



Swiped: Analyzing Ground-truth Data of a Marketplace for Stolen Debit and Credit Cards

Maxwell Aliapoulios, Cameron Ballard, Rasika Bhalerao, Tobias Lauinger, and
Damon McCoy, *New York University*

<https://www.usenix.org/conference/usenixsecurity21/presentation/aliapoulios>

This paper is included in the Proceedings of the
30th USENIX Security Symposium.

August 11-13, 2021

978-1-939133-24-3

Open access to the Proceedings of the
30th USENIX Security Symposium
is sponsored by USENIX.

Swiped: Analyzing Ground-truth Data of a Marketplace for Stolen Debit and Credit Cards

Maxwell Aliapoulios, Cameron Ballard, Rasika Bhalerao, Tobias Lauinger, Damon McCoy
New York University

Abstract

This paper presents the first empirical study of ground-truth data from a major underground shop selling stolen credit and debit cards. To date, there is little quantitative knowledge about how this segment of the underground economy operates, despite it causing fraud losses estimated at billions of dollars a year. Our analysis of four years of leaked transactional data allows us to characterize this shop's business model, sellers, customers, and finances. The shop earned close to \$104 M in gross revenue, and listed over 19 M unique card numbers for sale. Around 97% of the inventory was stolen magnetic stripe data, commonly used to produce counterfeit cards for in-person payments. Perhaps surprisingly, customers purchased only 40% of this inventory. In contrast, the shop sold 83% of its card-not-present inventory, used for online fraud, which appeared to be in short supply. Demand and pricing were not uniform, as buyers appeared to perceive some banks as having weaker countermeasures against fraud. Even multiple years into the U.S. EMV chip deployment, the supply of stolen magnetic stripe data continued to increase sharply. In particular, we identified a continuing supply of newly issued cards not equipped with EMV chips, especially among prepaid cards. Our findings suggest that improvements to EMV chip deployment in the U.S., combined with a limited supply of stolen card-not-present data, could be avenues to decreasing the revenue and profitability of this shop.

1 Introduction

Fraud due to counterfeit credit and debit cards is a growing problem, estimated at 20 billion dollars worldwide for 2018 [10]. These losses were not distributed evenly. E.U. countries experienced some of the lowest levels of fraud, and the U.S. some of the highest [10]. This is largely attributed to the E.U.'s early adoption of anti-counterfeit chip technology *EMV*. The U.S. introduced EMV only relatively recently, and has not yet achieved comprehensive deployment. In October 2015, liability for counterfeit card payments shifted to

merchants failing to process chip payments [6, 19]. Yet, the U.S. Federal Reserve estimated that in 2018, 43.3% of *card-present* (in-person) payments were still processed by reading the magnetic stripe instead of a chip [8].

For card-present payments, counterfeit cards are typically produced by encoding magnetic stripe data stolen from authentic cards. Magnetic stripe data may be stolen through breaches of merchants' Point of Sale (PoS) terminals [26], or using skimmers installed in ATMs [36] and gas pumps [16,35]. Often, the data is then resold in forums and marketplaces.

Much of the prior academic work has focused on the communities behind this fraud [11, 21, 24, 31, 41, 44], and on developing methods to detect physical skimming devices [16,35,36] or cloned magnetic stripe cards [37]. Several industry studies have been able to provide insights into pricing based on data scraped from *carding shops* dedicated to selling stolen credit and debit card data [9, 12, 22]. A key limitation of these studies was that they were based on external measurements with limited visibility into internal operations. For example, to date there is little understanding of the financial aspects and profitability of such shops, and we do not know which parts of a shop's inventory are actually purchased.

This paper presents the first empirical case study based on ground-truth data of a major shop selling stolen credit and debit cards with a focus on the U.S. market. The data was leaked in a breach of the carding shop. When we received a copy of the database, we needed to reverse engineer its schema, and conducted internal and external validity checks to assess its authenticity.

The leaked database covers the period from January 2015 to January 2019. During this time, the shop earned close to \$104 M in gross revenue, and listed over 19 M unique card numbers (stolen *accounts*) for sale. The majority (97%) of the inventory was stolen magnetic stripe data. Perhaps counterintuitively, magnetic stripe supply increased after the EMV liability shift; in 2018 and 2019, the shop added an average of 93,600 stolen magnetic stripe accounts per week. This supports reports that large-scale breaches of PoS systems are fairly common [26–28].

We conducted an analysis of EMV deployment in the U.S. from the perspective of the carding shop. In the last two years of the leaked data, 85% of the stolen magnetic stripe data originated from EMV chip-enabled cards. This suggests that current incentives might be insufficient to reduce risky use and acceptance of magnetic stripe transactions. Furthermore, even three years after the liability shift, there still was a small but persistent supply of newly issued cards without chips, especially among prepaid cards.¹ Such non-EMV accounts saw much greater demand than EMV accounts, and made up 30.4% of the shop's gross revenue after the liability shift.

Out of the over 19 M accounts listed in the shop, 60% did not sell, despite prices starting at only 21 cents. We investigated what made such a large fraction of stolen accounts apparently undesirable for carders, and found that they preferred to purchase magnetic stripe accounts issued by certain banks but not others. In particular, carders appeared to prefer accounts from medium-sized and smaller banks. This suggests that buyers perceived differences in anti-fraud measures.

We estimate that the shop earned \$24 M before labor and infrastructure costs, with profits growing consistently over the years. Revenue from stolen magnetic stripe data was flat in 2017 and 2018, but it still accounted for 92.2% of gross revenue in 2018. The top 5 magnetic stripe buyers in 2018 spent over \$100k each on stolen magnetic stripe accounts, indicating that they were likely able to evade EMV and transactional risk-based anti-fraud measures.

Around 3% of the shop's inventory was *card-not-present* data used for online fraud. Supply and demand were increasing, but these accounts only made up 7.8% of gross revenue in the last year of the leaked data. The shop paid sellers higher commission rates for stolen card-not-present accounts, yet it appeared unable to attract supply at the same level as magnetic stripe accounts. Based on the perspective of this one shop, it appeared to be more difficult in the U.S. to steal large amounts of card-not-present accounts as opposed to magnetic stripe accounts.

This paper makes the following contributions:

- We characterize the behavior of buyers, and show on the basis of pricing and demand that buyers had clear preferences among card issuers and card types.
- We study the state of U.S. EMV deployment through the lens of this shop. While effects of EMV were visible, deployment had no major impact on the shop's prosperity.

2 Background

The sale of stolen payment cards online has a long history. It has been conducted in public Internet Relay Chat (IRC) channels [21] and underground forums [11, 24, 31, 34]. More recently, payment card sales have migrated away from these

¹Prepaid cards are not linked to a credit or conventional bank account; only funds deposited into the associated prepaid account can be spent.

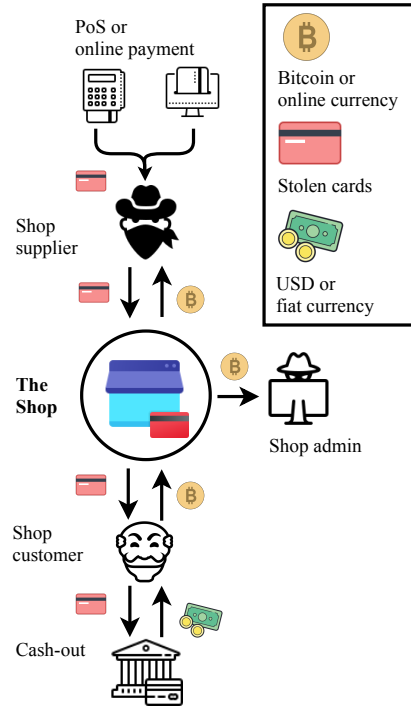


Figure 1: Stakeholder overview. The carder shop receives *accounts* (payment card data) from suppliers who presumably steal them by compromising PoS (Point of Sale) or online payment systems. Shop customers buy accounts for use in fraudulent purchases or for resale. This paper focuses on the shop; how accounts are stolen or cashed out is out of scope.

ad-hoc channels to more streamlined, dedicated *carder shop* websites [29]. These carder shops have functionality and a degree of automation similar to their legitimate e-commerce counterparts. For example, carder shops feature automated customer account creation, a search interface for available payment cards, shopping carts, automated checkout and payment, automated refunds, and ticket-based customer support.

This new carder shop structure has created specializations and likely efficiencies in the ecosystem. As Figure 1 illustrates, it is no longer the case that the person stealing the cards has to deal with the rest of the sales process. Instead, carder shops, including the one that we study, are often market platforms where multiple sellers provide stolen payment cards on consignment to the shop operators. The carder shop pays the sellers a commission for each sale, and handles tasks such as vetting of sellers, advertising of the shop on underground forums, building and maintaining the website infrastructure, payment processing, and customer support. This platform structure is also used in other illegitimate markets such as drug sales [20, 38] and bullet proof hosting [33].

2.1 “The Shop”

The shop that we study in this paper has been in operation since 2015, and can be reached at a regular Internet domain name. As indicated in the leak, it was run by two administrators. To become a customer, users create an account and make a deposit. While paid in Bitcoin, the balance (and the prices in the shop) are marked in U.S. dollars. Customers can search the shop’s inventory according to features such as the card network (e.g., MasterCard or Visa), type (e.g., business or gold), bank, zip code, or price. This enables buyers to identify cards they believe to be more likely to complete the intended fraudulent transaction. Upon purchase, they have the option to check the validity of a card for \$0.50. The card “checker” presumably requests authorization for a small test purchase using the card data. If authorization fails, the buyer obtains an automatic refund for the card. (Certain card types are not eligible for refunds, and the shop places time restrictions and upper limits on how many refunds a single buyer can obtain.) It is worth noting that the shop sells each card only once, to give the buyer confidence that it has not been previously used for fraud. The shop is continually restocked with new releases of stolen cards. The most loyal customers enjoy exclusive access to new cards for the first hours after release, and receive a discount on all their purchases. Based on a review of conversations in the customer support ticketing system, it appears that prices, and discounts for older stock, are set by the shop. Supply is provided by external sellers, who negotiate their commission with the shop.

2.2 Payment Cards, Authorization and Fraud

When we refer to stolen payment cards in this paper, we do not mean cards that have been physically stolen. Instead, they have had some of their data copied, which would allow their use or cloning without access to the original card. Payment cards are typically identified by their number and expiration date. We refer to this data as a stolen *account*. The first six digits of the card number are the *Bank Identification Number (BIN)*, indicating the card network and the issuing bank. In some cases, the card network issues all or most of the cards; in other cases, the network acts as an umbrella for independent banks and their technical service providers.

In order to prevent fraud, a payment usually cannot be processed with knowledge of the card number and expiration date alone. What exactly is needed to authorize a payment depends on various circumstances, including whether the payment is processed with the *card present* or the *card not present*, the physical security features of the card, such as magnetic stripe, contact chip or contactless, the capabilities of the device reading the card, and the policies and anti-fraud measures implemented by the involved parties such as the merchant, the merchant’s payment processor, the card network, and the issuing bank.

A fragmented ecosystem and differing anti-fraud mechanisms imply that opportunities for fraud are not uniform. Carders frequently share strategies for monetizing stolen cards and avoiding anti-fraud systems. These tips include attempting fraudulent transactions close to where the legitimate card owner likely lives, selecting banks perceived to have less effective anti-fraud systems, and making specific types of purchases that target gaps in these systems. To wit, underground carding forum members frequently sell lists of BINs that are thought to be more vulnerable to fraudulent transactions.

In the following, we briefly discuss the differences between card-present and card-not-present payment authorization, and which card data is typically required during the authorization process. We base these definitions on the manner in which stolen accounts were marketed on the shop.

Card Present (Magnetic Stripe). In this scenario, the payment card is physically read by the merchant’s point of sale (PoS) system. Traditionally, this meant “swiping” the card to read the track data encoded in the magnetic stripe, which includes the card number, expiration date, cardholder name, and CVV1. The card verification value CVV1 is only encoded on the magnetic stripe but not visible or typically known to the cardholder; when authorizing transactions, it is used as an indicator that the magnetic stripe of a card has been read. Since magnetic stripe data is static, it is trivial to clone cards for fraudulent use. Countermeasures include the introduction of contactless and contact chip (EMV) technology [1–4]. However, adoption of EMV technology in the U.S. has lagged behind other countries [43] with only 57% of 2018 U.S. card present transactions processed using EMV [8]. Although vulnerabilities have been discovered in the EMV specification [17, 32], we did not find any indications of these being exploited in the leaked data.

The shop calls data stolen during card-present transactions “dumps;” we refer to them as *magnetic stripe* accounts to distinguish from chip-based attacks. The shop also contains data stolen from chip cards, but we believe that this data originates from reading the cards’ magnetic stripe despite the presence of the chip. (Chips are protected against cloning through private data that cannot be read by the terminal, and the publicly readable chip data contains an iCVV instead of the CVV1 that is encoded in the magnetic stripe data.) Magnetic stripe data is commonly stolen from compromised PoS systems [26] or with physical card skimmers installed in ATMs and gas pumps [16, 35, 36]. Some skimmers additionally record PINs needed for ATM withdrawals, but we do not consider this category in our analysis since only 5,801 (0.03%) accounts with PINs were for sale in the shop.

Card Not Present (CNP). This authorization method is used in scenarios where the card cannot be physically present, such as in online shopping. It requires information visually present on the card, or known to the cardholder. All transactions require the card’s account number, and many online payment processors also verify the expiration date, CVV2, and billing

zip code. Some banks may request additional second-factor authentication before authorizing online payments, such as a one-time code sent to the cardholder's phone. This requirement appears to be most common among international banks but not in the U.S., the main market of the shop that we study. In carder slang, this type of stolen payment data is called "cards" or "CVV2s," we refer to it as *CNP* accounts. The CVV2 is a card verification value printed on the card, but not electronically encoded and therefore not automatically read during payments in brick-and-mortar stores. The billing zip code is not present on the card. By design, *CNP* authorization data is chosen such that it cannot be used for card present authorization, and vice versa. As a result, stolen *CNP* data typically originates from compromised online merchants, their payment processors, or end host malware infections that steal card information entered into a website. If an online payment processor does not verify the CVV2 and billing zip code, scammers can also attempt to use magnetic stripe data for online purchases.

2.3 Related Work

The ecosystem of carder shops includes hundreds of sellers, and thousands of buyers. While the mechanics of the older IRC and underground forum business models are described in prior work [11, 16, 21, 24, 31, 41, 44], the dynamics of the actors and the underlying constants that define the present economics are not well understood. Prior work [12, 22, 38, 42] has used limited scraped data to describe some of what is offered for sale on carder marketplaces, but lacked finer grained data. In this paper, we empirically explore interesting questions such as the impact of EMV chip deployment by focusing in depth on the leaked data from one of the major carder shops.

A study on the impact of EMV chip deployment in the U.K. found a strong displacement effect of fraud moving from card-present to card-not-present transactions [13, 14]. We observed far less pronounced displacement in this shop. Another study cited carder shop prices of \$20–30 per account [16]. While this range matched initial offer prices, we show that buyers were often able to purchase accounts at discounted prices.

Prior work has developed methodologies and performed analyses of leaked or seized back-end data of for-profit cybercrime enterprises including bullet-proof hosting [33], DDoS attack services [18, 25], fake AV [39], illicit pharmaceuticals [30], reshipping scams [23], and spam campaigns [40]. In our study, we adapted some of these methods to the leaked dataset, and developed new analysis methodologies and metrics tailored to understanding marketplaces selling stolen accounts. This has resulted in a deeper understanding of these marketplaces, and an analysis framework that could be reused on similar datasets by stakeholders such as law enforcement.

3 Authenticity and Ethics

The data that we analyze in this study was not collected by us, but initially obtained by an unknown third party through a presumably unauthorized means of access. We do not speculate about the motives behind this hack, but it is clear that as in prior studies, our use of "found data" creates concerns of authenticity and ethics [30, 39, 40].

3.1 Ethics

The first concern is whether it is acceptable to analyze third-party data likely obtained through unlawful means, and without the usual standard of "informed consent." We make a utilitarian argument in line with prior studies of criminal backends [18, 23, 30, 40]. Better understanding the economic forces and common strategies in financial fraud provides benefits to society by helping to improve countermeasures and reduce future losses. The authors also believe that making this knowledge more widely known does not contribute to future fraud, as fraud strategies are already well documented in underground forums.

We additionally took a number of steps to ensure that our analysis does not create further harm. Firstly, we do not name the shop in order not to provide any validation or benefit to the perpetrator of the hack. We also wish to avoid publicity for the shop, which is still open for business. Secondly, when we received a copy of the leaked data, we were assured that law enforcement and other stakeholders such as card networks and banks had already been notified about the affected accounts. Thirdly, we only report aggregate or pseudonymous data. Personally identifiable information was either removed from the dataset before it was shared with us, or hashed to allow detection of duplicates while avoiding identification of involved parties. Most notably, the analyzed database contained no identifying information about cardholders, as names had been removed, and account numbers had been hashed, except for the first six digits identifying the bank. We submitted our protocol to NYU's Institutional Review Board (IRB) and our study was deemed to not be human subjects research and it was exempted.

3.2 Authenticity

Given the circumstances of how the dataset was obtained, we needed to assess its authenticity and accuracy. As the first step, we confirmed with security companies that information in the database matched information they had previously scraped from the shop. Furthermore, we received confirmation that test purchases done on behalf of banks were indeed present in the database.

Second, we considered the internal and cross consistency of the data. There were direct concordances between the different elements of the database schema. For example,

when (re)computing seller commissions for batch releases of accounts, we encountered only 16 (0.2%) of over 8,505 seller-release combinations where our calculations did not match the shop's calculations. The median disagreement was roughly \$30, or the price of about two accounts. We also compared the number of sales per release in the orders table to the recorded number of sales in the release statistics table. There were only 11 (0.1%) of over 8,505 seller-release combinations where the numbers differed. Disagreements between order and inventory tables were minimal as well. Out of over 7 million purchased magnetic stripe accounts, only 46 did not appear in the inventory table. For CNP accounts, 59 out of 300 k purchases were not listed in the inventory table. We also found concordances between transaction and customer support ticket data. For example, refund tickets had a corresponding shop transaction. Finally, we determined that 96.2% of 260 k unique Bitcoin wallet addresses from the database were present on the public blockchain. This makes us confident that the data we analyze is indeed authentic.

4 Data

While we received a copy of the leaked database, it did not include any code. Consequently, we needed to reverse engineer the database schema and resolve ambiguities in the data.

Shop User Accounts. The database contains a table with information related to individual shop users, such as user identifiers and current account balances. Some entries were ostensibly test accounts, others were created to vent frustrations of being banned. For the purpose of this study, we consider only active user accounts that have successfully bought or sold accounts, which reduces the number of user accounts from 89,196 to 75,109. We also deduplicated user accounts based on information the shop administrators kept in the table. Duplicate user accounts were created when a user was locked out of their original account, since the site did not appear to support password resets. Some users appeared to create new accounts when the reputation of their old one was such that they were no longer issued refunds. This further reduced the number of accounts to 67,812 buyers and 121 sellers.

Card Account Information. Two tables contained a complete record of all stolen accounts listed on the shop, split into magnetic stripe and CNP. This data included the card number (hashed in our case), BIN, timestamp of when the card was initially available for purchase, zip code (of the billing address for CNP, and presumably where the card was stolen for magnetic stripe), and additional metadata such as the card type. The shop database also listed the issuing banks' names, which we manually mapped to a uniform representation in case a single bank's name appeared with variations.

Shop Transactions. The database contained a table with a row for each completed order. We utilized several data points in this table to perform our analysis of card pricing, buyer activity, and revenue. We were able to compute the fees

collected by the shop for each sale by joining two internal tables. Based on orders and fee percentages, we computed the seller's commission for each release.

To identify whether an order had been subsequently refunded, we combined the orders table with another table that indicated user balances before and after a transaction, where a refund can be identified by the after balance being larger.

Customer Support Tickets. The database also contained customer support tickets, which we used to contextualize and anecdotally support some of our findings. Site users, both buyers and sellers, created support tickets through a dedicated "support" tab on the shop. In their responses to these private inquiries, site admin often used boiler plate language, indicating that many shop users had similar questions.

5 Analysis

We begin our analysis of the shop with an overview of key statistics over the four-year span of the dataset, as summarized in Table 1. A total of 19 million unique accounts issued by 7,092 different banks were listed for sale. The shop had accumulated \$103.9 M in gross revenue over the 4 years of the data. The vast majority (95%) of this revenue was from sales of stolen magnetic stripe accounts, over 20 times more than the \$4.8 M in gross revenue for CNP. The relative demand for CNP accounts, however, was far greater than for magnetic stripe accounts. Indeed, the shop sold 84% of all CNP inventory, in contrast to only 40% of available magnetic stripe accounts. With an inventory 42 times smaller than magnetic stripe, the CNP market appeared to be limited by supply.

After deduplication, we counted 67,813 unique buyers and 121 sellers who had completed at least one purchase or sale on the marketplace. There were 11 dual-role accounts that both bought and sold. The markets for magnetic stripe and CNP were fairly segmented, with only eight sellers (7%), and 21,718 buyers (32%) active in both. These "universal" shop users were more prolific than users operating in only one domain. For example, their median net spend was 9.7 times higher than CNP-only buyers (3.7 times higher than magnetic stripe-only), and their median commissions were 4.4 times that of specialized sellers. This suggests that only the most skilled carders operated in both markets. In contrast, most small-to-mid level actors were active in only one domain, and may have specialized due to different skill sets and strategies necessary for stealing and cashing out magnetic stripe and CNP accounts. Most of our analysis explores these two largely disparate markets separately.

5.1 Sellers

The shop depends on a consistent supply of freshly stolen accounts, which are offered by external sellers. Gross revenue of the shop was heavily biased towards the top sellers. The largest seller alone contributed over 28% of the entire

Type	Sellers	Buyers	Releases	Inventory	Sold	Purchases	Revenue
CNP	11 (9%)	31 K (46%)	523 (6%)	448 K (3%)	374 K (4%)	278 K (11%)	\$4.8 M (5%)
Magnetic stripe	118 (97%)	59 K (86%)	7,821 (94%)	19 M (97%)	7.5 M (96%)	2.4 M (89%)	\$99.1 M (95%)
Overall	121	68 K	8,349	19.45 M	7.83 M	2.69 M	\$103.9 M

Table 1: Summary of the carder shop data, January 2015 – January 2019. Sellers and buyers listed after de-duplication, with at least one sale or purchase. Releases are batches of stolen accounts. Inventory and Sold refer to the total number of accounts available and purchased, respectively. Purchases are sales transactions. Revenue is the total gross sales, before refunds.

shop’s gross revenue. When considering only CNP sales, the largest seller in that domain was responsible for almost 60% of revenue. Furthermore, as shown in Figure 2, just 18 of the highest-revenue sellers (15%) accounted for 81% of the shop’s gross revenue. From an intervention perspective, this indicates that undermining the activities of these few top sellers could significantly decrease the supply of stolen accounts, and consequently the shop’s revenue. Especially the more supply-constrained CNP market relies on its top seller. Given the almost two times larger fraction of unsold magnetic stripe inventory, however, it is possible that demand in that market could be satisfied by the remaining sellers and inventory (albeit it is unclear whether their supply is as *attractive* as that of the top sellers).

Accounts listed on the shop were grouped in releases, each originating from a single seller. These releases had median sizes of 791 magnetic stripe accounts, and 564 accounts for CNP. Each release was assigned a seller commission rate based on negotiation with the shop. These seller commissions varied considerably, with a minimum of 45% and a maximum of 90%. Perhaps due to the more restricted supply, the average commission for a CNP release was 81%, whereas an average magnetic stripe release yielded only 55%. As shown in Figure 3, the median seller earned \$47 k in commissions, and the highest earning seller almost \$16.9 M.

Based on the perspective of this single shop, it appears that many sellers are capable of sourcing stolen magnetic stripe data, and that competition among sellers has resulted in reduced commissions. On the other hand, few sellers seem to be able to steal CNP account data, which may have allowed them to negotiate higher commissions.

5.2 Buyers

The shop did not vet buyers, and provided a fully automated account creation and purchasing process. As a result, 21,209 users in the database never made any purchase. The shop attracted buyers of varying sophistication levels, ranging from “amateur” to “professional” fraudsters. From a financial point of view, nearly all spending in the shop was due to repeat customers (99.1% for magnetic stripe, and 91.9% of CNP). Loyal buyers with higher spending and lower refund rates were given discounts and access to new releases of accounts

before the rest of the customers. On the other end of the customer spectrum, we observed support tickets from likely inexperienced buyers requesting assistance in selecting accounts for purchase. The willingness of the shop to give advice to novice buyers highlights the potential of automated shops to facilitate both professional and amateur fraud.

As shown in Figure 6, buyers of both magnetic stripe and CNP accounts tended to spend more than buyers who focused on one account type, but all kinds of buyers exhibited large variation in the amount they spent. Even though the shop had over 67 k active buyers, 81.3% of total spending was concentrated in the top 9.3% (6,296) of buyers, as shown in Figure 5. The highest spending buyer accounted for only 0.48% (\$495 k) of the shop’s total revenue. This implies that an intervention targeting only the highest spending buyers would have a limited effect on the shop, unless it can disrupt thousands of buyers.

5.3 Pricing

When customers purchased accounts, the prices they paid ranged from \$0.21 to \$256.76 for magnetic stripe data, with a median price of \$13.91. CNP accounts ranged from \$0.93 to \$48.50 with a median of \$12.61. (These purchase prices do not account for refunds, which we discuss in Section 5.5.)

Purchase prices of accounts changed over the course of the dataset. The median purchase price of a magnetic stripe account decreased from \$15.66 in 2015 to \$12.75 in 2018 (–18.6%), whereas it increased from \$5.46 to \$14.55 for CNP (+166%). (This is also shown in Figure 12 in the appendix.) During this time period, the shop witnessed an accumulating oversupply of magnetic stripe accounts, and increasing demand for CNP accounts in short supply.

Customers frequently paid less than full price for their purchases. Around 31% of purchased magnetic stripe, and 11% of CNP accounts were advertised and sold at a discount. The shop offered such discounts for older stock or bulk account packages. In addition to these discounts that were available to all customers, the shop also granted discounts to loyal customers with a high purchase volume. Loyalty discounts applied to 75% of magnetic stripe and 90% of purchased CNP accounts, potentially overlapping with advertised discounts. Across all sold accounts (including non-discounted sales),

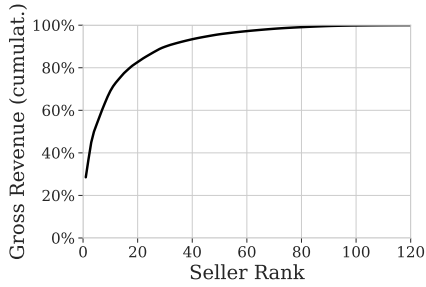


Figure 2: Rank plot of sellers' cumulative gross revenue. The top 18 sellers contributed 81% of the shop's revenue.

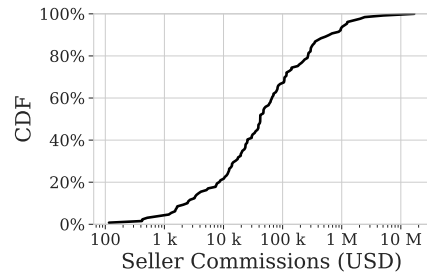


Figure 3: CDF of seller commissions (log scale). The median seller earned \$47k; the top seller \$16.9M (net of refunds).

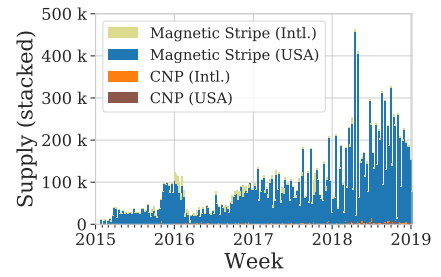


Figure 4: New accounts added weekly (stacked). Supply generally grew over the lifetime of the shop.

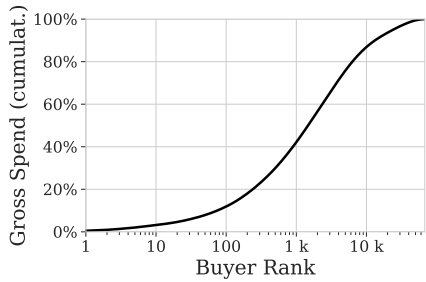


Figure 5: Rank plot of buyers' cumulative gross spend (log scale). The top 6,296 buyers (9.3%) spent 81.25% of the total.




Figure 6: CDF of buyer spending (log scale, net of refunds). Buyers active in both areas spent most (maximum \$495k).




Figure 7: Weekly gross revenue (stacked). Magnetic stripe dominated, but CNP's share was larger than its share of supply.

both discounts combined amounted to an average of \$4.49 per magnetic stripe account and \$1.33 for CNP, corresponding to 25% and 9% of the total listing prices, respectively. We provide detailed discount information to highlight a challenge for researchers monitoring carder shops by scraping their websites; advertised discounts require scraping the entire inventory regularly in order to make accurate revenue estimates. Furthermore, the quantity of loyalty discounts and refunded purchases may not be visible externally at all. It is an interesting question for future work how accurately internal shop data could be inferred from externally visible data, and whether increased breach data sharing among banks could potentially allow for more accurate inferences.

5.4 Supply & Demand

Weekly supply of accounts increased consistently over the lifetime of the shop, with a few significant spikes. While the shop added an overall average of 38,800 accounts per week in 2015, Figure 4 shows a temporary increase to more than 70,000 accounts per week in late 2015 and early 2016 (peaking at 123,929 accounts in one week). After that, average weekly supply increased each year, with 48,400 accounts added weekly in 2016, 64,600 in 2017, and 93,600 in 2018 and 2019. Many of the spikes were due to large releases from one or a few sellers, including the temporary increase in

late 2015 and early 2016, which was caused by an influx of accounts from Australia.

Weekly revenue, driven largely by magnetic stripe accounts and shown in Figure 7, mirrors some but not all trends of supply. Perhaps not surprisingly, revenue grew only slowly in the first six months of the shop's existence. The temporary spike in inventory of late 2015/early 2016 was mirrored in increased revenue, reaching more than \$700k per week for nine weeks, with a maximum of \$904k. Most notably, revenue during this period increased disproportionately for international magnetic stripe accounts, largely due to an influx of accounts from Australia. However, this effect subsided after a few months. (We investigate international accounts in Section 7.2.) While supply increased throughout the rest of the shop's lifetime, revenue never again exceeded \$800k per week. This suggests that the shop may have reached saturation in the magnetic stripe market, or at least an excess of less desirable accounts.

The market for CNP accounts appeared to be in a different situation. The first such accounts were added to the shop in October 2015, and never made up more than 1.7% of the accounts available for purchase on the store. Sales of CNP, however, accounted for up to 11% of the store's weekly revenue. The number of CNP accounts added to the store each week grew at an average rate of 22.7% per week, over four times as fast as the accumulation of unsold inventory (4.8%), meaning that as supply of CNP accounts increased, demand

grew along with it. For comparison, magnetic stripe accounts added to the store and unsold inventory grew at nearly identical rates of 4.0% and 3.7%, respectively. This indicates potential latent demand for CNP accounts, in contrast to the large back stock of less desirable magnetic stripe accounts.

Both the supply of accounts and spending of buyers exhibited regional differences. We based regional data on the zip code where the account was stolen for magnetic stripe accounts, and the billing zip code for CNP. Figure 8 shows a heatmap of this data in each state, normalized by capita. (We excluded Washington, D.C., because its unique status as a single city resulted in much higher per-capita spending than any state.) For magnetic stripe sales, South Carolina was by far the most popular state, with nearly one dollar spent per inhabitant (60% more than the next highest state). This finding is consistent with customer support conversations encouraging buyers to perform fraudulent magnetic stripe transactions in the southeastern U.S., where anti-fraud measures were perceived to be weaker. Colorado and Nevada stood out as hot spots for accounts added, but not for spending. In 2018, one seller added over 700k accounts from Colorado and 230k accounts from Nevada. In Colorado, this amounted to nearly eight times more accounts than all other years combined. However, per-capita spending was only three times that of other years, a much smaller proportional increase. In Nevada, per-capita spending in 2018 decreased by 23% compared to the previous year, despite the large increase in regional supply. These findings show that there may be factors other than supply driving regional demand of magnetic stripe accounts. Understanding these factors, whether they are possible security deficiencies of banks based in the region, a local presence of carder networks and other supporting criminal infrastructure, or simply a myth among carders, might illuminate what makes specific accounts more susceptible to fraud attempts.

For CNP accounts, the “home” region of a stolen account seemed to have little effect on purchasing habits. State-by-state per-capita supply and spending were nearly identical, suggesting that regional differences in demand were mainly a consequence of availability. Cashing out an account online requires less attention to the “home” region of the account, since an online purchase can be placed using a proxy IP address geolocated in the billing zip code, whereas in-person transactions would require physical travel. Kansas stood out as a hot spot, with 69% of accounts coming from a single seller, and 93% of those accounts added in 2018.

The shop that we studied had difficulty supplying more stolen CNP data. This is counter to prior research [13, 14] and therefore an important point for future work. Open questions include whether this finding is limited to this specific shop, due to the current stage of EMV transition in the United States, or a more general effect hinting at CNP data being more difficult to steal at a larger scale.

5.5 Refunds

When customers purchase accounts, they can have their validity checked by the shop’s checker services. A declined authorization could mean that the account has been flagged as stolen by the card issuer, and would allow the buyer to receive a refund from the shop, subject to certain restrictions. One of these restrictions is that the shop marks certain categories of accounts as non-refundable.

Shop policies regarding refundability evolved over time. The supply of refundable CNP accounts, for instance, decreased significantly from 96% in 2015 to 17% in 2018, and none of the CNP accounts added in January 2019 was listed as refundable. In responses to support tickets, the shop admin justified not granting refunds for most CNP accounts with the difficulty of accurately checking their validity. Supply of refundable magnetic stripe accounts, however, steadily increased from 46% in 2015 to 84% in January 2019. Over the entire dataset, 46% of purchased magnetic stripe accounts, and 55% of purchased CNP, were sold as non-refundable. Overall, the shop granted refunds for 1.9 M magnetic stripe accounts amounting to \$33.5 M of sales (34% of gross magnetic stripe revenue), and 49.5 k CNP accounts worth \$597 k (12% of gross CNP revenue).

Due to frequently changing refund rules, these trends do not allow us to draw any direct conclusion about issuers’ anti-fraud performance. However, they illustrate the scale and potential impact of this customer-friendly policy on the shop’s revenue margin, as we will further discuss in Section 9.

6 Pricing Strategies

Based on a review of support tickets, it appears that the shop, not sellers, were responsible for setting the prices of accounts. These prices were not uniform, and ranged from \$0.21 to \$256.76. The shop’s pricing strategy had two components, the initial asking price, and a possible discount that could be added at a later point for older back stock. In the following, we present a preliminary exploration of initial and discounted account prices. Our goal is to identify factors that may have influenced pricing, which in turn provides us with indicators for features that make stolen accounts more valuable.

Initial Asking Price. Through a random forest of decision trees, we were able to predict the initial asking price of accounts (irrespective of whether they were purchased) with an R^2 of 0.74 for magnetic stripe accounts. The average validity of a batch of accounts, as indicated by the shop upon release, explained 54% of pricing. According to customer support tickets, the shop computed the validity of a release based on 20 random accounts verified through one of the shop’s checkers when obtaining the release from the seller. In the dataset, the average initial price was \$56.75 in magnetic stripe releases with more than 95% validity, as opposed to \$24.23 in releases advertised as having less than 40% validity. Com-




Figure 8: Seller supply (# of accounts) and gross buyer spending (\$) per capita in each state. South Carolina stood out as a popular area for both supply and demand of magnetic stripe accounts. Colorado and Nevada’s large supply was due to isolated breaches in those states. CNP accounts showed little difference in how supply and demand were distributed across the country.

ments in the ticketing data further support the finding that releases with lower percentages of valid accounts had lower prices. Noteworthy additional features and their importance included debit versus credit (11.4%), type (such as prepaid vs. corporate cards, 10.4%), issuing bank (10.4%), and location (7.1%). The average initial price of magnetic stripe debit cards in the dataset was \$15.33, whereas credit cards cost an average of \$24.49.

For CNP accounts, we did not encounter any pricing features of significant importance. While the decision tree analysis determined that release validity was the most important feature for CNP pricing, the R^2 was only 0.33, and the average price difference for releases of different validity (segmented as above) was less than one dollar.

Sale Price. A similar analysis on the price at which accounts were purchased (after discounts) yielded similar results (R^2 of 0.85 for magnetic stripe, and 0.34 for CNP). Again, the most important feature was the average validity of the release (53% for magnetic stripe, and 86% for CNP), followed by the time during which the account had gone unsold on the shop (15% and 10%, respectively). Figure 9 plots the “shelf time” of accounts before they sold against the median price buyers paid for them, as well as the number of items sold. Around 47% of magnetic stripe sales, and 76% of CNP sales happened during the first 4 weeks of the account being added to the shop. Presumably due to the much more limited supply, CNP sales of older stock declined faster than for magnetic stripe accounts. At the same time, purchase prices of CNP remained relatively stable and did not appear to be correlated with the account’s age. Magnetic stripe buyers, in contrast, tended to purchase higher-priced accounts quickly within the first few weeks of being added. Median purchase prices of magnetic stripe accounts initially started out higher than for CNP (at \$18.48), and gradually decreased to \$2.91 for accounts 20 weeks and older. In customer support conversations, the shop operators indicated that the validity rate of magnetic stripe accounts decreases over time due to banks detecting the common point of purchase² for breaches, while the validity rate

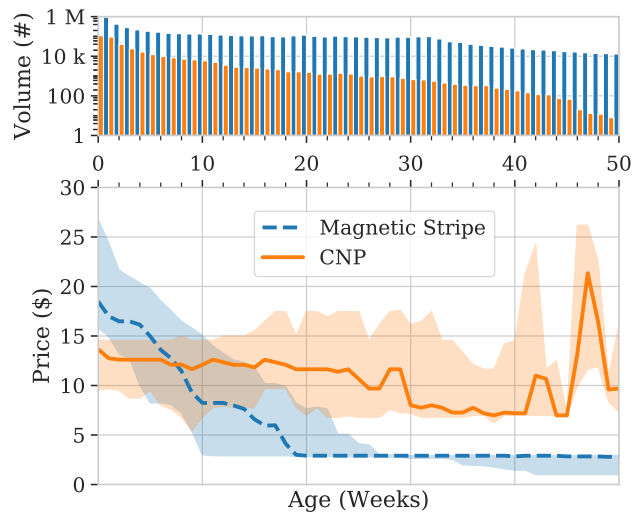


Figure 9: Median purchase price of magnetic stripe and CNP accounts relative to when they were added to the shop (aggregated in weeks); the shaded area corresponds to the range between the 25th and 75th percentiles. Above, sales (#) in log scale. The older magnetic stripe accounts, the less buyers pay for them; CNP prices remain more stable.

of stolen CNP accounts tends to remain more constant since they often do not have a common point of purchase. However, we cannot verify these claims.

Summary. The advertised validity rate of a release had the most impact on listing price, while its time on the shop had the largest impact on sale price. While our preliminary results shed light on factors that appear to influence pricing, we also note that we did not account for possible interactions between features, or longitudinal trends such as the shop optimizing the pricing strategy, or card issuers improving their anti-fraud measures. Furthermore, it is an open question to which extent these findings generalize to other carding shops.

²Common point of purchase is an anti-fraud technique aiming to identify a breached merchant. Starting from accounts reported for fraud, issuers look

for any overlap in their purchase history to infer a common location as the likely source of the breach. This allows issuers to flag additional accounts used at this location as potentially compromised before fraud is attempted.

7 Account Attractiveness

We investigate the varying *attractiveness* of magnetic stripe accounts by segmenting them based on the issuer, network, and type.³ We considered sets of accounts with a characteristic to be more attractive if 1) customers purchased a higher percentage of available accounts, or 2) customers purchased accounts for a higher price. We analyze U.S. and international accounts separately since U.S. accounts comprised 93% (17.4M) of available magnetic stripe accounts, whereas international accounts appeared to be more attractive overall.

7.1 U.S. Accounts

Shop customers bought stolen U.S. magnetic stripe accounts from a total of 6,929 issuers. 43% of this spending was concentrated in the top 10 issuers, which were all larger national U.S. banks (Table 2). While overall spending was in the millions of dollars for each of these banks, there were notable differences in how much of their inventory sold. Customers purchased 83.2% of USAA Savings Bank accounts, for instance, but only 27.0% of American Express-issued accounts.

We considered the next 104 entities medium-size issuers; they accounted for 25% of spending. It is noteworthy that in aggregate, the medium-size issuers had a higher fraction of their inventory bought than any of the top 10 issuers, except for USAA Savings Bank. Similar trends hold for the remaining 6,815 small issuers, which accounted for 22% of spending. For 10% of U.S. accounts we could not determine the issuer;⁴ they accounted for 9% of spending.

There were notable differences in terms of the card network. For the four major card networks, those with a higher rate of purchase also had a higher median price. Carders paid the lowest median price and purchased the smallest portion of American Express cards, then Discover, Mastercard, and Visa. Carders appeared to perceive American Express as having stronger defenses against fraud. However, 28% of American Express-branded accounts were issued by third-party issuers, and it is unclear whether these perceived defenses were at the issuer or card network level.

When segmenting by account type, prepaid debit accounts stand out at a purchase rate of 67.6%, more than 25 percentage points higher than non-prepaid debit or credit accounts. At the same time, prices paid were highest for credit accounts, followed by prepaid debit, and lowest for other debit accounts. The high purchase rate suggests that prepaid debit cards may be perceived as having the weakest anti-fraud measures. The slightly lower prices, in turn, might indicate that carders expect available balances to be lower than those of credit cards.

³We do not explore CNP accounts in this section since our prior analysis suggested that account features did not significantly affect their attractiveness.

⁴Resolving an account to an issuer is done based on lists of BIN-to-issuer mappings. No complete authoritative list of these mappings was publicly available, thus we used the resolution provided by the shop.

Segment	Spend (\$)	Median	Supply	Sold	Refunded
Chase Bank	8.58M	3.00	3.21M	27.4%	14.4%
Capital One Bank	6.02M	21.42	786k	42.1%	31.7%
Wells Fargo Bank	3.74M	4.12	1.54M	27.3%	17.4%
Citibank	3.51M	9.89	736k	37.4%	4.67%
Bank Of America	3.28M	2.91	1.54M	27.9%	16.6%
USAA Savings Bank	3.00M	22.17	169k	83.2%	37.9%
FIA Card Services	2.86M	17.85	381k	45.0%	29.5%
U.S. Bank	2.39M	12.37	494k	40.1%	29.0%
American Express	1.65M	2.91	829k	27.0%	10.7%
TD Bank	1.45M	10.18	289k	46.5%	29.5%
Top 10 Issuers	36.5M	7.89	9.97M	32.1%	19.2%
Medium Issuers	21.2M	14.28	2.92M	53.4%	30.9%
Small Issuers	19.3M	14.45	2.65M	55.2%	33.8%
Unknown Issuer	7.94M	10.50	1.82M	36.9%	23.0%
Card Networks					
Visa	49.8M	10.00	12.1M	36.6%	24.5%
Mastercard	30.9M	15.66	3.72M	54.1%	30.7%
American Express	2.69M	2.91	1.15M	28.8%	10.9%
Discover	1.47M	9.76	355k	33.0%	1.05%
Account Types					
Credit	45.9M	15.28	7.99M	38.7%	23.3%
Non-Prepaid Debit	36.2M	10.20	9.05M	39.6%	26.6%
Prepaid Debit	2.80M	14.45	321k	67.6%	31.6%

Table 2: Magnetic stripe accounts from U.S. issuers, segmented by the top 10 issuers, issuer size (both in terms of total spend), the four major card networks, and card type. Gross spend and median price of purchased accounts are in USD. Supply corresponds to all accounts available for purchase. Refunds relative to number of accounts purchased; some were non-refundable. Certain features appear to make accounts more valuable to carders. For example, buyers purchased a higher fraction of the available inventory for accounts from small and medium-size issuers compared to all but one of the top 10 issuers, and paid a higher median price.

7.2 International Accounts

Non-U.S. accounts made up 7% of magnetic stripe accounts available in the shop, but accounted for a disproportionately higher 14% of magnetic stripe revenue. Support tickets indicated that these international magnetic stripe accounts were likely being used to commit fraud within the U.S. Table 3 shows aggregate statistics for magnetic stripe accounts from issuers outside of the U.S. Compared to the U.S., carders bought a higher fraction of the inventory for 7 of the 10 most popular international countries. Carders also purchased more expensive accounts, with the median price of international magnetic stripe accounts almost twice that of the U.S. This points to anti-fraud measures of international accounts being perceived as weak when used within the U.S.

The supply of international accounts was more restricted than that of the U.S. A large portion of these accounts were added in a few distinct releases, indicating that many of the international accounts likely came from large isolated breaches. For international banks, discovering and responding to these breaches would significantly limit fraud.

Country	Spend (\$) Median		Supply	Sold	Refunded
U.S.A.	84.9M	12.64	17.4M (92.7%)	39.7%	25.3%
All Intl.	14.2M	20.37	1.37M (7.30%)	41.8%	26.6%
Canada	2.42M	10.18	505k (2.70%)	42.5%	19.8%
China	1.27M	52.21	33.2k (0.18%)	59.4%	10.3%
Australia	942k	38.12	41.3k (0.22%)	56.9%	36.5%
Spain	865k	35.37	32.2k (0.17%)	83.3%	47.9%
U.K.	849k	28.17	170k (0.91%)	17.4%	40.7%
Korea	635k	34.80	53.3k (0.28%)	39.3%	30.7%
Germany	587k	38.08	24.0k (0.13%)	65.1%	43.3%
Aruba	546k	34.80	32.6k (0.17%)	56.3%	28.5%
France	402k	34.80	18.1k (0.10%)	65.1%	34.8%
Brazil	350k	31.77	35.8k (0.19%)	29.1%	26.9%
Other	5.33M	30.16	423k (2.26%)	42.8%	27.2%

Table 3: Magnetic stripe accounts segmented by country of the issuer. Gross spend and median price of purchased accounts are in USD. Refunds relative to number of accounts purchased; some were non-refundable. Purchase rates and inventory varied considerably, such as 17.4% of 170k accounts in the U.K., 39.7% of 17.4M in the U.S., and 83.3% of 32.2k in Spain. Median purchase prices for accounts from the U.S. and Canada were lower than all other countries.

Chinese accounts were consistently popular throughout the duration of the dataset. They were almost 50% more expensive than the next highest priced country, and over four times more expensive than U.S. accounts. In late 2018 and early 2019, a single seller added more than 15,000 Chinese magnetic stripe accounts to the shop, a large majority of which were non-EMV cards (82%) and priced at over \$100 per account. As we will discuss in Section 8, magnetic stripes extracted from non-EMV cards might be easier to monetize than magnetic stripes from EMV cards, driving up buyer demand and cost. The shopkeeper also mentioned in support tickets that Chinese accounts typically had poor anti-fraud protections when used within the United States. Australian accounts were the second most expensive, and drove part of the revenue spike in late 2015 and early 2016 when a seller uploaded over 25,000 such accounts to the shop. Canada was the second most popular country in terms of spending, and exhibited similar demand and pricing characteristics as U.S. accounts, probably due to their close proximity.

In summary, we found that carders had apparent preferences for certain issuers and card types. These preferences suggest that carders expected different fraud gains, potentially based on perceived differences in anti-fraud measures. Increased sharing of expertise and intelligence could help the “weaker” issuers improve their defenses, especially in the case of international banks targeted for fraud in the U.S.

8 U.S. EMV Chip Deployment

In order to reduce fraud from counterfeit payment cards, issuers have begun equipping their cards with an EMV chip in

addition to the magnetic stripe. These chips, in contrast to magnetic stripes, are thought to be more secure against duplication attacks. To discourage merchants from processing magnetic stripe transactions, card networks imposed a liability shift for card-present transactions involving counterfeit cards. In the U.S., it took effect on October 1, 2015 [19]. (Other major markets had already implemented a similar liability shift prior to that date.) Since this date, merchants, not banks, have been responsible for fraud losses when a card equipped with an EMV chip is processed as a magnetic stripe swipe instead of reading the chip. We study the impact of increasing EMV adoption on the carder shop, especially with regard to the availability and pricing of magnetic stripe data, which is required to produce counterfeit cards for in-store purchases. From a supply perspective, this data can only be stolen when merchants read the magnetic stripe instead of the chip. This may occur when merchants are unwilling or unable to process chip transactions. For example, at the time of writing, there is still an exception from the liability shift for gas pump transactions [5], and there were reports of merchants disabling chip transactions during peak holiday shopping periods in 2015 to shorten checkout times [15]. Our analysis in this section is limited to U.S. magnetic stripe accounts, since EMV is not used for CNP transactions, and other countries had already completed their transition to EMV.

Supply of EMV and non-EMV accounts. The liability shift occurred 10 months after the start of the leaked data. Overall supply of magnetic stripe accounts continued to grow significantly until the end of the dataset, as discussed in Section 5.4. This suggests that breached merchants were still processing large numbers of magnetic stripe transactions, despite the incentive to read the chip. Figure 10 shows that most of this magnetic stripe data was stolen from cards equipped with a chip. New supplies of chipless accounts decreased after the liability shift, but never went to zero. Instead, supply remained at a relatively stable level during the last years of the dataset.

In terms of stock available for purchase in the shop, there was an oversupply of EMV accounts. The supply of EMV accounts nearly doubled every year, such as a 93% increase from 2017 to 2018. However, only 35% of EMV accounts added after the liability shift were purchased. In contrast, 84.2% of non-EMV accounts were purchased, and available stock was effectively shrinking from 359,351 accounts after the liability shift to 192,078 accounts after 2018. As Figure 11 shows, new non-EMV supply was added at a pace similar to the purchase rate, but existing older stock was lost due to the cards reaching their expiration dates.

Sales of non-EMV accounts made up 67.1% of the shop’s gross revenue before the liability shift, but only 30.4% afterwards. EMV sales increased from 32.7% to 68.7%. This increase appeared to be driven mostly by volume, not prices. Around the time of the liability shift, buyers paid a median price of \$20.37 for EMV compared to \$14.45 for non-EMV accounts. In early 2016, however, the median purchase price of

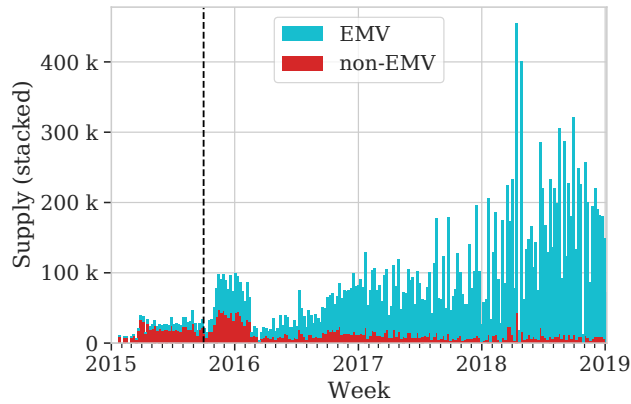


Figure 10: Weekly supply of U.S. magnetic stripe data. Even three years after the EMV liability shift (of October 2015), there was significant increase in supply, suggesting that breached merchants were still processing large numbers of magnetic stripe transactions. Most of the magnetic stripe data was stolen from cards equipped with a chip, but there appeared to be a relatively steady supply of chipless cards, suggesting that non-EMV cards were still being issued.

EMV accounts sharply decreased to \$9.76, and has remained consistently lower than non-EMV accounts since then. Buyers were willing to spend more for non-EMV accounts, but their supply was limited. (Figure 15 in the appendix shows longitudinal pricing for EMV and non-EMV accounts.)

Continued supply of non-EMV accounts. Even three years after the liability shift, the shop continued to receive new supplies of non-EMV accounts. One hypothesis to explain this phenomenon is that these accounts were mostly invalid. However, the 29.8% refund rate for non-EMV accounts was not much higher than the 23.8% refund rate for EMV accounts.

Another hypothesis is that these non-EMV accounts might have been issued before the liability shift, and the issuer was waiting for them to expire before reissuing them with EMV. In this case, we would expect non-EMV accounts to have relatively little time remaining until their expiration date when they were added to the shop. Indeed, in 2016, the median remaining lifespan of non-EMV accounts was 1.2 years shorter than for EMV accounts. From 2016 to 2018, however, the median remaining lifespan of non-EMV accounts increased by about 100 days; the non-EMV population was getting *younger*, whereas EMV accounts aged by the same amount. (Figure 13 in the appendix shows a box plot of the remaining lifespans from 2016 to 2018.) This suggests that new non-EMV cards continued to be issued after the liability shift.

Table 4 shows the percentage of EMV support among the accounts of the ten largest U.S. issuers. For accounts added to the shop before the liability shift, EMV capability ranged from 18% (USAA Savings Bank) to 72% (American Express). Issuers progressed at a different pace during the next




Figure 11: U.S. magnetic stripe inventory by its cumulative weekly purchase status, for EMV (above) and non-EMV accounts (below). Weekly EMV supply was significantly larger than purchases, leading to an accumulation of unsold accounts. In contrast, non-EMV accounts were purchased at a rate similar to new supply; available inventory effectively shrunk due to older accounts reaching their expiration dates.

14 months, with EMV support between 58 and 94%. In 2018 and 2019, around three years after the liability shift, each of these large issuers had reached levels of 91 to 100% of EMV capability. Small and medium-size issuers collectively started from lower EMV capability rates of 5.8% and 13%, respectively, and reached 89% and 91% in 2018.

There was a clear difference in card types among EMV and non-EMV accounts added to the shop after the liability shift. The EMV-capable accounts were composed of 52.8% credit, 46.5% non-prepaid debit, and 0.45% prepaid debit cards. In contrast, only 23% of non-EMV accounts were credit cards. The remainder were debit cards: 67.8% non-prepaid and 9% prepaid debit. Prepaid debit cards appeared to be issued predominantly without EMV capability; 77.4% of prepaid debit cards added after the liability shift had no chip. Anecdotally, we confirmed that many prepaid debit cards issued today do not support EMV. For example, we found that 99.8% (9,559) of Green Dot-issued prepaid cards added to the shop after the liability shift do not support EMV.

In summary, the transition to EMV progressed at a different pace depending on the card issuer and card type, and had not completed at the end of the dataset. Despite the progress in deployment, EMV did not lead to a decrease in the availability of stolen magnetic stripe data in the shop.

Issuer	2015-01 – 2015-09	2015-10 – 2017-12	2018-01 – 2019-01
Chase Bank	58.0k (38%)	1.29M (94%)	1.66M (99%)
Capital One Bank	10.4k (23%)	317k (91%)	400k (99%)
Wells Fargo Bank	17.0k (23%)	535k (80%)	777k (97%)
Citibank	24.4k (59%)	314k (91%)	345k (98%)
Bank Of America	44.9k (57%)	637k (89%)	716k (96%)
USAA Savings Bank	2.81k (18%)	53.2k (58%)	60.9k (100%)
FIA Card Services	21.3k (70%)	168k (94%)	171k (100%)
U.S. Bank	5.36k (22%)	137k (67%)	239k (91%)
American Express	30.1k (72%)	379k (94%)	405k (98%)
TD Bank	5.39k (20%)	159k (70%)	99.8k (93%)
Medium Issuers	20.3k (13%)	1.01M (71%)	1.30M (91%)
Small Issuers	8.36k (5.8%)	632k (54%)	1.21M (89%)
Unknown Issuer	40.9k (52%)	743k (83%)	1.04M (94%)

Table 4: Top 10 U.S. magnetic stripe issuers in terms of total spend, and their proportion of added accounts that were EMV-capable. Major issuers progressed at a different pace, reaching EMV support levels between 91 and 100% around three years after the liability shift.

9 Marketplace Finances

Deterring profit-motivated attackers is akin to disrupting a business process. Thus, it is important to understand the cost structure of this shop, and how the payment industry could increase the shop’s operating costs to reduce profitability. Table 5 depicts the shop’s yearly sales revenue, seller commissions, buyer refunds, and profit margins as they are visible in the leaked database. We do not have information on operating costs such as website infrastructure, employees of the shop, or advertising costs related to customer acquisition. Therefore, the margins we compute are an upper estimate.

Operating Costs. The two main operating costs are commissions paid to sellers, and refunds provided to buyers. The shop monitored refund rates and adjusted refund policies to maintain them at an average of 33% of gross revenue.

Commissions were paid to the seller as a percentage of each non-refunded account sale. These commission rates were negotiated individually with every seller. While the average commission was 65.9%, larger sellers and especially CNP sellers were able to negotiate higher commission rates of up to 90%. Over time, the shop improved its bargaining position with sellers, reducing average commission rates on a total sales basis from 78.6% in 2015 down to 55.7% in 2018.

Margins. We estimate that the shop made an overall profit of up to \$23.8 M, at a 23% margin. From 2016 to 2017, revenue grew by 34.2%, which resulted in a profit increase of \$1.3 M (23.2%). At the same time, the margin fell from 23% to 21% due to higher operating costs from increased refunds. In 2018, gross revenue increased by only 4% (\$1.3 M). All of this increase was from growing CNP sales; magnetic stripe sales decreased slightly by 0.3% (\$92 k). Despite this small increase in revenue from 2017 to 2018, the shop was able

Year	Revenue	Commissions	Refunds	Margins
2015*	13.4M	7.7M (57%)	3.6M (27%)	2.1M (16%)
2016	24M	10.8M (45%)	7.6M (32%)	5.6M (23%)
2017	32.2M	13.6M (42%)	11.8M (37%)	6.8M (21%)
2018	33.5M	13.6M (41%)	10.8M (32%)	9.1M (27%)
2019*	770K	313K (41%)	241K (31%)	217K (28%)
Total	103.9M	46M (44%)	34.1M (33%)	23.8M (23%)

Table 5: Yearly finances of the shop, in USD. *Partial data for 2015 and 2019. The shop earned \$23.8M before costs such as advertising, employees and infrastructure.

to increase profits by \$2.3 M (33.8%) as a result of reducing costs through lower seller commissions and refund rates.

The growth opportunity for the shop appeared to be in CNP sales, while magnetic stripe sales remained the primary source of revenue. A stagnant supply of CNP accounts, and a steep decline in magnetic stripe supply or demand, possibly by improved EMV adoption, might force the shop to reduce costs further, and could ultimately erode its profitability.

10 Discussion and Implications

The shop has created a scalable and lucrative model for selling stolen accounts. Several measures were aimed at maintaining the reputation of the marketplace and the loyalty of customers. During the last months of the dataset, the shop added hundreds of thousands of stolen accounts per week. Despite the introduction of EMV chip cards, the shop accumulated an oversupply of newly stolen magnetic stripe data in the three years following the payment industry’s liability shift towards merchants processing magnetic stripe transactions; the supply was so large that only 40% of the shop’s inventory was purchased. Breaches of PoS systems appear to have become common events, and the risk of magnetic stripe data being stolen when a card is swiped is non-trivial. It appears that the liability shift alone has not been sufficient to disincentivize merchants from swiping EMV-enabled cards and curb the supply of stolen magnetic stripe data. Further disincentives could include increased fees for processing magnetic stripe payments, and liability for breached merchants. If the payment industry fails to agree on more effective self-regulation, government regulation might ultimately be necessary.

A 2018 study from the U.S. Federal Reserve estimated a decline of 20.9% (\$770 M) in card-present fraud, and a 34.4% (\$1.2 B) increase in card-not-present fraud after the U.S. adoption of EMV chip technology [7]. Yet, from the perspective of the shop’s finances, EMV had not caused a major impact (yet). Carders continued to spend millions of dollars to purchase magnetic stripe accounts in the years after the U.S. deployed EMV, suggesting that EMV had not (yet) significantly impaired their ability to conduct fraud with stolen magnetic stripe data.

Buyer preferences, however, did exhibit a noticeable impact from EMV. Magnetic stripe data stolen from cards equipped with a chip appeared to be less desirable than data from chip-less cards. This suggests that carders perceived fraudulent magnetic stripe transactions as less likely to succeed when the data was stolen from EMV-enabled cards. Similarly, carders appeared to perceive several U.S. banks, and many international banks, as having weaker anti-fraud measures than other banks. It is unfortunate that this “folk wisdom” is available to carders, but not to the banks’ legitimate customers. This lack of transparency enables each bank to make independent cost-benefit determinations for their anti-fraud measures, whereas customers are left in the dark about the likelihood of encountering fraud with their payment cards. Requiring banks to disclose quantitative fraud data could help customers make informed decisions, and might incentivize underperforming banks to improve their anti-fraud measures. More generally, better information sharing among banks, and potentially more centralized initiatives at the card network level, could also help make fraud deterrents more uniform.

Carders also appeared to have a preference for prepaid debit cards, which include gift cards, electronic benefit transfer (EBT) cards such as SNAP food benefit cards, and payroll cards for unbanked populations. Transaction risk scoring for prepaid debit cards may be less accurate because issuers have less information about account owners and their transaction history. Furthermore, prepaid account owners may be less likely to regularly check their statements and notice fraudulent transactions. Despite these systemic difficulties in preventing fraud, prepaid debit cards in the carder shop’s inventory had a particularly high fraction of 77.4% not equipped with an EMV chip. The data furthermore suggests that new cards continued to be issued without a chip. For instance, nearly all (99.8%) of the prepaid debit cards issued by Green Dot did not support EMV, despite being added to the shop after the liability shift. Issuers may consider the cost of EMV to be higher than the fraud losses they and their customers might have to bear. As an illustration, Scaife et al. quote a manufacturing cost of \$2.00 for an EMV chip card as opposed to \$0.08 for a magnetic stripe card [37]. Yet, the continued issuance of non-EMV cards makes the goal of disallowing magnetic stripe transactions elusive, and arguably holds back anti-fraud progress in the entire industry. It would be worthwhile studying whether holders of prepaid debit cards, which include low-income populations, are less successful in detecting and reverting fraudulent charges on their prepaid debit cards than holders of regular debit and credit cards. Such a finding could justify regulatory mandates for improved anti-fraud measures in cases where economic considerations prevent issuers from taking these measures voluntarily.

The EMV transition, and the prospect of magnetic stripe supply eventually drying up may be a (distant) threat to the shop. However, three years after the liability shift, the shop still appeared to be prospering, and it is unclear on which

time horizon EMV might cause more pressing issues. Prior work has reported on a trend of payment fraud migrating from card-present to card-not-present transactions, presumably due to EMV [7, 13]. The shop we studied had a relatively limited, but highly demanded inventory of card-not-present accounts, and appeared unable to secure a larger supply. This constraint may eventually become a threat to this particular shop.

Limitations. Our study encompassed data leaked from a single carder shop. While the scale of the data, such as the inventory of over 19 M stolen accounts, and the gross revenue of almost \$104 M suggest that the shop played a significant role, it remains unclear how representative it was at the time of the leak. In particular, we do not know whether supply constraints for CNP accounts translate to the entire ecosystem.

Another limitation is that the leaked data does not include information about how accounts were initially stolen, or how carders attempted to monetize them after purchase. We cannot measure, for instance, whether the introduction of EMV had an influence on expected fraud returns, or on the effort necessary to cash out stolen magnetic stripe accounts.

Lastly, we have no absolute certainty that the leaked data is authentic. However, due to our consistency checks along with vetting by other companies, we are fairly confident that it is authentic.

Future Work. We found that certain card issuers and types of cards commanded higher underground prices or sold in higher numbers; they were disproportionately being targeted for fraud. However, as issuers update their anti-fraud measures, these trends are likely to evolve. While we were able to observe these trends in “ground-truth” back-end data, such data is not commonly available, and it is an open question how future trends could be identified without access to such data. Future work could study whether such trends can be inferred from partial data, such as scraped from the shop’s public front end. Another interesting question is how metrics could be derived for high-level consumer advice, e.g. for customers to compare banks by their aggregate fraud risk.

11 Conclusion

We have presented the first inside analysis of an underground marketplace for stolen credit and debit cards. We found that most of the supply (97%) and revenue (95%) originated from stolen magnetic stripe data. The shop’s supply and profits continued to increase even three years after the U.S. shifted liability for fraud using counterfeit cards to merchants who failed to read the chip of EMV-enabled cards.

The shop accumulated an inventory of 19 M stolen magnetic stripe accounts, but the majority (60%) of them did not sell. Buyers had clear preferences for accounts issued by certain banks, and for cards that did not support EMV, likely because buyers perceived them as more vulnerable to fraud. While we do not know whether these perceptions were accurate, the reality for affected account holders was that their

stolen accounts incurred a disproportionate amount of fraud attempts compared to accounts stolen from issuers perceived as more secure. Our hope is that by further studying and understanding these marketplaces, we can inform new and more effective directions for mitigating this threat.

Acknowledgments

This work was funded by the NSF through grants 1717062 and 2039693. Our research lab has also received gifts from Google. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the view of our funders.

References

- [1] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 1: Application Independent ICC to Terminal Interface Requirements, Version 4.3 ed.* EMVCo, LLC, November 2011.
- [2] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 2: Security and Key Management, Version 4.3 ed.* EMVCo, LLC, November 2011.
- [3] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 3: Application Specification, Version 4.3 ed.* EMVCo, LLC, November 2011.
- [4] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 4: Cardholder, Attendant, and Acquirer Interface Requirements, Version 4.3 ed.* EMVCo, LLC, November 2011.
- [5] EMV at the pump. <https://usa.visa.com/visa-everywhere/security/emv-at-the-pump.html>, 2016.
- [6] US Payments Forum – Understanding the U.S. EMV liability shifts. <https://www.uspaymentsforum.org/wp-content/uploads/2017/07/EMV-Fraud-Liability-Shift-WP-FINAL-July-2017.pdf>, 2017.
- [7] Changes in U.S. payments fraud from 2012 to 2016: Evidence from the Federal Reserve payments study. <https://www.federalreserve.gov/publications/files/changes-in-us-payments-fraud-from-2012-to-2016-20181016.pdf>, October 2018.
- [8] The 2019 Federal Reserve payments study. <https://www.federalreserve.gov/newsevents/pressreleases/files/2019-payments-study-20191219.pdf>, 2019.
- [9] The Armor 2019 black market report. <https://cdn.armor.com/app/uploads/2019/11/2019-Q3-Report-BlackMarket-SinglePages-1.pdf>, 2019.
- [10] Nilson Report – Card fraud losses reach \$27.85 billion. <https://nilsonreport.com/mention/407/1link/>, 2019.
- [11] Sadia Afroz, Vaibhav Garg, Damon McCoy, and Rachel Greenstadt. Honor among thieves: A common’s analysis of cybercrime economies. In *eCRS*, 2013.
- [12] Max Aliapoulios and Ian Gray. BSides Las Vegas 2019 preview: Visualizing Joker’s Stash. <https://www.flashpoint-intel.com/blog/bsides-las-vegas-2019-visualizing-jokers-stash/>, 2019.
- [13] Ross Anderson, Chris Barton, Rainer Bohme, Richard Clayton, Carlos Ganan, Tom Grasso, Michael Levi, Tyler Moore, and Marie Vasek. Measuring the changing cost of cybercrime. In *WEIS*, 2019.
- [14] Ross Anderson, Chris Barton, Rainer Bohme, Richard Clayton, Michel J.G. van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In *WEIS*, 2012.
- [15] Kate Ashford. Chip cards take so long, some retailers disabled them for the holidays. <https://www.forbes.com/sites/kateashford/2015/12/27/chip-cards-take-too-long/>, 2015.
- [16] Nishant Bhaskar, Maxwell Bland, Kirill Levchenko, and Aaron Schulman. Please pay inside: Evaluating Bluetooth-based detection of gas pump skimmers. In *USENIX Security*, 2019.
- [17] M. Bond, O. Choudary, S. J. Murdoch, S. Skorobogatov, and R. Anderson. Chip and skim: Cloning EMV cards with the pre-play attack. In *Security & Privacy Symposium*, 2014.
- [18] Ryan Brunt, Prakhar Pandey, and Damon McCoy. Booted: An analysis of a payment intervention on a DDoS-for-hire service. In *WEIS*, 2017.
- [19] Christopher Budd. October 1, 2015: Happy EMV day! What it means for you. <https://blog.trendmicro.com/october-1-2015-happy-emv-day-what-it-means-for-you/>, 2015.
- [20] Nicolas Christin. Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In *WWW*, 2013.
- [21] Jason Franklin, Vern Paxson, Adrian Perrig, and Stefan Savage. An inquiry into the nature and causes of the wealth of Internet miscreants. In *CCS*, 2007.

- [22] Ian Gray and Vlad Cuiujuclu. Giving credit where it's not due: Visualizing Joker's Stash. <https://go.flashpoint-intel.com/webinar/Jokers-Stash>, 2019.
- [23] Shuang Hao, Kevin Borgolte, Nick Nikiforakis, Gianluca Stringhini, Manuel Egele, Michael Eubanks, Brian Krebs, and Giovanni Vigna. Drops for stuff: An analysis of reshipping mule scams. In *CCS*, 2015.
- [24] Andreas Haslebacher, Jeremiah Onaolapo, and Gianluca Stringhini. All your cards are belong to us: Understanding online carding forums. In *eCrime*, 2017.
- [25] Mohammad Karami and Damon McCoy. Understanding the emerging threat of DDoS-as-a-service. In *LEET*, 2013.
- [26] Brian Krebs. Sources: Target investigating data breach. <https://krebsonsecurity.com/2013/12/sources-target-investigating-data-breach/>, 2013.
- [27] Brian Krebs. Dairy Queen confirms breach at 395 stores. <https://krebsonsecurity.com/2014/10/dairy-queen-confirms-breach-at-395-stores/>, 2014.
- [28] Brian Krebs. Zip codes show extent of Sally Beauty breach. <https://krebsonsecurity.com/2014/03/zip-codes-show-extent-of-sally-beauty-breach/>, 2014.
- [29] Brian Krebs. Russians shut down huge card fraud ring. <https://krebsonsecurity.com/2020/03/russians-shut-down-huge-card-fraud-ring/>, 2020.
- [30] Damon McCoy, Andreas Pitsillidis, Jordan Grant, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey Voelker, Stefan Savage, and Kirill Levchenko. PharmaLeaks: Understanding the business of online pharmaceutical affiliate programs. In *USENIX Security*, 2012.
- [31] Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage, and Geoffrey M Voelker. An analysis of underground forums. In *IMC*, 2011.
- [32] S. J. Murdoch, S. Drimer, R. Anderson, and M. Bond. Chip and PIN is broken. In *Security & Privacy Symposium*, 2010.
- [33] Arman Noroozian, Jan Koenders, Eelco van Veldhuizen, Carlos H. Ganan, Sumayah Alrwais, Damon McCoy, and Michel van Eeten. Platforms in everything: Analyzing ground-truth data on the anatomy and economics of bullet-proof hosting. In *USENIX Security*, 2019.
- [34] Sergio Pastrana, Daniel Thomas, Alice Hutchings, and Richard Clayton. CrimeBB: Enabling cybercrime research on underground forums at scale. In *WWW*, 2018.
- [35] Nolen Scaife, Jasmine Bowers, Christian Peeters, Grant Hernandez, Imani N. Sherman, Patrick Traynor, and Lisa Anthony. Kiss from a rogue: Evaluating detectability of pay-at-the-pump card skimmers. In *Security & Privacy Symposium*, 2019.
- [36] Nolen Scaife, Christian Peeters, and Patrick Traynor. Fear the reaper: Characterization and fast detection of card skimmers. In *USENIX Security*, 2018.
- [37] Nolen Scaife, Christian Peeters, Camilo Velez, Hanqing Zhao, Patrick Traynor, and David Arnold. The cards aren't alright: Detecting counterfeit gift cards using encoding jitter. In *Security & Privacy Symposium*, 2018.
- [38] Kyle Soska and Nicolas Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *USENIX Security*, 2015.
- [39] Brett Stone-Gross, Ryan Abman, Richard A. Kemmerer, Christopher Kruegel, and Douglas G. Steigerwald. The underground economy of fake antivirus software. In *WEIS*, 2011.
- [40] Brett Stone-Gross, Thorsten Holz, Gianluca Stringhini, and Giovanni Vigna. The underground economy of spam: A botmaster's perspective of coordinating large-scale spam campaigns. In *LEET*, 2011.
- [41] Kurt Thomas, Danny Huang, David Wang, Elie Bursztein, Chris Grier, Thomas J. Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. Framing dependencies introduced by underground commoditization. In *WEIS*, 2015.
- [42] Rolf Van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Gañán, Bram Klievink, Nicolas Christin, and Michel Van Eeten. Plug and prey? Measuring the commoditization of cybercrime via online anonymous markets. In *USENIX Security*, 2018.
- [43] Josephine Wolff. Why is the U.S. determined to have the least-secure credit cards in the world? <https://www.theatlantic.com/business/archive/2016/03/us-determined-to-have-the-least-secure-credit-cards-in-the-world/473199/>, 2016.
- [44] M. Yip, N. Shadbolt, and C. Webber. Structural analysis of online criminal social networks. In *ISI*, 2012.

A Appendix

In this appendix, we present additional statistics computed from the carder shop's database.

Sellers. The largest seller of magnetic stripe accounts had a commission 25 percentage points higher than the average of all other magnetic stripe sellers. The largest releases contained 2.5 M magnetic stripe, and 8,611 CNP accounts. The smallest release for magnetic stripe was only 1 account, whereas for CNP it was 14.

Buyers. Magnetic stripe buyers tended to spend more than CNP buyers, as shown in Figure 6. The median net spend of magnetic stripe-only buyers was \$89.35 for 7 non-refunded purchases, as opposed to \$34.25 for 3 non-refunded purchases in the case of CNP-only buyers. Buyers active in both domains exhibited the highest median net spend of \$333.67 on 28 non-refunded purchases. Among them was the overall highest spending buyer, who, over the course of 3 years and 76 days, bought 16.2k magnetic stripe and CNP accounts (plus 752 that were refunded) for a net total of \$495k after refunds.

Pricing. The shop database contained 1.3k accounts (0.02% of sales) with a purchase price higher than the initial listing price.

Supply and Demand. When the shop opened in the first half of 2015, sellers added between 9,255 and 41,005 unique accounts to the shop's inventory each week, as shown in Figure 4. Nearly all supply was in U.S. magnetic stripe accounts. Supply increased drastically from November 2015 to February 2016, staying above 70,000 unique accounts per week for fifteen weeks and reaching a maximum of 123,929, before returning to the prior rhythm for another seven months. Overall, there was an average of 38,800 accounts added per week in 2015. Weekly supply increased each year, with 48,400 accounts added weekly in 2016, 64,600 in 2017, and 93,600 in 2018 and 2019.

Refunds. For refundable magnetic stripe and CNP accounts, customers checked 57.8% and 100% of their purchases, respectively. Actual refunds of refundable magnetic stripe purchases increased slightly from 31% in 2015 to 38% in 2018. The overall refund rates were 25% and 13% of sold accounts, for magnetic stripe and CNP accounts, respectively. The lower rate for CNP might partially be due to a higher fraction of non-refundable sales. On a per-buyer basis, the median

refund rates were 32% of magnetic stripe accounts, and 0% of CNP (averages: 31% and 11%). High refund rates had negative consequences for buyers, as they gradually reduced the time window during which purchases were eligible for refund. Customers with over 40% refunds lost eligibility for refunds.




Figure 12: Weekly median price of purchased accounts. The shaded area represents the range between the 25th and 75th percentiles. Median purchase prices for CNP were initially lower than for magnetic stripe accounts but increased over time, whereas magnetic stripe prices decreased.




Figure 13: Box plot of time remaining until the expiration date when U.S. magnetic stripe accounts were added to the shop. From 2016 to 2018, the median remaining lifetime of non-EMV accounts increased by 110 days, suggesting that new non-EMV accounts were issued after the liability shift.




Figure 14: Weekly availability of U.S. non-EMV and EMV magnetic stripe accounts in the shop, computed as the number of accounts added minus those that were purchased or expired. The shop accumulated an oversupply of EMV accounts.




Figure 15: Weekly median price of purchased non-EMV and EMV magnetic stripe accounts from the U.S. The shaded area represents the range between the 25th and 75th percentiles. Before 2016, EMV accounts tended to be more expensive than non-EMV accounts, but the trend later inverted.