iTuring.ai

HOT MONEY COLD TRUTH

We analyze how to accelerate deposit growth for banks using net new money predictive models and relationship based pricing recommendations

By SUMAN KUMAR SINGH

ABSTRACT

Finding True Growth in a World of Hot Money

Every financial institution is in a relentless race to grow deposits. But what if one of your most critical growth metrics is misleading you?

Banks spend millions on campaigns to attract new funds, yet a significant portion of that "growth" is often an illusion. It's the result of "hot money": existing funds merely shuffled between accounts to capture a promotional rate.

This phenomenon creates a costly blind spot. It looks like progress on a dashboard, but it silently inflates marketing costs and distorts strategic planning.

The true measure of deposit growth, the lifeblood of a bank's balance sheet, is Net New Money (NNM): genuine, incremental capital brought in from outside the bank. The challenge is that the line between illusory "hot money" and real NNM is nearly invisible to traditional methods.

How can you be certain you are rewarding the *right* customers, the ones who drive sustainable growth, and not just those chasing short-term offers?

The answer is hidden in plain sight: within the vast, complex patterns of your own customer data.

This whitepaper reveals a proven methodology for solving the Net New Money puzzle, to bring in new capital and for building an automated system that turns insight into measurable ROI.

A NEW FRAMEWORK FOR DEPOSIT GROWTH

By the end of this paper, you will understand the framework required to unlock a powerful new capability for your institution. This data-driven approach is designed to achieve three core business objectives:

Pinpoint True Growth Potential

Move beyond simplistic metrics to precisely identify the specific customers who are most likely to bring genuine Net New Money to the institution.

Engineer Intelligent Offers

Replace one-size-fits-all campaigns with a dynamic, relationship-based pricing and offer strategy, ensuring that incentives are targeted to drive profitable behavior and long-term value.

Automate and Deepen Relationships

Implement a scalable, AI-driven contact strategy that grows deposits and systematically enhances the customer experience, fostering loyalty while reducing attrition.

Insights and Objectives

This whitepaper seeks to explore and validate a novel, data-driven approach to accurately identify and harness Net New Money (NNM) as a key driver for sustainable deposit growth within financial institutions.

Through these objectives, this paper aims to equip banking leaders and data practitioners with actionable insights and an evidence-based blueprint to radically improve their deposit acquisition strategies.

Specifically, the primary objectives are:

- To define and isolate true deposit growth (Net New Money) from misleading metrics such as internal fund transfers or "hot money" movements.
- To develop and validate a predictive model that can precisely identify customers with a high propensity to contribute sustainable new capital.
- To formulate a dynamic pricing optimization framework that tailors offers for maximum impact while preserving institutional margins.
- To demonstrate the operational integration of predictive analytics and pricing engines into core banking workflows for real-time, automated campaign execution.
- To validate the methodology through a real-world case study, quantifying tangible financial outcomes including increased deposits and enhanced marketing efficiency.

ANALYTICAL FRAMEWORK

This study leverages a comprehensive dataset comprising transaction and account data for over 450,000 customers at a leading North American financial institution.

The research methodology encompasses a multi-stage, integrated approach:

Data collection & preprocessing

Aggregation of customer-level transaction histories, account balances, and demographic attributes, followed by rigorous data cleaning to remove statistical noise and inconsistencies.

Feature engineering

Creation of advanced behavioral variables and financial ratios (e.g., growth rates, interaction variables, trend variables) designed to translate raw data into meaningful features for predictive modeling.

Predictive modeling

Application of a champion-challenger framework to evaluate multiple machine learning algorithms, including logistic regression, gradient boosting models (GBM, LightGBM), and artificial neural networks (ANN), to classify customers based on their likelihood to contribute Net New Money.

Price elasticity modeling

Development of non-linear elasticity models to quantify customer sensitivity to interest rate changes across various dimensions, including region, occupation, and product segments.

Optimization & validation

Integration of predictive scores and elasticity models within an optimization framework to determine individualized offers. All models are rigorously validated using out-of-time datasets to ensure stability and real-world accuracy.

Deployment & monitoring

Illustration of the automated execution engine that operationalizes offer delivery and establishes a closed-loop system for continuous learning and model refinement.

> By combining rigorous data science with a focus on practical deployment, this methodology ensures that the resulting framework is both predictively accurate and commercially applicable.

TABLE OF CONTENTS

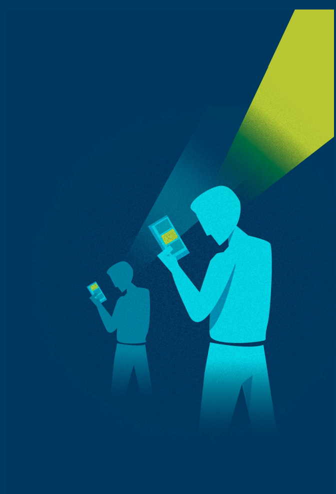

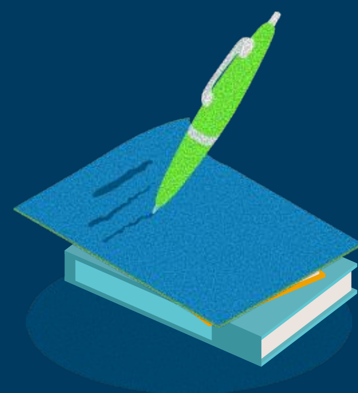

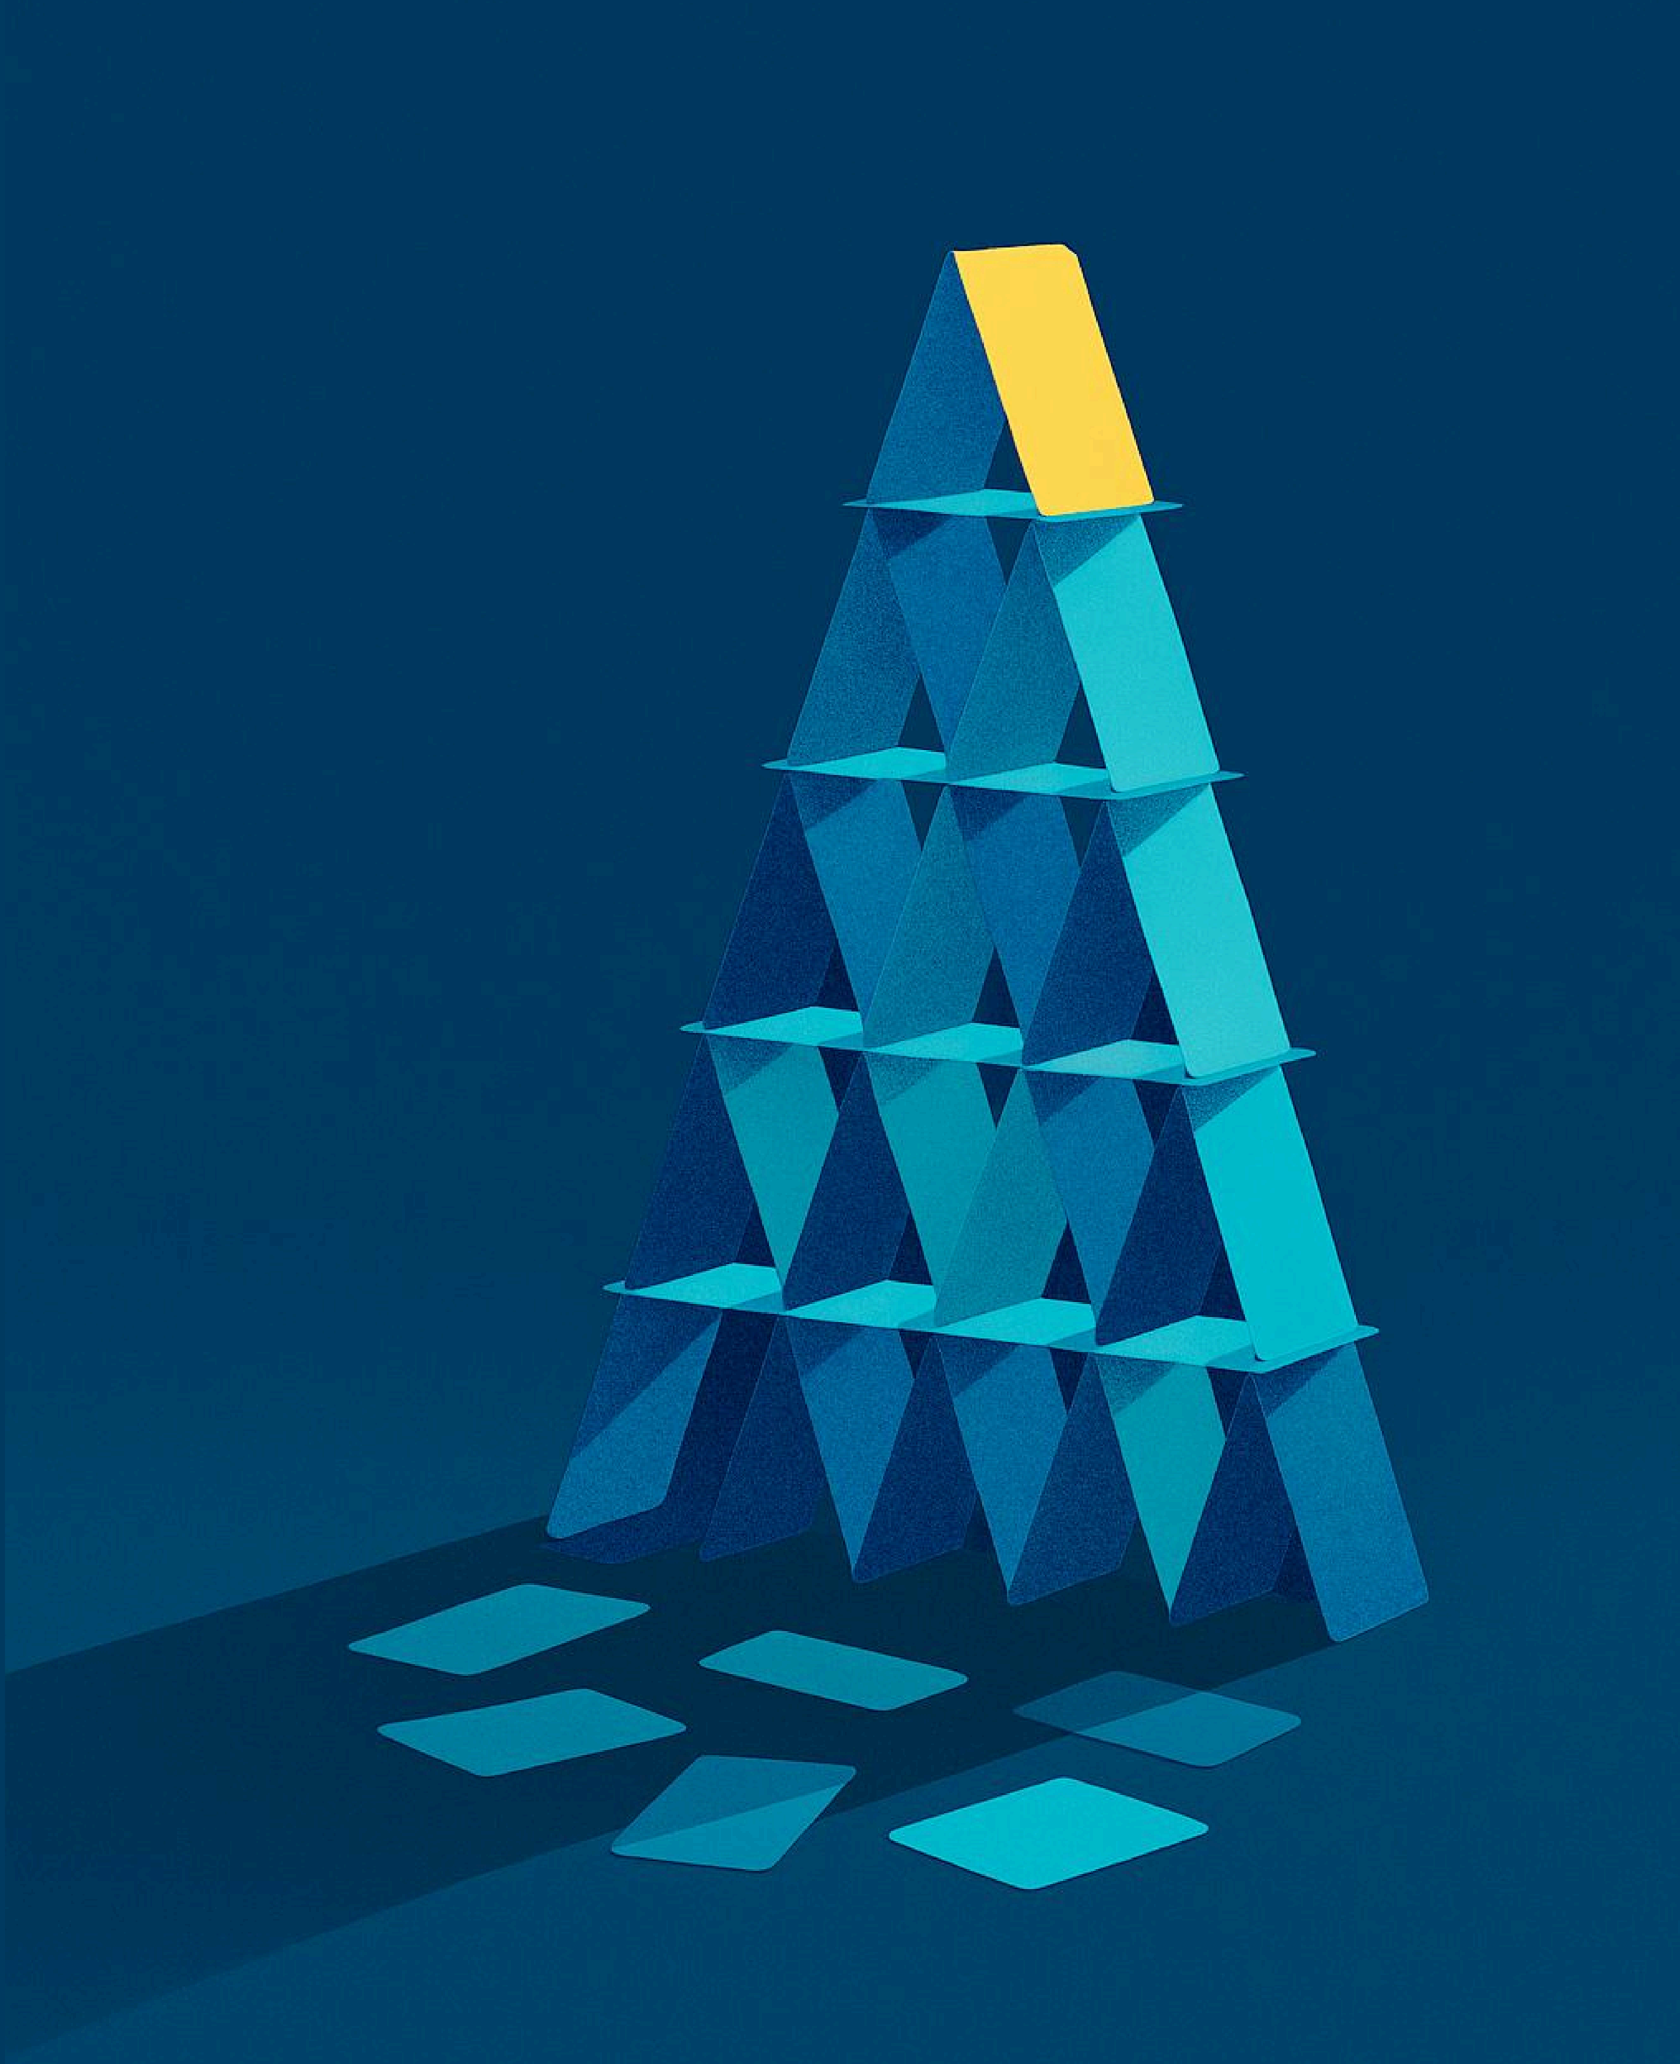

01

Compound Interest, Not Problems

Why attracting the right deposits changes everything

THE MONEY THAT MATTERS

To isolate the true signal of deposit growth, we must first establish a precise definition that cuts through the ambiguity of daily transactions. This signal is Net New Money (NNM), which represents the cumulative new money a customer brings into the bank. Formally, NNM is defined as:

> The total additional deposit inflows across all accounts together under a customer after deducting the previous balance from a benchmark date (period). Including all regular fees, charges and adjusting any account cancellations within a given time period.

LIFEBLOOD OF A BANK

The strategic value of accurately tracking NNM cannot be overstated. Since Net New Money represents the actual additional money from a customer, it arrives as a lump sum capital that is risk-free and immediately strengthens the bank's balance sheet.
A higher NNM growth rate directly enables a bank to decrease its cost of funds and thereby increase the margin on its lending products.

From a regulatory perspective, this is one of the most efficient levers for improving the cost-income ratio (CIR) and managing capital adequacy requirements. The benefits are clear and significant. Why, then, does identifying it remain such a persistent challenge?

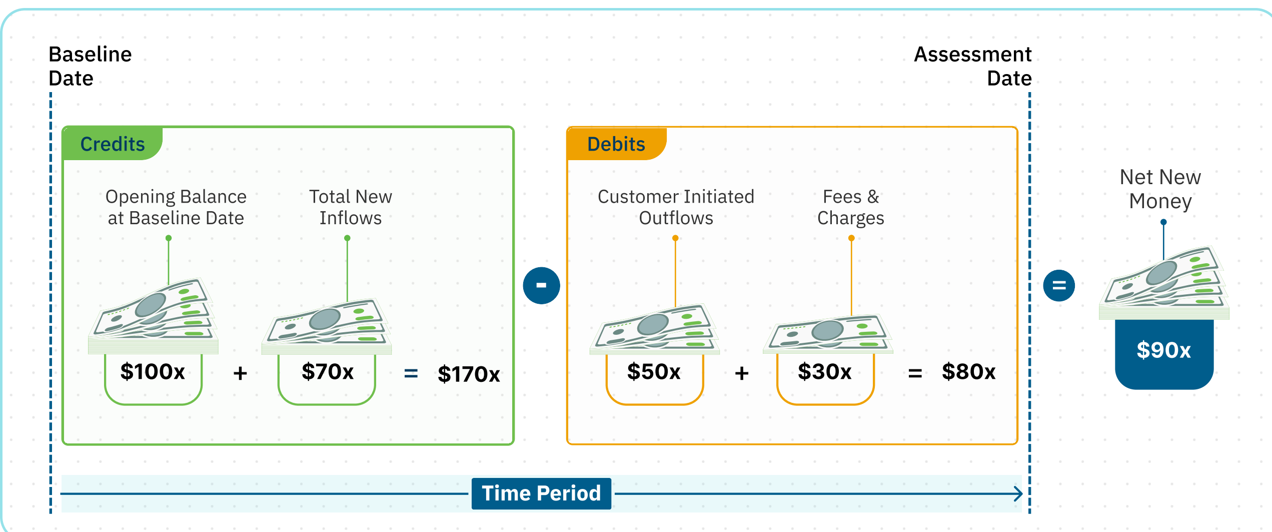

This calculation provides the foundation for accurately measuring sustainable capital growth.

THE HIGH COST OF THE "HOT MONEY" ILLUSION

Herein lies the central problem. While the definition of NNM is clear, the computation is complex. The primary difficulty is that the true signal of cumulative new money is often obscured by "hot money": funds that customers shuffle between their own accounts within the same bank. This creates a significant gap between perceived and actual growth.

For instance, in the example provided below, a bank might calculate an inflow of $80X and consider it Net New Money.

However, this figure is misleading because it fails to exclude internal transfers. Without a system to automatically identify and discount this "hot money," the institution is making critical decisions based on inflated and inaccurate data.

> **The fundamental question then becomes:** How can a bank move from a manual, error-prone process to a reliable, automated system that can accurately identify true Net New Money at scale?

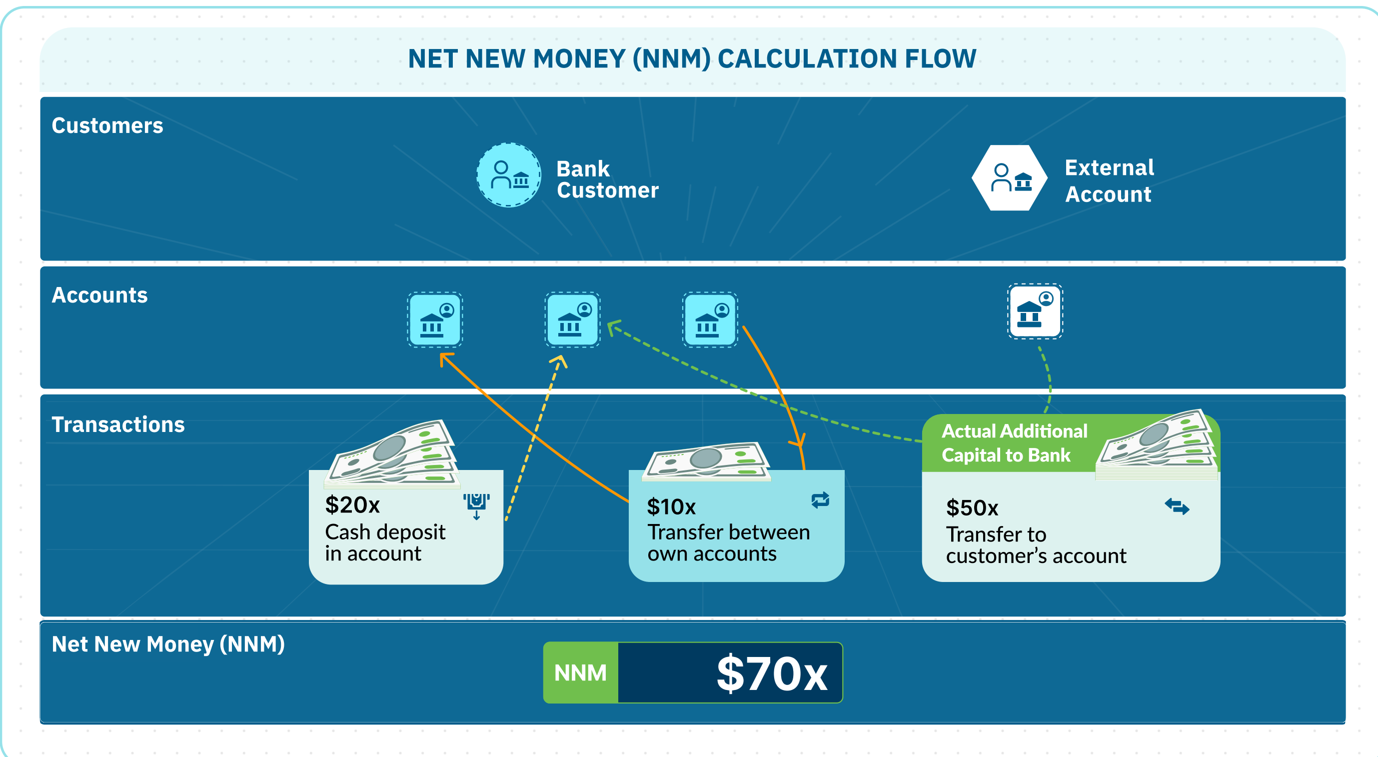

02

The Moneyball Moment For Banking

How to stop swinging for the fences and start scoring real growth

PUTTING GROWTH ON AUTOPILOT

The solution to the Net New Money puzzle requires moving beyond manual analysis to a fully automated system. One that can distinguish "hot money" from genuine inflows to empower business users to act on that intelligence.

The system must be sophisticated enough to handle complex fund-movement calculations, yet intuitive enough for marketing staff to configure and deploy campaigns without specialized technical skills.

That is precisely the challenge that the Turing Decision AI is designed to solve. It empowers a bank to maximize Net New Money by enabling them to:

- Design and configure a contact strategy using a prioritized list of customers who are most likely to bring in net new money.
- Configure and define specific rules for what constitutes Net New Money for individual customers, ensuring calculations are precise and tailored.
- Execute dynamic business rules for customer-centric pricing, offers, and incentives, allowing the bank to automatically deliver the "next best product" to the right customers.

PREDICTING YOUR NEXT BEST CUSTOMER

At its core, this approach requires a fully automated data science and machine learning platform, complete with an integrated MLOps system. An environment of such nature is designed to enable real-time, personalized customer strategies by equipping banks to develop and gather powerful predictive intelligence.

Such a platform makes the process of identifying high-potential customers both straightforward and intuitive. It moves beyond simple rules-based calculations to generate actionable predictive scores. These scores, combined with the right KPIs, allow for the creation of a highly effective contact strategy, often built on sophisticated cross-sell/up-sell models.

Furthermore, by deploying advanced machine learning techniques, banks can apply rigorous scientific methods like champion-challenger testing and Design of Experiments. This allows them to test multiple offers, scientifically sharpening their customer targeting and maximizing the ROI of every campaign.

With the combined intelligence of a predictive AI platform and a decision-making engine, a bank can finally upgrade from executing costly "blanket offers" to deploying a more targeted and intelligent basket of offers for specific customers. This approach simultaneously keeps costs low and increases revenues, all without compromising the customer experience.

THAT'S WHAT WE DO AT iTURING.

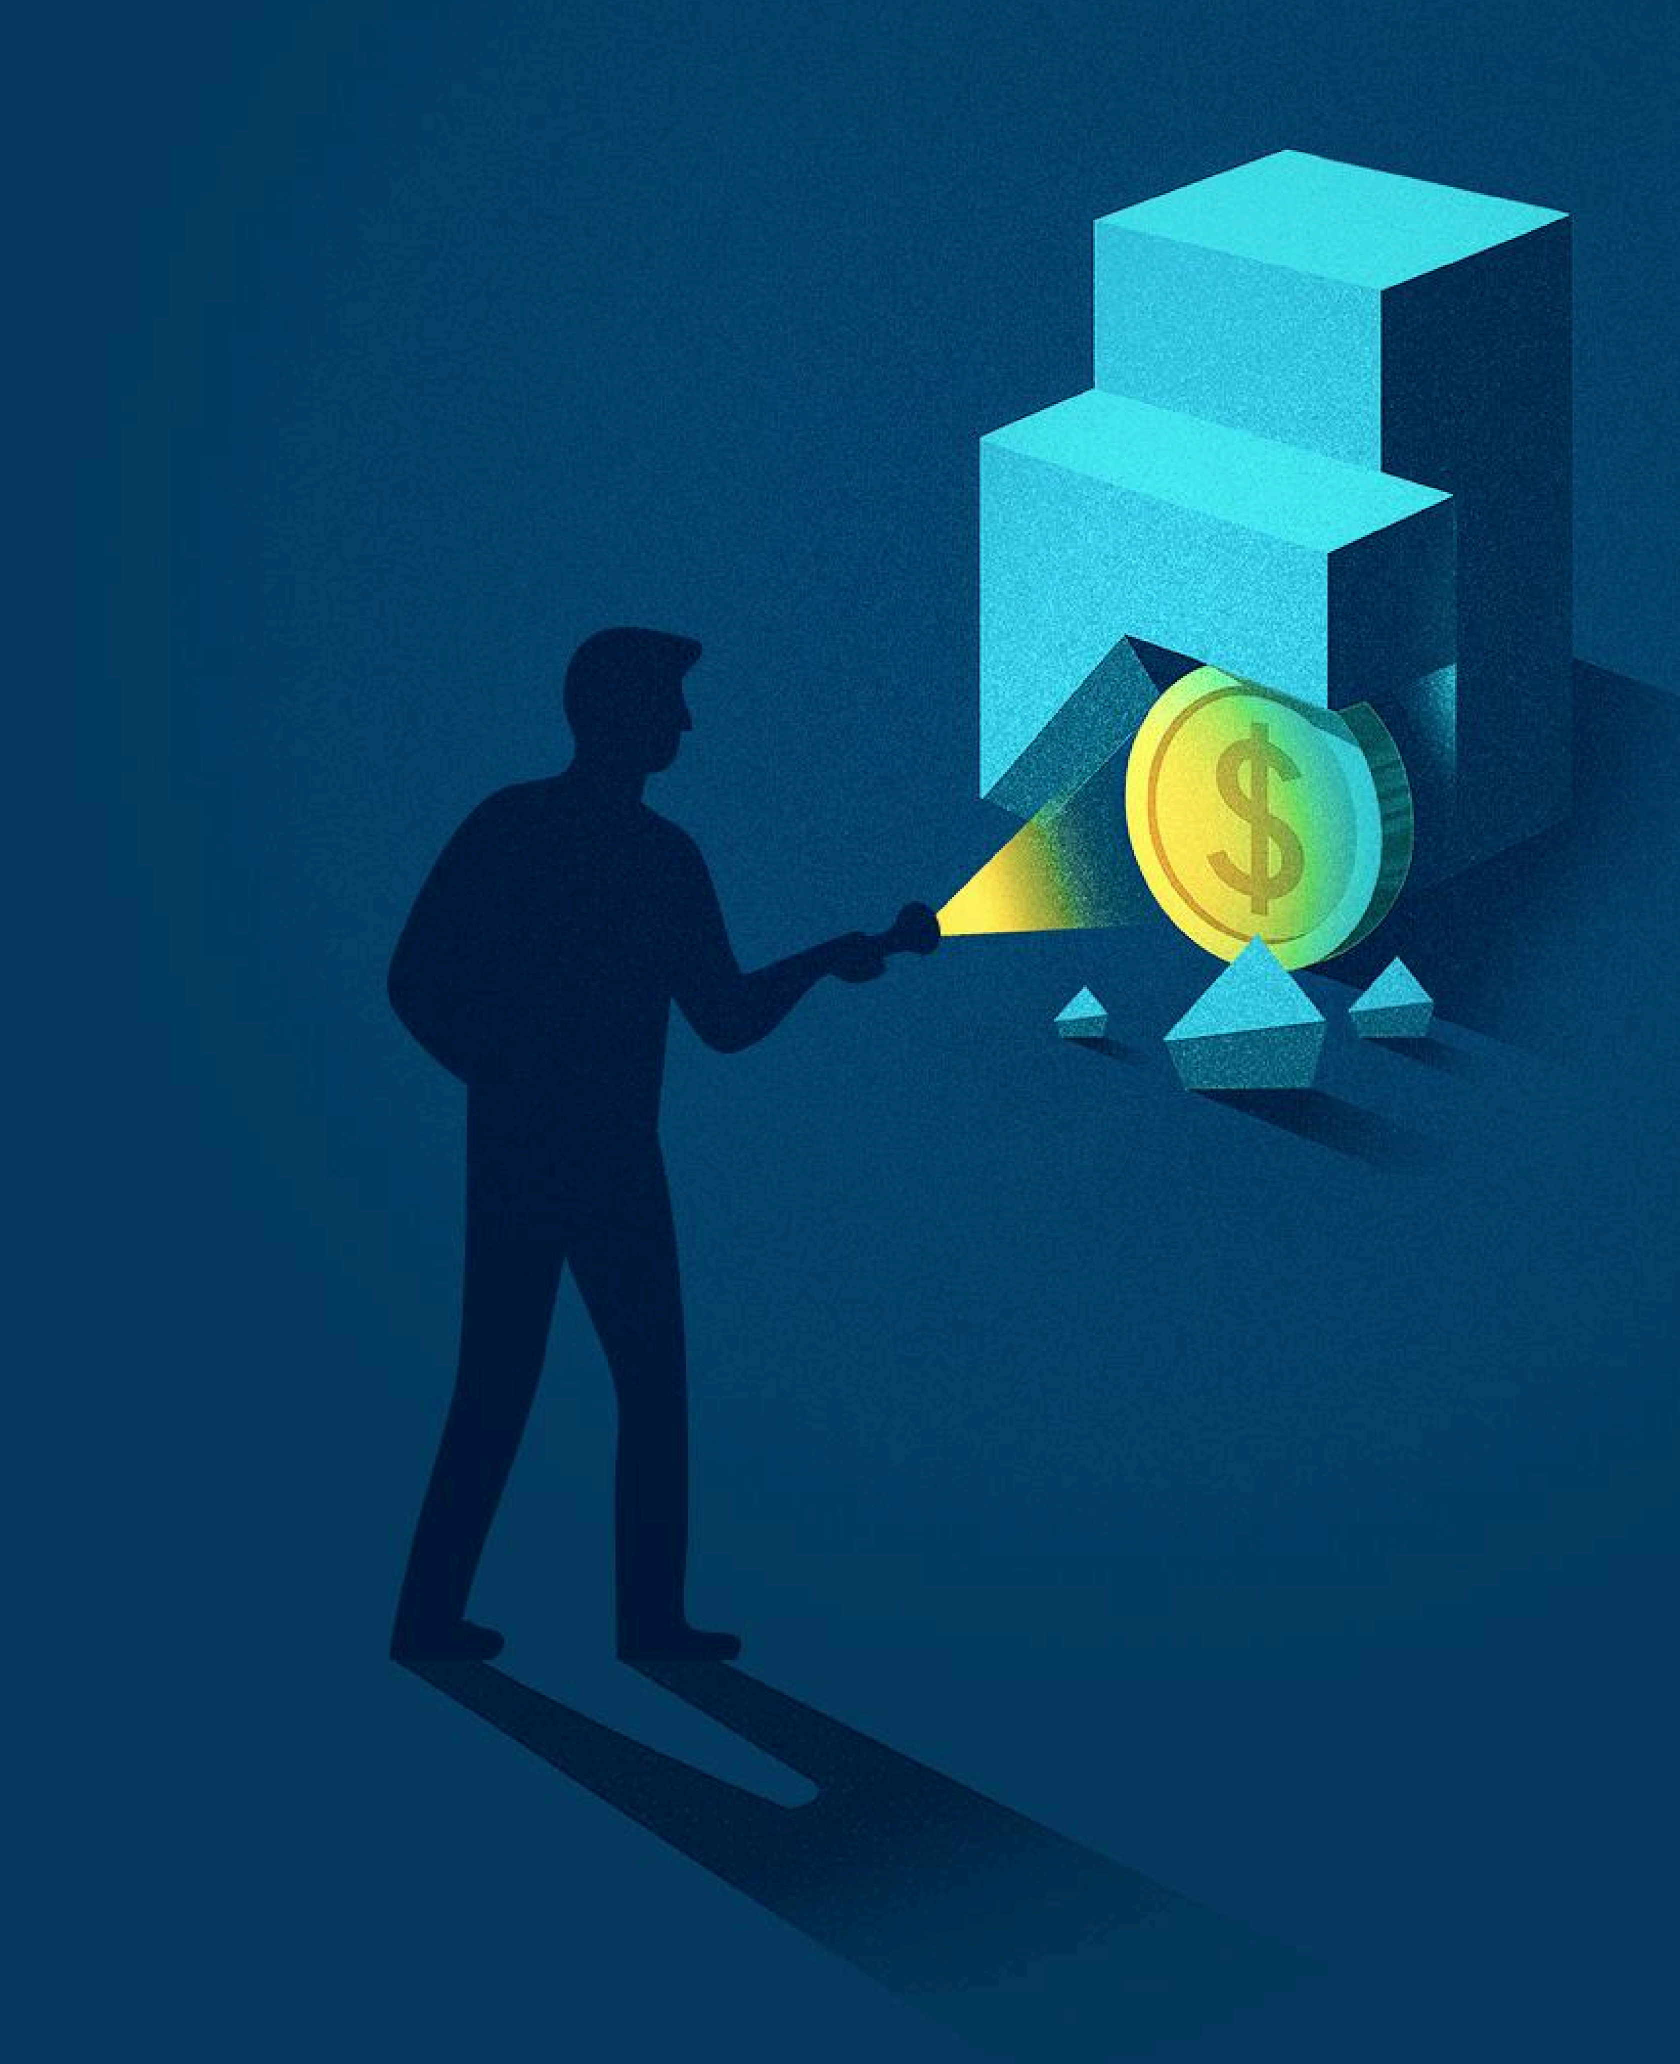

03

How to Find $45 Million in Hidden Growth

How a leading North American bank drove millions in new capital

Setting the Stakes: Escaping the "Blanket Offer" Trap

A leading North American bank sought to move beyond conventional marketing and focus its resources on customers with a high propensity to bring in genuine Net New Money.

The business rationale was clear: if the bank could accurately predict which customers would drive real deposit growth, it could create targeted, cost-effective incentive programs to secure those funds.

FROM A MILLION CUSTOMER POINTS TO A SINGLE, SHARP SIGNAL

The project was built on a comprehensive dataset covering 458,698 customers within a performance window. The data was rolled up to the customer level, encompassing all deposit accounts (Savings, Checking, and Term Deposits).

To ensure model stability and validity, the analysis used a three-month observation period and a separate out-of-time dataset of 464,029 customers for validation.

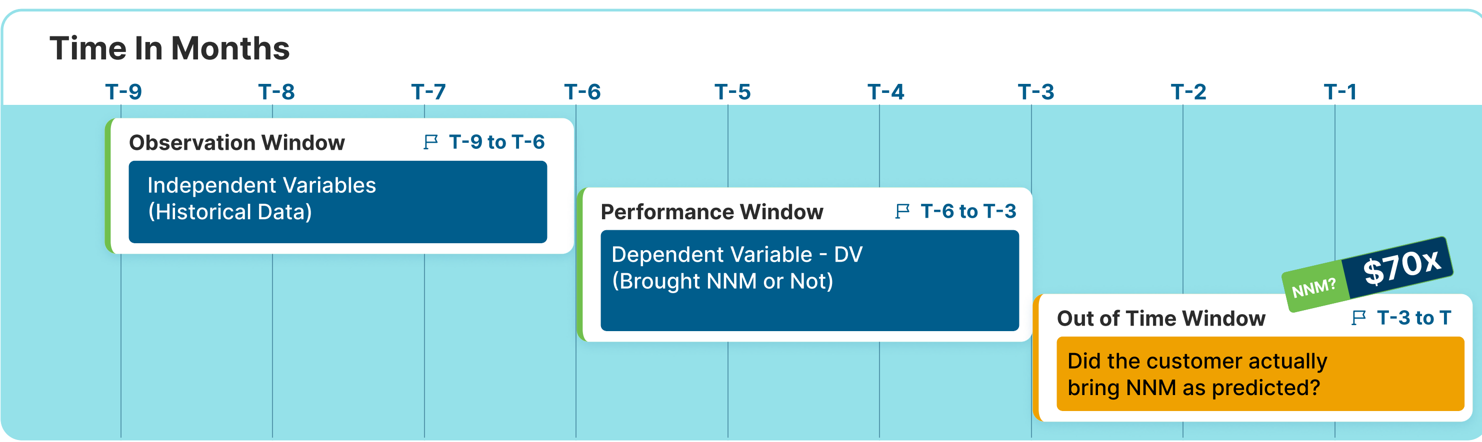

The Dependent Variable (DV), the target outcome for the model, was rigorously defined. A positive case was a customer who met all of the following criteria:

- **Purchased a new Term Deposit within the observation window.**
- **Increased their total Term Deposit balance.**
- **Increased their total balance across all deposit accounts.**
- **Achieved a total growth of >= $500 in their deposit and/or Term Deposit balances.**

The data was split into a 60% training set and a 40% test set to build and evaluate the models.

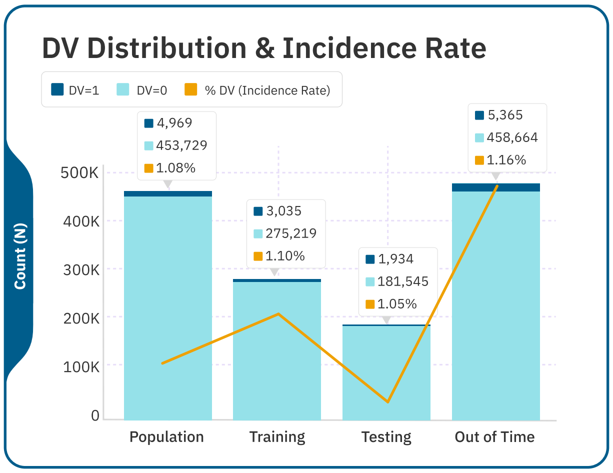

The initial dataset captured a wide range of customer behaviors, demographics, and value parameters across 1,224 raw variables. However, to build a powerful predictive model, this raw data first needed to be meticulously refined to separate the signal from the noise. This involved several critical data science techniques:

Data Cleaning

Inadmissible variables were removed due to business considerations or statistical issues like a high percentage of missing values.

Feature Reduction

To identify the most impactful variables, Fisher's score was utilized, retaining the top 70%. Subsequently, a correlation matrix was analyzed to remove variables with a correlation greater than 0.8, and the Variance Inflation Factor (VIF) was used to eliminate variables with a VIF greater than 10 to tackle multicollinearity.

Data Outlier and Missing Value Treatment

Missing values were imputed using the median for continuous variables (like age) or replaced with 0 for transactional variables (where a missing value indicated no transaction). Outliers were handled by flooring at the 1st percentile and capping at the 99th percentile. For the specific case of age, any value below 20 was imputed with 20.

WHEN DATA STARTS TO TALK

With a clean and robust dataset established, the investigation shifted from data preparation to active discovery. The objective of this phase was to move beyond the raw variables and begin engineering new, more powerful features that could capture the subtle signals of customer intent.

Leveraging deep business and domain knowledge, a series of new features were created, including financial ratios, growth rates, interaction variables, and trend variables. This process is designed to translate raw data points into meaningful business concepts that describe customer behavior over time.

To visually analyze the potential of these newly created features, bar graphs and histograms were utilized. This allowed for an initial analysis of the correlations between the potential explanatory variables and the dependent variable (the likelihood of bringing in Net New Money). These pre-model analyses provided the first concrete clues in the investigation:

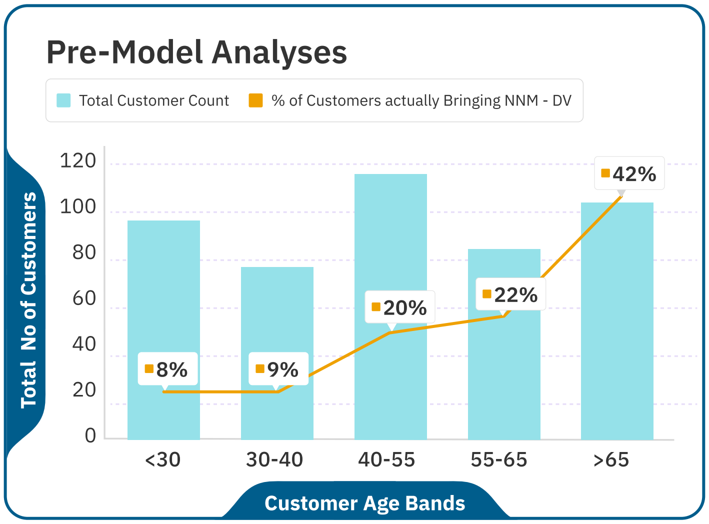

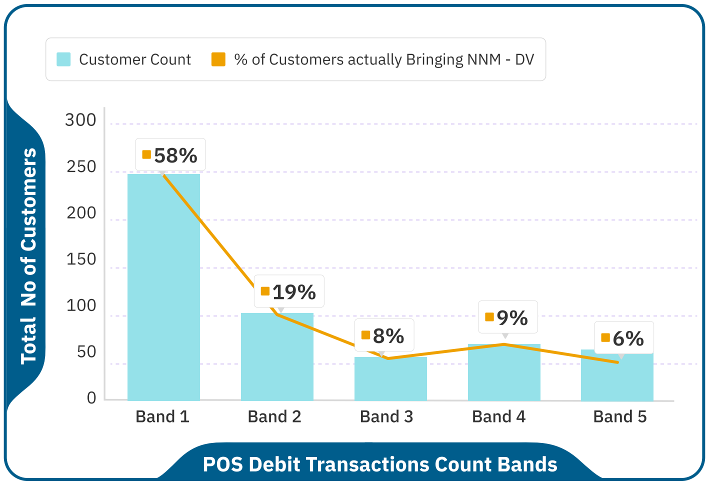

As these simple bivariate graphs demonstrate, a clear and strong relationship began to emerge between the key explanatory variables and the target outcome. This initial analysis confirmed that there were, indeed, powerful predictive signals hidden within the data, paving the way for the next phase: formal model development.

SURVIVAL OF THE FITTEST: A BATTLE FOR PREDICTIVE SUPREMACY

With the initial clues from the exploratory analysis in hand, the next step was to determine which modeling technique could most accurately and reliably predict the likelihood of a customer bringing in Net New Money. To achieve this, a rigorous champion-challenger approach was employed, leveraging a variety of platforms and tools including R, Python, and the Scikit-learn and Keras libraries for advanced machine learning.

Six distinct models were developed and pitted against each other to identify the top performer:

- Logistic Regression
- GLM Model
- ANN with RMSprop as the optimizer
- GBM Model
- LightGBM model
- ANN with ADAM as the optimizer

The core objective was to identify the algorithm with the highest predictive power, while also ensuring that its performance was stable over time and across different data samples. To validate this, each model was evaluated not only on the test data but also on a separate out-of-time dataset.

Performance was measured by looking at the Area Under the Curve (AUC) score, a key metric that provides an aggregate measure of a model's performance across all classification thresholds. As is evident from the results, the LightGBM model emerged as the clear champion. It achieved an excellent out-of-time AUC score of 0.8413. Just as importantly, the difference between the training AUC and the out-of-time AUC was very low, revealing that the model's performance is highly stable over time.

The classification matrix of the LightGBM model further confirmed its strength, demonstrating that the model achieved excellent discrimination and remained stable in both the testing and out-of-time samples on multiple matrices.

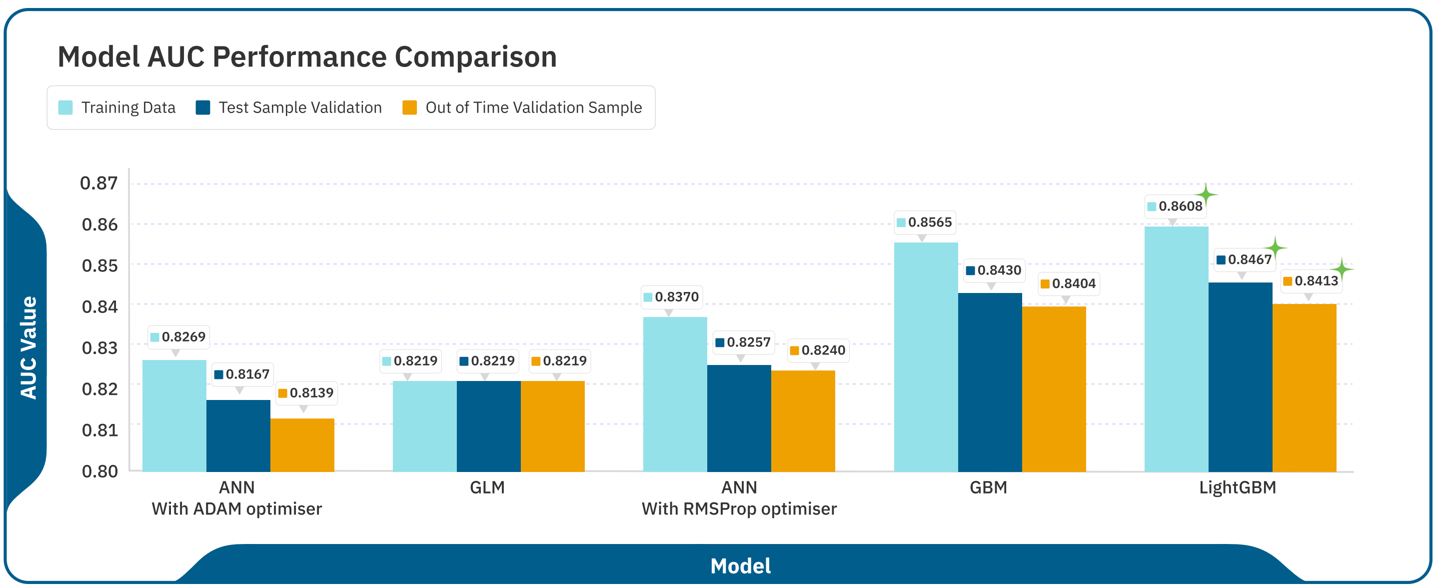

CLASSIFICATION TABLE FOR LIGHTGBM MODEL

Metrics	Precision	Specificity	Sensitivity	Accuracy	FPR	FNR	AUC Score	Gini	Brier Score	KS	Discriminant_slope
Training	3.85%	78.17%	78.32%	78.17%	21.83%	21.68%	0.86078971	0.72	0.01	0.57	0.051
Validation	3.56%	78.21%	75.44%	78.18%	21.79%	24.56%	0.84668527	0.69	0.01	0.54	0.045
oot	3.91%	78.86%	73.61%	78.80%	21.14%	26.39%	0.84129676	0.68	0.01	0.53	0.045

FEATURE IMPORTANCE FOR THE LIGHTGBM MODEL

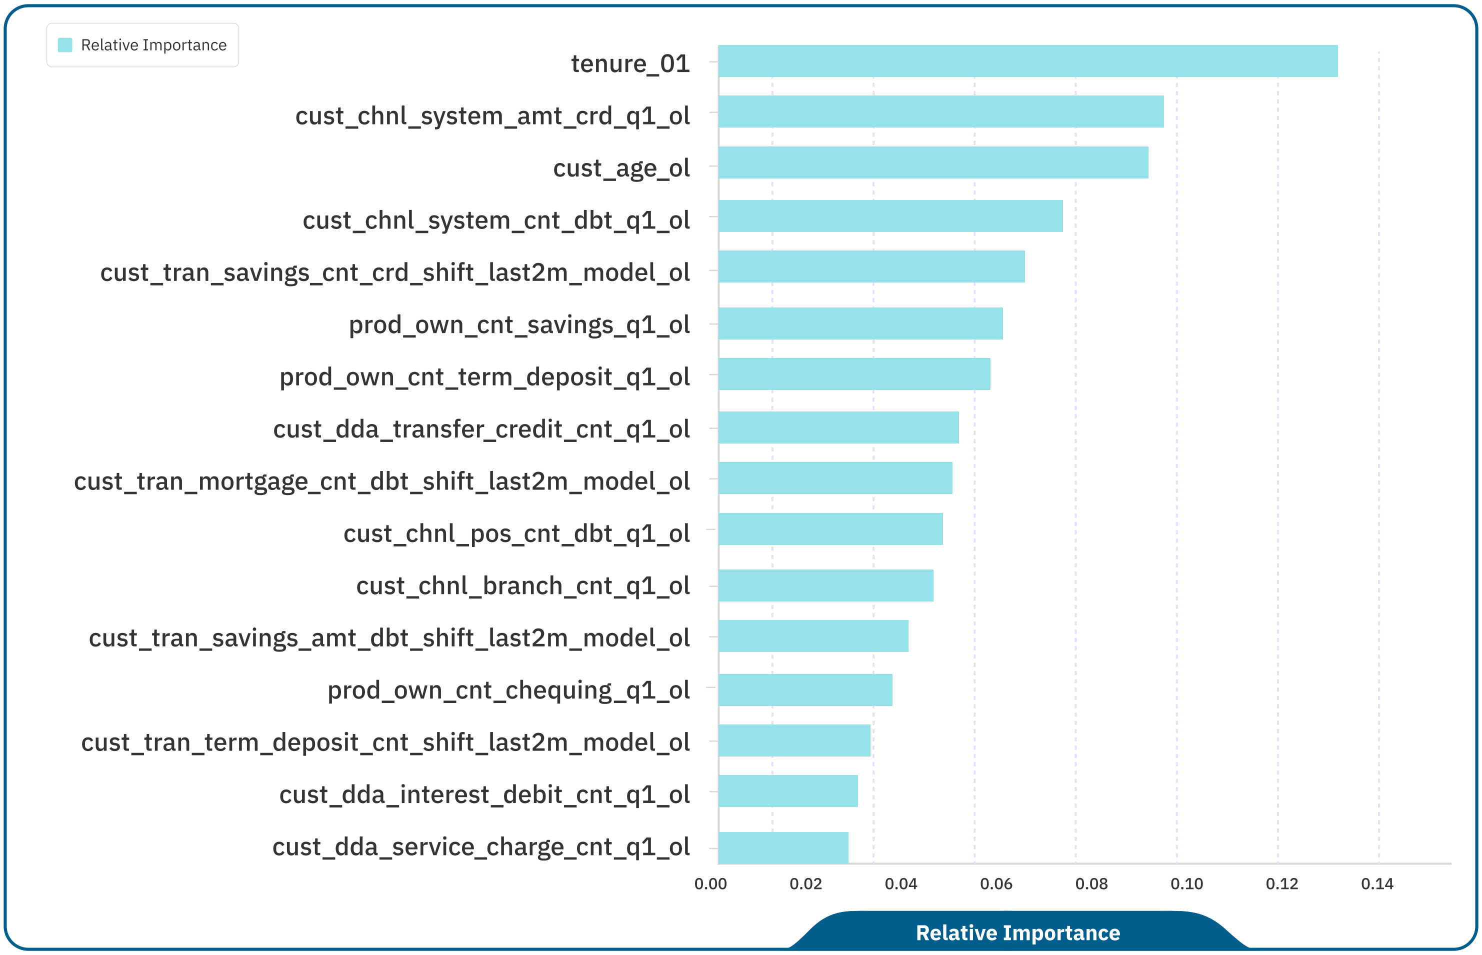

With a clean and robust dataset established, the investigation shifted from data preparation to active discovery.

The objective of this phase was to move beyond the raw variables and begin engineering new, more powerful features that could capture the subtle signals of customer intent.

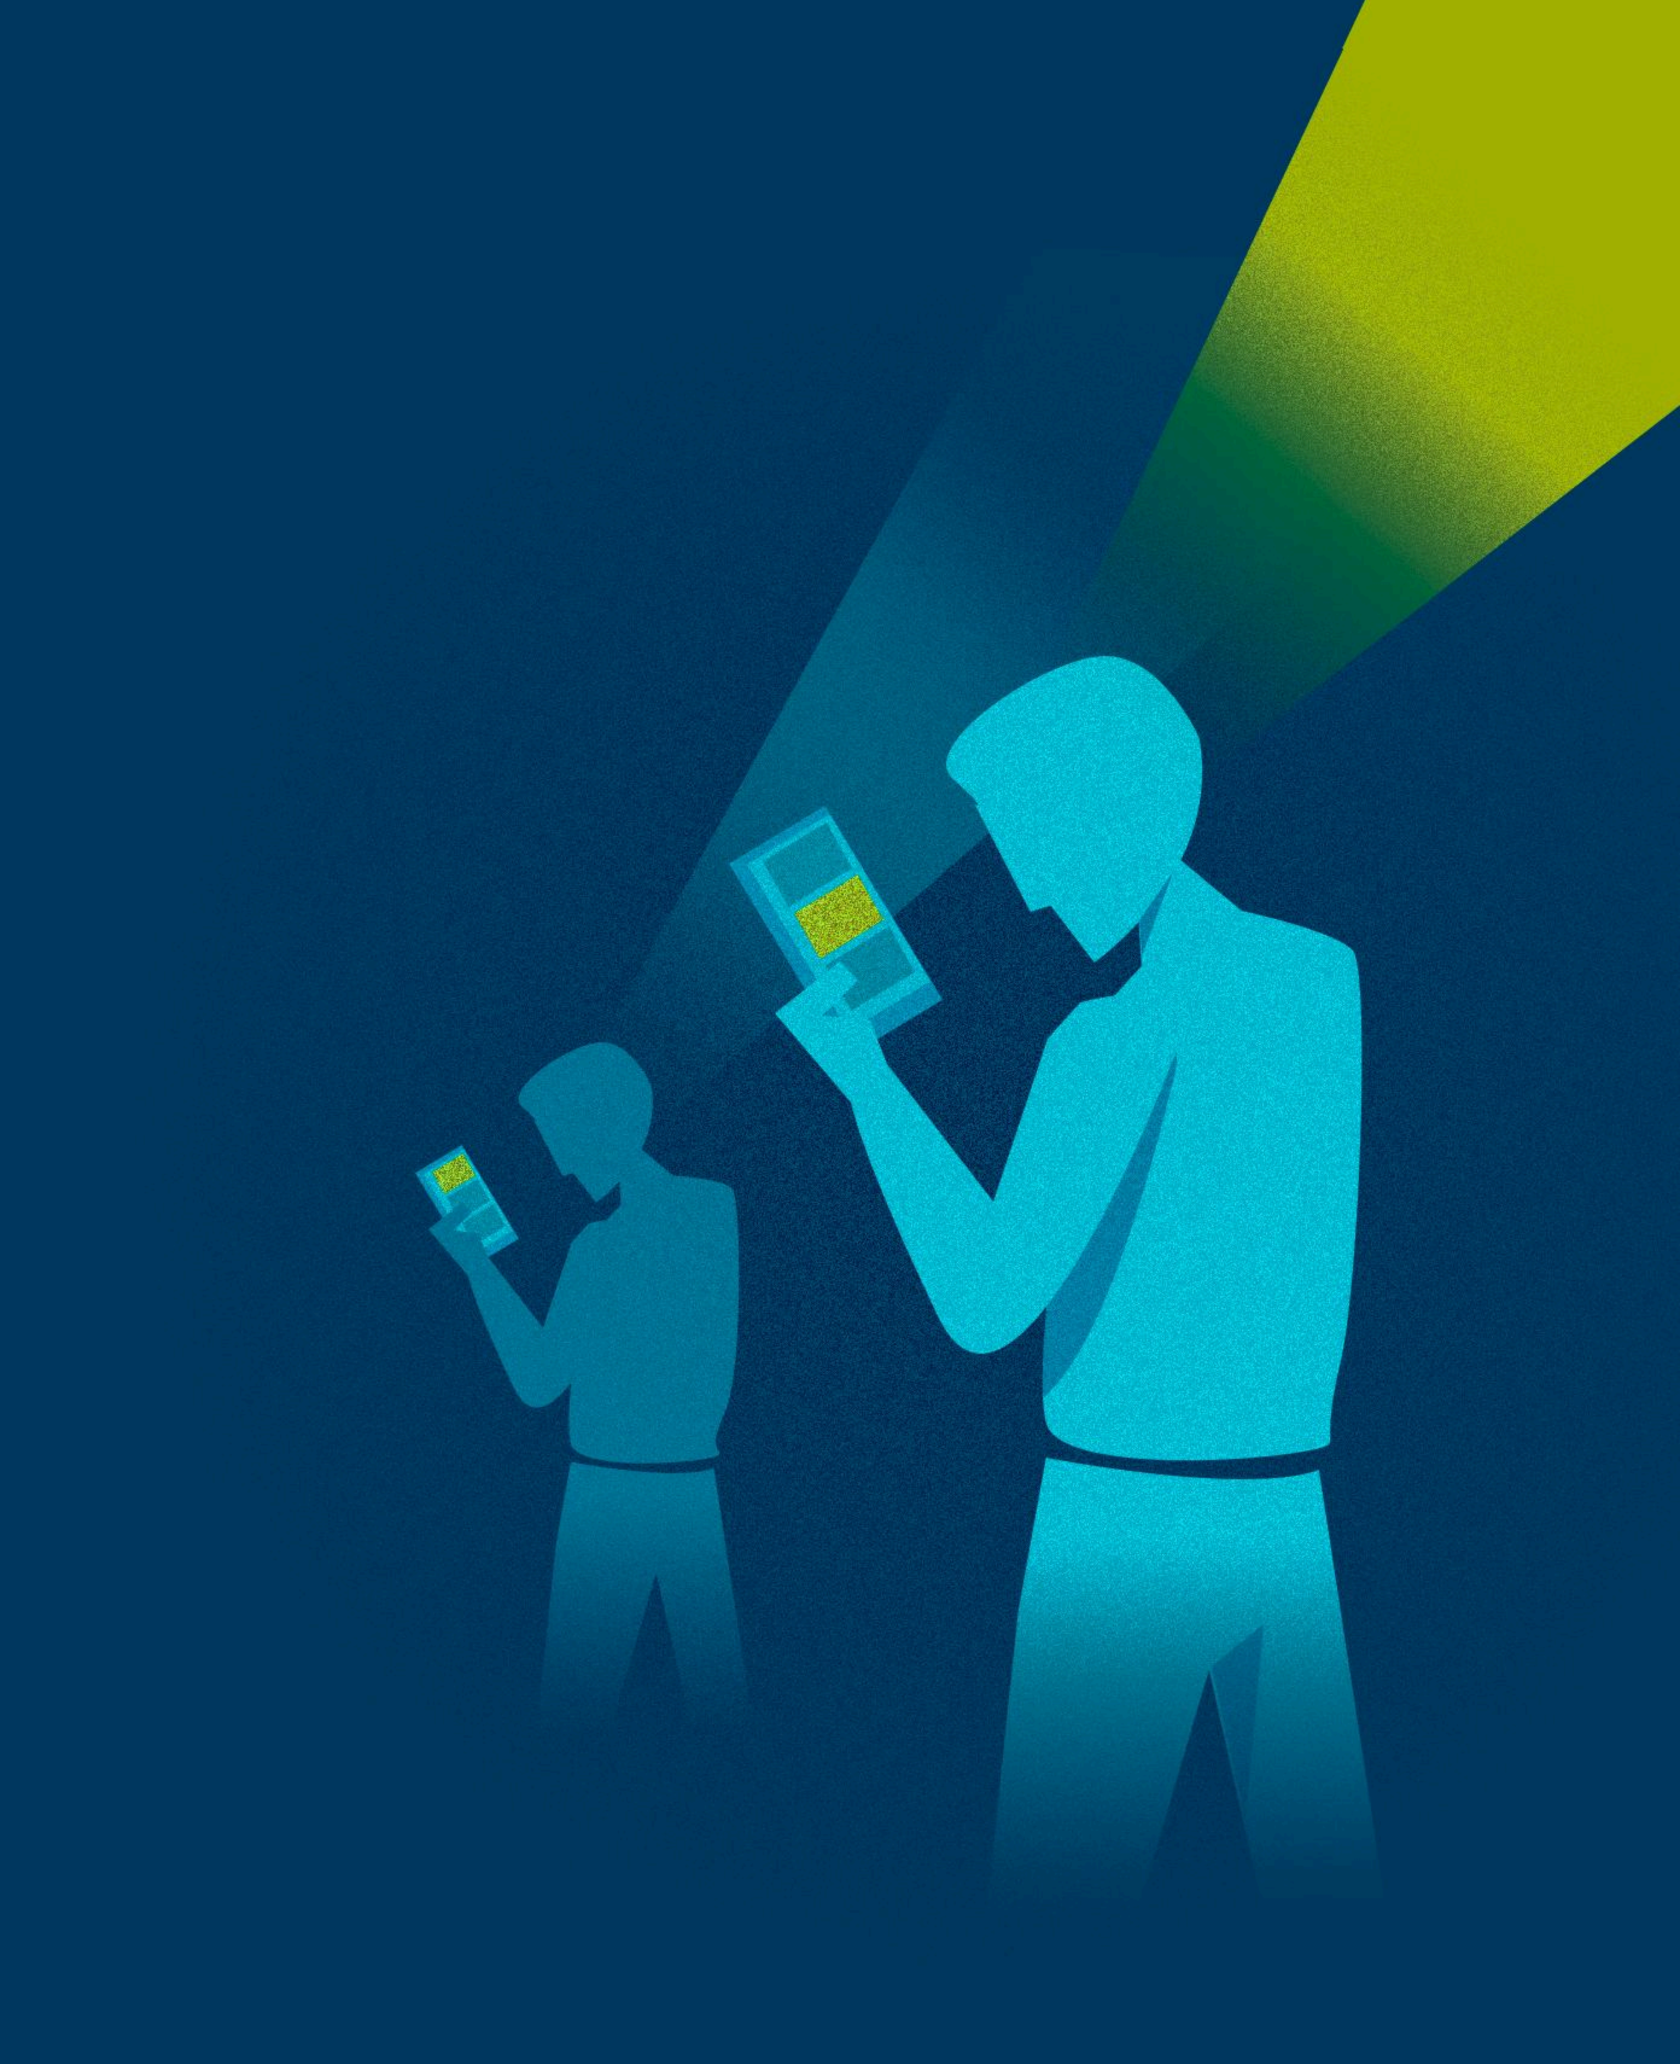

04

The Price is Right: Science of The Perfect Offer

A blueprint for dynamic, personalized, and profitable pricing

How to Motivate Without Giving Away Margin

Identifying customers with a high propensity to bring in Net New Money is a critical first step. However, a predictive score alone, no matter how accurate, does not answer the next crucial question:

What is the optimal offer to make to each customer?

How can a bank design pricing that motivates the right behavior without unnecessarily giving away margin?

The answer lies in moving from a one-size-fits-all pricing strategy to a dynamic, data-driven one. This requires a deep understanding of how different customers will react to different offers—a concept known as price elasticity.

READING THE CUSTOMER'S MIND

A price elasticity model for deposit growth quantifies how sensitive different customer segments are to changes in deposit interest rates. This enables a financial institution to forecast how rate changes will impact deposit inflows and retention, taking into account competitive pressures and macroeconomic conditions.

To achieve this, a non-linear price elasticity model was developed to estimate price sensitivity across multiple dimensions, including:

- REGION
- PRODUCT
- OCCUPATION
- CUSTOMER SEGMENTS

This allows for the measurement of elasticity at a granular level, defining which customers have a positive or negative sensitivity to price changes.

ENGINEERING THE "JUST RIGHT" OFFER

With a clear understanding of price sensitivity, the bank can then move to optimize its pricing. By analyzing sensitivity across products and regions, it is possible to maximize profitability while keeping a close watch on competition and promotions.

This is achieved by using non-linear optimization techniques to calculate the optimal price for every customer, based on custom constraints defined for each segment.

LANDING THE PERFECT PRICE EVERY TIME

The final step is to translate this analysis into a concrete action. By performing a price positioning analysis, the system can analyze profitability, measure the impact of market competition, and recommend the optimal price for each customer accordingly. This ensures that every offer is not just personalized, but also strategically sound.

THE 1-2 PUNCH: WHERE PREDICTION AND PRICE CONVERGE

The goal is to identify customers who are going to bring Net New Money. And configure the right offer for the right price to the right customers.

This is where the two streams of analysis, prediction and pricing converge into a single, powerful capability. By integrating the Net New Money predictive model with the price elasticity model, a bank can move beyond simply identifying high-potential customers to knowing exactly how to motivate them with the right offer at the right price.

> The integrated solution is operationalized through an automated offer execution engine: the Price Recommendation Engine (PRE).
>
> The PRE dramatically reduces the time it takes to get a campaign from conceptualization to execution. It provides marketing teams with a simple, web-based interface where anyone can configure a new offer to be executed at set frequencies.

Crucially, the NNM calculation is built directly into the PRE, so it automatically discounts any "hot money" movement and ensures that offers are targeted only at generating real, incremental deposits.

The integration with a bank's core banking system is a one-time setup. From that point on, the PRE can automatically collect the required data to execute offers seamlessly.

It creates a powerful, closed-loop system for continuous improvement. As the PRE executes campaigns, all customer and offer data flows back into the iTuring platform.

The result is the rapid development and refinement of machine learning models. It also enables deep performance analysis, generating insights into which offers worked, for which customer segments, and what the precise dollar impact was.

You get a virtuous cycle of learning and optimization, where every campaign makes the next one smarter.

THE ART OF NUDGING: HOW TO MOTIVATE YOUR BEST CUSTOMERS

With a powerful predictive model in place, the bank can shift from a reactive to a proactive engagement strategy. Instead of waiting for customers to act, the bank can now proactively approach those with a high likelihood of bringing in Net New Money, armed with insights from their historical data.

There is an opportunity to educate and motivate these high-potential customers with personalized offers. The system also enables the deployment of intelligent, trigger-based offers, ensuring that the timing of the offer is perfectly suited for maximum customer conversion.

> **At the same time, this predictive capability allows the bank to "ring-fence" these valuable customers, proactively managing their experience and addressing any service issues that could trigger dormancy or attrition.**

Recognizing that different customers will respond differently to various offers, the platform provides a robust framework for scientific testing. To begin with, a bank may not know which offers will work best.

By using the Price Recommendation Engine (PRE) and the NNM propensity scores, the bank can sub-segment the customer base and test multiple, small-scale offers before committing to a full-scale launch.

The data-driven approach to offer design and testing is seamless and smooth, as new offers can be easily configured in the PRE, and their performance can be quickly analyzed.

It dramatically reduces the time-to-market for full-scale offers and ensures that when they are launched, they have already been tested and validated on a smaller sample, leading to much higher conversion rates.

Identifying customers with a high propensity to bring in Net New Money is a critical first step.

However, a predictive score alone, no matter how accurate, does not answer the next crucial question: What is the optimal offer to make to each customer?

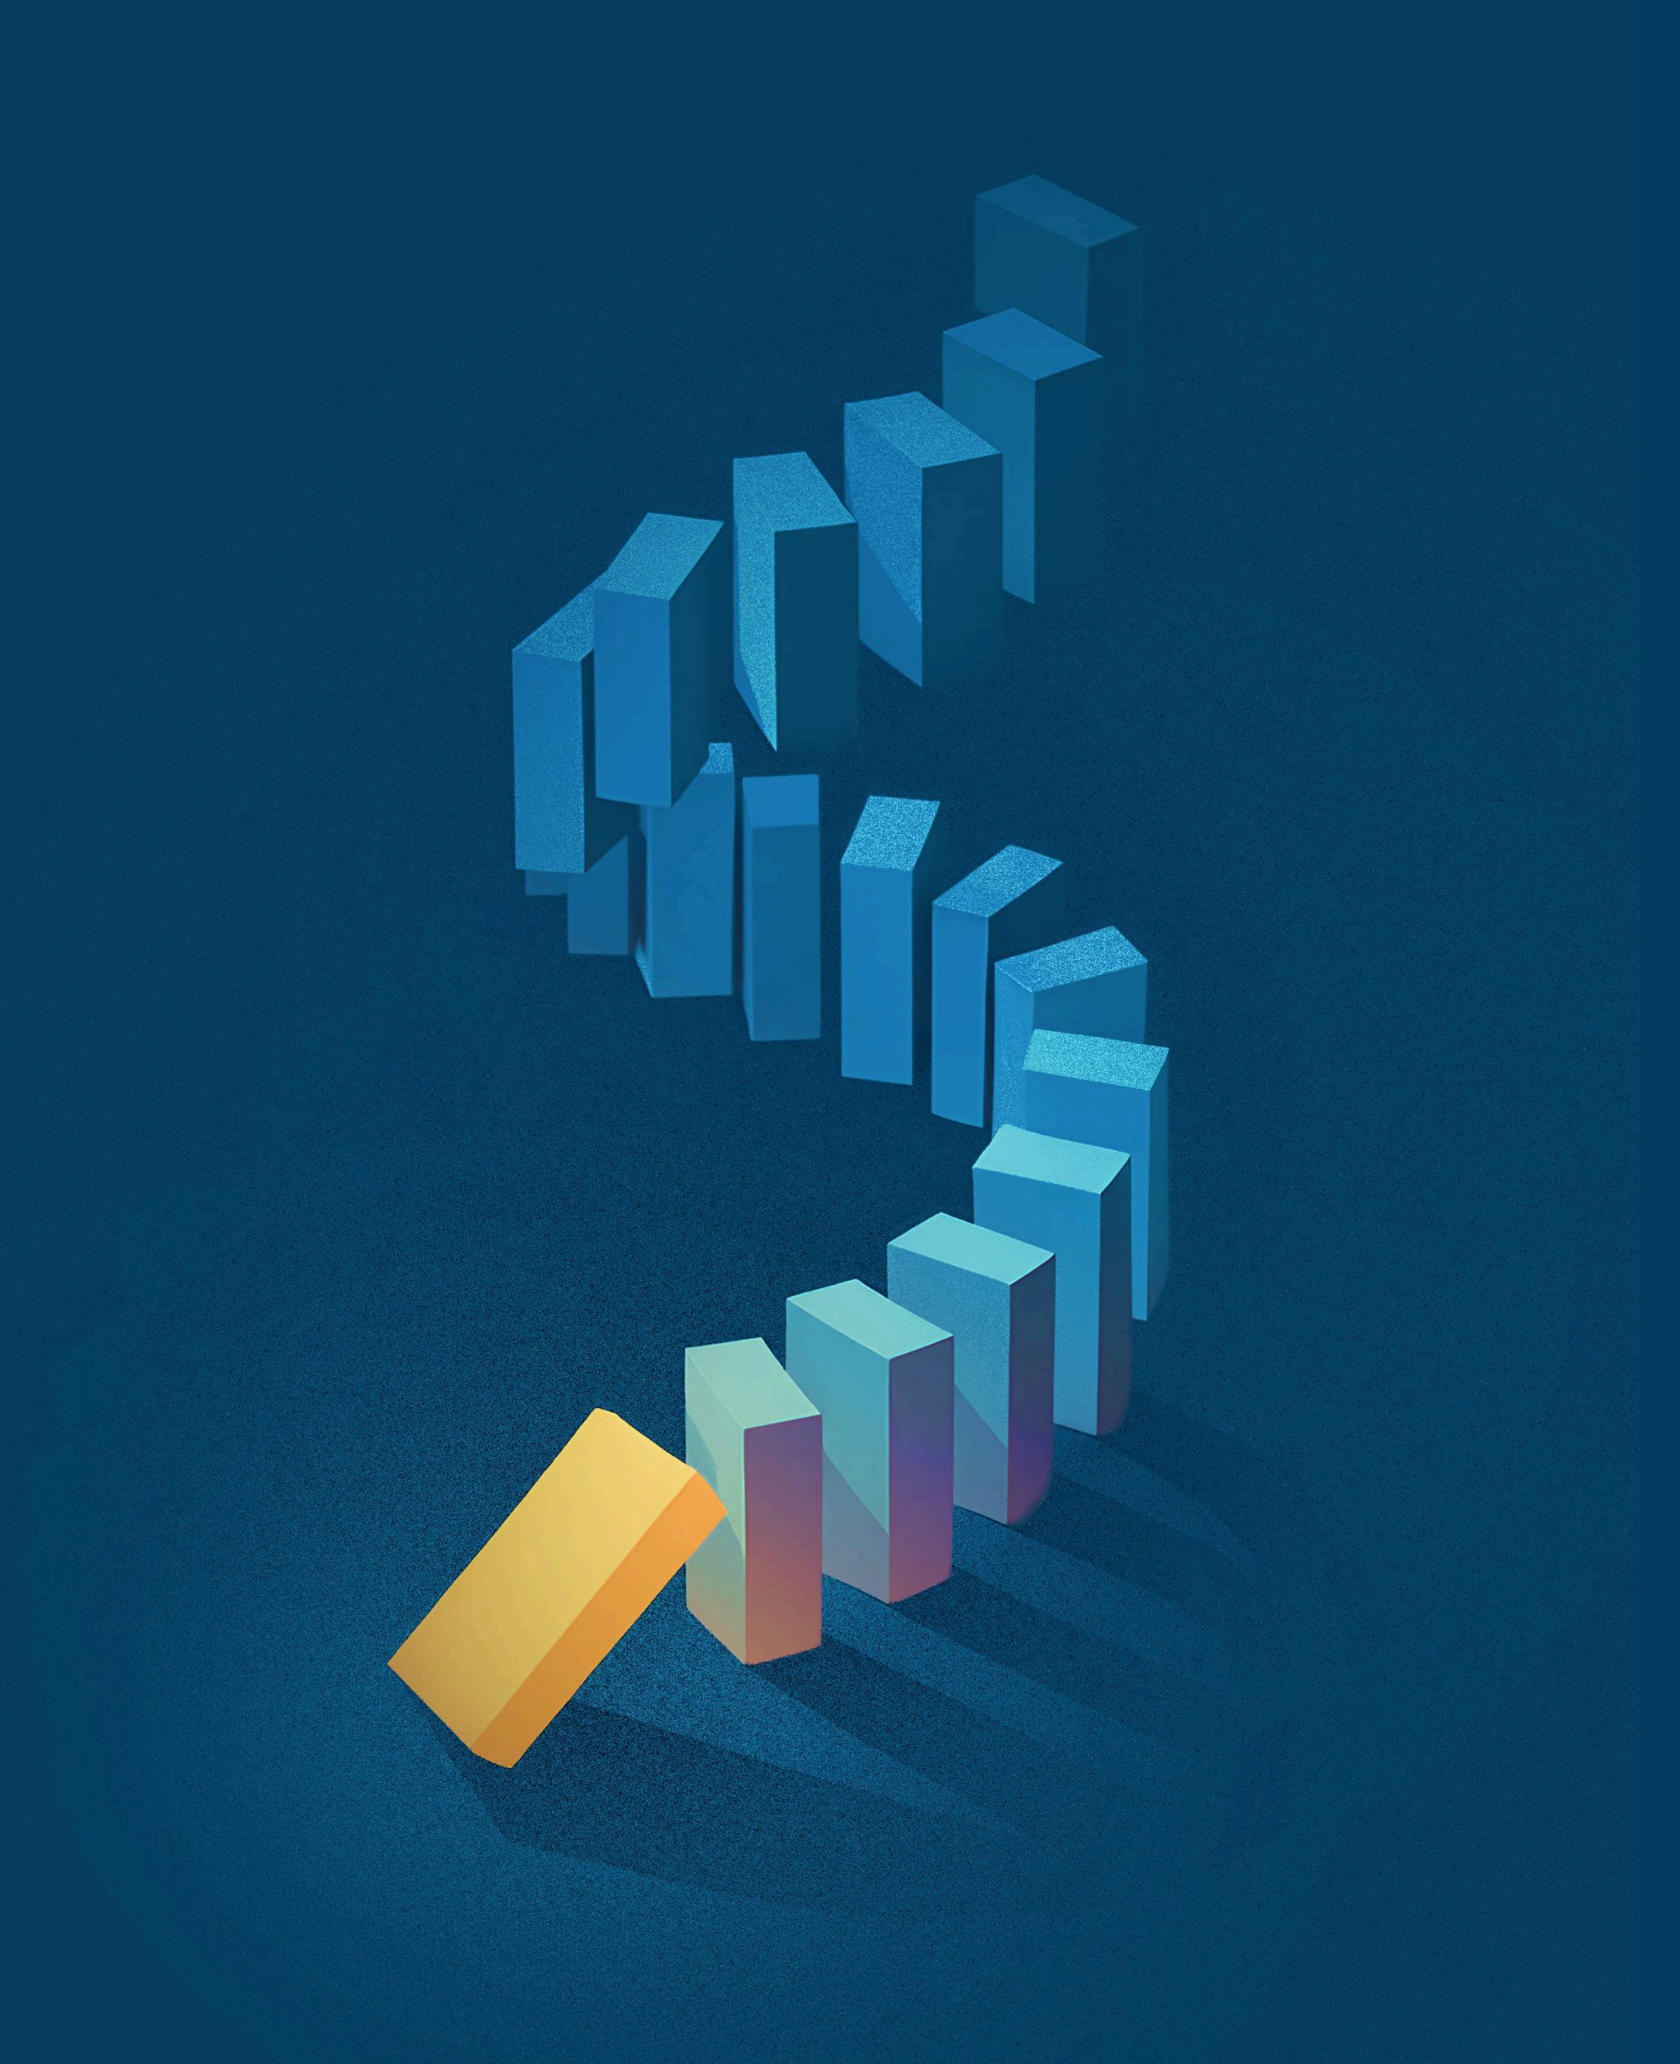

05

The Domino Effect

How one model ignited total growth

The Anatomy of a Win

The true measure of any analytical model is its impact in the real world. By overlaying the Net New Money predictive score with the Price Recommendation Engine, the bank was able to target its new Term Deposit (TD) product campaigns with unprecedented precision.

The results were immediate and significant:

21% increase in deposits, which translated to over $45 million in new capital in just 90 days.

82% engagement match between the customers predicted to respond and those who actually did, providing powerful validation of the model's accuracy.

We achieved it through a strategy of personalized outreach, targeted promotions, and the dynamic monitoring of campaign effectiveness all orchestrated to strategically grow NNM.

THE START OF A NEW STORY

Net New Money is the most important KPI for most banks. As this case study has clearly elucidated, the challenge of distinguishing genuine growth from the noise of "hot money" is no longer an unsolvable puzzle. By utilizing sophisticated machine learning algorithms, it is possible to predict which customers will bring in Net New Money with a high degree of accuracy.

The conclusion is clear: the data-driven approach is a powerful and essential tool for increasing efficiency, reducing marketing costs, and directly improving profits.

By using the combined power of the Price Recommendation Engine and predictive score-based prioritization, the bank was able to optimize its resource utilization and maximize its Net New Money inflows. This journey from a complex business problem to a measurable, high-impact outcome demonstrates the transformative power of an integrated AI strategy.

The same engine that drives deposit growth can also be used to develop and execute a variety of strategic marketing initiatives. This includes tapping cross-sell/up-sell opportunities, predicting customer attrition, and deploying trigger-based offers for a truly competitive advantage in the modern financial landscape.

06

About the Author

Suman Kumar Singh

ABOUT THE AUTHOR

Suman Singh is the Founder and Chief Executive Officer of iTuring.ai, a next-generation AI and machine learning platform designed to transform the banking and financial services industry.

With over two decades of experience at the intersection of AI, analytics, and technology leadership, Suman has been a pioneering force in delivering scalable, practical AI solutions to major financial institutions globally.

Prior to founding iTuring.ai, he held senior leadership roles including Chief Analytics Officer at Zafin and key positions at Fiserv, where he spearheaded initiatives to integrate AI-driven decision-making into complex financial services environments.

Under his leadership, iTuring.ai has emerged as a recognized innovator, earning multiple industry awards for its end-to-end, no-code platform that simplifies and accelerates the path from data to measurable business value.

Suman is passionate about bridging the gap between cutting-edge AI research and real-world business impact, empowering financial organizations to unlock sustainable growth for a true competitive advantage.

www.ituring.ai
iTuring.ai