



FIVE

AI

SOLUTIONS

EVERY RETAIL BANK NEEDS

THE FUTURE OF AI IS NOW



INTRODUCTION

FOR AS LONG AS PEOPLE HAVE WANDERED THE EARTH,

lending has been a part of the economy. That means that lenders have been trying to estimate risk for thousands of years. It is no surprise, then, that banks tend to be among the most sophisticated users of machine learning among all other industries. Nowhere is this more true than in the retail bank. Data is the fuel for risk modeling, and consumer data is plentiful and, over the past decade, has become substantially better curated than the data for commercial lines of business. One consequence of all this is that retail banks tend to be more invested in data science than other lines of business, and this is particularly true for larger institutions.

In recent years, banks have developed automated loan approval processes, credit scoring systems, targeted marketing capabilities, and a variety of other types of solutions. They have done this as a matter of necessity. Particularly in the last 5 years, the number of upstart fintech companies that have formed, particularly in the lending and payments spaces, has grown at a staggering pace. These fintechs are gobbling up market share. They've been successful, in part, because of their customer-centric user experience but also because of their ability to utilize data to optimize their business.



INTRODUCTION (CONT.)

Retail banks must dramatically expand their use of AI and machine learning to optimize every part of their business, from customer acquisition and relationship deepening to managing delinquent loans and portfolio management (buying and selling assets in the secondary market). Organizations that fail to harness this technology will certainly decline and ultimately fail. This article is intended to highlight just a few of the important AI solutions that all retail banks need to develop in order to be competitive. There are many more opportunities than there is space to cover in this e-book, but any bank that completes the use cases outlined below will dramatically boost revenue and improve their competitive position, while at the same time uncovering new ideas and opportunities, and building the capacity and capability to execute and deploy AI solutions.





TABLE STAKES

It's worth noting that predictive modeling is nothing new to the banking space—particularly retail banking. Credit scoring is not new. In fact, the ubiquitous FICO score was first developed in the 1950's, so it's not surprising that nearly all retail banks have some sort of automated credit scoring and approval process built into their system. The scale of the consumer lending market is too big to do business without it.

The largest retail banks likely have their own credit scoring models. These models predict the likelihood that a loan will default—i.e., Probability of Default (PD)—and how big of a loss the bank is likely to suffer in the event that the loan does default; i.e., Loss Given Default (LGD). Many lenders still rely on vendor solutions to risk rate their customers—certainly there is enough regulation in the lending space to make building a credit scoring solution sufficiently daunting. The accuracy improvement in internally-built credit models, though, is usually worth the effort and cost of building them—even if banks build and deploy them entirely by hand, as is most commonly the case.

The fact that automated underwriting and credit scoring models exist (and are table stakes) in the banking sector, though, does not mean that innovation is not needed. For lending, accuracy is king—particularly with such low margins and high competition in the space—the banks that can build the most accurate risk selection and pricing solutions will have a substantial edge in the marketplace, finding and serving the markets with the most attractive risk/reward.

Of course, building and maintaining accurate credit risk models is key not only from a new business acquisition perspective but also from a loss forecasting and stress testing perspective. New disclosure rules make loss forecast accuracy a differentiator for investors and analysts. While this article will not dive deeply into this area, the need for these types of models is critical and non-negotiable for any retail bank.





1

**CLIENTS WITH DEEPER
RELATIONSHIPS TEND TO BE BOTH
MORE LOYAL AND MORE PROFITABLE,
SO ANTICIPATING AND MEETING
CLIENT NEEDS IS CRITICAL.**



PREDICTIVE PRODUCT RECOMMENDATION

The basic nature of the customer relationship is fairly straightforward: All things being equal, the more products that a customer utilizes or buys, the more profitable and loyal that customer will be. Consequently, identifying and building deep customer relationships where the bank provides multiple services to the client—e.g., card, deposit, auto loan, lines of credit, mortgage refinance, overdraft protection, safety deposit boxes, etc.—is an important pillar for building a profitable business.

The most straightforward way of doing this is to use the purchasing and transactional history of your existing customers to determine what sorts of customers have needs for which of your product offerings. Machine learning is an obvious and natural way to solve this problem. Once built, these models will allow personal bankers, customer service representatives, and tellers to make relevant and personalized offers to existing customers. And marketing departments will be able to target the right customers with the right offers, meaning more efficient, productive campaigns.

The time component of the customer journey, though, adds complexity to this modeling process. Certain times of life or life events make certain products more appealing. For instance, a customer with a growing family might be interested in a college savings plan. Once that need has been met it will be practically impossible to sell that product to that customer.

Identification of life events then is a critical tool in making the most relevant product offers to an existing customer base. Transaction history and other customer data provide tremendous insight into when life events—e.g., weddings, births, moves, job changes, and so on—may have occurred and which products these may indicate an emerging need for.

Some organizations have developed rules to identify these life events, but machine learning is a much better solution for detecting them. By utilizing a bank's substantial historical experience with its customers and by tapping third party data providers, the opportunity to predict customers' changing needs has never been bigger. With these predictions in hand, the bank stands much better prepared to anticipate the financial needs of its customers and be ready with the right product at the right time.





KEEPING EXISTING
CUSTOMERS IS AT LEAST
AS IMPORTANT AS FINDING
NEW ONES.

2





KEEPING CUSTOMERS AND KEEPING THEM HAPPY

Acquiring new customers and deepening existing relationships is the offensive side of the customer battle. Competitors are all working to do precisely the same thing. That means that keeping existing customers is at least as important as finding new ones. The relationship deepening work described above will undoubtedly help with this problem since deeper relationships are more difficult to break. As in many things, though, there is no replacement for a good, strong defense.

Given the sheer number of interactions that a bank has with its customers, from transactions to online banking to branch visits and more, identification of unhappy customers or customers at high risk of attrition is relatively straightforward. Machine learning provides the perfect toolset for the detection of customers who may be shopping for a new bank, whether the reason is an unfavorable interest rate or an unpleasant customer support exchange. Once at-risk customers are identified, then specialized teams can be deployed to retain or win back the lost business.

The great untapped resource in the battle for customers is the call center—whether physical or virtual. Apart from irritating telemarketing calls, call centers are largely passive, fielding calls,

chat messages and emails from both happy and unhappy customers 24 hours per day. Very few organizations use intelligent routing to better handle incoming calls.

Given the sheer amount of information that a bank knows about its customers, it's straightforward to predict why many customers are calling. Has someone just had trouble logging into the website? Why not prompt them to be transferred to online support? More important, though, a predictive model can be used to identify which customers are likely upset. For example, their transaction history (including overdrafts, returned checks, or declined card charges) and all previous interactions with bank personnel should be a good start as input to this model. Then potentially frustrated customers and customers who have had difficulty with a transaction can be routed to more experienced personnel with more authority and a greater ability to resolve the issue quickly.

Using a combination of proactive outreach and intelligent routing in the call center will not only improve the customer experience, but it will reduce overall call volume, produce happier customers, and ultimately reduce churn.





3

RETAIL BANKS CAN USE
MACHINE LEARNING MODELS
TO MEASURE PRICE ELASTICITY
AND FINE TUNE THE BALANCE
BETWEEN VOLUME AND
PROFIT MARGIN.





PRICING OPTIMIZATION AND LIFETIME VALUE

In retail channels, the size and volume of individual loans preclude a one-by-one loan approval process. You have to decide *a priori* what acceptable levels of risk are - default risk, loss severity - and price for them in such a way that you balance volume and profit. Set credit standards too loosely and losses will eat into your profits. Set prices too low and you may gain market share, but at the cost of thinner margins. Set credit standards too tightly and volume may suffer even as prepayments rise.

Many banks use a score-carding process—if the borrower and loan terms meet these criteria (e.g. FICO score, LTV, term) then the rate to be charged is %X. Often the primary analytics, perhaps the only analytics, used in these scorecards are estimates of default probability and loss severity. Estimates of price elasticity of demand (how much more or less volume will I get for a change in price?) are often judgmental or empirical based on trial and error.

Many banks have used machine learning to leverage their own experiential data to build better loss models. But machine learning can also build models that predict application volumes for a given price based on prior volumes and market data. You can also build models that predict total profitability for a credit product, taking into account the client's probability of

default, loan structure, net interest margin and fees, and likelihood of prepayment. Prices can then be set that balance profitability and volume. Line of business heads can either “tune” their process for greater volume, or for higher profit margin, or some mixture of both.

In the past, building models for retail credit channels was a slow, expensive, laborious process requiring a critical mass of data science talent. In many cases, single models crossed different regions, client segments, even products. With automated machine learning, you can easily and quickly build hundreds of models. So you can learn from your own experience how behavior will differ across regions, channels, products, client segments or risk spectrum. You can make informed decisions about underwriting, risk acceptance and pricing—knowing how those decisions are likely to shape the composition of your future portfolio.

With precise models, informed decisions can be made that drive the right effort, for the right new business volume, at the right price point, with the right credit and profitability characteristics. That level of insight is a real advantage in the highly competitive market for retail credit and allows bankers to tailor their origination to achieve the desired portfolio mix.





**MACHINE LEARNING IS THE IDEAL
SOLUTION FOR FIGHTING FRAUD.**

4



FRAUD DETECTION/PREVENTION

Losses due to fraud seem to increase every year, with some estimates claiming worldwide losses to fraud as high as \$200B in 2017. Despite the cost, many banks are either fighting fraud with antiquated, rules-based systems or with expensive, black-box vendor models.

Running a successful fraud solution means not only minimizing losses due to fraud but also minimizing irritation and impact to customers. Blocking a legitimate transaction or placing excessive holds on a deposit may reduce direct losses to the bank, but at the cost of a tangible, substantial impact on customer satisfaction, retention and churn.

Machine learning is the ideal solution for fighting fraud. By the very nature of the business, banks record mountains of relevant information about all types of transactions and their counterparties, and whether or not these transactions are fraudulent. This historical data is the foundation of the machine learning approach.

Machine learning models can predict which checks are likely to be bad, which ATM deposit envelopes are likely to be empty, which loan applications are likely to be based on identity fraud and which point-of-sale transactions are likely to be fraudulent. Implementing these models can prevent millions of dollars in losses to fraudsters because fraudulent transactions

can be blocked in real-time to prevent fraud before it happens. Anomaly detection can also be used to spot unusual transactional patterns for a client, which provides a final safety net.

Since many legacy fraud detection mechanisms are application specific, implementing new AI based fraud prevention models can require modification of multiple core systems within a bank. Making changes to these systems may give even the most veteran CTO heartburn. In addition, models must be monitored for accuracy over time, as new types of fraud emerge and the models age. Automated machine learning allows models to be built and deployed quickly, to be retuned and retrained easily, and to be deployed into your production environment with minimal risk.

CREDIT CARD/TRANSACTIONAL FRAUD

BAD CHECK DETECTION

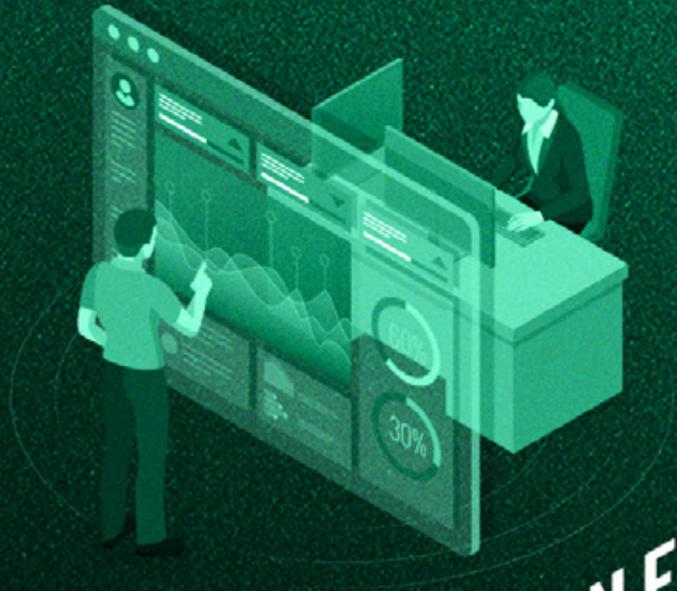
ATM/DEPOSIT FRAUD

FRAUDULENT LOANS CHECK KITING

STOLEN CHECK FRAUD

APPLICATION/IDENTITY FRAUD





WITH MACHINE LEARNING-BASED FORECASTING, YOU CAN CREATE BASELINE FORECASTS BY REGION, BRANCH, OR EVEN BANKER.

5



SALES GOAL SETTING

Setting sales and performance targets can be challenging: set goals unrealistically high, and they may not incent (might even dis-incent) your workforce, goals set too low may lead to paying out too much while leaving growth opportunities on the table.

When setting goals, it's hard to know how much past performance was attributable to natural growth or higher market opportunity. Comparably weak performance may have simply been the result of less promising local demographics and fewer opportunities. It's entirely possible that a branch with lower absolute performance metrics actually outperformed a branch with better performance metrics in a more affluent location or one with higher local growth rates or rising home values. It's also possible that a branch with a higher volume of new deposit accounts at the end of year 1 naturally experienced more deposit balance growth in year 2 with little additional effort.

With machine learning-based forecasting, also known as time-series modeling, you can create baseline forecasts by region, branch, or even banker. The algorithm learns from past performance along with indicative data—local demographics, year-over-year changes in house values or employment, beginning client and account mix—to predict future performance. The more explanatory data you provide on past

performance and local markets, the better able to differentiate performance the algorithm will be.

With projected weekly sales by region, branch, or even individual banker, performance targets can be set realistically so that they serve as strong incentives, you match pay with true performance, and you don't unfairly under-compensate frontline employees in areas with lower growth potential.

Machine learning algorithms can be trained with prior data, and tested and tuned with more recent actual data, to predict volumes (like the number of new card applications), to predict deposit growth (new and existing accounts), and to predict new loan dollar volume. If you have the historical data and enough local level explanatory market data—many of which are available from 3rd party data providers, often free of charge—you can build predictors for every KPI that will provide you with baseline forecasts at every level.

Guessing correctly which regions, branches, and bankers will (or should) perform better when setting goals requires enormous experience and insight, and even then may be more accurate in some places than others. With machine learning, you take the guesswork out. You set better goals with less time and effort and ultimately get better performance from your retail channels.





HONORABLE MENTIONS

Selecting which solutions to include in this guide and which to omit was no easy feat, because the number of high value solutions in this space is extensive. The challenge for every leader in every organization is to develop a roadmap of customized solutions that fits the appetite and needs of the organization. There is no one-size-fits-all solution, and the details of each implementation of any use case will vary from bank to bank.

These are just a few of the many ideas that did not quite make the cut:



Overdraft prediction and other alerts:

Customer-facing models such as alerts generated when customers are at risk of overdrawing their checking accounts represent a largely untapped category of AI solution. The first bank to make these a part of the standard retail offering will earn tremendous credibility and trust from their customers.



Manual vs. automatic account opening:

The process of opening accounts is largely automated and is great when it works. For some institutions, the rate of exceptions and delays in the account opening process reaches as high as 10-15%. Detecting, and correcting, these issues before these gum up the system would mean much great efficiency and a better customer experience.



Branch employee hiring optimization:

Employee attrition is a serious issue for most retail banking heads. High turnover roles, such as in branches, represent a significant investment in both recruiting and training. Identifying long-tenure candidates at the time of resume screening in the hiring process is just one of many recruiting and human resources-related use cases that could mean big savings and happier customers.



Delinquent asset management:

While the lending market is relatively benign now, the credit cycle will eventually turn. Being able to efficiently and optimally handle problem assets will only become more important. Delinquent asset valuation, collections probabilities and amounts, and even participation in the secondary market for distressed debt all consist of many opportunities to optimize the business using machine learning.





CONCLUSION

ARGUABLY, AUTOMATED MACHINE LEARNING IS THE MOST

significant advancement in AI for retail banks in history. Larger banks with sophisticated data science practitioners, can leverage their existing data and expertise across a much broader opportunity set with no incremental staff expense while delivering solutions and responding to market changes faster. Smaller banks can now leverage AI and machine learning without the massive upfront investment in tools, platforms, and people that was once the price of admission. Automated machine learning offers smaller or less sophisticated banks the opportunity to compete head to head with the largest most sophisticated players. That kind of opportunity does not come along often.

[DATAROBOT.COM](https://datarobot.com)

**SIGN UP FOR A
FREE TRIAL TODAY**
TO FIND OUT HOW DATAROBOT CAN HELP
YOUR ORGANIZATION