Stripe Snapshot

Online Fraud Trends and Behavior

December 2017

Introduction

In 2016, an estimated 1.61 billion people worldwide purchased goods online, and global e-retail sales amounted to $1.9 trillion. And recent projections show a growth of up to $4.06 trillion spent online by 2020, with traffic increasingly coming from mobile devices.

Unfortunately, online shopping isn't the only thing on the rise: As chip-enabled cards have made brick-and-mortar shopping safer, fraudsters are increasingly targeting online stores—and these internet businesses are responsible for not only detecting fraud, but also paying the associated costs. **On average**, every $1 of fraudulent orders costs an online store an additional $2.62 and a mobile store $3.34.

What can online businesses do about it? Stripe looked across several years' worth of fraud data to seek out patterns by country, time-of-day and other behaviors to help guide businesses' approaches to combatting fraud. While fraud can certainly be managed by setting sophisticated payment rules or deploying anti-fraud software, our hope is that this data will also help businesses better understand the underpinnings of fraudulent behavior to create specific strategies best suited for their businesses.

Fraud rates by country

Stripe data indicate that fraud rates in some card-issuing countries can be 2-3 times higher than fraud rates in other countries. Card purchases from consumers in Argentina, Brazil, India, Malaysia, Mexico, and Turkey are particularly fraudulent, although U.S., Canadian, and French cards are also susceptible. And it's important to note that these rates are still within a small percentage of overall shopping volume—so merchants should obviously be cautious about blocking legitimate transactions. Potential solutions to tackling country-specific fraud are testing geography-based rules, or seeking more information (CVV numbers, full addresses, etc.) from purchases on cards from those countries.

2016 Global Fraud Rates by Share of Transactions

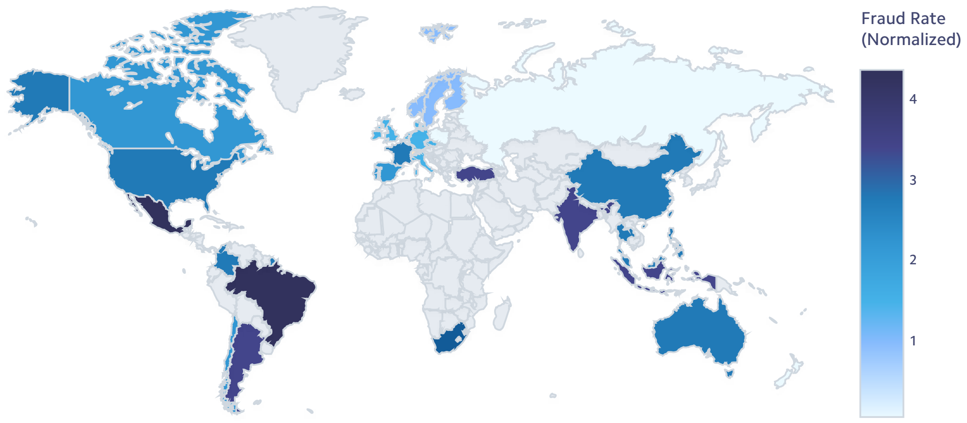

Notes: Some countries excluded based on lower transaction volume.

Fraud by day and time

Not surprisingly, Stripe found that fraud rates increase during the holidays and summer back-to-school season—but with some unusual characteristics. For example, fraud rates do not rise notably on heavy shopping days like Black Friday, but rather on days like Christmas when many people aren't shopping, yet fraudsters continue to operate. Stripe also found that fraud rates are lower for recurring payments, most likely because these transactions involve a subscription service or extended relationship that has been verified by businesses over time.

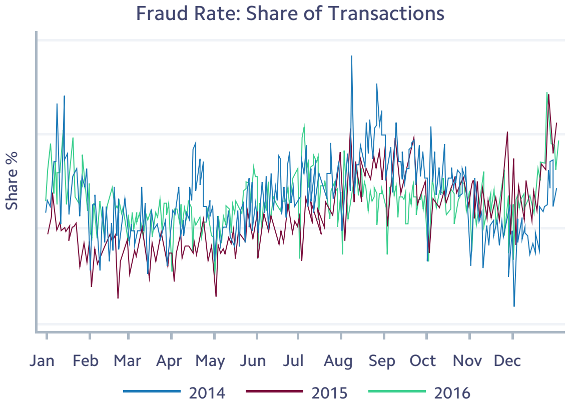

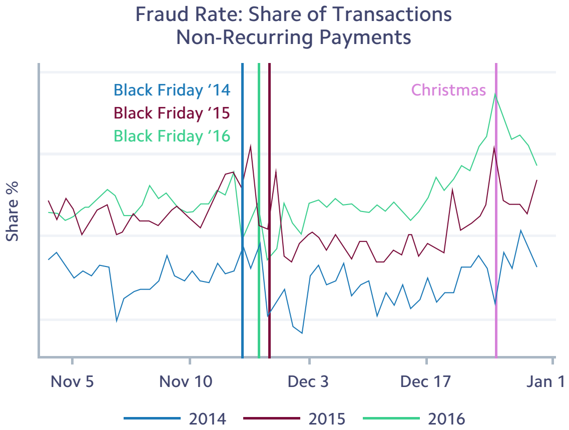

The finding that fraud rates increase during "quiet times" also applies when looking at time of day. Normalizing to local time zones in each country, we see that traffic peaks during workday hours and plummets during nocturnal hours. However, fraud rates follow a stark reversed pattern, peaking late at night and flattening out during the day. This likely reflects the wider geographic dispersion of fraudsters, who may well be operating remotely around the world, conducting their activities during hours when businesses' regular customers are asleep.

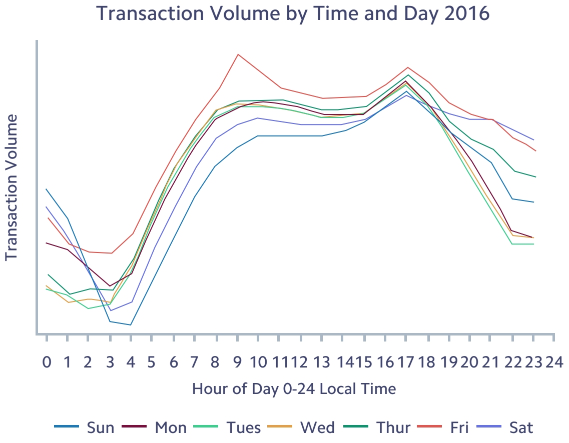

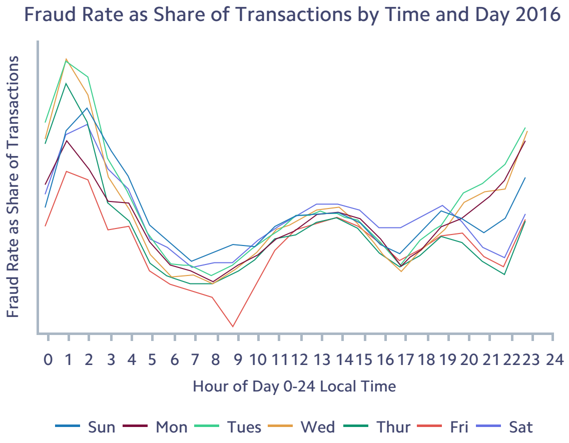

How should businesses think about battling this day/time phenomenon? Solutions might involve adding extra scrutiny for transactions outside of normal business hours, either through manual reviews or more stringent filters. This could help prevent fraudsters from striking while businesses are on their summer vacations or sitting down for their Christmas dinners.

Fraudster behavior vs. normal behavior

One interesting fact about fraudulent transactions is that they are often small. This is surprising given that fraudsters are not paying for the products they buy. In the United States, Stripe data shows that fraudulent transaction amounts are only slightly larger than regular transaction amounts. But in many other countries, fraudulent transactions are significantly larger than normal transactions—typically about twice as large, and in some countries more than five or even ten times as large.

Note: In all countries, fraud is a very small percentage of overall traffic; the figures below are normalized for comparison purposes; full-year 2016 data only.

Distribution of Transaction Amount in Australia

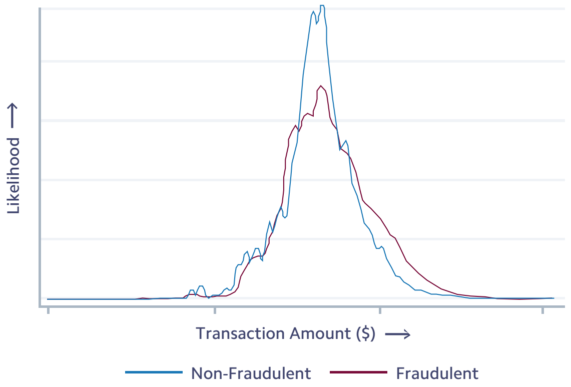

Distribution of Transaction Amount in Finland

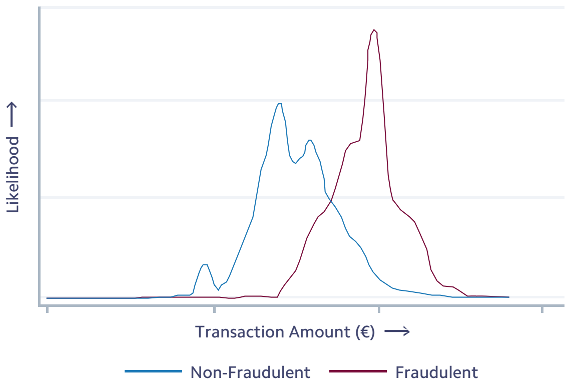

Distribution of Transaction Amount in France

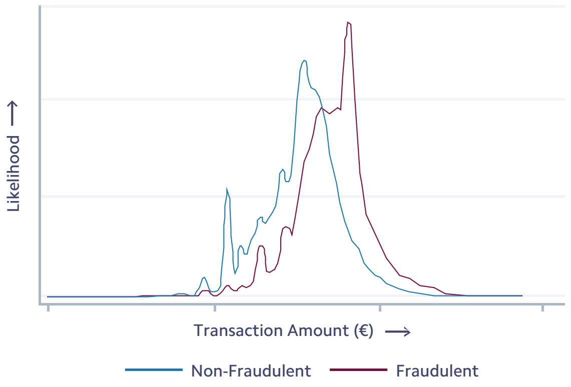

Distribution of Transaction Amount in Germany

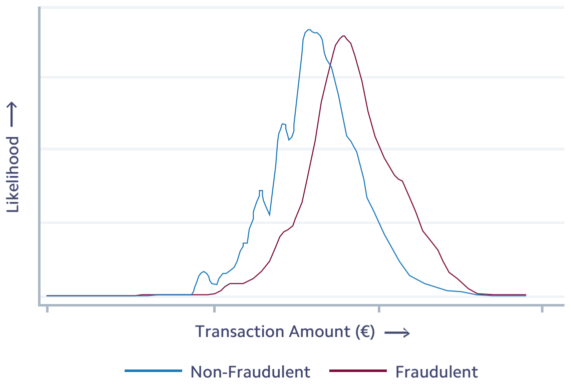

Distribution of Transaction Amount in Italy

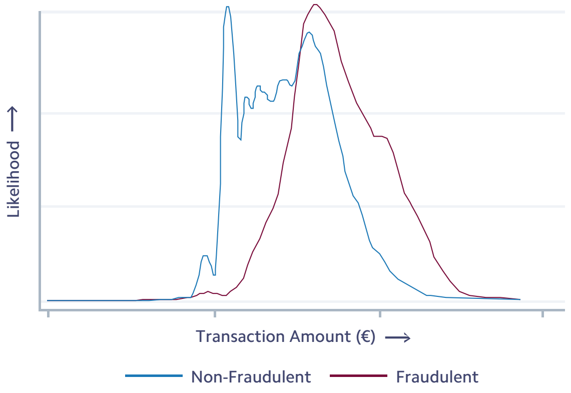

Distribution of Transaction Amount in Japan

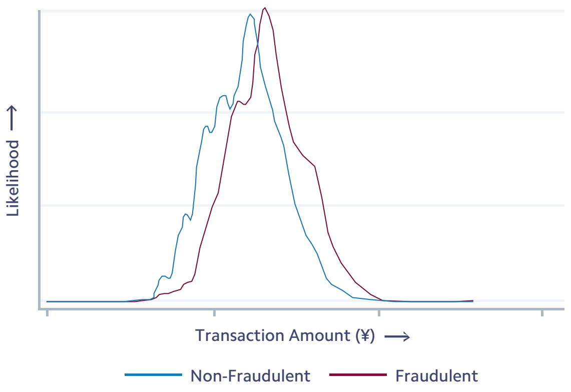

Distribution of Transaction Amount in Netherlands

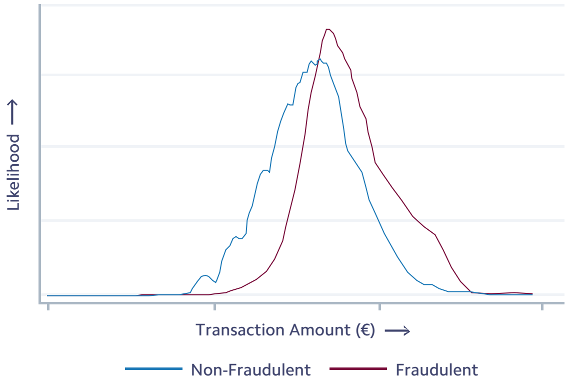

Distribution of Transaction Amount in Norway

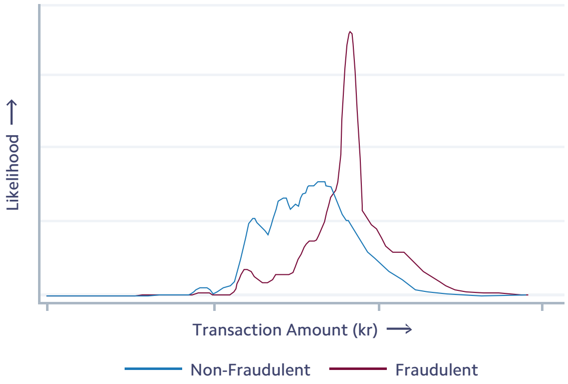

Distribution of Transaction Amount in Singapore

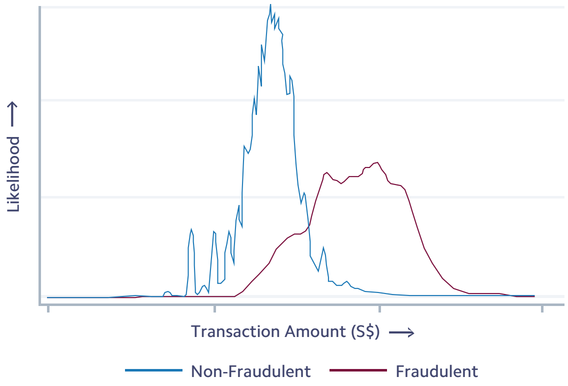

Distribution of Transaction Amount in Spain

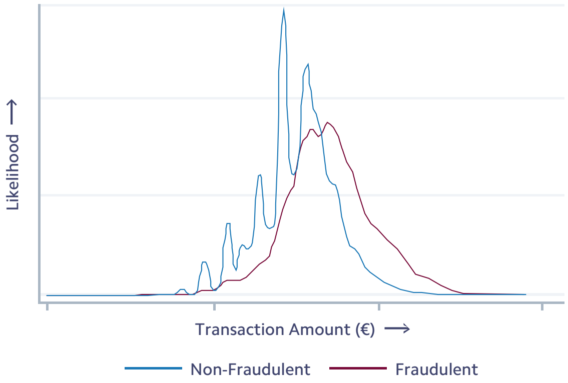

Distribution of Transaction Amount in Sweden

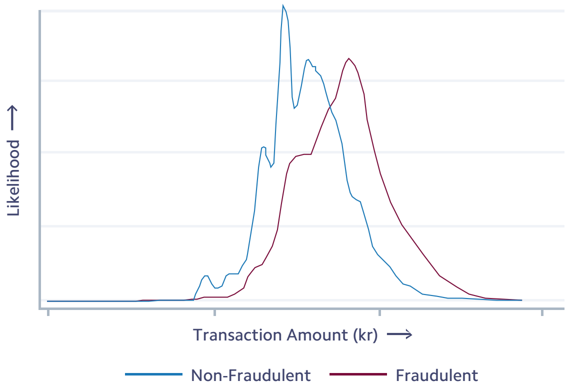

Distribution of Transaction Amount in U.K.

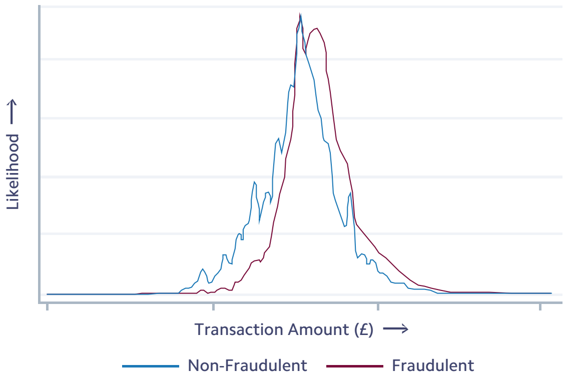

Distribution of Transaction Amount in United States

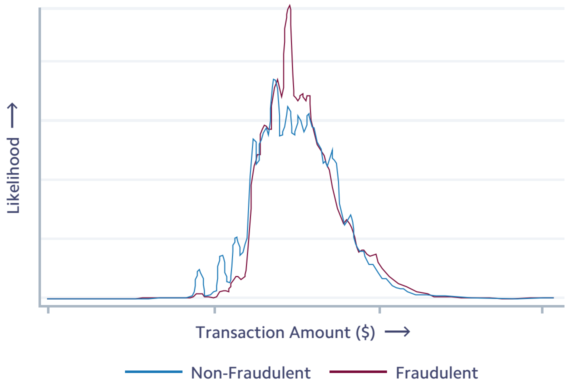

Fraudsters exhibit a more revealing signature when it comes to where and how often they shop, especially in repeat purchases on the same stolen card. Unfortunately, repeat fraud on a card is common: more than 40% of compromised cards are charged for more than one fraudulent transaction.

First, fraudsters shop repeatedly at the same merchant, rather than buying products from multiple businesses. For example, cards that have made four fraudulent charges typically have made all these charges at the same identical merchant, while four normal charges tend to be spread out across an average of two distinct businesses. Second, fraudsters make these repeat purchases much more quickly than normal transactions. In fact, these consecutive charges happen *ten times* more quickly than actual cardholders.

This conspicuous pattern of "high-velocity" purchases with a single merchant contradicts the notion of fraudsters trying to blend in with a cardholder's normal transaction patterns. One lesson here for businesses is to be cautious with many rapid-fire transactions from the same credit card—though, as always, it is important not to block good transactions with blanket rules.

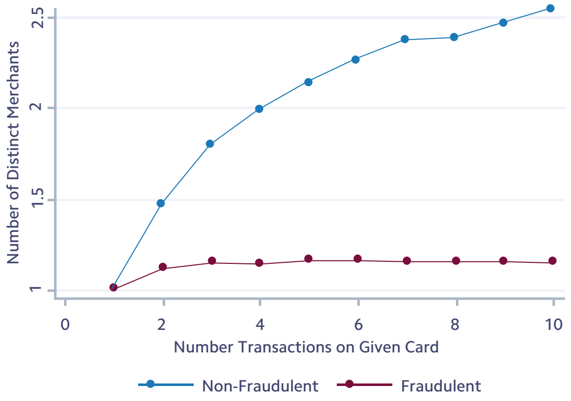
Buying from the Same Merchant 2016

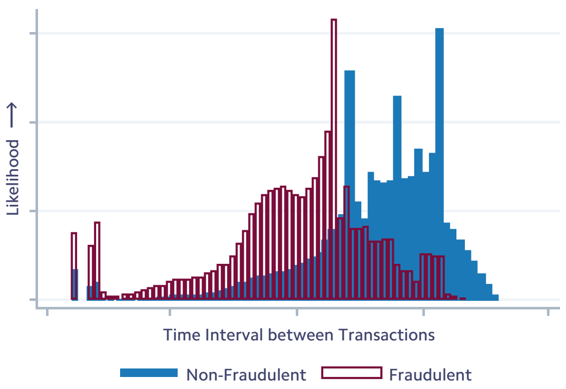
Distribution of Time Interval Between Transactions 2016

Fraudster behavior by business

Stripe data also reveals that fraudsters target certain kinds of businesses, and patterns here suggest that fear of arrest—not a desire to blend in with normal transactions on a cardholder's billing statement—offers a more compelling guide to behavior. A key challenge for online fraudsters is delivery. Delivery of physical goods to a home or office associated with the fraudster, or anyone in his or her social network, carries obvious risks.

To avoid these issues, fraudsters tend to buy products that do not need to be delivered, or products that are routinely delivered to places not linked to the buyer. Many services, as opposed to physical goods, do not require delivery. But the most appealing services are performed *immediately* before the charge has any chance of being detected and invalidated. One implication is that on-demand services pose an attractive target for fraud. Low-end consumer goods also stand out, perhaps because fraudsters assume authorities deprioritize low-stakes theft.

Conclusion

As consumer behavior and fraud schemes continue to evolve, businesses should be aware that these global trends scratch the surface of big data patterns that can be used to detect fraud. Effective fraud prevention takes into account the specific context of the business. Machine-learning models address this challenge by incorporating many context-specific nuances in order to reject only the most suspicious transactions, rather than putting in place blanket rules that can easily wind up blocking good transactions. Merchants should work with payment processors with machine learning and other technologies to optimize these complex tradeoffs between stopping fraud and maximizing profitability.

Methodology

In conducting the analysis in the report, Stripe examined transaction data across hundreds of thousands of its customers in 25 countries. The heat map on Global Fraud Rates by Share of Transactions (page 2) includes 2016 data only and excludes those countries with lower transaction volumes. The graph of fraud by day of the month (page 3) is aggregated data for 2014-2016 across all countries where Stripe operates, normalized by working hours for each timezone. The Fraud Rate as Share of Transactions by Time and Day (page 3, righthand figure) graph includes 2016 data only and is normalized for comparison purposes with Transaction Volume by Time and Day (page 3, lefthand figure). The country histograms (pages 4-5) include 2016 data only, and non-fraudulent and fraudulent distributions were normalized for comparison purposes. Finally, the Buying from the Same Merchant and Distribution of Time Interval Between Transactions figures (page 6) includes 2016 data only.