Managing Credit Risk
Through Embedded Intelligence in On-line Transaction Processing:
First Union National Bank, Charlotte, NC[1]

THE ORGANIZATION

First Union National Bank (FUNB) is the nation's sixth largest banking company with $9.3 billion in total stockholders' equity and $17.2 billion in market capitalization. FUNB offers a diverse array of products such as 401(k) plans, checking and savings accounts tailored to customer needs, investment banking, certificates of deposit, mutual funds, credit cards and other loan products. FUNB has presence in 12 eastern states and Washington, DC It recently acquired of First Fidelity Bancorporation on January 1, 1996, giving it a customer base of about 12 million customers from Connecticut to Florida.

On June 30, 1996, FUNB company had assets of $139.9 billion. The bank's network has 1,981 offices -- the nation's largest branch banking system. It also has the nation's fifth largest automated teller machine network. In addition, First Union is pioneering one of the first direct banks on the Internet.

Over the years, FUNB has been working towards developing a leading position in several markets, including deposits and credit cards, on the assumption that scale is an essential element of keeping unit costs low. First Union has the leading deposit share in its home state of North Carolina; it ranks second in Florida, New Jersey and the Washington, DC, and third in Connecticut. The company serves eight of the 10 richest counties in the US.

Since 1993 FUNB has nearly tripled the size of its credit card portfolio. This growth was achieved through targeted, national market solicitations aimed at providing geographic diversity and by trying to attract high quality, revolving credit customers. It was the 14th largest issuer of credit cards in the nation as of June 30, 1996. FUNB also offers on-line credit card applications over the Internet.

This case focuses on the credit card business of FUNB.

THE PROBLEM

FUNB's problem was simple: to minimize the dollar loss incurred as a result of credit card fraud. Why is this a difficult problem?

Traditionally, credit card issuers have been concerned about losses due to defaults -- borrowers not paying back their charges, but in the last few years, a new kind of loss has also arisen. Financial fraud is

[1] This case was prepared by Vasant Dhar for purposes of class discussion and not to illustrate appropriate or inappropriate handling of an administrative situation.

not a petty theft situation but a high stakes business that "employs" a significant number of individuals. It is organized crime that exploits several aspects of electronic commerce:

- a side effect of the "electronization" of commerce is that more and more information about individuals is publicly available without people even being aware of this
- individuals want as little red tape as possible; they don't want to be questioned about the legitimacy of every transaction
- control systems are never perfect. There are invariably "holes" in most systems that criminals try and find become they are plugged

Because of the large amounts of money involved, organized crime gets into the game. And this means sophisticated methods of theft. Criminals are constantly looking for innovative ways to exploit availability of credit.

To prevent losses, issuers engage in credit screening. Screening occurs at a number of levels. Issuers routinely have a "pre booking" process involving the screening of applicants, determining which ones are "safe" enough for credit cards to be issued.

Screening at the transaction level is a lot more difficult. Transaction level fraud occurs when a legitimate credit card or account is illegally used or taken over by a criminal. It must be detected in real time. This is the problem FUNB decided to focus on.

Why is detecting a fraudulent credit transaction a difficult problem?

First, credit is an easy medium for criminals to exploit because of the difficulty of discriminating between a legitimate and a fraudulent transaction. Criminals are capable of getting a lot of information about individuals such as social security numbers and other key pieces of data without much difficulty. When they take over a credit card or an account, this information provides them easy access to high priced goods and services such as jewelry, airline tickets, etc.

Once criminals know that they have an exploitable credit card, they first "test" it by engaging in seemingly "normal" activity, and then hit the card hard and fast, extracting as much of its available credit line as quickly as possible. Often, this behavior is not much different from that of the legitimate holder of the card. In other words, detecting fraud is a subtle and complex: there are no obvious patterns that stand out as pathological behavior.

Second, banks must minimize the risk of denying legitimate transactions, also referred to as the "type II cs. Embarrassment and inconvenience are a sure way of losing customers: you have to be virtually certain before intervening. This increases the likelihood of approving bad transactions, or the "type I error".

Third, customers are not always available to verify the legitimacy of a transaction even when the bank deems it important enough to do so. People are often not available by phone, sometimes for days.

Besides, criminals may also pose as the customer, especially when they are armed with sensitive customer information. Criminals have been known to completely "take over" a bank account, sometimes changing the billing address to their own without the real customer being aware of this situation!

Until a few years ago, banks were largely unable to combat the increasing fraud. The process of trying to unearth fraud was woefully inadequate. Mary Ann Miller, manager of fraud prevention in the Customer Direct Access Division at First Union described the problem as follows:

> "In 1991, a lot of your information came from Issuers Clearinghouse Reports. These had information such as social security numbers, phone numbers and addresses of known problematic cards in the last 30, 60, and 90 days. This information is useful for screening out transactions coming from cards that are known to be bad -- it is a useful pre-booking fraud tool. However, the information is not relevant in finding fraud that is under way *now*. To deal with transaction level fraud, we used to specify some crude "filters" that consisted of broad parameters such as "cash transactions over $1000", which an on-line system would use to flag transactions. But these kept our printers really busy. We would get reams of output. These reams would be taken over to analysts who would spend hours or days analyzing them, trying to judge which ones might be fraudulent. In the early 90s, days could go by without finding a fraudulent transaction. As a result, we devoted more analyst effort, but it was just that: more effort, a more labor intensive process that did not scale well. We needed a system to do this, a system whose costs could be amortized over a large number of transactions."

Figure 1 shows a general schematic of the situation. The problem was there was no good *model* of a fraudulent profile. Without such a model, the bank could at best put transactions through crude filters, which represented their best guesses the kinds of things you might see in risky transactions.

The trouble was that these filters were much too loose: while many fraudulent transactions might correspond to them, there were *a lot more* legitimate ones that also fit them. Without any way to discriminate, analysts were deluged with a lot of irrelevant data. They were required to exercise *too much* judgment, and their performance depended largely on their experience, expertise, motivation, and luck. The bottom line was that analysts simply could not handle the volume of output that flowed into them. The odds of catching fraudulent transactions were extremely low, reflected by the fact that they sometimes went for days without catching a single bad transaction. Clearly, this was not a scaleable approach to the problem.

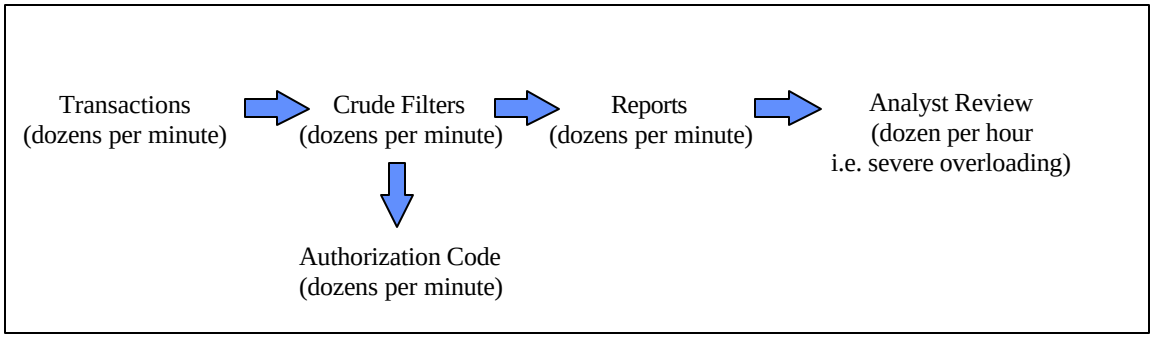

Figure 1

In a sense, the situation called for developing better filters. But not only did they have to be better, they also had to be *dynamic*, since the profile of bad transactions can keep changing: the bank might detect a pattern, criminals realize that it no longer works, and alter their strategy. By definition, the bank is always playing catch up. They could not take the approach of developing "the correct" filters. Rather, the filters needed to be flexible enough to adapt to new kinds of fraudulent activity, while "not
In a sense, they had to be capable of evolving in tandem with fraudulent behavior!

By the same token, the problem required that a solution scale well across different customers: the bank's customer base would keep changing, requiring that the system scale to handle different types and larger numbers of customers. For example, if FUNB decided to develop an system for handling each "profile" of customer (i.e. "urban professional"), they would need to be able to extend the system easily to handle new profiles.

While a system had to be flexible enough to handle the evolving subtleties of criminal behavior, FUNB managers felt that their approval process should also be able to express and test certain "profiles" that they might hypothesize as being potentially problematic and hence worth checking out statistically.

For example, if experts felt that things like "three month payment history trend" or "number of address changes in the last six months" might indicate something about consumer behavior (such as propensity to default), the system should be able to track this information and allow an analyst to query the existing database with such a filter. This would require a system to compute and store such statistics when directed by an analyst.

FUNB also realized that it would never be able to eliminate fraud completely either in principle or practically. The only way to do this would be to call customers to approve each transaction. Customers would not tolerate this. Practically speaking, it was important to determine when a transaction was deemed sever enough to get approval from the customer. The goal, therefore, was not have the most accurate system, but one that would keep fraud to "reasonable limits". Furthermore, it was important to limit the calls to the customer. Figure 2 shows the "desired" process envisioned by FUNB as the first step. The idea was that each transaction would be scored to reflect the probability of it being a bad transaction. The idea was that each transaction would be scored to reflect the probability of it being a bad transaction. It would then be passed on to a "Case Management System" that would decide

whether to put it into a “queue” for further analysis by an analyst. The outcome of this analysis would in turn affect the score assigned to the next transaction.

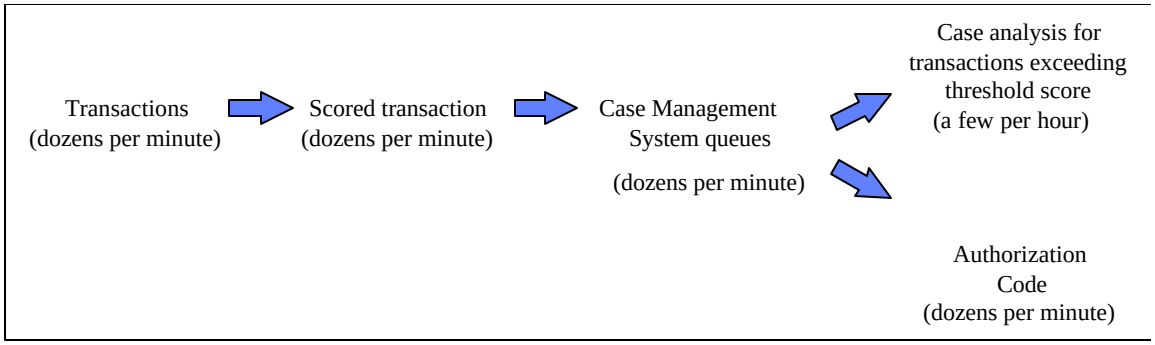

Figure 2

Finally, FUNB also realized that their data would always contain a certain number of errors. Some of these would be in the demographic data. More importantly, perhaps, there would always be some “bad” transactions that would have been classified as good, i.e. not detected as bad. Some of these might never be rectified.