN). The Issuer checks the payment data sent by the Merchant along with this SCA identification-data of the payer. Should the checks fail, the payment is declined by the bank. There are variations to card and PIN verification, especially if the payment is not made by card⁴. If the payment is contactless, derogation from the SCA rules can apply and single-factor authentication is required instead. SCA remains required every few payments or after a set of payments has exceeded set value (*e.g.*, EUR/GBP150 in EU/UK).

4. Payment Authorisation and Clearing. Funds are settled during the final phase, called *clearing*, as follows. (i) The Merchant, via their acquirer, requests payment from the card issuer. The issuer verifies details like transaction location, payer identity, and merchant information. (ii) If the cardholder has sufficient funds, the issuer deducts the amount from their account⁵. The final authorization is handled by the issuing

⁴*E.g.*, if the Payer uses a smartphone SCA verification by the issuing bank is replaced by *Consumer Device Cardholder Verification Method (CDCVM)* executed on the phone. That is, the payers fingerprint or face identification is read by the phone and used as customer authentication. The result of that is later checked by the issuing bank.

⁵This is for debit cards. For credit cards, this differs slightly.

bank, possibly in consultation with payment networks like Visa or Mastercard. Once approved, the funds are transferred to the Merchant’s/acquiring bank.

As explained in Appendix C, the information necessary for a payment authorisation varies based on the business model (*e.g.*, from Visa to Mastercard) and not all Merchant information is truly necessary. For example, Curve [8] operates in the following way (and *e.g.*, Visa incentivises it [62]): they over-submit Merchant data especially if their MRI is high, to increase the probability of authorisation and therefore maintain customer satisfaction. There may be leeway in provisioning M_{PS} from the viewpoint of the Issuer, since the minimal amount of Merchant data needed is not standardised.

5.2 Mobile-EMV Tokenisation

Mobile payment applications such as ApplePay [15], Google Pay [11] *etc.* allow registering plastic cards to pay via a mobile application. The onboarding requires an authorisation from the card’s issuing bank, and therefore KYC is observed. The Payer can then use the app for contactless payments. When a payment is made, the card’s long-term PAN is *tokenised*, and the payment transcript between the phone and the Merchant’s PoS looks different from one made with the physical card, with PAN-related data replaced by tokenised values. Mobile-payment transcripts (see Appendix G for an example) include the following payer-identifying data relevant to us:

- *one-time tokenised PAN* – an ephemeral account number that changes with each payment and each app: each payments made with card C through mobile app A_1 or app A_2 will each generate a different number.
- *long-term PAR* – a fixed value that is shared amongst various/all payment apps $A_1, A_2 \dots$ to refer to *any/all* payment made based on the same physical card C ; the PAR was introduced at the request of the Merchants and payment networks, so that mobile payments made with one card C , though showing varying tokenised PANs, can all be linked together.

The tokenised PAN and PAR are sent by the Merchant onto the payments networks, just as the “plastic” PAN was. However, before these reach the Issuer, the tokenised PAN is de-tokenised by entities in the EMV system who transform it back to the associated PAN. All the rest of the backend part of payment processing is as described in Section 5.1. Details of payment tokenisation and use cases are given in [46, 48]. So from the viewpoint of the Merchant, mobile EMV payments achieve a form of P_{PS} , via the PAN, but do not achieve Unlnk, due to the PAR.

Takeaway Message. In our designs, we carefully combine the ideas explained above: (1) from mobile payments – namely, we will make the PAN/PARs be one-time/removed (to achieve payer pseudonymity), and make all their long-term data randomised for each transaction (to achieve payment unlinkability); (2) from payment-authorisation – Curve already

sends selective payment-authorisation data to the Issuer; we will further filter or replace that data through a proxy to achieve merchant pseudonymity. On top, astute orchestration between a proxy and the Issuer, as well as careful protocol design, will lead our designs to achieve EMV-compliance and abiding by AML, KYC and SCA regulations. This intuition is developed further in Appendix C.

6 Anonymous EMV In-Shop Payments

Now, we propose two constructions, compatible with EMV contactless payments, providing privacy as per P_{PS} , M_{PS} , Unlnk, with provable guarantees, all the while being compatible with the aforementioned law and regulations. Legal frameworks may vary locally, our solutions may require adjustments to accommodate varying regulatory environments.

At the core of our first construction called *PrivBank*, given in Figure 1, there is a privacy-friendly issuing bank who provisions P_{PS} and Unlnk for its customers. To do this, this bank strongly partners with a Proxy who mediates and curates customers’ payments providing M_{PS} . Meanwhile, at the heart of our second construction called *PrivProxy*, given in Figure 2, there is no longer a bank, but rather a pseudonymity-friendly Proxy which aims to provide P_{PS} , M_{PS} and Unlnk of its own accord and at its own risks, to Payers who bank with whoever they chose to, independently of the Proxy.

The crux of our designs is to compose several standard, non-private EMV-payments or parts thereof, such as to obtain one mobile, contactless EMV payment which attains P_{PS} , M_{PS} , Unlnk. We realise this via the design and use of proxies (also called instant escrow), without modifying EMV elements in the original payments, and without cryptographic additions. As such, all the cryptography used in our schemes and all EMV building blocks can be treated as black-boxes inherited from EMV, and our only focus is going to be the design of the proxied systems, from an engineering perspective alone. Indeed, our proofs w.r.t. P_{PS} , M_{PS} , Unlnk follow from the proxied construction, and the cryptographic or inner protocol details (*e.g.*, Visa, Mastercard variations) are irrelevant therein, as they are in the descriptions that follow.

We will now describe the functionality of our proposals and answer our first research question RQ1 by describing the main aspects and intricacy of *PrivBank* and *PrivProxy*.

6.1 Construction *PrivBank*

In *PrivBank* (Figure 1) a Proxy, contractually committed with the bank, intermediates all in-store transactions done by a Payer with a Merchant. The Proxy gets to know who the Merchants are but not the long-term identifiers of the Payers, whereas the bank knows who the Payers are but not the Merchants. As is required by AML regulation, the two entities can work together to recover full knowledge on any transaction. They have an agreement in place, thus sharing

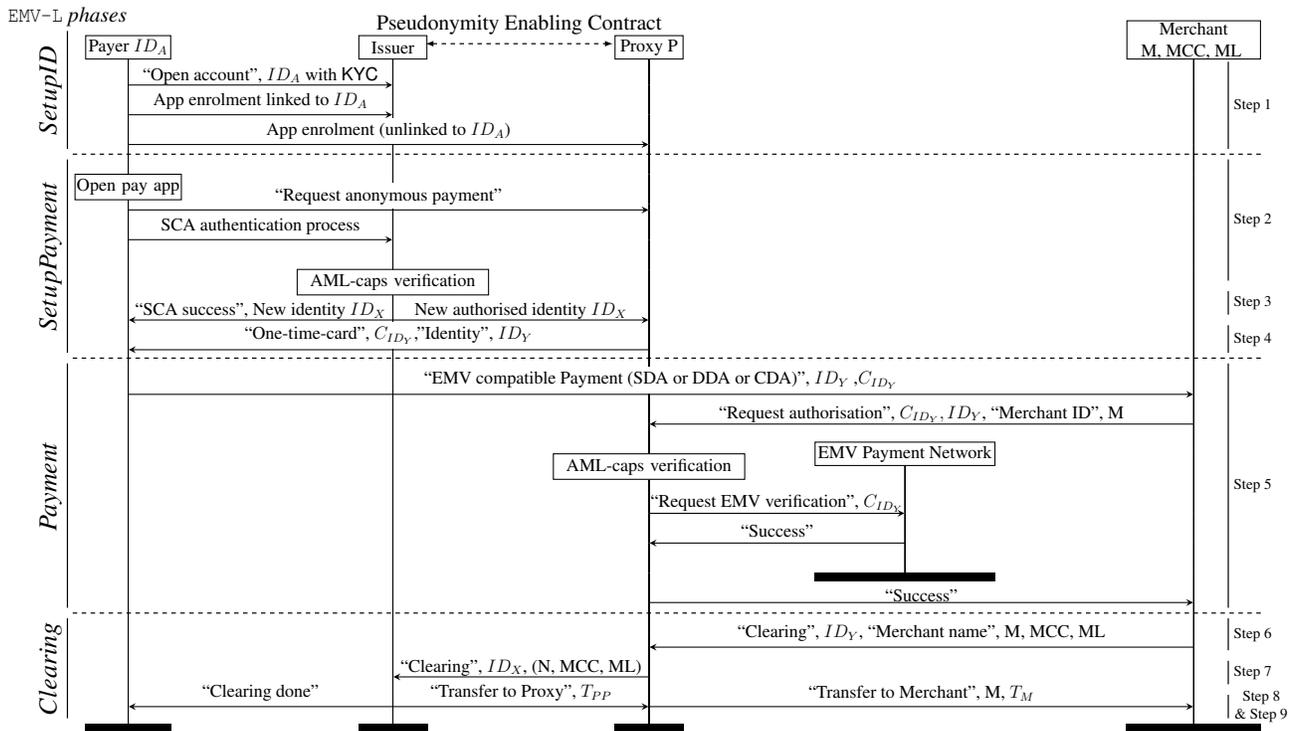


Figure 1: Protocol/Implementation flow of PrivBank. All communications apart from the payment from the Payer to the Merchant are assumed to be executed on a secure channel (encrypted and authenticated communications).

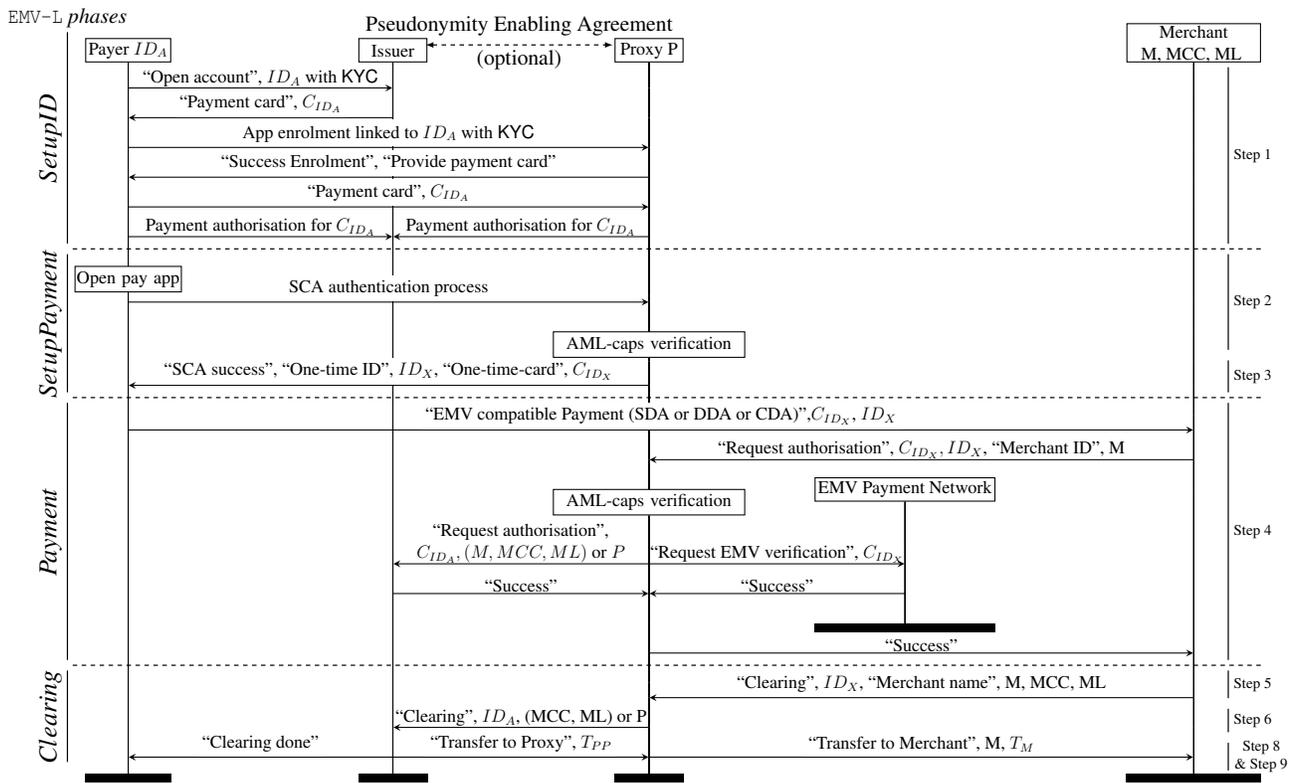


Figure 2: Protocol/Implementation flow of PrivProxy. All communications apart from the payment from the Payer to the Merchant are assumed to be executed on a secure channel (encrypted and authenticated communications).

risk and liability. It's important to note that an agreement can create civil liability, e.g., to indemnify the bank for a fine if the Proxy fails to do something it has promised but it cannot apportion criminal liability or liability to pay a fine. Civil liability to the customer would only be by the bank. Payments done via `PrivBank` are supported via one mobile app⁶ provided in partnership between the Issuer and the Proxy. We now describe, step by step, how a payment is made possible as well as carried through via `PrivBank`. These steps are also highlighted in Figure 1.

Step 1 (Registration, see *SetupID* in Figure 1). The Payer, whose identity is " ID_A " opens a bank account with the Issuer. This bank account comes with a "premium" option of support for privacy à la $P_{PS}, M_{PS}, Unlnk$, which is achieved via `PrivBank`. The Payer's banking account and their `PrivBank` account are with the Issuer, which handles KYC authentication, not the Proxy. The Issuer does not share the Payer's identification details with the Proxy. To allow this, the contract stipulates certain terms and conditions (T&C) as we will detail below.

As an account holder with the Issuer, the Payer accesses the `PrivBank` app, which connects the Payer, Issuer, and Proxy from an engineering standpoint. However, the app is provided by the Issuer and links only the app-store identifier to the Payer's identity ID_A on the Issuer's servers. The Proxy and other third parties identify the Payer through their app-store account, not the Issuer's method.

Under AML regulations, and similarly to other payment methods, the amounts spend using `PrivBank` may be capped (e.g., EUR/GBP1000 per month); we call these the *AML caps*. The T&Cs set this limit, which the Issuer and Proxy enforce.

Step 2 (Authentication, see *SetupPayment* in Figure 1). When a payment is to be made by the Payer to a PoS of a Merchant's, the Payer ID_A opens the `PrivBank` app. The opening of the app prompts both the Issuer and the Proxy on secure (e.g., HTTPS) channels:

(a) *App-Triggers on the Proxy's Side:* At this stage, the push by the app to the Proxy only says that someone, *with no specific identity revealed*, is intending making a payment.

(b) *App-Triggers on the Issuer's Side:* The SCA and AML checks are triggered through a request to the Issuer. SCA is a two-factor authentication, ensuring that Payer ID_A is making the payment. If needed, the Issuer also verifies that ID_A has not exceeded the AML caps for `PrivBank`. If the caps are reached or SCA fails, the protocol halts.

Upon successful SCA and AML checks, the Issuer creates a one-time virtual identity ID_X , pseudorandom and statistically independent of ID_A and any existing long-term card C_{ID_A} .

Step 3 (Pseudo-identity issuance, see *SetupPayment*). On the back-end (i.e., not via the `PrivBank` app), the Issuer sends the identifier ID_X to the Proxy, which, in the light of

Steps 1 and 2, only knows that an account holder with the Issuer using `PrivBank` wants to make a payment.

The Proxy expects this push⁷, having received an alert in Step 2 that someone intends to make a payment.

Step 4 (One-time card issuance, see *SetupPayment*). At this stage, the Proxy issues – for whom it knows as Payer ID_X – a one-time, virtual, EMV-compliant card C_{ID_Y} with all aspects (PAN, certificates *etc.*) being freshly generated for one-time use, including an attached one-time, virtual card-holder name of " ID_Y ". ID_Y/C_{ID_Y} are pseudorandom and statistically independent of ID_X and of ID_A/C_{ID_A} . The card issued is "loaded" onto the app.

Step 5 (Payment, see *Payment* in Figure 1). The Payer pays with it at the Merchant's PoS. Here, SCA authentication may have to be carried out offline using the CDCVM tag in the case of preloaded payment methods ID_Y/C_{ID_Y} . The Proxy (alternatively the Issuer) checks that it would not reach the AML caps through this specific amount being paid via `PrivBank` and executes the AML scrutiny. If any of these conditions fail, the protocol stops. AML caps were checked in Step 2 of `PrivBank`, but those checks did not account for the current payment.

Liability Shifts & Fraud-protection. Under its partnership with the Proxy, the Issuer accepts controlled *shift of liability* with respect to fraud protection. To this end, for *selected stores* – that are nominated based on MCC, MRI and ML, *etc.* – the partnership allows that the Issuer receives from the Proxy sanitised information. In practice, the *list of selected stores* can be large (e.g., all Merchants in a country with given MCCs), as is the case for the "*Ticket Restaurant*" services with Edenred [10] or Up-one [18]. The sanitised information does not reveal the original Merchants' full identity, instead it contains what we call *pseudo-merchant identities*. These are prescribed, such that the Issuer can check the Proxy's compliance to the agreement⁸.

In terms of fraud-detection disputes, the Proxy and the Issuer have to come together to resolve this, and the Proxy has to disclose to the Issuer the full Merchant data. This is reflected in the T&C of the contract that the Payer has with the Issuer on using the `PrivBank` product, i.e., the Payer knows that it can use `PrivBank`, in selected stores.

Step 6 (Merchant clearing operation, see *Clearing* in Figure 1). Using standard EMV mechanisms, the (Acquirer of the) Merchant begins to resolve the payer's payment by contacting the Proxy, which is the Issuer of C_{ID_Y} .

Step 7 (Proxy clearing operation, see *Clearing*). If the Merchant M is not on the "pre-selected" list, the protocol stops. Otherwise, using the `PrivBank`'s back-end, the Proxy goes to the Issuer to resolve a payment for Payer ID_X , and provides a *pseudo-merchant identity N* instead of the true

⁶For compliance with banking regulations, this app may require a Trusted Execution Environment (TEE), a chip designed for secure storage and cryptographic operations.

⁷The time between any push by the app in Step 2 and receiving ID_X is capped at 2 seconds due to EMV security constraints.

⁸E.g., a Merchant M in the country, with Merchant Location and Merchant Category Code getting pseudonymised as a fixed pseudo-merchant identity.

identity of the Merchant M.

Step 8 (Balance adjustment, see *Clearing*). The Issuer checks if Payer ID_X can pay to a pseudo-merchant N via PrivBank as per the pre-agreed list of merchants and as per the rules of PrivBank. Then, the Issuer further checks that Payer ID_X has funds to pay.

Step 9 (Balance adjustment, see *Clearing*). If step 8 went through, the payment is resolved towards the Proxy and then from the Proxy to the Merchant.

6.2 PrivBank Law & EMV Compliance

PrivBank align with the specification and all regulations applicable to the banking system. Firstly, note that when a payment is made with the PrivBank app, the transcript is the same as in one made with an EMV contactless payment card. The CDCVM is the only different aspect to standard EMV, and we detail this in the four paragraphs below; these also answer our research questions RQ1 and RQ2.

On Compliance with SCA. Strong customer authentication is checked by the Issuer when Payer ID_A intends to pay at Step 2 to provide ID_Y/C_{ID_Y} to Payer ID_A . Note that an SCA authentication (that may be carried out offline using the CDCVM tag) may be required for the payment if it is not executed within a few seconds after the first one.

On Timing Compliance w.r.t. the EMV System. EMV-compliant payments are required to set a maximum general processing time for each transaction. Allowed timings range between a few hundred milliseconds to a few seconds. In the stages defined in PrivBank, Step 1 corresponds to an initial setup independent of any payment. Steps 2 to 4 involve SCA authentication, up to the point where the card is loaded into the app. This process can be executed ahead of time for one or several one-time virtual cards. Steps 5 to 9 exactly correspond to a timed EMV process of payment. Over all, PrivBank results in a processing time within the range of current payment standards and similar to already deployed solutions such as tokenization or Curve. Thus, from a technical point of view, PrivBank complies with the standard.

On Compliance with KYC. KYC regulations are fulfilled via the Issuer, who checks Payers' identification documents upon them opening a bank account. A contract in between the Issuer and the Proxy mandate the Issuer for the identity verification. On the other side, there is a liability shift towards the Proxy on the verification of the Merchant's identity.

On Compliance with AML. The T&Cs of PrivBank subscribers are such that the amount of payments that any Payer ID_A makes via PrivBank will be capped to values as per AML regulations (e.g., EUR/GBP1,000). Upon the AML verification done by the Proxy, if a breach is found by the Proxy, then an alert would be sent to the Issuer. The Issuer's officers would investigate and generate a so-called "Suspicious Activity Report (SAR)" [55] would send this to the authorities, since ultimate AML liability in PrivBank sits with the Issuer.

The latter would need to collaborate with the Proxy to solve the case, and –for this– any Proxy and Issuer in PrivBank would need an initial legal agreement. The privacy of entities will be reverted in case of a SAR investigation.

6.3 Construction PrivProxy

PrivProxy (Figure 2) is based on the same three main parties, but while at the core of PrivBank there is a "pseudonymity-friendly" Issuer, now, in PrivProxy, it is a Proxy who provides a service to add pseudonymity on top of EMV payments. Note that there could be many such Proxies.

Step 1 (Registration, see *SetupID* in Figure 2). The Payer is known via their identity as " ID_A " by the Issuer and the Proxy, and it holds accounts with both. During the onboarding process, the Payer ID_A links the bank account they hold with the Issuer with the user account they hold with the Proxy.

EMV Compliance. At onboarding with PrivProxy, banking pre-authorisation is conducted, where the Proxy can take up to a fixed total from the Payer's bank account. In line with the EMV rules, this *pre-authorisation cap* can be a maximum of x amount per year/month/day (e.g., EUR/GBP1000/month). The Payer gets access to the PrivProxy app as a Proxy-provisioned service and has to comply with KYC procedure via their identity ID_A to use it. This time, the app, from an engineering perspective, has nothing to do with the Issuer.

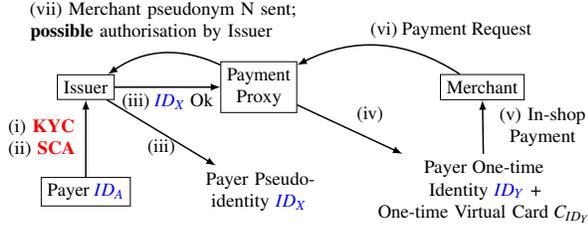
Step 2 (Authentication, see *SetupPayment* in Figure 2). When a payment is to be made by the Payer ID_A to a Merchant with identity M, the Payer opens the PrivProxy app. The opening of the app prompts the Proxy to do the SCA process, see *SetupPayment* in Figure 2. Upon successful SCA checks, the Proxy checks that ID_A has not reached their AML-cap. If they have, the protocol stops.

Step 3 (One-time card issuance, see *SetupPayment* in Figure 2). At this stage, for Payer ID_A , the Proxy creates a one-time virtual identity and a one-time EMV-compliant card, shown as ID_X and C_{ID_X} in Figure 2; they are pseudorandom and statistically independent of ID_A and C_{ID_A} . The card issued is automatically loaded into the PrivProxy app. The transcript produced between the app and the Merchant's PoS, when paying with this app, is that of the EMV physical card except for CDCVM.

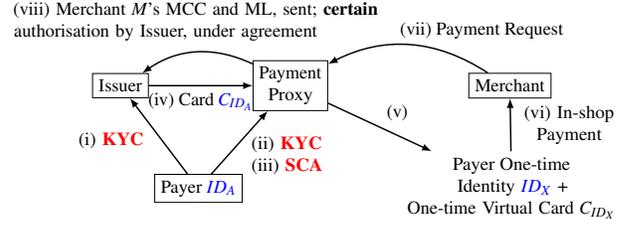
Step 4 (Payment, see *Payment* in Figure 2). The Payer goes to pay with its one-time EMV-compliant virtual card by the Merchant's PoS. Accounting for the value of the current payment, the Proxy checks that the Payer's pre-authorisation caps are not reached, and nor are the AML-caps and executes the AML scrutinising.

Step 5 (Merchant clearing operation, see *Clearing* in Figure 2). Using standard EMV mechanisms, the (Acquirer of the) Merchant begins to resolve the payer's payment by contacting the Proxy, which is the Issuer of C_{ID_X} .

Step 6 (Proxy clearing operation, see *Clearing* in Fig-



(a) PrivBank: EMV-compliant payments with pseudonymity provisioned by privacy-friendly Issuer, using a third-party Proxy.



(b) PrivProxy: EMV-compliant payments with pseudonymity provisioned by third-party Proxy.

Figure 3: Contrasting our proposals, especially on KYC, SCA, and identities.

Figure 2). The Proxy behaves differently if the Merchant is on their pre-vetted selected stores list to use `PrivProxy` or not. In either case, the Proxy aims to provide Merchant pseudonymity and proceeds as described below.

If the Merchant is on the “selected-stores” list, then, using the payment networks’ back-end, the Proxy matches ID_X with the payer’s identity ID_A and asks the Issuer to resolve a payment for Payer ID_A . It does not declare Merchant identity M but its own identity P as being the Merchant.

If the Merchant is not on the “selected stores” list, the Proxy still does not declare the Merchant’s identity M to the Issuer. Instead it provides restricted information (e.g., ML and MCC) using the payment network’s back end. In this case, the Proxy takes on risks in terms of their authorisation rates (i.e., it is possible that the Issuer will not approve the payment due to too little information on the Merchant).

Liability & Fraud-protection. In terms of fraud-detection disputes, the Proxy takes on the liability. This is stipulated in the contract it has with the Payer, meaning it will have to reimburse the payer in some cases as it is the payment provider. Should the Issuer need to be involved, this is done entirely by the Proxy, with no legal obligation on the Payer.

Step 7 (Balance adjustment, see *Clearing*). If the payment is not authorised by the Issuer (which is unlikely for Merchants on the selected-stores list), the protocol stops. Otherwise, the Issuer checks if Payer ID_A has funds to pay and approves the transaction if it does.

Step 8 (Balance adjustment, see *Clearing*). If all went through the payment is resolved towards the Proxy and then from the Proxy to the Merchant M .

6.4 PrivProxy Law & EMV Compliance

`PrivProxy` also align with the EMV specification (answers RQ1) and all regulations applicable to the banking system (answers RQ2). Arguments similar to those set out in Section 6.2 apply. We detail below.

On Compliance with EMV Specifications. EMV-compliant one-time cards are issued and delays are managed in the same way as in `PrivBank`. Here, it is also possible to pre-load the

pseudo-identity ID_X and the card C_{ID_X} in order to carry out EMV contactless payments with CDCVM authentication.

On Compliance with KYC. KYC regulations are fulfilled by the Issuer and the Proxy, who check Payers’ identification documents upon them opening accounts with each.

On Compliance with SCA. Since the Proxy did KYC on-boarding of the Payer, the Proxy can do the SCA step and checks that it is indeed Payer ID_A attempting to pay.

On Compliance with AML. Like in `PrivBank`, here we have AML Caps and “Suspicious Activity Reports (SAR)” [55] if they are broken. But, unlike in `PrivBank`, the Proxy has all the payer’s data and has ultimate responsibility in terms of AML; they will raise a SAR directly with the authorities, but the Issuer will be involved, as bound by KYC and other banking regulations. The privacy of entities will be reverted in case of a SAR investigation.

6.5 Comparing PrivBank & PrivProxy

Different Designs. The best way to see differences in designs between `PrivBank` and `PrivProxy`, as well as in the most-relevant protocol steps is Figure 3. A comparison between key properties of the two designs is given in Table 2.

Varying Features. For instance, `PrivBank` requires additional assumptions to enhance pseudonymity, such as a legal agreement of collaboration between the Proxy and the Issuer (Table 2, row 7). Conversely, `PrivProxy` could be offered by the Proxy independently of an Issuer, though legitimate payments may be rejected (row 10) due to the lack of such an agreement and the Proxy curating Merchant-related information of their own accord. Furthermore, `PrivBank` allows for multiple Proxies as partners of the Issuers, offering Payers a choice of providers; if it is the Payers who pay for the Proxy service, the balance of trust may shift, potentially making `PrivBank` more appealing to some Payers. Also, liability is shared in `PrivBank`, but not in `PrivProxy` (rows 3 & 4). `PrivBank` boasts stronger decentralisation of knowledge, but requires collaboration between Issuers and Proxies for the management of the app, as it necessitates a robust agreement (row 7). Furthermore, in reality, pseudonymity-friendly Issuers may be rare. Ultimately, the choice between the two

Criteria	PrivBank	PrivProxy
1. Payer pseudonymity	w.r.t. the Proxy and the Merchant 😊	w.r.t. the Merchant 😞
2. Merchant pseudonymity	w.r.t. the Issuer 😊	w.r.t. the Issuer 😊
3. Liability for legal compliance	Issuer and Proxy	Proxy
4. Liability for economic risk	Shared between Issuer and Proxy	Proxy
5. Payers' identities distributed	Yes 😞	No 😊
6. Payer's trust	In the Proxy and the Issuer	In the Proxy
7. System's assumptions	Trust between the Proxy and the Issuer 😞	None 😊
8. Security assumptions	Issuer's app not leaking identities 😞	None 😊
9. Feasibility	Privacy-friendly Issuers may be rare 😞	Immediately feasible by some companies 😊
10. False rejection rate	None 😊	Risks of low payment authorization rates 😞

Table 2: PrivBank and PrivProxy: Pseudonymity-provision, Advantages and Disadvantages.

proposals depends on priorities and incentives driving system deployment. Thus, we continue to present and analyse both PrivBank and PrivProxy, as they cater to different markets and operate under distinct assumptions.

Achieving P_{PS} , M_{PS} , Unlnk. PrivBank and PrivProxy offer different pseudonymity guarantees (see Table 1). Most of the relevant comparative aspects between our two protocols are detailed in Table 2 with the relevant design choices. illustrate that neither PrivBank nor PrivProxy can be deemed superior.

PrivBank achieves P_{PS} from the perspective of the Proxy, which is down to the Issuer and the Proxy having only partial access to identifying information. This is not the case in PrivProxy, so P_{PS} cannot be achieved there. All such results recounted in Table 1 are formalised in Section 7. The information essential for pseudonymity is divided between the Proxy (payment information) and the card Issuer (identification information). This is necessary to comply with legislation.

On Legal Compliance. For an overview of conformance with KYC and SCA regulations, see Table 1. In PrivBank, a legal agreement allows the Proxy to rely on the Issuer for KYC and SCA; thus, payments can proceed only when the Proxy has received SCA approval from the Issuer. This is unlike the case of PrivProxy where SCA is no longer required from the Issuer, as the Proxy has done KYC and there is also a pre-authorisation by the Payer on some of funds made to the Proxy during enrolment. In terms of AML, in PrivBank the main responsibility is with the Issuer, and in PrivProxy it is with the Proxy; but, in both cases, they share responsibility in the case of an AML official alert [55] and this is stipulated in their initial contracts.

On Using Proxies. The use of proxies in the banking system, such as Curve, and Revolut, is well-established, widely accepted, and utilised by millions of customers worldwide. Additionally, our design leverages the principle that banks do not require all merchant data for payment authorisation. This ensures that both the Issuer and the Proxy can potentially monetise the service, supporting the acceptability and feasibility of our approach. Meanwhile, the Issuer (*resp.* Proxy) does not have access to the Payer's full payment information (*resp.* ac-

count information), promoting greater data decentralisation and encouraging potential client adoption.

On Implementation Aspects. Both solutions/apps generate one-time cards (*i.e.*, single-use PANs) and run contactless EMV as for physical cards with the PoS. Loading a card mandates use of TEE to comply with the AML regulations. Alternatively, the server could load tokens associated with one-time cards. Tokens do not require to be securely stored. This second scenario generates a single-use PAR, but this element would lose its traceability purposes. Thus, we prefer the first solution. Apart from these minor considerations, our two solutions can be implemented on the basis of an adapted version of the services provided by existing proxies.

On Additional Demands. Our proposed solutions impose greater demands on the network: more communication and production of large volume of virtual cards. They also imply new shifts in liability for economic risk. Consequently, this form of payment may entail higher costs, which could be absorbed by merchants, as is the EMV specification, and/or by users opting for a "premium" privacy-preserving service.

7 Formal Treatment of Payments Anonymity

We introduce a formalism that allows us to reason formally about the privacy notions P_{PS} , Unlnk and M_{PS} . This formalism is accessible and answers RQ3. In Appendix E, we give a "traditional" cryptographic model and analysis of our schemes; we also show that the "simpler" definitions here fully capture the cryptographic definitions.

7.1 Execution Model

To study the security/privacy of (payment) protocols, we consider the parties in the protocol, denoted by *Parties*: *Payers*, *Issuers*, *Merchants*, and *Proxies*. These represent machines, devices or humans, associated with long-term identifiers, well-defined PPT (Probabilistic Polynomial Time) algorithms to execute, and they may hold cryptographic material. There can be any number of such parties, they are all executed concurrently and outputs of some are inputs for others, in the way

of Interactive Turing Machines (ITMs) [63].

7.2 EMV-L: A Language for EMV Protocols

To describe the main protocols executed by our payment parties, we define a language EMV-L that, similarly to an API (Application Programming Interface), describes the main procedures and sub-procedures of any EMV-compliant payment protocol. We omit the setup of the EMV system as it is already in place and focus on the registration and payment for the payer ID .

Definition 1 (EMV-L: A Language for EMV Protocols)

EMV-L is formed of the following procedures: $SetupID(ID) \rightarrow (\lambda_{ID}, C_{ID})$: sets up the execution environment denoted λ_{ID} , in which a Payer party with identity ID can make payments. The object λ_{ID} enables and encapsulates all payment information related to the Payer ID . It may also produce a long-term card C_{ID} .

$SetupPayment(ID) \rightarrow C$: based on a Payer’s identity ID and its execution environment λ_{ID} , creates an EMV-compliant payment device C (e.g., a physical card, or a mobile device with a card registered to it) that can transact with an EMV-compliant PoS.

$Payment((ID, C), M) \rightarrow \text{pay}$: based on an identity ID , a payment device C correctly set up as above, and a Merchant M , generates an EMV-compliant payment transcript pay .

$Clearing(ID, M, \text{pay}) \rightarrow T$: based on an identity ID , a Merchant M and a transaction pay produced as per the above, finalises the payment by balancing the account of the participants, and returns the terminating data T .

Each transcript tran resulting from these procedures is divided between the entities \mathcal{E} taking part. We later consider the restricted output of, e.g., the Merchant M in the payment pay as pay^M . A standard in-shop EMV payment, or an EMV-compliant one such as our PrivBank and PrivProxy using the EMV-L language, follows flow of EMV-L procedures:

$SetupID \rightarrow SetupPayment \rightarrow Payment \rightarrow Clearing$.

$SetupID$ is executed once per Payer, the other procedures can be run multiple times. The instantiation of PrivBank and PrivProxy in these procedures is detailed in Figures 1 and 2, as well as in Appendix D.

In this study, we examine the above operations from a network perspective rather than delving into cryptographic considerations. Indeed, actions like registering a Payer with the Issuer does not involve cryptography. Only $SetupPayment$ and $Payment$ can be viewed (partially) as cryptographic protocols, respectively card key generation and payment protocol [43]. Moreover, the clearing process ($Clearing$) occurs between Issuers and the network lacking any unique or public specifications; so, we treat these operations as black boxes.

7.3 Formalising Payments Privacy

Pseudonymity P_{Ps}^{ID} holds if, when a Payer ID makes a payment P , the adversary considered by our threat model (Section 3.3) cannot build a relation of the type “payment P is related to a Payer ID ”. Similarly, we define notions on mathematical relations to later describe P_{Ps} , Unlnk , and M_{Ps} . Consider R , a binary relation on pairs (x, y) .

One-Way Relation. We say that R is *one-way* if for all x , for all $y \in R_x$ (all y such that $(x, y) \in R$), given y , finding x is hard⁹.

Class Hiding Relation. We say that R is *class hiding* if for all x and y , given x and y , it is hard to determine if $(x, y) \in R$.

Intuitively, we formalise our property as the ability to correctly break these properties according to the transcript a subset of all entities $\mathcal{E} \subset \text{Parties}$ may have. The concept of a one-way relation is meant to reflect the inability to recover $x = ID$, the identity of a payer, from $y = \text{pay}$, the transcript of a payment made by ID . Meanwhile, the concept of a class hiding relation formalises the idea that, for two payments $x = \text{pay}_1$ and $y = \text{pay}_2$, one cannot determine whether a link exists between them—in this case, whether they were produced by the same payer. To this end, we formalise the relation for which the above properties should apply.

Definition 2 (Payer/Payments/Merchant Relation) *Let $[\text{Alg}]$ be the set of possible outputs of the algorithm Alg . Let π be an EMV protocol described in EMV-L. Let ID be a set of at least two long-term Payers’ identifiers. Let CARD be a set of at least two long-term Payers’ cards. Let M be a set of at least two Merchant identifiers. Let PAY be a set of transcripts outputted by $Payment$, generated by Payers with identifiers in ID toward some Merchants M in M . We define the relations:*

Payer Relation. *The payer relation $R_{\text{Pldt}} \subseteq ID \times \text{PAY}$ consists of all pairs (ID, pay) such that there exists $(\lambda, C_{ID}) \in [SetupID(ID)]$ and $C \in [SetupPayment(ID)]$, where $\text{pay} \in [Payment((ID, C), M)]$. Informally, a pair (ID, pay) is in R_{Pldt} if payment pay has been made by Payer ID .*

Card Relation. *The card relation $R_{\text{Cldt}} \subseteq \text{CARD} \times \text{PAY}$ consists of all pairs (C_{ID}, pay) such that there exists $(\lambda, C_{ID}) \in [SetupID(ID)]$ which encompasses a card C_{ID} and $C \in [SetupPayment(ID)]$, where $\text{pay} \in [Payment((ID, C), M)]$. Informally, a pair (C_{ID}, pay) is in R_{Cldt} if payment pay has been made with the long-term card C_{ID} .*

Payments Relation. *The payments relation $R_{\text{Payms}} \subseteq \text{PAY} \times \text{PAY}$ consists of all pairs $(\text{pay}, \text{pay}')$ such that there exist $ID, M, M', (\lambda, C_{ID}) \in [SetupID(ID)]$, and $C, C' \in [SetupPayment(ID)]$, where $\text{pay} \in [Payment((ID, C), M)]$ and $\text{pay}' \in [Payment((ID, C'), M')]$. Informally, a pair of payments is in R_{Payms} if it was produced by the same Payer.*

Merchant Relation. *The Merchant relation $R_{\text{Mldt}} \subseteq M \times \text{PAY}$ consists of all pairs (M, pay) such that there exist $ID, (\lambda, C_{ID}) \in [SetupID(ID)]$, $C \in [SetupPayment(ID)]$,*

⁹No PPT algorithm can compute it with non-negligible probability.

for which $\text{pay} \in [\text{Payment}((ID, C), M)]$. Informally, a pair (M, pay) is in R_{PAYMS} if payment pay has been directed to Merchant M .

Restricted Relations. For each relation, we define the restricted relation $R_*^{\mathcal{E}}$, for $\mathcal{E} \subset \mathcal{Parties}$, by considering the restricted view of the payments, i.e., the restricted view of the transcripts in all the algorithms of Definition 1.

Assuming that no two transactions within any of the protocol's views leads to identical data transcript (due to the timestamp amongst others), $\text{PAY}^{\mathcal{E}}$ is in bijection with PAY and the elements in a relation $R_*^{\mathcal{E}}$ are in bijection with the elements in R_* . Hence, the restricted relations are well defined. By requiring intractability of properties of these relations, we define our three payments-privacy properties.

Payer Pseudonymity. This property entails that from all one-time payment-transcripts which an attacker sees, included from corrupted parties in a set \mathcal{E} , the attacker cannot link back to a correct long-term payer identity ID_A for an honest payer ID_A . Mathematically the payer relation $R_{\text{Pldt}}^{\mathcal{E}}$ formed by tuples $(ID_A, \text{pay}^{\mathcal{E}})$ is one-way for the set \mathcal{E} hence, the view of $\text{pay}^{\mathcal{E}}$ of the payment pay . In the execution scenario considered the adversary also has access to the transcripts subsequent to the payment $\text{pay}^{\mathcal{E}}$, i.e., the clearing transcript and is an external observer of the other procedures. Payer's long-term cards C_{ID} can also be regarded as sensitive. The pseudonymity w.r.t. the Payer's long-term card, $P_{\text{Ps}}^{C_{ID}}$, is defined similarly and can be considered against the same corruption set \mathcal{E} .

Definition 3 (Pseudonymity - $P_{\text{Ps}}^{ID} \& P_{\text{Ps}}^{C_{ID}}$) Let R_{Pldt} be a Payer relation defined by an EMV protocol described in EMV-L and \mathcal{E} be a set of parties. We say that P_{Ps}^{ID} holds in front of a set \mathcal{E} of parties if the relation $R_{\text{Pldt}}^{\mathcal{E}}$ is one-way for the payments made by an uncorrupted Payer ID , i.e., it is unfeasible to create ID from $\text{pay}^{\mathcal{E}}$ and other transcripts seen by parties in \mathcal{E} . Similarly, pseudonymity $P_{\text{Ps}}^{C_{ID}}$ is attained in front of the corruption set \mathcal{E} if the relation $R_{\text{Cldt}}^{\mathcal{E}}$ is one-way i.e., it is intractable to yield C_{ID} from $\text{pay}^{\mathcal{E}}$ and other transcripts seen by parties in \mathcal{E} .

Unlinkability. Unlinkability refers to the capacity to ascertain whether two payments originate from the same Payer. This property is formalised by saying that the attacker cannot form the relation whereby the pairs are two payments emitted by the same Payer; our notion for this is called 'class-hiding', i.e., the attacker cannot form equivalence classes over payment transcripts, hence requiring the class-hiding property for the relation $R_{\text{PAYMS}}^{\mathcal{E}}$, given again for a corruption set \mathcal{E} .

Definition 4 (Unlinkability - Unlnk) Let R_{PAYMS} be a payments relation defined by a EMV protocol and \mathcal{E} be a set of parties. We say that the protocol attains unlinkability Unlnk for a set \mathcal{E} of parties if $R_{\text{PAYMS}}^{\mathcal{E}}$ is class hiding for uncorrupted Payers.

Merchant Pseudonymity. Merchant pseudonymity is defined similarly to Payer pseudonymity P_{Ps} . It is also based on the one-way property, but, this time, of a Merchant relation R_{Mldt} .

Definition 5 (Merchant Pseudonymity - M_{Ps}) Let R_{Mldt} be a Merchant relation defined by a payment protocol and \mathcal{E} be a set of parties. We say that the protocol attains Merchant pseudonymity M_{Ps} for a set \mathcal{E} of parties if the relation $R_{\text{Mldt}}^{\mathcal{E}}$ is one-way for payments made to an uncorrupted Merchant.

7.4 Provable Anonymity

We state our properties P_{Ps} , Unlnk, and M_{Ps} against PrivBank and PrivProxy . As per our threat model, any observer between the Payer and the Merchant is at most as strong as a corrupt Merchant. Similarly, someone who breaks the back-end channel between the Merchant and the Issuer is at most as strong as a corrupt Issuer. So, we state our results below only w.r.t. corrupt parties.

We start with the pseudonymity of the Payer, which differs between our constructions. Indeed, a KYC procedure is required in PrivProxy , where the pseudonymity-friendly Issuer provides a one-time identity for the Payer to present to the Proxy in PrivBank . This KYC-based difference indirectly leads to:

Proposition 1 (PrivBank respects P_{Ps}^{ID}) Consider an arbitrarily picked honest Payer with identifier ID , and PrivBank in the threat model given, where the Issuer which gives service to Payer ID is not corrupted. Then, PrivBank attains P_{Ps}^{ID} in front of the Proxy and the Merchant.

No long-term card C_{ID} is ever produced by $\text{SetupID}(ID)$ in PrivBank . Hence, the pseudonymity $P_{\text{Ps}}^{C_{ID}}$ cannot be defined for PrivBank . Pseudonymity of cards is still guaranteed for any long-term card C_{ID} held by a Payer outside of PrivBank . Indeed, neither this card nor any data related to it would ever need to be provided by the Payer in PrivBank .

Proposition 2 (PrivProxy - P_{Ps}^{ID} and $P_{\text{Ps}}^{C_{ID}}$) Consider an arbitrarily picked honest Payer with identifier ID , and PrivProxy in the threat model given, where the Issuers and Proxies which give joint service to Payer ID are not corrupted. Then, PrivProxy attains P_{Ps}^{ID} and $P_{\text{Ps}}^{C_{ID}}$ in front of the Merchant.

We move to the attainment of payment unlinkability. In PrivProxy , the Proxy can link payments by virtue of controlling all sides of identifiers of the payers. In PrivBank , even if the Proxy does not know the Payer's ID_A , the Proxy can link their payments by using the Android/Apple account on the phone of the Payer for which is receives requests. We change this w.r.t. the app (especially since it is provisioned by the Issuer); however, this would complicate the resolution/authorisation of payments by the

Proxy towards the Issuer. This leads us to the result below.

Proposition 3 (PrivBank and PrivProxy– Unlnk)

Consider an arbitrarily picked honest Payer with identifier ID , and $PrivBank$ and $PrivProxy$ in the threat model given, where the Issuers and Proxies which give joint service to Payer ID are not corrupt. Then, $PrivBank$ and $PrivProxy$ attain Unlnk in front of the Merchant.

We move to Merchant pseudonymity. Our result is:

Proposition 4 (PrivBank and PrivProxy – M_{PS})

Consider an arbitrarily picked honest Merchant with identifier M , and $PrivBank$ and $PrivProxy$ in the threat model given, where the Proxies and Payers which jointly pay to Merchant M are not corrupted. Then, $PrivBank$ and $PrivProxy$ attain M_{PS} in front of the Issuer.

All proofs are included in Appendix D.

7.5 On Privacy Treatments

Our privacy properties are specific to payment systems; moreover, pseudonymity is a restricted form of anonymity. However, our notions follow commonplace privacy definitions in the cryptographic setting. In Appendix E, we indeed show that our definitions can be re-cast and re-proven via standard game-based, cryptographic proofs [60]. Concretely, P_{PS} and Unlnk are both recast, in the game-based formalism, via indistinguishability-style games. As such, their nature is to protect identities strictly regarding protocol data and fields. Orthogonally, well-known studies [39] link payments via behaviour (time, GPS, shopping patterns), using data science. Our notions, even if lifted to temporary identifiers, cannot distinguish such “hidden” links, as they focus solely on protocol transcript-based relations. Profiling attacks [39] are not addressed by our protocols.

Our formalism resembles [33], where privacy is defined using attacker-infeasible relations between identifiers and secure-layer messages. Unlike our one-wayness for P_{PS} and M_{PS} , [33] assumes direct relation formation. Privacy relations in [33] use ephemeral identifiers, like protocol data. Similarly, we could redefine our privacy by targeting short-term payer identifiers (ID_X or ID_Y). [33] aligned relational privacy on ephemeral IDs with *weak unlinkability* and the work on long-term IDs with *strong unlinkability* by Arapinis *et al.* [26].

Our attacker is any PPT algorithm, which is known to be stronger than a Dolev-Yao attacker [22]. So, any Dolev-Yao attack translatable to our PPT relational definition applies here. In some cases, an attacker’s “if-then-else” tests cause errors, enabling them to mount privacy attacks, by distinguishing to whom a message belongs; this is akin to reverting payment relations and break P_{PS} (or stronger ephemeral-ID versions). For example, [37] show (an implementation of the) e-passport failing privacy, due to unmasked errors. However, if $PrivBank$

and $PrivProxy$ randomise all transcripts as specified, no such leakage occurs, and P_{PS} holds against such attacks. Orthogonally, some attacks do not apply here: counter-based one (e.g., [33] on IoT), as all counters are randomised

8 Discussions and Conclusions

On Legal Adoption. A main hurdle around adoption of privacy-enhancements in EMV is that they must not infringe the laws and regulations that are relevant: AML, KYC, PSD2 and SCA. Regarding KYC, PSD2, SCA, the concerns are around dealing with “common” impersonation fraud; for this, in our proposals, the agreements put in place between the Issuers and the Proxies would have to abide by laws such as the Proceeds of Crime Act 2002, hinging on their shared/sole liability; this is legally straightforward. Managing regulatory enforcements for anti-money laundering at scale is more complex. In Sections 6.2 and 6.4, we discussed how each proposal will deal with this, in accordance to the AML regulation and the guidance by the Financial Conduct Authority.

On Practical Adoption. Perhaps surprisingly, there is no technical or legal barrier to adoption. The practical hurdle is the sheer volume of virtual cards to be inserted in the payment systems; in fact, if $PrivBank$ or $PrivProxy$ would be entirely feasible, if offered only to a select few. To this end, a product director at Curve states: “*PrivBank and PrivProxy appear feasible in practice. These schemes make considerable progress on how to deal with AML compliance, and it seems feasible to make agreements with banks on how to deal with liabilities and the exact merchant data that proxies would send to issuing banks to balance ‘merchant anonymity’ vs. good payment-authorisation rates. The key practical challenge for Curve would be the significant cost of managing the high volume of one-time virtual cards, especially while complying with the business rules and regulations, across different countries and their Payment Services.*”

So, we proposed $PrivBank$ and $PrivProxy$: EMV-compliant, law-abiding solutions achieving payment pseudonymity, unlinkability, and merchant pseudonymity. We formalised and proved these, using a new privacy model. Our proposals are industry-implementable, and we will pursue practical studies.

Acknowledgements. We thank Robin Savage, Head of Payment Products at Curve, for useful insights into real-life payments systems and valuable feedback on our constructions.

Ethical Considerations. We rigorously adhere to ethical standards; for instance, here we used no sensitive or personal financial data. We focused solely on publicly available information and theoretical models. We carefully considered the potential impact of our findings on security and privacy, striving to enhance the payment ecosystem responsibly without introducing risks to users or financial institutions.

Open-Science Considerations. While our research does not involve the generation of data, code, or related materials, we remain committed to the principles of open science. We will ensure that our findings and methodologies are thoroughly documented and openly shared through the publication, including on open repositories, in some format. By doing so, we aim to contribute valuable insights to the community and facilitate further research and discussion in the field.

References

- [1] The electronic money regulations 2011, 2011. <https://www.legislation.gov.uk/uksi/2011/99/2016-03-01>, Last accessed on 11-16-23.
- [2] Central Bank Digital Currency: functional scope, pricing and controls, 2021. <https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op286~9d472374ea.en.pdf>, Last accessed on 11-16-23.
- [3] Central bank accounts are dangerous and unnecessary, a critique of two papers, 2022. <https://taler.net/papers/accounts-dangerous-2022.pdf>, Last accessed on 11-16-23.
- [4] The ECASH Act, 2022. <https://ecashact.us/>, Last accessed on 11-16-23.
- [5] Overview of EMVCo, 2022. <https://www.emvco.com/about/overview/>, Last accessed on 11-16-23.
- [6] Paypal: Account limits, 2022. <https://www.paypal.com/uk/smarthelp/article/faq1253>, Last accessed on 17-05-22.
- [7] Visa prepaid reloadable personal cards, 2022. https://www.visa.co.uk/content/VISA/usa/englishlanguage/master/en_US/home/pay-with-visa/cards/prepaid-cards/all-purpose-reloadable.html, Last accessed on 01-06-22.
- [8] Curve: Rule your money, 2023. <https://www.curve.com/en-gb/>, Last accessed on 11-16-23.
- [9] ebay payment terms of use, 2023. <https://pages.ebay.co.uk/payment/2.0/terms.html>, Last accessed on 11-16-23.
- [10] Edenred, 2023. <https://www.edenred.fr/ticket-restaurant>, Last accessed on 02-24-24.
- [11] Google pay help: United kingdom: Supported payment methods, 2023. <https://support.google.com/pay/answer/7351534>, Last accessed on 11-16-23.
- [12] Lyf, 2023. <https://www.lyf.eu/en/>, Last accessed on 02-23-24.
- [13] Mastercard prepaid | just load and pay, 2023. <https://www.mastercard.co.uk/en-gb/personal/find-a-card/general-prepaid-mastercard.html>, Last accessed on 11-16-23.
- [14] Otac: A new paradigm for user authentication and device authentication, 2023. <https://www.swidch.com/technology/otac>, Last accessed on 11-16-23.
- [15] Set up apple pay, 2023. <https://support.apple.com/en-us/HT204506>, Last accessed on 11-16-23.
- [16] Shop with amazon pay, 2023. <https://pay.amazon.co.uk/using-amazon-pay>, Last accessed on 11-16-23.
- [17] Spend safely online with single-use cards, 2023. <https://www.revolut.com/cards>, Last accessed on 11-16-23.
- [18] Up-one, 2023. <https://up-one.up.coop>, Last accessed on 02-24-24.
- [19] Mastercard payment account reference inquiry, 2024. <https://developer.mastercard.com/payment-account-reference-inquiry/documentation/>, Last accessed on 08-29-24.
- [20] Takepayments: What are contactless payments and how do they work?, 2024. <https://www.takepayments.com/blog/product-information/what-are-contactless-payments-and-how-do-they-work/>, Last accessed on 08-29-24.
- [21] Visa scan to pay, 2024. <https://developer.visa.com/innovation-corner/example-projects/scan-and-pay>, Last accessed on 02-27-24.
- [22] M. Abadi and P. Rogaway. Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption). *J. of Cryptology*, 20(3):395–395, 2007.
- [23] UK Public General Acts. The money laundering and terrorist financing (amendment) regulations 2019. <https://www.legislation.gov.uk/uksi/2019/1511/contents/made>, Last accessed on 08-30-24.
- [24] UK Public General Acts. Proceeds of crime act 2002. <https://www.legislation.gov.uk/ukpga/2002/29/contents>, Last accessed on 08-30-24.
- [25] Nicholas Akinyokun and Vanessa Teague. Security and privacy implications of nfc-enabled contactless payment systems. In *Proceedings of the 12th international conference on availability, reliability and security*, 2017.

- [26] Myrto Arapinis, Tom Chothia, Eike Ritter, and Mark Ryan. Analysing unlinkability and anonymity using the applied pi calculus. In *2010 23rd IEEE computer security foundations symposium*, pages 107–121. IEEE, 2010.
- [27] Financial Conduct Authority. Financial crime guide: a firm’s guide to countering financial crime risks (fcg). <https://www.handbook.fca.org.uk/handbook/FCG.pdf>, Last accessed on 09-02-24.
- [28] David Basin, Ralf Sasse, and Jorge Toro-Pozo. The emv standard: Break, fix, verify. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021.
- [29] Mike Bond, Omar Choudary, Steven J Murdoch, Sergei Skorobogatov, and Ross Anderson. Chip and skim: cloning emv cards with the pre-play attack. In *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014.
- [30] Ioana Boureanu, Liqun Chen, and Sam Ivey. Provable-security model for strong proximity-based attacks: With application to contactless payments. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, 2020.
- [31] Ioana Boureanu, Tom Chothia, Alexandre Debant, and Stéphanie Delaune. Security analysis and implementation of relay-resistant contactless payments. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, 2020.
- [32] Christina Brzuska, Nigel P Smart, Bogdan Warinschi, and Gaven J Watson. An analysis of the emv channel establishment protocol. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 2013.
- [33] Ksenia Budykho, Ioana Boureanu, Stephan Wesemeyer, Daniel Romero, Matt Lewis, Yogaratnam Rahulan, Fortunat Rajaona, and Steve Schneider. Fine-grained trackability in protocol executions. In *30th Annual Network and Distributed System Security Symposium, NDSS 2023, San Diego, California, USA, February 27 - March 3, 2023*. The Internet Society, 2023.
- [34] Jeffrey Burdges, Florian Dold, Christian Grothoff, and Marcello Stanisci. Enabling secure web payments with gnu taler. In *Security, Privacy, and Applied Cryptography Engineering: 6th International Conference, SPACE 2016, Hyderabad, India, December 14-18, 2016, Proceedings 6*. Springer, 2016.
- [35] Sergiu Bursuc, Ross Horne, Sjouke Mauw, and Semen Yurkov. Provably unlinkable smart card-based payments. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023.
- [36] David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology*, Boston, MA, 1983. Springer US.
- [37] Tom Chothia and Vitaliy Smirnov. A traceability attack against e-passports. In Radu Sion, editor, *Financial Cryptography and Data Security*, pages 20–34, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [38] Véronique Cortier, Alicia Filiipiak, Jan Florent, Said Gharout, and Jacques Traoré. Designing and proving an emv-compliant payment protocol for mobile devices. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2017.
- [39] Yves-Alexandre de Montjoye, Laura Radaelli, Vivek Kumar Singh, and Alex “Sandy” Pentland. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, 347(6221):536–539, 2015.
- [40] Joeri De Ruiter and Erik Poll. Formal analysis of the emv protocol suite. In *Theory of Security and Applications: Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31-April 1, 2011, Revised Selected Papers*. Springer, 2012.
- [41] Florian Dold. *The GNU Taler system: practical and provably secure electronic payments*. PhD thesis, Université Rennes 1, 2019.
- [42] Danny Dolev and Andrew Yao. On the security of public key protocols. *IEEE Transactions on information theory*, 1983.
- [43] LLC EMVCo. Emv integrated circuit card specifications for payment systems, book 2, security and key management, version 4.3, 2011.
- [44] LLC EMVCo. Emv next generation. next generation kernel system architecture overview., 2014.
- [45] LLC EMVCo. Emv payment tokenisation frequently asked questions (faq) — general. <https://www.emvco.com/specifications/emv-payment-tokenisation-faqs/>, 2022.
- [46] LLC EMVCo. Emv payment tokenisation specification technical framework v2.3. *EMVCo: Foster City, CA, USA*, 2022.
- [47] LLC EMVCo. Emv annual report 2023: Enhancing emv technologies to supporting emerging payments. *EMVCo: Foster City, CA, USA*, 2023.
- [48] LLC EMVCo. Emv payment tokenisation – a guide to use cases. *EMVCo: Foster City, CA, USA*, 2023.

- [49] EUR-Lex European Union. Regulatory technical standards for strong customer authentication (sca) and common and secure open standards of communication. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02018R0389-20230912>, Last accessed on 02-21-24, 2018.
- [50] EUR-Lex European Union. Directive on payment services in the internal market. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02015L2366-20151223>, Last accessed on 02-21-24, 2021.
- [51] EUR-Lex European Union. Directive on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02015L0849-20210630>, Last accessed on 02-21-24, 2021.
- [52] L.-A. Galloway and T. Yunusov. First contact: New vulnerabilities in contactless payments. In *Black Hat Europe*, 2019.
- [53] Ross Horne, Sjouke Mauw, and Semen Yurkov. Unlinkability of an improved key agreement protocol for emv 2nd gen payments. In *2022 IEEE 35th Computer Security Foundations Symposium (CSF)*. IEEE, 2022.
- [54] ICANN (Internet Corporation for Assigned Names and Number). Information for Privacy and Proxy Service Providers. <https://www.icann.org/resources/pages/pp-services-2017-08-31-en>, 2004.
- [55] AML Legislations. Suspicious Activity Report Documentation. <https://www.lawsociety.org.uk/topics/anti-money-laundering/suspicious-activity-reports>.
- [56] Steven J Murdoch, Saar Drimer, Ross Anderson, and Mike Bond. Chip and pin is broken. In *2010 IEEE Symposium on Security and Privacy*. IEEE, 2010.
- [57] Satoshi Nakamoto and A Bitcoin. A peer-to-peer electronic cash system. *Bitcoin*.—URL: <https://bitcoin.org/bitcoin.pdf>, 2008.
- [58] Andreea-Ina Radu, Tom Chothia, CJ Newton, Ioana Boureanu, and Liqun Chen. Practical emv relay protection. In *Proc. 43rd IEEE Symp. Security Privacy*, 2022.
- [59] Rosehaslina Razali. Overview of e-cash: Implementation and security issues, 2002. <https://www.giac.org/paper/gsec/1799/overview-e-cash-implementation-security-issues/103204>, Last accessed on 11-16-23.

EMV	Europay Mastercard Visa
EMV-L	EMV language
EUR	Euro
AML	Anti-Money Laundering
KYC	Know Your Customer
SCA	Strong Customer Authentication
PSD2	Payment Services Directive (version 2)
MCC	Merchant Category Code
MN	Merchant Name
MRI	Merchant Risk Index
ML	Merchant Location
PAN	Application Primary Account Number
PAR	Payment Account Reference
TEE	Trusted Execution Environment
CVM	Cardholder Verification Method
CDCVM	Consumer Device CVM
CBDC	Central Bank Digital Currency
PSP	Payment Service Provider
T&C	Terms and Conditions
CBDC	Central Bank Digital Currency

Table 3: List of the Most Used Acronyms.

- [60] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. <http://eprint.iacr.org/2004/332>, 2004.
- [61] Aleksei Stennikov Timur Yunusov, Artem Ivachev. New vulnerabilities in public transport schemes for apple pay, samsung pay, gpay. *White Paper*, 2021.
- [62] Visa. Minimum data requirements for merchants, 2021. <https://www.elavon.ie/content/dam/elavon/en-gb/documents/customer-centre/customer-news/merchant-best-practice-infographic.pdf>, Last accessed on 02-21-24.
- [63] Peter Wegner. Interactive foundations of computing. *Theoretical Computer Science*, 1998.
- [64] Timur Yunusov. Hand in your pocket without you noticing: Current state of mobile wallet security. *Black Hat Europe*, 2021.

Auxiliary Supporting Material

A Acronyms

In this paper, we frequently use various acronyms to streamline the presentation of complex terms and concepts. For the reader’s convenience, we provide a list of these acronyms along with their full forms below in Table 3.

B Traditional Payment Systems & Their Attainment of Payment-Privacy

The following are well-known, “traditional”¹⁰, in-shop payment solutions¹¹. We look at them and consider aspects linked to their provision of KYC of pseudonymity for the payer. A summary of their properties has been given in Section 4, Table 1. In this section, we provide a description of each of them and give reasons for the claimed properties.

1. Cash (banknotes and coins). This works for in-shop purchases only. There is no KYC and strong pseudonymity; although shops may use of security cameras and subsequent review may allow the payer to be identified, this is beyond the scope of the payment methods. Banknotes might be traced, based on their serial numbers; but that requires a complicated set of step taking by different banks and the Merchant, which again makes the matter beyond into a complex type of identifiability. Coins which do not feature serial numbers. Thus, we can say that cash has no KYC and provides pseudonymity and unlinkability (within reasonable/normal measures). For the same reason, namely that cash cannot be linked to the payer, the SCA cannot be required before payments.

2. Cheques. They provide a payment mechanism which involved banks as intermediary between the payer and the merchant or the entity paid. In addition, cheques require the name of the payer and the merchant to be written on them, as well as an bank account number of the payer. The account number is also divulging to the bank issuing the cheque to the payer. This method therefore has KYC in place, and does not provide pseudonymity in any of the possible ways. In many countries, it has become common practice to require proof of identity in order to make a payment. In such cases, the payer is strongly authenticated and SCA holds.

3. E-cash. This was originally proposed by Chaum in his 1983 paper on “Blind Signatures for Untraceable Payments” [36]. Payers have an e-wallet topped-up with e-cash, from which they spend e-cash like they would real cash, in principle; limits on the expenditure exist though, e.g., for AML reasons. Thus, this digital mechanism aims to provide the user with the same level of pseudonymity that they achieve when using cash. If the link between the e-cash wallet and the owner’s bank account is just for applying limits to the amount of e-cash stored on a wallet, then the merchant should not obtain information about the customer and even the bank should gain little information. How it is regulated and implemented will determine the outcomes which may differ from Table 1.

Now, we give further details on e-cash, since it is an interesting variant to what we do. It needs a separate network to EMV, but it has supporters, still, today. Chaum commercialised this idea, founding DigiCash in 1994. This and other early e-cash developments are described in [59]. Although none of these have been commercially successful, there is still interest in developing e-cash systems [34].

In the US, the “Electronic Currency And Secure Hardware Act” [4] proposes that an electronic dollar should be created with the same privacy properties as the dollar itself. If achieved there would be no *Static Data Authentication* (SDA) requirement and this could also be used for online purchases, although for full pseudonymity the purchaser may need to use a VPN to hide their IP address, use a separate e-mail account and a delivery locker for their purchase.

In contrast, the European Central Bank’s document discussing e-cash (Central Bank Digital Currency, CBDC) [2] recommends that to avoid too much money being stored in consumer’s digital wallets (and not in the banks) limits should be applied and to enable this the consumer would need to be identified and their CBDC wallet linked to their bank account. Others [3] opposed it, and it is not clear what might ultimately be decided.

4. Plastic/physical credit or debit cards. This are to be used in-shop or online, to make payments in ways we all know. In Section 5.1, we already explained what KYC, SCA and pseudonymity they provide and why.

5. Google Pay [11] and Apple Pay [15]. These are two of the most used methods of mobile payment, *i.e.*, payment via a mobile app “inside” which a physical card is registered. Like in physical cards, KYC and SCA are the norm here. In mobile-payments SCA, when making the payment, the customer may be asked to confirm their identity onto the payment device too, via PIN, fingerprint or face recognition. But, unlike payments by physical cards, making and authorising mobile payments need more intermediaries in the payment networks, and tokenisation (as outlined in Section 5.1). Due to this tokenisation mechanism, the merchant can link purchases by the same payer using their long-term PAR – created once during the onboarding of their card onto the app and used in all their payments thereafter. The PAR replaces the PAN, which is then hidden from the merchant, making payments anonymous but not unlinkable.

6. Pre-paid cards. These are cards which may have no bank account associated with them and one tops up with a set amount or one buys already topped up and uses. We divide pre-paid cards into two categories:

- a) Top-up cards – cards offered with services such as those by Revolut [17]. Other providers exist: for example, UK pre-paid Mastercard cards’ providers, listed on the

¹⁰They do not use crypto currencies.

¹¹Online payments are included here as they may give some insights into what might be possible

Mastercard website, for general use and as gift all behave as per the below. The Payer may not have to undergo credit checks, but must satisfy identity and address checks and have money available in their account to cover any payments made. So, KYC is generally done; indeed, most of the issuers of such pre-paid cards act as electronic money institutions and are regulated by the Electronic Money Regulations [1]. Since these are cards, their pseudonymity properties are the same as plastic cards, or that mobile apps – if they are loaded therein.

- b) Gift cards – card that can be bought is store cards, with set amounts preloaded onto them. There is generally no KYC done. Related to this, their use is restricted (to specific merchants) and the amount on each card is small, to satisfying the AML requirements [51]. If they are purchased with cash and not linked to a bank account (for re-charging, for example), then subject to the same caveats about IP addresses, e-mail accounts and deliveries, payer pseudonymity can be achieved.

7. Virtual or “one time” cards (VC) These provide pseudorandom card details (card number, expiry date, CVV), for each transaction. You can remain anonymous to the merchant, but the virtual card is linked to your ‘real’ card to enable payment to be made and so the issuer knows who the payer is and from whom they are purchasing. These are marketed for online use, in general. One example is a card offered the Revolut card [17]. Another company, Swiich [14] offer a range of services based on One Time Access Codes (OTAC) and this includes *ephemeral* cards; as for other virtual cards these are linked to a registered *real* payment card (requiring KYC). In this context, an intermediary may be acting as a payment proxy although how individual providers handle the payments differs from one company to another. In addition, unless strict usage limits are applied, SCA is required for all payments.

8. Payment service providers There are a number of providers in this category, for example:

PayPal PayPal offers a range of services. In terms of the discussion here, PayPal accounts can be used to make payments. Figure 4 shows how PayPal is used for making payments. The stages are:

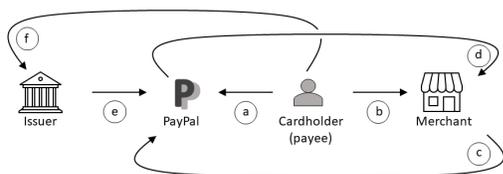


Figure 4: Making payments with PayPal.

1. The cardholder opens an account and registers a card to be used for payment. SCA for the Issuer is carried out at this point, but is not necessary afterwards. Unless the account holder confirms their identity and address (KYC) accounts are restricted and have limits placed on amounts that can be sent, received, or withdrawn [6]. PayPal is acting as the *payment service provider* (PSP) and knows who you are and gets to see who you are buying from and when.
2. The cardholder purchases an item from a merchant and pays using their PayPal account. Purchases can be made online, or in person.
3. The merchant receives their payment *from PayPal*. And not directly from the payer, thus the purchaser may use a pseudonym and separate e-mail account and a delivery locker for their purchase to obtain full pseudonymity.
4. PayPal pays the merchant.
5. PayPal charges the issuer. The issuer knows that something was purchased, but not who from.
6. The cardholder pays the issuer.

In this case, payment or identity information has been provided to PayPal, but has not been made known to the merchant because PayPal has filtered most of it. Other information, such as e-mail, may be pseudonyms, not linked to the payer’s identity.

Amazon Pay. Amazon Pay [16] offers a similar service and allows paying online with a credit card, debit card or by direct debit. They make it clear that the merchant does not receive your payment details: “We do not share your full credit card, debit card, or bank account number with sites or charitable organisations that accept Amazon Pay. The merchant only receives information that is required to complete and support your transaction. This information may include your name, email address, and shipping address.”

Curve. Curve [8] provides a payer with a card and a payment application. Curve users must satisfy the KYC rules. The payer registers multiple bankcards issued by one or several banks in the Curve app. When the payer pays with their Curve card or with the Curve app, one of their Curve-registered bankcards is charged in principle, but the payment is not “settled” entirely (*i.e.*, the payment onto the classical card is “pending”). Curve pays the merchant on behalf of the payer on the spot, but Curve also provides the payer a period of 30-days to potentially move the transaction to another of their registered, classical cards. So, there is an intricate process of what is called “payment authorisation” between Curve, the issuers of classical bankcards and other bodies (*e.g.*, Visa,

Mastercard). During the payment authorisation, all the information required is shared without filtering, so that all the entities know each other’s identities.

9. Online marketplaces. Examples of these are provided by Amazon and eBay. We view these as merchants here. So, from that viewpoint, clearly, here is generally no KYC or SCA needed to open accounts with them, as a payer for their goods. However, eBay, for example, states in their terms [9] that they may require “any other data about the buyer which the buyer’s payment service provider or we may require”. Aside from that, if goods bought from them are sent to a pickup locations, then some degree of pseudonymity can be achieved.

Further Payment Alternatives. Cryptocurrencies [57] are alien to EMV. But, non-EMV payments close to EMV exist. For instance, Lyf [12] and Visa [21] propose payment services which rely on their own payment network and QR codes. Or, a large-scale, EU-funded project tries to push new payments based on the GNU Taler initiative [34] and using the well-known e-cash idea by Chaum [36]. They perform online transactions and are compliant with online-payments’ regulation; they do not use the card-to-PoS-merchant payment networks like us. Attaining privacy via online transactions is easier – *e.g.*, via one-time cards without the worry of “in-shop” SCA but relying on 3D secure, without having to share credentials over an app between different-domain entities.

C Our Solution At A Glance

We took inspiration from existing payment systems: plastic card and mobile EMV, disposable EMV cards, proxying of EMV payments by Curve [8] and machinations during EMV-payment authorisation.

(A). On Payers’ Pseudonymity and Payments’ Unlinkability. The main inspiration for our designs here come from mobile EMV-payments and one-time/disposable cards for online shopping, and we bring the later into the space of “in-shop” payments. As a result of enhanced security, mobile payments are already more privacy-preserving than plastic cards as they hide the main identifier of the physical card, the *Primary Account Number* (PAN), via an ephemeral card-like number called *tokenised PAN*, which contributes to payers’ pseudonymity. Yet, mobile-payments made via the same bankcard still contain the fixed card-identifying *Payment Account Reference* (PAR); this leads to payments being linkable. All of our designs will revert to tokenisation and PAN-PAR-based constructions in mobile payments. Instead, our mobile apps will utilise one-time disposable cards which will produce transactions as per plastic cards, which is akin to having a one-time PAN.

(B). On Merchants’ Pseudonymity. Here, we take inspiration from EMV-payment proxies such as Curve [8] (see

Section 5). We add an intermediary in the interaction between the payer and the merchant, which also relays the payment to the issuing bank, but stripped of certain merchant-related data. In more detail, based on an agreement between the issuer and the proxy *w.r.t.*, *e.g.*, certain categories of merchants with sufficiently low *Merchant Risk Indicators*, the proxy omits sending the merchant name to the issuer, while still providing the latter with some merchant identification data.

However, there is one last hurdle to our designs, chiefly the sets of regulations, as follows.

(1) AML and counter-terrorism financing regulations require payments’ auditability by certain payment-system parties, therefore, for any transaction, the payer and the merchant must be traceable.

(2) SCA/*Payment Services Directive* (PSD2) require identification of payers prior to using a payment service, including opening bank-accounts and making payments.

So, we carefully combined the ideas in (A) and (B) above to achieve payers and merchants pseudonymity as well as payments unlinkability, while still achieving EMV-compliance and abiding by regulations (1) and (2) above.

To achieve this, some entities retained some of the identity information required and, when all combined, the systems obtained are in accordance with regulations (1) and (2) above.

Our solution is given at a glance in Figure 5 with `PrivBank` in Figure 5a and `PrivProxy` in Figure 5b. There the knowledge of long-term and one-time identities is described.

D Proofs of our Main Results

In this section, we give proofs of the privacy preserving properties, given in Section 7.4, and obtained by our two constructions `PrivBank` and `PrivProxy`. These properties are also summarised in Table 1.

To this end, we first instantiate the four procedure of EMV-L language of `PrivBank` and `PrivProxy` according to their step descriptions in Section 6.1 and Section 6.3.

Instantiation of `PrivBank` in EMV-L. The `PrivBank` protocol instantiated in the EMV-L language is as follows:

SetupID(ID) $\rightarrow (\lambda_{ID}, C_{ID})$: The three participating entities are the Payer, identified by an identity ID , an Issuer, and a Proxy. This algorithm corresponds to step 1 of `PrivBank`, as described in Section 6.1. The output $\lambda_{ID} = (\lambda_{ID,Payer}, \lambda_{ID,Issuer}, \lambda_{ID,Proxy})$ represent the enrolment outputs of the three respective entities. The output is C_{ID} is set to \perp , indicating that no long-term card is generated within `PrivBank`.

SetupPayment(ID) $\rightarrow C$: The three participating entities are the Payer, identified by an identity ID , an Issuer, and a Proxy. This procedure corresponds to steps 2, 3 and 4 of `PrivBank`. The Payer’s output is $C = (ID_Y, C_{ID_Y})$, consisting of a one-time virtual cardholder name “ ID_Y ,” and a one-time

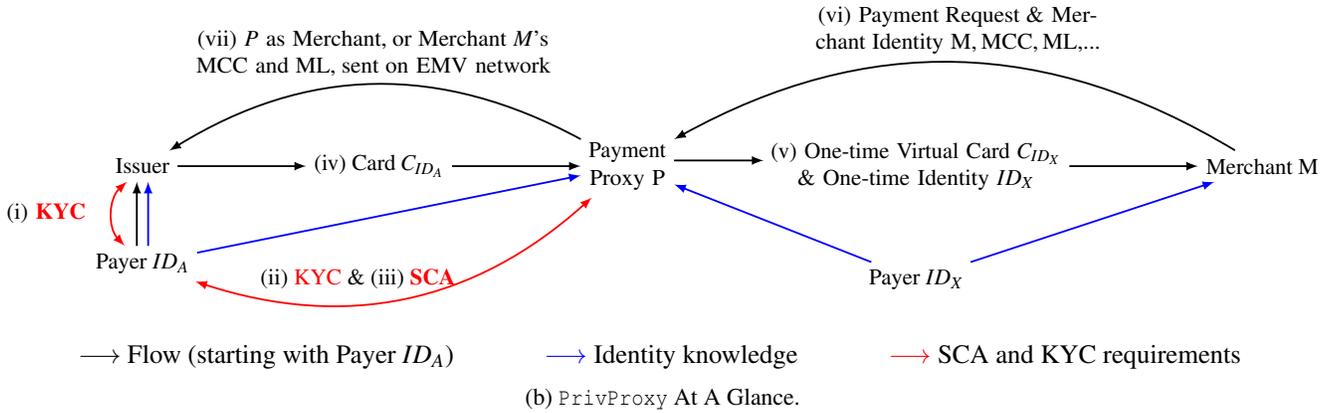
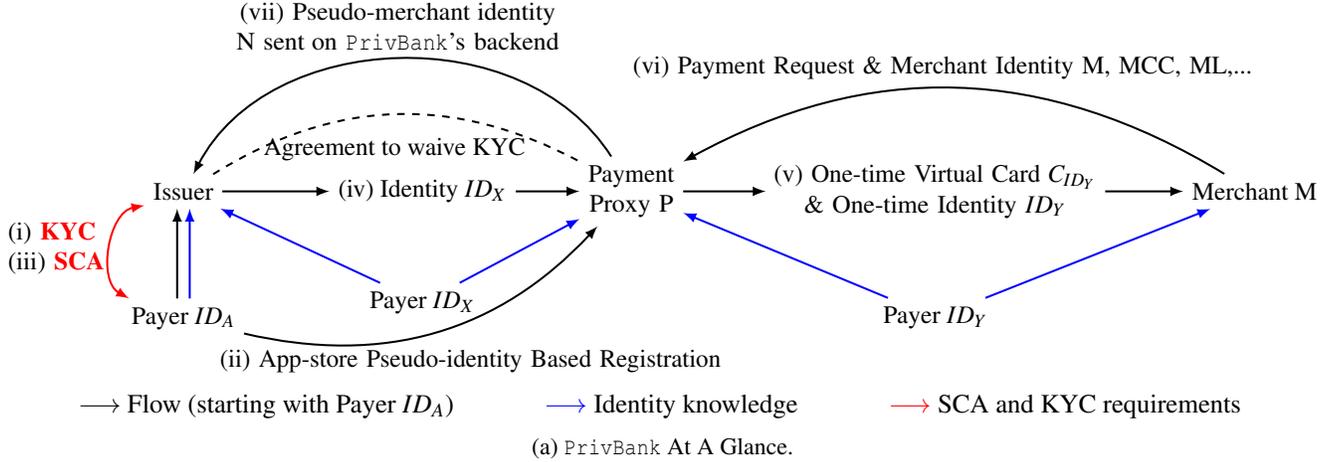


Figure 5: Succinct description of the our two proposals: Sketch of Steps regarding KYC, SCA and Identification.

virtual, EMV-compliant card C_{ID_Y} . The other entities produce their respective procedure transcripts.

$Payment((ID, C), M) \rightarrow \text{pay}$: The four participating entities are the Payer holding (ID, C) , an Issuer, a Proxy, and a Merchant identified by M . This procedure corresponds to step 5 of PrivBank. The output is an EMV-compliant *payment transcript* $\text{pay} = (\text{pay}_{Payer}, \text{pay}_{Issuer}, \text{pay}_{Proxy}, \text{pay}_{Merchant})$, with each entity retaining its respective portion of the transcript.

$Clearing(ID, M, \text{pay}) \rightarrow T$: The four participating entities are the Payer identified by ID , an Issuer, a Proxy, and a Merchant identified by M . Each entity also inputs its view of the payment pay . This procedure corresponds to steps 6, 7, 8 and 9 of PrivBank. The output is the *terminating data*, $T = (T_{Payer}, T_{Issuer}, T_{Proxy}, T_{Merchant})$.

Instantiation of PrivProxy in EMV-L. The PrivProxy protocol instantiated in the EMV-L language is as follows:

$SetupID(ID) \rightarrow (\lambda_{ID}, C_{ID})$: The three participating entities are the Payer, identified by an identity ID , an Issuer, and a Proxy. This algorithm corresponds to step 1

of PrivProxy, as described in Section 6.3. The outputs $\lambda_{ID} = (\lambda_{ID, Payer}, \lambda_{ID, Issuer}, \lambda_{ID, Proxy})$ represent the enrolment outputs of the three respective entities. The Payer's long-term card C_{ID} is also returned by all the participating entities.

$SetupPayment(ID) \rightarrow C$: The three participating entities are the Payer, identified by an identity ID , an Issuer, and a Proxy. This procedure corresponds to steps 2 and 3 of PrivProxy. The Payer's output is $C = (ID_Y, C_{ID_Y})$, consisting of a one-time virtual cardholder name "ID_Y," and a one-time virtual, EMV-compliant card C_{ID_Y} . The other entities produce their respective procedure transcripts.

$Payment((ID, C), M) \rightarrow \text{pay}$: The four participating entities are the Payer holding (ID, C) , an Issuer, a Proxy, and a Merchant identified by M . This procedure corresponds to step 4 of PrivProxy. The output is an EMV-compliant *payment transcript* $\text{pay} = (\text{pay}_{Payer}, \text{pay}_{Issuer}, \text{pay}_{Proxy}, \text{pay}_{Merchant})$, with each entity retaining its respective portion of the transcript.

$Clearing(ID, M, \text{pay}) \rightarrow T$: The four participating entities are the Payer identified by ID , an Issuer, a Proxy, and a Merchant identified by M . Each entity also inputs its view

of the payment pay. This procedure corresponds to steps 5, 6, 7 and 8 of `PrivProxy`. The output is the *terminating data*, $T = (T_{\text{Payer}}, T_{\text{Issuer}}, T_{\text{Proxy}}, T_{\text{Merchant}})$.

These step descriptions are visually represented in Figures 1 and 2.

In our designs, we avoid delving into the exact cryptographic details. Instead, we treat these as black boxes, particularly in the card generation, payment and clearing processes. The payment mechanisms in our proposals remain unaltered compared to the extensively analysed EMV payment procedures between the card-issuing entity (which is the Proxy in our cases), a merchant's PoS and a Payer with its card (see Section 2 for related protocol analysis). Furthermore, the backend infrastructure lies outside the public specification of EMV standards and varies among different entities. As the emphasis is on the non-disclosure of certain identity-related information, we also treated them as black-boxes, only formulating a few related and direct hypotheses on these, in our threat model (see Section 3.3).

Given the EMV-L description of `PrivBank` and `PrivProxy`, we now show that our privacy properties hold on our constructions: for Property 1 and Property 2 – that R_{Pldt} and R_{Cldt} are one-way against any attacker in our model; for Property 3 – that R_{Payms} is class hiding against any attacker in our model; Property 4 – that R_{Mldt} is preimage resistant one-way against any attacker in our model.

Proposition 1 (PrivBank respects $P_{\text{Ps}}^{\text{ID}}$) Consider an arbitrarily picked honest Payer with identifier ID , and `PrivBank` in the threat model given, where the Issuer which gives service to Payer ID is not corrupted. Then, `PrivBank` attains $P_{\text{Ps}}^{\text{ID}}$ in front of the Proxy and the Merchant.

Proof 1 (Proposition 1: PrivBank respecting $P_{\text{Ps}}^{\text{ID}}$) We need to prove that payer pseudonymity $P_{\text{Ps}}^{\text{ID}}$ holds for `PrivBank`: relation R_{Pldt} as generated by `PrivBank` is one-way against any PPT adversary in our model (specifically, including corrupted Proxies and corrupted Merchant).

According to our threat model (see Section 3.3), at least two Payers, with identities referred to as ID_1 and ID_2 , remain honest, as do their Issuers. As stated in our model, the apps of the uncorrupted Payers are also uncorrupted.

Let \mathcal{A} be any PPT adversary controlling a corrupted Proxy and Merchant, with whom the honest Payers ID_1 and ID_2 and their Issuer interact. The adversary \mathcal{A} can also control other entities, with all corrupted entities being part of the set \mathcal{E} .

We do our proof of property $P_{\text{Ps}}^{\text{ID}}$ against \mathcal{A} in two steps: (a). we first demonstrate that, given a payment pay made by Payer ID_1 , we can replace it by Payer ID_2 in all the process of the EMV-L language without affecting the transcript $\text{pay}^{\mathcal{E}}$ that has been observed by the adversary \mathcal{A} .

(b). then, we prove that it is possible only if the relation R_{Pldt}

is one-way against the adversary \mathcal{A} , i.e., that $P_{\text{Ps}}^{\text{ID}}$ holds for `PrivBank`.

Step (a). We do this step, by reasoning, on each EMV-L procedure inside `PrivBank`, at one time, i.e., we show that in each EMV-L procedure, in fact ID_1 could be indistinguishably replaced (in terms of the procedure's output) by ID_2 , and vice-versa.

SetupID: there, a Payer executes the KYC procedure with the uncorrupted Issuer and enrolls in the `PrivBank` app using ID . Subsequently, it enrolls to the Proxy using their app-store account, which is, by hypothesis, unlinked to any Payer's identity. Hence, this authentication to the Proxy is independent of the identity of the Payer. It could have been either Payer ID_1 or Payer ID_2 participating in the KYC process with the Issuer for the app-store account presented to the Proxy. The Proxy thus obtains a pseudo-identities for the Payers, which it can link to all its actions, but these pseudo-identities are completely independent of the Payers' actual identities ID_1 and ID_2 in this process, and the output of the procedure is interchangeable over ID_1 and ID_2 .

SetupPayment: the procedure begins with the Payer opening the `PrivBank` app. As explained previously for *SetupID*, app enrolment does not reveal the Payer's identity to the Proxy, the only corrupted party interacting with the Payer and the Issuer in *SetupID* and *SetupPayment*.

Next, the Payer is authenticated via SCA with the honest Issuer. An identity ID_X , independent from the Payer's actual identity ID_1 or ID_2 , is then sent to the Proxy and used in subsequent interactions in *SetupPayment*. Any pseudo-identity ID_X remains completely independent of the Payers' actual identities and could belong to ID_1 or ID_2 in the same way; the allocation of ID_X is pseudorandom and statistically independent of the long-term identities ID_1 or ID_2 . Thus, the notification pushed by the app to request for a payment to the Proxy is the same of ID_1 as it is for ID_2 .

Payment: this procedure is an interaction between the honest Payer and the adversary, controlling the Proxy and the Merchant. Here, the Payer only provides the one-time identity ID_Y and pays using the one-time card C_{ID_Y} . Both were provided to the Payer by the adversary. Thus, nothing more is revealed to the adversary about the real identity of the Payer, regardless of the outcome of the transaction: it being for ID_1 or for ID_2 , it is just the same from the viewpoint of \mathcal{A} .

Clearing: the clearing of the payment is a process where the Merchant and the Proxy, i.e., here the adversary, contact the Issuer to obtain the due payment. The relevant step is the adversary sending ID_X , the Merchant's pseudo-identity, and other transaction-related data to the Issuer and its answer. Once again, the adversary's claim for payment, as well as the honest Issuer's response, rely solely on the pseudo-identity ID_X with all elements sent being independent of the Payer's long term identity: i.e., it would be the same process for ID_1 , as it would be for ID_2 .

With all the above, we have shown that whether it is

ID_1 or ID_2 , a pseudo-identity ID_X is consistently used in *SetupPayment*, *Payment*, and *Clearing*, and all further sub-transcripts obtained by the adversary \mathcal{A} rely solely on ID_X , without any link to the long-term identities ID_1 or ID_2 . Hence, the same *PrivBank* transaction/transcript, call it pay^E , can be derived for Payer ID_1 or for Payer ID_2 , and seen by \mathcal{A} unknowingly of which is the case.

Step (b). We now proceed to step (b) in our proof, i.e., we prove that if the conclusion of step (a) is the case, then the relation R_{Pldt}^E yielded by *PrivBank* is one-way against the arbitrary adversary \mathcal{A} .

Assume, by *reductio ad absurdum*, that there exists a PPT algorithm \mathcal{B} capable of breaking the one-wayness of the payment relation R_{PAY}^E . Then, the attacker \mathcal{A} would call \mathcal{B} on the transcript pay^E yielded in step (a), and \mathcal{A} would then know if it belongs to ID_1 or ID_2 . This contradicts the conclusion of step (a), built on the execution of *PrivBank* based on its assumptions (e.g., ID_X s being pseudorandom and statistically independent of the long-term identities ID_1 and ID_2 , etc.).

So, the relation R_{Pldt}^E yielded by *PrivBank* has one-wayness: i.e., *PrivBank* has payer pseudonymity P_{Ps}^{ID} .

Proposition 2 (*PrivProxy* – P_{Ps}^{ID} and P_{Ps}^{CID})

Consider an arbitrarily picked honest Payer with identifier ID , and *PrivProxy* in the threat model given, where the Issuers and Proxies which give joint service to Payer ID are not corrupted. Then, *PrivProxy* attains P_{Ps}^{ID} and P_{Ps}^{CID} in front of the Merchant.

Proof 2 (Proposition 2: *PrivProxy* respecting P_{Ps}^{ID})

We aim to prove that payer pseudonymity, P_{Ps}^{ID} holds for *PrivProxy* in the presence of a corrupted Merchant. According to our threat model (see Section 3.3), at least two Payers (and their app), with identities referred to as ID_1 and ID_2 , remain honest, as do their common Issuers and their Proxy. As stated in our model, the apps of the uncorrupted Payers are also uncorrupted.

Let \mathcal{A} be a PPT adversary controlling a Merchant, with whom the honest Payers ID_1 and ID_2 , their Issuer and their Proxy interact. The adversary \mathcal{A} can also control other entities, with all corrupted entities being part of the set \mathcal{E} . We begin by proving the property P_{Ps}^{ID} against \mathcal{A} .

Once more, we do our proof of property P_{Ps}^{ID} against \mathcal{A} in two steps:

(a). we first demonstrate that, given a payment pay the Payer ID_1 , resp. the card $C_{ID,1}$, can replace by Payer ID_2 , resp. the card $C_{ID,2}$ in all the process of the EMV-L language without affecting the transcript pay^E that has been observed by the adversary \mathcal{A} .

(b). then, we prove that is possible only if the relation R_{Pldt}^E , resp. R_{CIDt}^E is one-way against the adversary \mathcal{A} , i.e., that

P_{Ps}^{ID} , resp. P_{Ps}^{CID} holds for *PrivBank*.

We start by the proof for the Payer’s Pseudonymity P_{Ps}^{ID} .

Pseudonymity P_{Ps}^{ID} . Step (a). We do this step, by reasoning, on each EMV-L procedure inside *PrivProxy*, at one time, i.e., we show that in each EMV-L procedure, in fact ID_1 could be indistinguishably replaced (in terms of the procedure’s output) by ID_2 , and vice-versa.

SetupID: this procedure, consisting of two KYC authentications, is executed by the Payer, their Issuer, and the Proxy. The adversary controlling the Merchant does not participate. As all communications are assumed to be secure, nothing is leaked to \mathcal{A} at this point: an execution for a different Payer would result in a transcript that is computationally indistinguishable to the observer \mathcal{A} .

SetupPayment: this procedure, consisting of an SCA authentication of the Payer to the Proxy, is executed solely between these two honest entities. Once again, all communications are assumed to be secure, ensuring that nothing is leaked to \mathcal{A} at this point. Through this process, the Payer obtains a one-time identity ID_X and a one-time payment card C_{ID_X} . These two elements are pseudorandom and statistically independent of the Payer’s identity.

Payment: the same argument as in the proof *PrivBank* P_{Ps}^{ID} apply: the Payer only provide one-time identity and only the success or failure is retruned by the other honest entities after that.

Clearing: once again the same argument as in the proof *PrivBank* P_{Ps}^{ID} apply: the clearing rely solely on the pseudo-identity ID_X and the transaction executed based on the one-time card C_{ID} with all element sent being independent of the Payer’s long term identity.

To sum up, *SetupID* and *SetupPayment* are executed by honest parties and, following the second procedure, produce a pseudorandom and statistically independent one-time identity ID_X and one-time card C_{ID_X} for the Payer. *Payment* and *Clearing* are subsequently executed, with all exchanges relying solely on elements ID_X and C_{ID_X} . As both elements are pseudorandom and statistically independent of the Payer’s identity, then, if, without loss of generality, we assume that Payer ID_1 has received ID_X and C_{ID_X} , which lead to the payment transcript pay^E , from the adversary’s perspective, the same transcript pay^E could have resulted from Payer ID_2 being provided ID_X and C_{ID_X} .

Step (b). This part is similar to step (b) of the proof of P_{PsID} for *PrivBank*. This allows us to conclude the proof of the property P_{Ps}^{ID} of *PrivProxy*.

Pseudonymity P_{Ps}^{CID} . We now focus on the pseudonymity P_{Ps}^{CID} for *PrivProxy*. The adversary remains unchanged. We demonstrate, in step (a), that the long-term card data is interchangeable and completely hidden behind the one-time card used for the payment.

Step (a). During the execution of $\text{SetupID}(ID)$, an element λ_{ID} and a long-term card $C_{ID,1}$ are generated. SetupID is executed between the Payer, the Issuer, and the Proxy. This is the same as in SetupPayment . There, a one-time payment card C_{ID_X} is derived from the Payer's card $C_{ID,1}$ in SetupPayment . C_{ID_X} is a completely new card which is pseudorandom and statistically independent of $C_{ID,1}$. We now analyse the two other procedure where the adversary is involved:

Payment: the payment is executed with the Merchant using C_{ID_X} , which is completely independent of the card $C_{ID,1}$. The long-term card $C_{ID,1}$ is then used for the payment between the Proxy and the Issuer, both of whom are honest entities. No information about this second payment is transferred to the Merchant, except for its approval. Hence, in the payment process, all elements transmitted to the Merchant are entirely independent of the long-term card $C_{ID,1}$.

Clearing: the clearing of the transaction executed using C_{ID_X} is performed between the Merchant and the Proxy. The clearing of the other transaction, based on $C_{ID,1}$, occurs between the honest Proxy and the honest Issuer. Once again, all information transmitted to external corrupted entities is free of the card $C_{ID,1}$ (i.e., statistically independent), with only the approval of the second clearing being relayed to \mathcal{A} .

With all the above, a payment pay leading to pay^E and its subsequent clearing T , leading to T^E , are both completely free of the card $C_{ID,1}$. Indeed, the view of \mathcal{A} executing all entities in \mathcal{E} depends only on C_{ID_X} , and as we have just shown, C_{ID_X} is statistically independent of $C_{ID,1}$. If the same ephemeral card C_{ID_X} were produced based on a card $C_{ID,2}$, then pay^E , T^E , and all further transcripts obtained by \mathcal{A} would remain identical.

Step (b). We now prove that the step (a) of the proof prevent from having an adversary breaking one-wayness of the $R_{C_{IDt}}^E$ relation. For that, assume, by reductio ad absurdum, that there exists a PPT algorithm \mathcal{B} capable of breaking the one-wayness of the payment relation $R_{C_{IDt}}^E$. Then, the attacker \mathcal{A} would call \mathcal{B} on the transcript pay^E yielded in step (a), and \mathcal{A} would then know if it is produced based on $C_{ID,1}$ or $C_{ID,2}$. This contradicts the conclusion of step (a), built on the analysis of PrivProxy based on its assumptions (e.g., C_{ID_X} s being pseudorandom and statistically independent of the long-term card $C_{ID,1}$ and $C_{ID,2}$, etc.).

So, the relation $R_{C_{IDt}}^E$ yielded by PrivProxy has one-wayness: i.e., PrivProxy has payer pseudonymity P_{Ps}^{CID} .

Proposition 3 (PrivBank and PrivProxy– Unlnk)

Consider an arbitrarily picked honest Payer with identifier ID , and PrivBank and PrivProxy in the threat model given, where the Issuers and Proxies which give joint service to Payer ID are not corrupt. Then, PrivBank and PrivProxy attain Unlnk in front of the Merchant.

Proof 3 (Proposition 3: Our proposals respecting Unlnk)
We need to prove that Payer unlinkability Unlnk holds for PrivBank and PrivProxy in front of a dishonest Merchant: relation R_{Payms}^E as generated by both proposals is class hiding against any PPT adversary in our model (specifically, including corrupted Merchant).

According to our threat model (see Section 3.3), at least two Payers, with identities referred to as ID_1 and ID_2 , remain honest, as do their Issuers. As stated in our model, the apps of the uncorrupted Payers are also uncorrupted.

Let \mathcal{A} be any PPT adversary controlling a corrupted Merchant, with whom the honest Payers ID_1 and ID_2 , their Issuer and Proxy interact. The adversary \mathcal{A} can also control other entities, with all corrupted entities being part of the set \mathcal{E} .

We start by treating the case of PrivBank .

Payer Unlinkability Unlnk of PrivBank. *Step (a).* Consider two payment transcripts, pay_1 and pay_2 , and their respective transcripts view by the adversary \mathcal{A} : pay_1^E and pay_2^E .

These payments were obtained from two executions of Payment , each preceded by a respective execution of SetupPayment , which produces for each payment a unique one-time identities $ID_{Y,1}$ and $ID_{Y,2}$ and one-time cards $C_{ID_{Y,1}}$ and $C_{ID_{Y,2}}$. These one-time cards are generated pseudorandomly (and are thus independent of one another) and remain statistically independent of the respective identities of the Payers. Consequently, linking $C_{ID_{Y,1}}$ or $ID_{Y,1}$ to $C_{ID_{Y,2}}$ or $ID_{Y,2}$ is unfeasible. Furthermore, since these one-time cards and identities are used in only one transaction, no link between transactions based on them can be established.

The SetupPayment procedure is performed exclusively between honest entities: the Payer, its Issuer, and its Proxy. As a result, only the outputs subsequently used in the payment and clearing processes might reveal a potential link, and only when executed by the Payer.

Payment: An EMV-compatible payment is executed using the one-time card $C_{ID_{Y,1}}$ and the one-time identity $ID_{Y,1}$. The response from the corrupted Merchant must be a transaction authorisation request based on $C_{ID_{Y,1}}$ and $ID_{Y,1}$. If this request is based on any other one-time card and identity not currently in use for the payment, the transaction will fail. Subsequently, a "Success" or failure message concludes the transaction.

AML caps are shared among honest Payers, ensuring they do not trigger differently. The validation based on $C_{ID_{Y,1}}$ is determined solely by the format of the request, and the profile of the Payers does not influence the transaction's acceptance. As a result, the Proxy's response does not reveal any information to the adversary.

Clearing: the clearing is based on the elements sent by the Merchant to the honest entities. These entities then respond with a money transfer to the Merchant's bank account, provided everything checks out. Since there is no distinction between the Payers, the same outputs are produced for both

transactions. Thus, the adversary is left with a deterministic answer based on its request.

Based on the above, we can conclude that any two payments, pay_1 and pay_2 , are statistically independent, regardless of whether they are executed by the same Payer or two different Payers.

Step (b). We now prove that the step (a) of the proof prevent from having an adversary breaking the class hiding property of the relation R_{Payms}^E . For that, assume, by reductio ad absurdum, that there exists a PPT algorithm \mathcal{B} capable of breaking the class hiding of the payment relation. Then, the attacker \mathcal{A} would call \mathcal{B} on the transcript pay_1^E and pay_2^E yielded in step (a), and \mathcal{A} would then know if they were produced both by one a the Payer ID_1 or ID_2 or by multiple Payer. This contradicts the conclusion of step (a), built on the analysis of PrivBank based on its assumptions (e.g., C_{ID_x} s being pseudorandom and statistically independent of the long-term card $C_{ID,1}$ and $C_{ID,2}$, etc.).

So, the relation R_{Payms}^E yielded by PrivBank has class hiding: i.e., PrivBank has Payer unlinkability Unlnk .

Payer Unlinkability Unlnk of PrivProxy . We demonstrate Payer unlinkability Unlnk using the same strategy as for the payer unlinkability Unlnk of PrivBank .

The procedures SetupID and SetupPayment are executed for the two honest Payers ID_1 and ID_2 , involving only honest entities (the Issuer and Proxy). In PrivProxy , just as in PrivBank , the procedure SetupPayment results in one-time identities $ID_{X,1}$ and $ID_{X,2}$, along with one-time cards $C_{ID_{X,1}}$ and $C_{ID_{X,2}}$. These identities are directly derived from the original identities of the Payers, yet they are generated pseudorandomly (and are thus independent of one another) and remain statistically independent.

The investigation is similarly applicable to Payment, which operates in PrivProxy much like in PrivBank . The only addition is the payment authorisation that is relayed to the Issuer. As the Payers' identities do not influence the acceptance of the transaction, there is no way to distinguish between the behaviours of payments made by different Payers.

Subsequently, the clearing Clearing of the transaction remains unchanged, so the analysis for this step is identical to that in PrivBank .

Hence, for the same reasons as above, any two payments, pay_1 and pay_2 , are statistically independent, regardless of whether they are executed by the same Payer or by two different Payers. Moreover, step (b) performed for PrivBank also applies here, allowing us to conclude that PrivProxy achieves Payer unlinkability Unlnk .

Proposition 4 (PrivBank and $\text{PrivProxy} - M_{\text{Ps}}$)

Consider an arbitrarily picked honest Merchant with identifier M , and PrivBank and PrivProxy in the threat model given, where the Proxies and Payers which jointly pays to Merchant M are not corrupted. Then, PrivBank and PrivProxy attain M_{Ps} in front of the Issuer.

Proof 4 (Proposition 4: Our proposals respecting M_{Ps})

We need to prove that Merchant pseudonimity M_{Ps} holds for PrivBank and PrivProxy in front of the Issuer: relation R_{Mldt}^E as generated by both proposals is one-way against any PPT adversary in our model.

According to our threat model (see Section 3.3), at least two Merchant, with identities referred to as M_1 and M_2 , remain honest, as do the Payer and the Proxy for the considered transactions. As stated in our model, the apps of the uncorrupted Payers are also uncorrupted.

Let \mathcal{A} be any PPT adversary controlling a corrupted Issuer, with whom the honest Payer their Issuer, Proxy and Merchants M_1 and M_2 interacts. The adversary \mathcal{A} can also control other entities, with all corrupted entities being part of the set \mathcal{E} .

We start by treating the case of PrivBank .

We do our proof of property M_{Ps} against \mathcal{A} in two steps: (a) we first demonstrate that, given a payment pay made to a Merchant M_1 , we can replace it by Merchant M_2 in all the process of the EMV-L language without affecting the transcript pay^E that has been observed by the adversary \mathcal{A} .

(b) then, we prove that it is possible only if the relation R_{Pldt} is one-way against the adversary \mathcal{A} , i.e., that P_{Ps}^{ID} holds for PrivBank .

Merchant pseudonimity M_{Ps} of PrivBank . *Step (a)* We do this step, by reasoning, on each EMV-L procedure inside PrivBank , at one time, i.e., we show that in each EMV-L procedure, in fact Merchant M_1 could be indistinguishably replaced (in terms of the procedure's output) by Merchant M_2 , and vice-versa.

SetupID and SetupPayment: these procedures are between the honest Payer and Proxy with the corrupted Issuer. They correspond to the production of the payment mean opening prior to any payment and before any interaction with a Merchant is made. For any other entity, an execution for a different Payer would produce a transcript that is computationally indistinguishable.

Payment: this procedure involves the Merchant, who is contacted by the honest Payer. The Merchant then communicates with the honest Proxy, providing its identity M along with the Payer's one-time card C_{ID_X} and one-time identity ID_X . Once again, the Issuer is not involved and does not receive any information about the Merchant behind this transaction.

Clearing: the Issuer becomes involved during the clearing process. At this stage, the Proxy has received a clearing

request from the Merchant containing the payment reference based on the Payer’s identity ID_Y and the Merchant’s identity M . The request forwarded by the Proxy to the Issuer includes the Payer’s identity ID_X and a pseudo-merchant identity N , i.e., a fixed pseudo-merchant identity associated with this Merchant. By hypothesis this pseudo-merchant identity does not reveal the original Merchant identity. Hence the clearing request sent by the Proxy does not reveal it neither.

We have shown that throughout all processes of the EMV-L language, none of the data sent to the Issuer reveals the original identity of the Merchant. With all the above, we have shown that whether it is M_1 or M_2 , a pseudo-identity N or sanitisation with the Proxy identity P is consistently used in Clearing, and all further sub-transcripts obtained by the adversary \mathcal{A} rely solely on N or P , without any link to the original merchant identities M_1 or M_2 . Hence, the same $PrivBank$ transcript, call it pay^E , can be derived for Merchant M_1 or for Merchant M_2 , and seen by \mathcal{A} unknowingly of which is the case.

Step (b). We now proceed to step (b) in our proof, i.e., we prove that if the conclusion of step (a) is the case, then the relation R_{Mldt}^E yielded by $PrivBank$ is one-way against the arbitrary adversary \mathcal{A} .

Assume, by *reductio ad absurdum*, that there exists a PPT algorithm \mathcal{B} capable of breaking the one-wayness of the payment relation R_{Mldt}^E . Then, the attacker \mathcal{A} would call \mathcal{B} on the transcript pay^E yielded in step (a), and \mathcal{A} would then know if it is to Merchant M_1 or M_2 . This contradicts the conclusion of step (a), built on the execution of $PrivBank$ based on its assumptions.

So, the relation R_{Mldt}^E yielded by $PrivBank$ has one-wayness: i.e., $PrivBank$ has Merchant pseudonymity M_{Ps} .

Merchant pseudonymity M_{Ps} of $PrivProxy$. *Step (a)* We do this step, by reasoning similarly as for $PrivBank$: on each EMV-L procedure inside $PrivProxy$, we show that in each EMV-L procedure, in fact Merchant M_1 could be indistinguishably replaced (in terms of the procedure’s output) by Merchant M_2 , and vice-versa.

SetupID and SetupPayment: these procedure still correspond to the production of the payment mean appening prior to any payment and before any interaction with a Merchant is made. Moreover, an execution for a different Payer would still result in a transcript that is computationally indistinguishable to any entity not involved in the process.

Payment: this procedure involves the Merchant, who is contacted by the honest Payer. The Merchant then communicates with the honest Proxy, providing its identity M along with the Payer’s one-time card C_{ID_X} and one-time identity ID_X . In $PrivProxy$, the Issuer is involved in validating the transaction. A validation request is sent by the Proxy to the Issuer, including the Payer’s payment details C_{ID_A} and either the Proxy’s identity P (as the recipient of the payment to be made by the Issuer) or a pseudo-merchant identity N , i.e., a

fixed pseudo-merchant identity associated with this Merchant. As in the case of $PrivBank$, by hypothesis, this pseudo-merchant identity does not reveal the original Merchant’s identity. Hence, the payment request sent by the Proxy does not reveal it either.

Clearing: the Issuer is also involved in the clearing. At this stage, the Proxy has received a clearing request from the Merchant containing the payment reference based on the Payer’s identity ID_Y and the Merchant’s identity M . The request forwarded by the Proxy to the Issuer undergoes similar sanitisation of the Merchant’s identity as in Payment. For the same reasons, this sanitisation ensures that the request is independent of the original Merchant’s identity. Hence, the transcript observed by the Issuer is completely independent of the original Merchant’s identity.

We have shown that throughout all processes of the EMV-L language, none of the data sent to the Issuer reveals the original identity of the Merchant. With all the above, we have shown that whether it is M_1 or M_2 , a pseudo-identity N or sanitisation with the Proxy identity P is consistently used in Payment and Clearing, and all further sub-transcripts obtained by the adversary \mathcal{A} rely solely on N or P , without any link to the original merchant identities M_1 or M_2 . Hence, the same $PrivBank$ transcript, call it pay^E , can be derived for Merchant M_1 or for Merchant M_2 , and seen by \mathcal{A} unknowingly of which is the case.

Step (b). We conclude, as in the proof of Merchant pseudonymity for $PrivBank$, that $PrivProxy$ achieves Merchant pseudonymity M_{Ps} .

E Game Based Formalisation

Below, we recast our security definitions in the widely-accepted game-based paradigm. These definitions demonstrate the security of our proposal in the cryptographic game based standard. Multiple formalisations for game-based properties, with varying degrees of strength are possible. The game based definitions which the first definition is for Payer’s pseudonymity and unlinkability and the second for Merchant’s pseudonymity are directly implied by the security introduced through relations in Section 7. De facto showing that our protocols matches the security defined in this section.

We have chosen to present only one game for the Payer’s privacy preserving properties, as in general unlinkability $Unlnk$ imply pseudonymity P_{Ps} (i.e., $Unlnk \implies P_{Ps}$). In all the upcoming games, an EMV compatible payment procedure is formalised through our EMV language (see Definition 1) and executed by a challenger following the games against an adversary \mathcal{A} . This adversary is assumed to corrupt and execute entities participating in the protocols in the experiments following the threat model of Section 3.3. All corrupted entities are encompassed in a set called \mathcal{P}_{cor} .

Payer's Payments Privacy: Pseudonymity + Unlinkability.

Payer Pseudonymity models that entities does not retrieve the Payer's long-term card data (card number, expiry date, CVV, etc.) nor its identity during the execution of payments. *Payments' unlinkability* models a stronger requirement against an adversary trying to match payments by distinguishing between cases where they were made by the same entity or not.

Below, we present an experiment encompassing both notions of pseudonymity and unlinkability. The adversary has access to a payment oracle and a clearing oracle which allows him on one side to link the payment under consideration or to infer the identity of the Payer based on the other executions it sees. Another mean of attack would be to only base itself on the payment in consideration. The latter relates to pseudonymity, while the first refer to the unlinkability property. Before introducing the experiment, let us set up the oracles in full details.

Payment. The oracle $Payment((\cdot, C), \cdot)$, is queried by an adversary to the challenger on the basis of a Payer identity ID and a Merchant identity M . The identity of the Payer must have been initialised prior to this call, meaning that: the $SetupID(ID)$ and $SetupPayment(ID)$ protocols, which returned a C_{ID} card, have been executed. It executes the protocol $Payment((ID, C_{ID}), M)$ with the adversary in the corruption frame under consideration.

Clearing. The oracle $Clearing(\cdot)$, is queried by an adversary for a payment pay . The payment pay must have been previously produced on the basis of the *Payment* oracle. It executes the $Clearing(pay)$ protocol with the adversary in the corruption frame under consideration.

Definition 6 (Game Based Unlinkability with Pseudonymity)

Consider the set of oracles $O = (Payment((\cdot, C), M), Clearing(\cdot, M, pay))$, each of them operating as defined above. Consider a PPT adversary \mathcal{A} participating in all executed protocols under the role of the corrupted entities \mathcal{P}_{cor} against a challenger executing the actions prescribed by $Exp_{\mathcal{P}_{cor}}^{PRIV}$ for uncorrupted entities.

$Exp_{\mathcal{P}_{cor}}^{PRIV}$

```

1:  $\{ID_i\}_{i=1}^n, M \leftarrow \mathcal{A}$ 
2: if  $n = 1$ , return  $\perp$ 
3: Let  $T = \emptyset$  the set of initialised cards.
4: for  $i \in \{1, \dots, n\}$ ,
5:    $(\lambda_i, C_{ID}) \leftarrow SetupID(ID_i)$ 
6:    $C_i \leftarrow SetupPayment(ID_i)$ 
7:    $T \leftarrow T \cup \{ID_i, C_i\}$ 
8: Sample an identity with the associated card
    $(ID, C) \leftarrow_s T$ 
9:  $pay \leftarrow Payment((ID, C), M)$ 
10:  $f \leftarrow Clearing(pay)$ 
11: return  $ID^* \leftarrow \mathcal{A}^O$ 

```

We say that a payment scheme has game based unlinkability with pseudonymity for a set of corrupted entities \mathcal{P}_{cor} if for adversaries \mathcal{A} controlling entities in \mathcal{P}_{cor} ,

$$|\Pr[Exp_{\mathcal{P}_{cor}}^{PRIV}(\mathcal{A}) = ID] - \frac{1}{n}| \leq \epsilon,$$

for an negligible ϵ ¹².

An alternative definition would require $Exp_{\mathcal{P}_{cor}}^{PRIV}$ to return a long term card. This definition is analogue but apply for Pseudonymity with respect to long-term card $P_{\mathcal{P}_{cor}}^{CID} \subset CARD \times PAY$ and require the opponent to return the Payer's long-term card C_{ID} instead of the Payer's identity ID . We call this experiment $Exp_{\mathcal{P}_{cor}}^{PRIV-CID}$, it is associated with the same winning probability.

Theorem 1 Assuming that the relation $R_{PAY}^{\mathcal{P}_{cor}} \subsetneq ID \times PAY$ is preimage resistant and that the relation $R_{Unlnk}^{\mathcal{P}_{cor}} \subsetneq PAY \times PAY$ is class hiding for a payment protocol described by an EMV language and for a set of corrupted entities, \mathcal{P}_{cor} then game base unlinkability is achieved for \mathcal{P}_{cor} .

Proof 5 Let \mathcal{A} be an adversary executing the corrupted entities from \mathcal{P}_{cor} against the challenger of the experiment $Exp_{\mathcal{P}_{cor}}^{PRIV}$.

\mathcal{A} has access to oracles $Payment((ID, C), M)$ and $Clearing(ID, M, pay)$ for all registered entities and for the identities it has chosen.

The execution scenario defined by the adversary's calls to the oracles and the prescribed execution of $SetupID$, $SetupPayment$, $Payment$ and $Clearing$ define a Payment Relation R_{PAY} and a linkability relation R_{Unlnk} . For both relation, the adversary's against these relations are their restricted counterpart: $R_{PAY}^{\mathcal{P}_{cor}}$ and $R_{Unlnk}^{\mathcal{P}_{cor}}$.

In the experiment, the adversary has two ways of guessing the Payer's identity. It can either try to retrieve the identity

¹²A function $\epsilon: \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every positive integer c , there exists an integer N_c such that for all $x > N_c$, $|\epsilon(x)| < 1/x^c$.

ID from its view of the payment transcript associated with pay, or try to deduce it on the basis of any links it might find with the other executions called through the oracles and then infer the identity based on the knowledge of the inputted identities. As \mathcal{A} chooses the identity *ID* for its calls to the oracles, the relation between the transcript it sees and the Payer's identities is not hidden from it.

The first case refers to an adversary recovering the identity from the payment pay:

Consider an adversary \mathcal{A} executing entities in \mathcal{P}_{cor} and interacting with the challenger of Exp^{PRIV} . \mathcal{A} has access to oracles $\text{Payment}((ID, C), M)$ and $\text{Clearing}(ID, M, \text{pay})$ in addition to the view it gets from the executions of SetupID , SetupPayment , Payment and Clearing . Based on the above algorithms, two sets can be populated: *ID* and *PAY*. The link between the elements of these sets, obtained based on the challenger's view of the execution allows defining the associated payment relation R_{PAY} . Based on the corruption set \mathcal{P}_{cor} and the relation, the restricted relation $R_{\text{PAY}}^{\mathcal{P}_{\text{cor}}} \subseteq \text{ID} \times \text{PAY}^{\mathcal{P}_{\text{cor}}}$ is defined for the adversary's execution scenario.

Assume \mathcal{A} returns a response ID^* for its view $R_{\text{PAY}}^{\mathcal{P}_{\text{cor}}}$ of the relation R_{PAY} and has probability significantly different from $1/n$ to win the experiment. When the adversary wins the experiment and returns the right answer $\mathcal{A}(\text{Exp}^{\text{PRIV}}) \rightarrow ID^*$ for the execution scenario, this means that $(ID^*, \text{pay}) \in R_{\text{PAY}}$. As this is assumed to be with non-negligible probability, hence we invert $R_{\text{PAY}}^{\mathcal{P}_{\text{cor}}}$ with non-negligible probability.

If the adversary provides a correct ID^* , for the payment pay, we can build a simulator algorithm based on the adversary to break the one-wayness of R_{PAY} . This lead to a contradiction under the assumption that the relation $R_{\text{PAY}}^{\mathcal{P}_{\text{cor}}}$ is difficult to invert.

We may now consider the second case: assume that \mathcal{A} made a link between pay and one of the payment it has sees during it call to the oracles. As the adversary knows the link between the second payment and the Payer making it, this lead to a link between pay and the identity ID^* . If the adversary provides a correct ID^* , which is the input of a Payment oracle query or a Clearing oracle query, we can build a simulator based on \mathcal{A} breaking the class hiding property Unlnk . This also leads to a contradiction, as we have assumed $R_{\text{Unlnk}}^{\mathcal{P}_{\text{cor}}}$ to be class hiding resistant.

Both potential attacks are covered, hence, under one-wayness of $R_{\text{PAY}}^{\mathcal{P}_{\text{cor}}} \subseteq \text{ID} \times \text{PAY}$ and class hiding resistance of $R_{\text{Unlnk}}^{\mathcal{P}_{\text{cor}}} \subseteq \text{PAY} \times \text{PAY}$ then game based unlinkability is achieved.

Remarks on the Proof. One can view this proof as a game hop. First, a random permutation is introduced over the identities that \mathcal{A} queries to the oracle, while still expecting the original identity as a response. This constitutes the first step of the proof, where we reduce security to pseudonymity \mathcal{P}_{PS} . The second part of the proof involves a direct reduction from

the modified experiment to unlinkability Unlnk .

Long-term card data $\mathcal{P}_{\text{PS}}^{CID}$. The same theorem and proof can be given for *Pseudonymity with respect to long-term card data* associated with the experiment $\text{Exp}^{\text{PRIV}-CID}$ by requiring the opponent to return a long-term card instead of an identity. In this latter case, the same implication holds and the security based on the basic game model is also ensured for our two protocols under the associated corruption sets \mathcal{P}_{cor} .

Merchant Pseudonymity (\mathcal{M}_{PS}). Merchant pseudonymity guarantees the Payers that the identity of the Merchant they are paying is not disclosed to other entities during the payment or its clearing. First, let us set up the oracles useful for our definition.

SetupID. The oracle $\text{SetupID}(\cdot)$ is queried by an adversary from the challenger on the basis of the identity of a Payer *ID*. The challenger executes the protocol $\text{SetupID}(ID)$ interacting with the adversary in the considered corruption frame.

SetupPayment. The oracle $\text{SetupPayment}(\cdot)$, is queried by an adversary from the challenger on the basis of an identity Payer's identity *ID* which should have been initialised based on SetupID . The challenger executes the protocol, $\text{SetupID}(ID)$ interacting with the adversary in the considered corruption frame.

Payment & Clearing. These oracles are the same as before.

Definition 7 (Game Based Merchant Pseudonymity)

Let \mathcal{P}_{cor} be a set of identities describing the corruption setup and consider the set of oracles $\mathcal{O}_{\mathcal{M}_{\text{PS}}} = (\text{SetupID}(\cdot), \text{SetupPayment}(\cdot), \text{Payment}(\cdot, C), M), \text{Clearing}(\cdot, M, \cdot))$, each of them operating as defined above. Consider a PPT adversary \mathcal{A} participating in all executed protocols under the role of the corrupted entities \mathcal{P}_{cor} against a challenger executing the actions prescribed by $\text{Exp}^{\mathcal{M}_{\text{PS}}}$ for uncorrupted entities.

$\text{Exp}_{\mathcal{P}_{\text{cor}}}^{\mathcal{M}_{\text{PS}}}$

- 1: $ID, M_0, M_1 \leftarrow \mathcal{A}$
- 2: $(\lambda, DC_{ID}) \leftarrow \text{SetupID}(ID)$
- 3: $C \leftarrow \text{SetupPayment}(ID)$
- 4: $b \leftarrow_{\$} \{0, 1\}$
- 5: $\text{pay} \leftarrow \text{Payment}((ID, C), M_b)$
- 6: $f \leftarrow \text{Clearing}(ID, M_b, \text{pay})$
- 7: **return** $M^* \leftarrow \mathcal{A}^{\mathcal{O}_{\mathcal{M}_{\text{PS}}}}$

A payment scheme has game Merchant pseudonymity for a corruption set \mathcal{P}_{cor} if for any adversaries \mathcal{A} executing entities in \mathcal{P}_{cor} ,

$$|\Pr[\text{Exp}_{\mathcal{P}_{\text{cor}}}^{\mathcal{M}_{\text{PS}}}(\mathcal{A}) = M_b] - \frac{1}{2}| \leq \epsilon,$$

for an negligible ϵ .

The experiment work as follows: attacker participate in the protocol for which it controls at least one entity involved; there are two Merchants; one gets paid with an associated transcript pay; clearing is executed for the payment, finally attacker guess the chosen Merchants based on its view.

Theorem 2 Assuming that the relation $R_{M_{PS}}^{P_{cor}} \subseteq M \times PAY$ is preimage resistant for a payment system and for a set of entities P_{cor} then Game M_{PS} is achieved for the corrupted set P_{cor} .

Proof 6 This proof is essentially the same as the first part of the proof of Theorem 1.

Assume for contradiction that an adversary \mathcal{A} can break $Exp_{P_{cor}}^{M_{PS}}$ for a corruption set P_{cor} while the relation $R_{M_{PS}}^{P_{cor}}$ still achieves preimage resistance. Under these assumptions, the protocol utilised achieves Pseudonymity but not Game Based Merchant Pseudonymity. Based on the adversary's interactions and its calls to the available oracles, two set $M = \{M_0, M_1\}$ and $PAY^{P_{cor}}$ can be defined, as well as a relation $R_{M_{PS}}^{P_{cor}} \subseteq M \times PAY^{P_{cor}}$.

In order to win against the game based Merchant pseudonymity, the adversary \mathcal{A} returns the identity M_b of a payment pay_b . When \mathcal{A} 's answer is correct, assumed with non negligible probability, we have found a preimage for one of the elements in $R_{M_{PS}}^{P_{cor}}$. Hence, based on this the adversarial algorithm we can construct a sequence of execution determining a relation $R_{M_{PS}}^{P_{cor}}$ for which it is possible to recover the Merchant's identity. This contradicts our hypothesis and shows that one-wayness of $R_{M_{PS}}^{P_{cor}}$ imply security against $Exp_{P_{cor}}^{M_{PS}}$.

With this proof, we have reduced our mathematical relationship-based model to the game-based model presented above. This gives us two formalisms for the security of our proposals. In addition, because we have reduced the security of our relational model to this model and prove the security of our proposals for the relationship-based model, our protocols meet the security requirements of both models.

F Sample Real Card Traces

Here we give a number of traces from card based payments to illustrate the points made earlier about the data included in card transactions. Here is part of a trace from a MasterCard card:

```
5A | len:8 Application Primary Account
      Number: 5521573039705376
5F24 | len:3 Application Expiration Date
      YYMMDD: 240430
5F25 | len:3 Application Effective Date
      YYMMDD: 200401
5F28 | len:2 Issuer Country Code: 0826
```

```
5F34 | len:1 Application Primary Account
      Number Sequence Number: 01
```

The card record contains the PAN, and the expiration date. We can also find information determining the currency in which the card is issued. Note: In the same trace we also see:

```
9F02 | len:6 Amount, Authorised
      (Numeric): 000000004600
9F03 | len:6 Amount, Other (Numeric):
      000000000000
9F1A | len:2 Terminal Country Code:
      0826
 95 | len:5 Terminal Verification
      Results: 0000008001
5F2A | len:2 Transaction Currency Code:
      0826
 9A | len:3 Transaction Date:
      210318
9F35 | len:1 Terminal Type:
      22
9F34 | len:3 Cardholder Verification
      Method Results: 1F0302
1F No CVM required;
```

This part shows the amount to be paid and other information about the terminal, such as the currency and the country where it is located. The cardholder verification method is also indicated. Merchant information comes from the terminal and are only sent to the acquirer in the backend, which is why it is not shown in the trace.

The same type of trace can be observed for other EMV-compatible card brands (Visa, AmericanExpress, etc.).

G Sample Mobile Application Traces

Here we give a number of traces from mobile phone applications to illustrate the points made earlier concerning tokenisation and the PAR. Here is part of a trace from an iPhone transaction:

```
70 | len:37 Record Template
5F28 | len:2 Issuer Country Code:
      0826
9F07 | len:2 Application Usage Control:
      C000
9F19 | len:6 Token Requestor ID:
      040010030273
5F34 | len:1 Application Primary Account
      Number (PAN) Sequence Number:
      00
9F24 | len:29 Payment Account Reference
      (PAR): 5630303130303133303136
      313936333535353639313035303937383933
```

The Token Requestor ID and the Payment Account Reference are part of the tokenisation process.

Note: In the same trace we also see:

```
70 | len:81 Record Template
5A | len:8 Application Primary Account
    Number (PAN):
    4831920272059474
5F24 | len:3 Application Expiration Date
    YYMMDD:
    231231
9F46 | len:176 ICC Public Key Cert:
149DD6A920995B05A5146C6ABEE823AFD2E2CBE91C
701C2648E395EAD23F5AD04C8C6B2D2DA4CD271B15
4339C1AB342E683964F812CA7C672E15F0407E6D3A
9253E064F9ECD01A49DD7D4C5B22388367F9C26108
FCC4AE2D94B169A322E29F65B02438FC0EC648AB94
9BAE006E270C4F17E52B40D11E2CDC8782C7AB873F
D625119DB250AED39E3CFF60F52635708BB36ED60C
8FEA5EC4
9F47 | len:1 ICC Public Key Expo:
    03
```

However, the PAN in this case identifies the Payment Token and is not the primary account number (see EMV Payment Tokenisation Specification Technical Framework v2.2, Page 74). This is a "shared token".

A payment from a Samsung phone with the same card loaded into the payment application gives the following trace:

```
70 | len:81 Record Template
9F69 | len:7 Card Authentication
    Related Data:
    01DC7BA93B0000
9F4B | len:128 Signed Dynamic Application
    Data (SDAD):
80892716925C6DBD4AB0817A929D40A6D56DBE5853
5ACC74B05C491BAA28E62D6951FFC5F49DE9FAB973
89DE800AFD04D391DEF44152C212F6959100B479BE
204124F847A4C005481D3998EDD5C2349F50274900
13DE56D77F98084EC49D748B2F45680075EF7F7868
13E6AF17851F26DF92392363F85AEF8ED225F9E46
2C41
9F07 | len:2 Application Usage Control:
    C000
5F28 | len:2 Issuer Country Code:
    0826
5F24 | len:3 Application Expiration Date
    YYMMDD: 231231
5A | len:8 Application Primary Account
    Number (PAN):
    4831920272190329
9F19 | len:6 Token Requestor ID:
    040010043095
```

```
9F24 | len:29 Payment Account Reference
    (PAR) :
563030313030313330313631393633353535363931
3035303937383933
```

We see that while the Token Requestor ID and Payment Token are different, the PAR values are the same. According to the EMV FAQs on tokenisation [45] the PAR "was introduced to resolve the challenges faced in the broader acceptance community including Merchants, Acquirers and Payment Processors, in regards to linking Payment Token transactions with each other or transactions initiated on the underlying PAN. This supports a variety of payment processes and value added services". So, in this scenario the payments can be linked both by the Payment Token and more widely by the PAR.