DRIVING IMPROVEMENT WITH OPERATIONAL RESEARCH AND DECISION ANALYTICS

IMPACT

SPRING 2015

Image © Rufus Stone / Alamy

REPLACING THE LONG HAUL FLEET

When British Airways needed a major fleet replacement, an O.R. model utilised hugely complex data to give an informed decision

ANALYTICS AND YOUR CREDIT CARD

Every decision a bank makes about your credit card is guided by analytics

WORKFORCE PLANNING IN THE NHS

System dynamics helps determine our need for doctors and dentists, and deliver substantial savings

THE OR SOCIETY

A DAY IN THE ANALYTIC LIFE OF A CREDIT CARD

A $13 Trillion Success Story for Operational Research

ANDREW JENNINGS

CREDIT CARDS TODAY seem so blandly utilitarian that most of us give them little thought — other than, of course, whether that next purchase will put us over our credit limit, or whether we've paid the bill!

We certainly don't credit the humble credit card as an analytic product. But in fact, that's exactly what it is.

Every decision a bank makes about your credit card — and there are more decisions than you may imagine — is guided by analytics. Every transaction you make with your card is monitored and enabled by analytics. As a success story in operational research, the credit card is right up there with airline bookings and supply chain optimization.

This is more than an analytics bragging story, though. Easier credit has helped fuel economies worldwide. And much of banks' operations are funded by credit card profits. In fact, about half of the net income from consumer banking at such global banks as Citi and Chase comes from credit cards, and while the ratios vary, cards are huge profit drivers for most major banks.

Cards' importance is growing, not diminishing. Card spending was estimated at $13 trillion worldwide in 2012, and is expected to hit $27 trillion by 2018, fuelled largely by growth in Asia Pacific.

We all know what makes cards so convenient – but what makes them so profitable? This is the story of analytics at work in the day of a life of a credit card.

AN ANALYTIC MELANGE

First, it's useful to note that there are three broad categories of analytics that drive credit cards' success. These move along the analytic continuum from understanding and reporting to action, to answer multiple questions that can improve performance.

These analytics are in turn applied across what we think of as the cardholder lifecycle, from card marketing to origination of new accounts to customer management decisions to fraud and collections. We'll look at some of these decisions shortly.

DESCRIPTIVE ANALYTICS (E.G., BUSINESS INTELLIGENCE, DEMOGRAPHIC CLUSTERING)	PREDICTIVE ANALYTICS (E.G., CREDIT SCORING, FRAUD MODELLING)	PRESCRIPTIVE ANALYTICS (E.G., DECISION MODELLING, OPTIMIZATION)
What is my delinquency rate by vintage? What is the average risk of new business? Are the usage trends stable?	What is the risk of this person? Will this person carry forward some of their balance (revolver) or pay it off each month (transactor)? Is this transaction fraudulent?	On what terms do I accept this application? What credit limit do I assign? What action will result in a payment of an overdue balance?

THE SCIENCE OF SAYING "YES"

The original application of analytics to credit cards began in the late 1950s, when two O.R. practitioners named Bill Fair and Earl Isaac decided to build their new business around a relatively new concept: credit scoring. Their work – originally in personal loans – led to the widespread adoption of credit scoring for all kinds of credit, from auto loans and mortgages to credit cards.

The concept is fairly simple: gather all the data known about cardholders when they applied for a card, compare it to their subsequent performance, add some reject inference, and you can determine which characteristics and attributes relate to credit risk, and how to weight them. Credit risk models are built to separate future "good" payers from "bad" payers. (The performance definition for "bad" is generally something like a payer who has ever been more than 2 cycles (60 days) delinquent.)

Originally, these credit risk models were developed on lenders' own data, and were specific to their business. In 1989, the company Fair and Isaac founded (now known as FICO) released the first FICO Score. These scores are now used in billions of credit decisions a year. While custom origination models are based on application data, FICO Scores are based solely from credit history data contained in an individual's credit bureau report. An extract from an example FICO Score model is shown in Figure 1.

CATEGORY	CHARACTERISTIC	ATTRIBUTES	POINTS
Payment History	Number of months since the most recent derogatory public record	No public record	75
		0–5	10
		6–11	15
		12–23	25
		24+	55
Outstanding Debt	Average balance on revolving trades	No revolving trades	30
		0	55
		1–99	65
		100–499	50
		500–749	40
		750–999	25
		1000 or more	15
Credit History Length	Number of months in file	Below 12	12
		12–23	35
		24–47	60
		48 or more	75
Pursuit of New Credit	Number of inquiries in last 6 months	0	70
		1	60
		2	45
		3	25
		4+	20
Credit Mix	Number of bankcard trade lines	0	15
		1	25
		2	50
		3	60
		4+	50

FIGURE 1 –SAMPLE FICO SCORE MODEL (SIMPLIFIED)

While this sample has been simplified, an actual FICO Score development involves the analysis of more than 600 candidate variables. For a typical consumer, well over 30,000 floating point operations are executed in order to return a valid FICO Score.

Scores do not predict an individual's specific credit performance or profitability. Rather, they rank-order individuals. People scoring 720 will as a whole perform better than people scoring 680, for instance. The actual ratio of good:bad accounts in any given score range, known as the odds ratio, will vary based on the economy and other factors.

KEEPING THE CREDIT MOVING

Provided you scored high enough, you now have the credit card in your hand. What happens next?

Once you start using your card, you become part of a network of analytic calculations designed to keep the payments system safe – and protect you from fraud.

Every single card transaction goes through a real-time authorization process, where information on the prospective purchase is communicated to the merchant acquirer (who has relationships with a set of merchants), a card processor (which processes transactions for the card type), a card network (Visa, MasterCard, Discover, AmEx) and the card issuer.

Focusing on the customer management decisions, one thing that's being checked during this process is whether the transaction keeps your card within its credit limit or not. A transaction that doesn't is referred to as an overlimit transaction, and your card issuer has options for handling those.

- Allow the transaction and raise your credit limit
- Allow the transaction at your current credit limit
- Allow the transaction at your current credit limit, but charge a fee for going overlimit
- Decline the transaction

How do they decide which action to take? Your credit score is involved — and now that you are a cardholder, there is another kind of score that can be used, your "behaviour score," which is based on your payment performance with that issuer. But other factors may also be considered. It's in the issuers' interest to approve the transaction if possible — you'll be happier, and they will get transaction revenue — but if you're going to default on your card payments they need to shut the spending down.

Card issuers make these decisions using strategies defined as decision trees, which are a handy way to apply several variables in order to segment a population into very specific groups for targeted decisions. In the 1980s, FICO introduced adaptive control technology

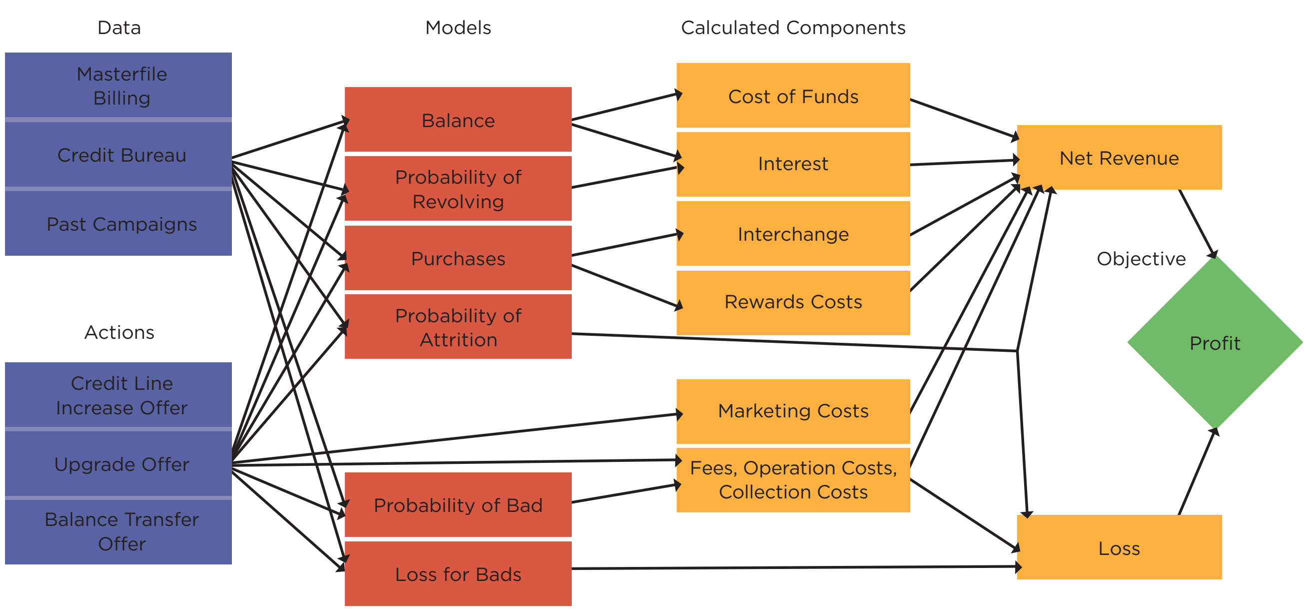

FIGURE 2 – INFLUENCE DIAGRAM FOR DECISION OPTIMIZATION

to credit card customer management. This enabled card issuers to run "champion/challenger" tests of decision strategies, changing a variable such as the cut-off score at which an overlimit transaction is approved and running the new strategy in parallel with the current "champion." This closed-loop learning process is now used almost ubiquitously in card account management, as issuers incrementally evolve their strategies to make them more profitable.

Of course, the complication in these strategies can become overwhelming. And it can be quite daunting to determine which variable to alter, and by how much, in order to make your next challenger strategy more successful. For about 15 years now, card issuers have been adopting optimization (prescriptive analytics) to derive their card account management strategies based on data. This involves influence diagrams (see Figure 2) and decision modelling to map all the components of a decision and their relationships with the target variable. These relationships can then be represented as a set of equations with an objective function and constraints which can be solved and implemented into the same operational systems that make the transaction authorisation decisions.

This in turn leads to some quite scientific analyses of how different strategies or scenarios move an issuer toward the "efficient frontier" — that min/max point where profit is maximized subject to other constraints. The example in Figure 3 illustrates how this evolution of credit strategies can produce substantially stronger profits.

PROTECTED BY NEURAL NETWORKS

If you use credit cards, you have probably received at least one frantic call from a card issuer, asking you to confirm a recent purchase was yours. You may even have imagined a cavernous room full of fraud analysts, scanning each transaction to see if it looks legit — occasionally an analyst strokes her chin and says, "Funny, I don't remember Marge buying that many cat toys before. Perhaps I'll give her a ring."

That's nearly how it works, too — except cut the number of fraud analysts down to a handful of case managers, who do the actual transaction monitoring with an ingenious analytic system. Ever since 1992, when HNC Software introduced its Falcon software, the chief card fraud defence around the world has been based on artificial intelligence.

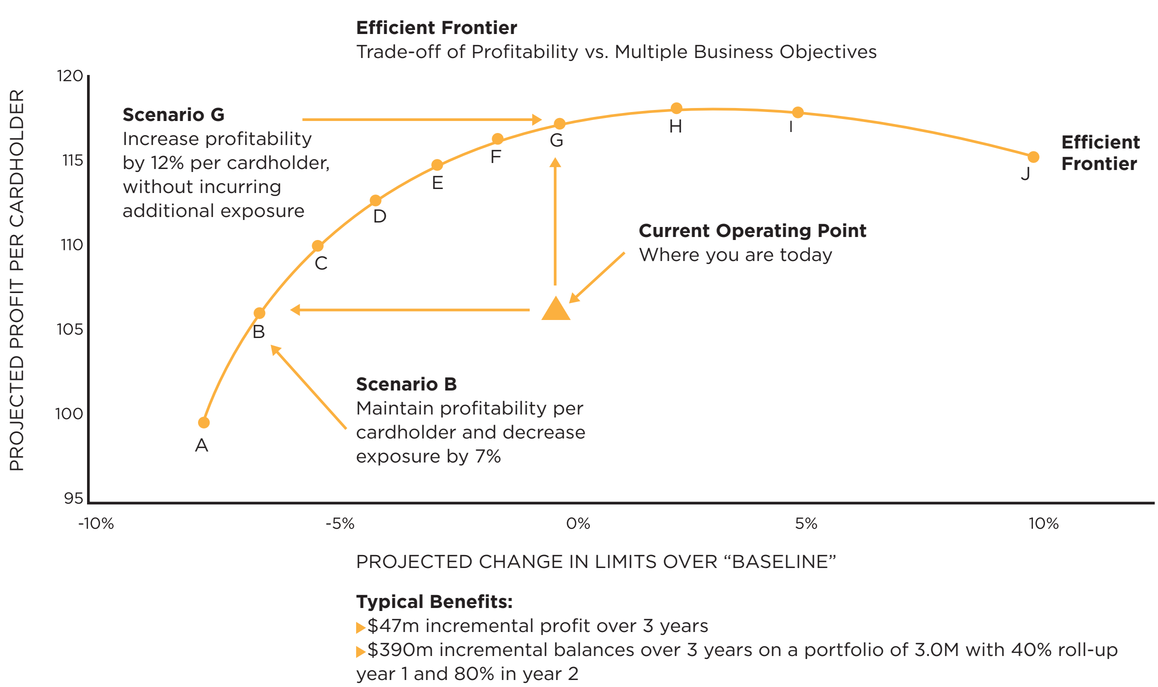

FIGURE 3 – HOW BETTER DECISIONS YIELD MORE PROFIT

Falcon introduced neural networks to financial services from military applications. Neural networks are modelled on the way the human brain works, and find connections between, for example, credit card transaction characteristics and the probability of fraud. The kinds of risk models described above are largely built by testing different combinations of data for their predictive power: transparency is important, since ultimately a lender may be required to tell a consumer why they didn't meet a score cut-off. In fraud, this kind of transparency is not as important, and neural networks are "trained" on massive datasets, and can even adapt to changes in patterns, which is critical in a realm where attackers are constantly changing their schemes to fool the system.

The models are developed in a supervised learning scenario – the input layer consists of features of transactions that are randomly weighted to start, and fed into neurons in the "hidden layer" of the network, which identifies relationships. The neural network makes predictions based on these weights, shifting the weights around to improve predictions until it reaches the best predictions. The neural nets output a score that identifies the fraud risk of a given transaction. This process is illustrated in Figure 4.

The neural networks are complemented by cardholder profiles,

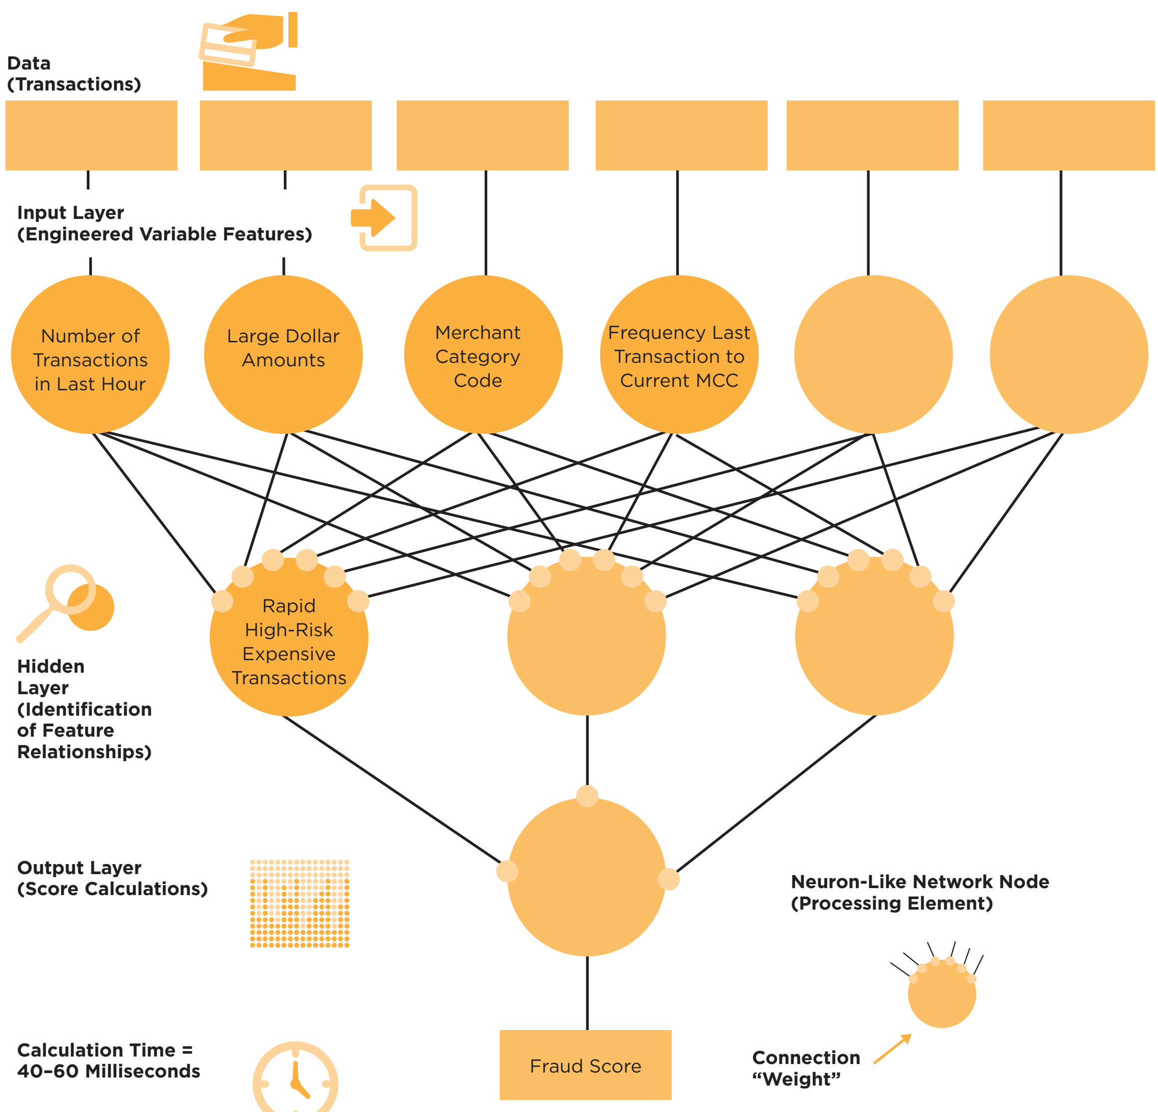

FIGURE 4 – NEURAL NETWORK FOR FRAUD DETECTION

Source: FICO

which identify normal and abnormal behaviour for a specific cardholder (not a type of cardholder). Because cardholder transaction histories can become quite massive, it is often impractical to retrieve the entire history in order to evaluate a single transaction. Transaction profiles contain recursive variables to summarize the relevant predictors.

For fraud, time is of the essence. The analysis of a transaction happens in about 40-60 milliseconds – about one-fifth the time it takes you to blink. In that time, the Falcon analytics perform 15,000 calculations.

The power of these analytics is incredible. While card fraud losses have multiplied globally over the last 20+ years, this is a function of card transaction growth. In fact, fraud losses as measured in basis points (1/100th of 1%) of card sales have declined dramatically during this period. In 1992, when Falcon was introduced, fraud losses in the U.S. stood at 18 basis points. They dropped to 8.4 by 2000, and have been below 6 since 2007. A very conservative estimate puts the savings for U.S. card issuers, based on these numbers, at more than $10 billion.

Falcon today, protecting more than 2.5 billion payment cards worldwide, has also contributed to dramatic reductions in losses in other countries. In the UK, card losses dropped from a high of £725 million in 2008 to £535 million in 2013, as the industry adopted chip and PIN technology and stepped up its use of analytic fraud detection.

The phenomenal rise of the credit card worldwide has been fuelled by analytics

Fraud protection is more than just a crackdown on criminals. Consumers want to be protected, but they also don't want to be inconvenienced by blocked transactions or unnecessary phone calls. To improve fraud detection and reduce the number of "false positives," data scientists have invented creative ways to boost performance. Some of the recent additions to the technology are:

- **Profiling of merchants and devices** to recognize patterns of fraud at specific point-of-sale terminals and ATMs.
- **Self-calibrating technology** that allows a system to fine-tune itself in real-time in response to shifts in transaction trends.
- **Adaptive analytics** that adjust the weighting in the neural networks as fraud patterns change.
- **Global intelligent profiles** that identify the riskiest ATMs, merchants and regions so extra scrutiny can be applied where the risk is greatest, without delaying the processing.
- **Behaviour-sorted lists**, which use a cardholder's most frequent transaction locations to reduce false positives.
- **Proximity correlation**, which compares the location of a cardholder's mobile phone to the place where the card transaction is occurring.
- **Behavioural archetypes**, which use Bayesian analysis to find soft clusters of cardholders with similar patterns, and can therefore determine that a rare transaction for a cardholder (say, buying a TV) is not unusual for someone in their cluster.

O.R. EVERY STEP OF THE WAY

The phenomenal rise of the credit card worldwide has been fuelled by analytics — in the UK all of this is now common practice and in APAC, the market with the greatest growth outlook, issuers are quickly adopting the latest technology to keep their profits strong and their payments safe. Every time you use your card, remember: it's not just a piece of plastic or a gateway to debt, it's a supercharged analytic engine!

Dr. Andrew Jennings is chief analytics officer at FICO, a leading analytic software company. He blogs on the FICO Blog at www.fico.com/blog.