Mitch Muroff

CLEAN Getaway

Responding to a highly sophisticated fraud attack

Most major merchants that accept payment cards from consumers deploy controls to manage payment acceptance fraud, particularly online. These controls typically take the form of rules or scores that flag or reject transactions with characteristics deemed risky. However, these rules and scoring factors tend to be standard across merchants and easy for criminals to anticipate.

The method that criminals use to evade these controls is simple: Learn the criteria merchants use to flag or reject transactions and then design fraudulent transactions that don't look like the bad transactions that rules and scoring models are designed to stop. Fully 46 percent of merchants participating in the Cybersource 2012 Online Fraud Report said fraudulent orders are getting "cleaner"—meaning they look more like good orders.

The following case study describes how we helped a finance executive at a Fortune 100 company, whose name cannot be disclosed due to confidentiality agreements, respond to an exceptionally sophisticated CNP fraud attack using advanced analytics.

The approach

The merchant had an established online presence and followed best practices for managing e-commerce payment fraud acceptance risk. All of their transactions passed through a fraud screening system for review. That system—and the rules embedded in it—was configured and maintained by a leading fraud solution provider. Furthermore, the company had a team of analysts responsible for reviewing suspicious orders, monitoring overall performance and investigating emerging fraud patterns.

Despite doing all the right things, this merchant's fraud rules and manual review processes were not able to detect a lot of the fraudulent transactions being presented, and therefore, the merchant's fraud rate—and subsequent fraud losses—were unacceptably high. To tackle this problem, the first step was to review the rules, procedures, and policies in place to ensure that nothing obvious was being missed. The rules in place were found to be reasonable and appropriate given the available data. The second step was to review the data for clues that might be useful in understanding the characteristics of the fraud and helpful in developing rules that could more effectively flag fraudulent transactions.

In more than 20 years of experience, I had never before seen fraud data that looked this clean.

While reviewing the data, I was surprised to observe that it did not contain typical indicators of fraud. There were few indications of high velocity purchase behavior or orders from high-risk countries. Names and addresses appeared to be valid, the order sizes were not unusually high, the addresses and security codes matched those on the credit cards and so forth. There was nothing obvious on first glance that appeared different when comparing good orders with fraudulent orders. In more than 20 years of experience, I had never before seen fraud data that looked this clean.

It quickly became apparent that solving this problem would require both an atypically rigorous approach and a deeper and broader data set than was being used by the current fraud screening system.

The screening

To analyze the data and find a strategy for reducing the fraud rate, we worked with the practitioner to obtain the data that was used by its fraud screening system and additional data that was not incorporated into his company's fraud screening process. This additional data request included information about account history and certain behavioral characteristics.

The data set used for fraud screening included the standard data elements collected on a typical checkout page: name, address, product ordered and IP address. The fraud vendor provided additional information about the IP address, including country, region and connection type.

The first step in solving this problem was to determine what variables might be most helpful to analyze. While it was obvious to use the data values provided by the company, additional variables were created by parsing those data elements, examining relationships between them, retrieving more information from third-party data sources and placing observations into categories that were more suited to analysis.

Several hundred additional variables were created to derive more meaningful insights that could help distinguish between good and bad orders, particularly when combined with existing rules.

Each variable was then tested to determine the strength and consistency of its relationship with both good and fraudulent orders. By studying each variable across time, we discovered that the company was not suffering from a single fraud attack but, in fact, had experienced several unique fraud attacks. In other words, against a backdrop of consistent fraud, there were some fraud attacks that had different characteristics from the general level of fraud. Since those attacks had abated and the objective was to achieve the fastest possible reduction in current fraud, the approach was refined to remove from further consideration any factors that would serve only to identify the characteristics of the specific historical attacks. Variables were selected for inclusion when they consistently related to the current fraud incident, with fairly stable characteristics over time.

With this foundational work in place, my firm evaluated all potential fraud-reduction strategies using the variables and values that survived the first cuts. For each potential strategy, we measured both the fraud reduced and the good orders that would have been flagged for review or rejected if the strategy were chosen. Then we selected the strategy that yielded the target fraud reduction at the lowest cost.

The results

Working with this simple data set, a solution was found to cut fraud by 25 percent by combining observations in specific ways that might not have been obvious using other methods. In particular, a very specific relationship was discovered between sets of locations where criminals were operating from, the products they were purchasing, and when those purchases were made that—combined with other account-level observations—allowed the company to distinguish between good and bad orders more effectively than before, using exactly the same data.

Although many strategies can achieve a desired fraud reduction, the costs of deploying these strategies can vary widely. We found more than 15 strategies that could each achieve the target reduction—using different sets of fraud rules. While all 15 solutions achieved the same fraud reduction, the delta between the lowest cost and highest cost option was 85 percent, meaning 85 percent more good orders would be flagged for review or rejected if the least favorable solution

were chosen, of the top 15 solutions, compared with the most favorable. Our ability to identify the most profitable strategy contrasts starkly against a common practice of developing rules based only on their ability to reduce fraud, without considering impact on good orders. This analysis shows that without measuring and optimizing for impact on good customers when developing rules to stop fraud, a merchant risks adopting a strategy that unnecessarily turns away millions of dollars in good sales, possibly without even recognizing it. This, ultimately, is the justification for doing the work to generate and rank a wide range of alternative strategies for responding to a fraud attack.

When one additional variable was added, using an internal data source from the company that wasn't included in the data set used for routine fraud screening, a further 39 percent cost-reduction occurred, meaning that if $0.50 worth of orders needed to be rejected or reviewed for each $1 of fraud reduced, the same fraud reduction could be achieved by reviewing and/or rejecting only $0.31 in good orders, simply by adding this 1 additional variable.

This experience shows that performing an analysis of a wider range of variables, and doing it in more depth than is typical can not only allow merchants to effectively combat increasingly clean fraud, but it can actually reduce the cost of managing fraud overall.

Three key principals

As fraud continues to become cleaner, there are three key principles that will drive successful fraud-mitigation strategies:

>>> Examples: Scoring factors to monitor

Address or security code does not match card.

High risk product.

Size of purchase.

Number of purchases in given period of time.

Age of account.

High risk IP country.

Distance between IP country, address, card issuer.

Source: Cybersource 2012 Fraud Report

- Merchants need to dig deeper into the transaction to gather information that describes behavior before and after the charge transaction and incorporate that information into their models. Currently, many merchants collect a very thin band of data that describes what happened at a single moment, when the checkout button was pressed. Merchants increasingly need to tie that data with information about behavior before the checkout page to develop a profile that allows them to effectively distinguish between good and bad orders because the limited data available at time of checkout is so effectively being manipulated by criminals to look like good orders. For example, where did the users come from? What did the users do and how did they behave before they hit the checkout page? What is known about prior transactions conducted by the users?
- Cleaner data increases the need to augment internal data with external data. External services can provide important information about the characteristics of key data elements used in fraud screening. Data elements for which third parties can provide additional data include: name, address, IP address, email address, phone number, payment card issuer, and the hardware used for the transaction. The goal is to learn more about each data element provided with a transaction and how those data elements relate to each other, then correlate those more complex transaction characteristics with fraud to build better rules and models that are harder for criminals to circumvent.
- Advanced analytical techniques are necessary to find powerful strategies. Firms need to study more data, in more creative ways, than is typically done, to better understand fraud characteristics so they can build more complex rules that are harder for criminals to evade. And, they need to ensure that those rules don't turn away too many good orders or cause an excessive volume of orders to be reviewed manually, since manual review is an expensive variable cost.

Mitch Muroff is president of Curaxian.