



**Causal Inference and Hyperpersonalization on a Scale: A
CONVERGENCE OF STOCHASTIC ARCHITECTURES AND INTELLIGENCE
Artificial Generative Engineering in Maximizing *Customer Lifetime
VALUE IN FINANCIAL ECOSYSTEMS***

CAUSAL INFERENCE AND HYPER-PERSONALIZATION AT SCALE: THE
CONVERGENCE OF STOCHASTIC ARCHITECTURES AND GENERATIVE AI IN
MAXIMIZING CUSTOMER LIFETIME VALUE WITHIN FINANCIAL ECOSYSTEMS

Author: Willian Gouveia de Aguiar

Bachelor's degree in Information Systems from Bandeirante University.

Product Management – PM3

SUMMARY

This study presents a comprehensive analysis of the application of Causal Inference and Generative Artificial Intelligence (GenAI) models in behavioral segmentation within large financial institutions. It investigates the transition from monolithic data architectures to event *-driven ecosystems*, demonstrating how reducing information latency directly impacts engagement and revenue metrics. The article details the engineering behind a real-world case study that resulted in a 46% increase in transactional message engagement and the generation of R\$19 million in incremental revenue, proposing a new paradigm for managing *Customer Lifetime Value* (LTV). The research concludes that the orchestration of stochastic algorithms with real-time data governance is the determining factor for customer retention in the digital economy.

Keywords: Hyperpersonalization, GenAI, Causal Inference, Data Architecture, LTV.

ABSTRACT

This study presents an exhaustive analysis of the application of Causal Inference models and Generative Artificial Intelligence (GenAI) in behavioral segmentation within large-scale financial institutions. It investigates the transition from monolithic data architectures to Event-Driven Architectures, demonstrating how the reduction of informational latency directly impacts

engagement and revenue metrics. The paper details the engineering behind a real-world case study that resulted in a 46% uplift in transactional message engagement and generated R\$19 million in incremental revenue, proposing a new paradigm for Customer Lifetime Value (LTV) management. The research concludes that the orchestration of stochastic algorithms with real-time data governance is the determining vector for customer retention in the digital economy.

Keywords: Hyper-personalization, GenAI, Causal Inference, Data Architecture, LTV.

1. INTRODUCTION

In the third decade of the 21st century, the global financial industry faces an unprecedented operational and strategic paradox: an abundance of transactional data *versus* a systemic inability to convert this volume of information into contextual relevance in real time.

Historically, banking institutions have based their customer relationship *management* (CRM) strategies on static demographic paradigms, segmenting individuals based on low-volatility variables such as declared income, age range, and residential geolocation. However, the rise of *Open Finance* and the ubiquity of mobile devices have radically altered the consumer journey, which has ceased to be linear and predictable, becoming fragmented, multichannel, and instantaneous. In this scenario, traditional propensity scoring models, based on simple logistic regressions applied to batch-processed data, prove mathematically insufficient to capture the stochastic complexity of modern human behavior.

The obsolescence of traditional methods lies not only in the statistical technique, but also in the underlying data architecture that supports decision-making. The latency inherent in legacy *data warehouses*, which often operate on weekly or daily update windows (D-1 or D-7), creates an insurmountable time gap between the event that generates the need (the behavioral "trigger") and the marketing action. By the time the offer reaches the customer, the "micro-moment" of interest has already dissipated, resulting in marginal conversion rates and, worse, a deterioration in the user's perception of brand value, who begins to see the bank's communications as noise or *spam*. The inefficiency of these generic campaigns imposes a gigantic opportunity cost, measurable not only by the wasted media budget, but also by the silent erosion of *Customer Lifetime Value* (LTV).

Faced with this problem, the imperative need arises for a paradigmatic reformulation that transcends simple descriptive ("what happened") and diagnostic ("why it happened") analysis, moving towards predictive ("what will happen") and, crucially, prescriptive ("what we should do") analysis. The technical answer to this challenge lies in the convergence of three vectors.

Cutting-edge technologies: Event -*Driven Architecture* (EDA), which allows for real-time signal processing; Causal Inference, which seeks to understand cause-and-effect relationships beyond spurious correlations; and Generative Artificial Intelligence (GenAI), which enables massive-scale semantic personalization. The central hypothesis of this work is that the synchronous orchestration of these technologies is capable of restoring the relevance of financial institutions in the daily lives of their clients.

This article, based on the author's empirical and academic experience in Information Systems and Product Management, aims to dissect the engineering and mathematical modeling necessary to implement this transformation. It is not merely a theoretical discussion, but a rigorous analysis of implementations that aim to unite data, artificial intelligence, and management platforms to create innovative solutions. The objective is to demonstrate how data governance, combined with timeliness, generates actionable *insights* that transform the customer experience.

2. THE TRANSITION OF DATA ARCHITECTURE: FROM MONOLITHS TO MESH ECOSYSTEMS

Traditional data infrastructure, predominant in the banking sector until the mid-2000s, was based on centralized monolithic architectures where data extraction, transformation, and loading (ETL) occurred within rigid time windows. While robust for financial accounting and regulatory reporting, this model is inadequate for the dynamics of modern digital products that demand immediate interactivity. Excessive centralization creates operational bottlenecks, where data engineering teams become the stranglehold for business innovation, hindering the agility needed to respond to market events or user behavior. The seven-day (D-7) latency, common in these environments for complex analytical data, makes any hyper-personalization strategy unfeasible, as the customer's context changes in seconds, not weeks.

In this context, the concept of *Data Mesh* emerges as a disruptive architectural response, proposing the decentralization of data ownership based on business domains. Instead of a monolithic central repository, the architecture distributes responsibility for the data to the teams that produce it (the domains), treating "data as a product." For a *Product Manager* focused on CRM and 360° Vision, this change is fundamental, as it allows multidisciplinary *squads* to access and consume data events autonomously and in a standardized way. This eliminates dependence on IT *ticket* queues and democratizes access to intelligence, allowing the business strategy to be fueled by fresh and reliable data.



Implementing an Event-Driven Architecture (EDA) is the technical component that operationalizes speed within the *Data Mesh*. Unlike synchronous requests (REST APIs) that couple services and can generate chain latency, EDA uses asynchronous message *brokers* (such as Apache Kafka) to propagate state changes in real time throughout the ecosystem. This means that when a client performs a transaction, simulates financing, or navigates a specific area of the application, this "event" is immediately published to a topic, becoming available for consumption by multiple subscribing services. This ability to react to events at the exact moment they occur is what allows for a drastic reduction in data availability time, going from days to milliseconds, or real-time (*online*).

The transition to real - time , however, requires a profound revision of data consistency paradigms. While traditional transactional systems (OLTP) demand strong consistency (ACID), real-time analytical and engagement systems often operate under the eventual consistency model (BASE). For the data product manager, this implies designing solutions that are resilient to replication delays and that know how to handle duplication or disorder of events. The technical robustness of this architecture is the foundation that supports the promise of a "360° View" of the customer, integrating fragments of identity dispersed across various legacy systems into a unified and coherent profile.

In addition to speed, horizontal scalability is imperative in financial ecosystems that process billions of transactions daily. Architectures based on microservices and containers (Kubernetes), orchestrated in a hybrid or public cloud, allow the computational capacity to process AI models to be dynamically allocated according to demand. This is crucial for seasonal campaigns or unforeseen market events (such as sharp stock market fluctuations or retail dates like Black Friday), where the volume of events skyrockets exponentially. The elasticity of the cloud ensures that the decision engine does not become a bottleneck, keeping the user experience fluid and responsive.

Data governance in a distributed, high-speed environment becomes exponentially more complex and critical. Ensuring the quality, lineage, and security of data in motion requires the implementation of rigorous data contracts *between* producers and consumers. The role of the *Product Manager* evolves to include defining SLOs (*Service Level Objectives*) and SLIs (*Service Level Indicators*) not only for software availability but also for data reliability and freshness. Without this structured governance, the *Data Lake* risks becoming a "Data Swamp," where information exists but is inaccessible or unreliable for strategic decision-making.

Finally, the integration of this modern architecture with market platforms, such as Salesforce, represents the final link in the value chain. It's not enough to process the data; it needs to be activated within the relationship channels. The architecture must provide robust and secure connectors that allow the *insights* generated by the AI models to be injected into automation tools.

Instant marketing. It is this efficient integration between proprietary technology, analytics, and SaaS platforms that closes the intelligence loop, transforming bits and bytes into meaningful and monetizable human interactions.

3. Causal Inference Versus Correlation: The New Frontier of Data Science

The statistical maxim that "correlation does not imply causation" is widely known, but frequently ignored in banking practice. Traditional predictive *machine learning* models are, in essence, extremely sophisticated correlation machines; they identify associative patterns in large volumes of data (e.g., "customers who buy diapers tend to buy beer"), but fail to explain the direction or cause of this association. For strategic product management, basing investment decisions or retention campaigns solely on correlations can lead to costly errors, such as offering discounts to customers who would buy the product anyway, wasting financial margin without generating real incrementality.

Causal Inference, grounded in the seminal work of statisticians and computer scientists like Judea Pearl, proposes a more rigorous approach, introducing the modeling of interventions and counterfactuals. In the banking context, the question should not simply be "what is the probability of this customer canceling their card?", but rather "what is the probability of this customer canceling their card *if I don't* offer a credit limit increase, compared to the probability *if I do*?". This nuance is fundamental for calculating the *Uplift*, that is, the true incremental gain of a marketing action. *Uplift Modeling* models segment the customer base not only by risk, but also by sensitivity to intervention, identifying the "persuadable" and avoiding the "lost cases" or "guaranteed customers".

The application of Directed Acyclic Causal Graphs (DAGs) allows data scientists and product managers to explicitly map assumptions about how variables interact in the financial ecosystem. This helps identify confounders *that* can bias the results of algorithms. For example, a campaign may appear successful in increasing credit card spending, but a causal analysis may reveal that the increase was due to external seasonality (such as Christmas) and not to the effectiveness of the communication. Without the rigor of causal inference, the bank risks attributing undue credit to its actions, perpetuating inefficient strategies.

The implementation of "event-based data strategies" takes on a new dimension when analyzed from a causal perspective. An event is not merely a trigger for an "if-then" rule, but a variable in a causal model that estimates the effect of the treatment in real time.



For example, the event "transaction denied due to insufficient funds" is a strong predictor of frustration and potential *churn*. A causal model can determine which intervention (offering an *overdraft*, suggesting installment payments, or simply sending an explanatory notification) best maximizes long-term customer retention, considering their specific history and behavioral profile.

The computational complexity of causal inference in large *datasets* is a significant challenge, requiring advanced processing power and optimized algorithms. Unlike standard supervised learning, which seeks to minimize prediction error, causal inference often grapples with the "fundamental problem of causal inference": we never observe the counterfactual outcome for the same individual (we cannot see what would have happened if we had not sent the email to the same person at the same time). Techniques such as *Double Machine Learning* and *Causal Forests* are employed to estimate these heterogeneous treatment effects, requiring a highly skilled data science team and a robust infrastructure.

Validating these models requires a culture of rigorous experimentation, with continuous A/B and multivariate testing running at scale. The *Product Manager* must lead this culture, ensuring that the product *backlog* prioritizes not only visible *features* but also experiments that refine the causal understanding of customer behavior. Evidence-based *backlog* prioritization ensures that development resources are allocated to initiatives with the highest proven probability of generating a return on investment (ROI).

Ultimately, the transition to causal inference allows the financial institution to move from a reactive to a proactive and surgical stance. By understanding the underlying mechanisms that govern customers' financial decisions, the bank can design journeys that not only sell products but also solve real and latent problems, building a relationship of trust and usefulness. It is this deep intelligence that allows for significant results, such as a 46% increase in engagement, as communication ceases to be an interruption and becomes an anticipated solution.

4. The Genai Revolution and Semantic Hyperpersonalization

Generative Artificial Intelligence (GenAI), powered by deep neural network architectures like *Transformers*, represents a technological disruption in how we interact with data and customers. Until recently, personalization at scale was limited to inserting a customer's name into an email *template* or recommending a product from a predefined list. GenAI breaks this barrier by enabling the generation of original content—

Text, image, audio — in real time, adapted not only to the demographic profile, but to the state.



Emotional, semantic context, and user interaction history. For a product manager focused on innovation, this opens up a range of unprecedented possibilities for value creation.

The application of *Large Language Models* (LLMs) in the banking context allows for the creation of dynamic "copy" (persuasive texts) that adapt to the customer's preferred tone of voice. A young, digitally savvy customer might receive a notification with informal language and emojis, while a high-income, conservative customer receives the same offer with formal, technical language. This semantic adaptation reduces the cognitive load required for message processing, increasing the likelihood of engagement and conversion. AI not only chooses the product but also "packages" the offer in the most attractive way possible for that specific individual at that specific moment.

Beyond stylistic variation, GenAI allows for the summarization and explanation of complex financial concepts in a personalized way. In an investment scenario, for example, AI can analyze a client's portfolio and generate a unique explanatory report, detailing why a particular asset performed in a certain way and suggesting rebalancing based on the user's stated life goals. This democratizes access to high-quality financial advice, previously restricted to *Private Banking clients*, scaling the service through technology. The ability to transform raw data into understandable narratives is a crucial competitive advantage in the information age.

Integrating GenAI with event-driven architecture enhances timeliness.

Imagine a scenario where a customer has a purchase declined during an international trip. The event system detects the failure, the causal model identifies the cause as a preventative security block, and GenAI immediately generates a *push* notification or a WhatsApp message explaining what happened, reassuring the customer and offering a quick unlock button with biometric authentication. This orchestrated, empathetic, and immediate response transforms a frustrating experience into a moment of "delight," increasing retention.

However, implementing GenAI in regulated environments requires extreme care to avoid model "hallucination" (generation of false or inaccurate information). Techniques such as RAG (*Retrieval-Augmented Generation*) are essential, where the generative model is restricted to consulting only internal and validated knowledge bases of the institution before formulating a response. The *Product Manager* must work closely with *prompt* engineers.

and *compliance* experts to define the *guardrails* that ensure AI operates within the ethical and legal boundaries of the financial system.

The economic efficiency of GenAI should also be monitored. The computational cost of inference in large language models can be high. Optimization strategies, such as the use of smaller, distilled models (*Small Language Models*) for specific tasks, or *caching*, should be considered.

Frequently asked questions are needed to ensure that the cost of customization does not exceed the...

Incremental revenue generated. Managing this *trade-off* between performance, quality, and cost is a vital competency for digital product leadership in the AI age.

Finally, GenAI doesn't replace human creativity, but enhances it. It frees marketing and product teams from the repetitive tasks of creating campaign variations, allowing them to focus on strategy, defining value propositions, and analyzing market trends. The symbiosis between human intuition and the processing and generation capacity of AI is what defines high-performance organizations. It is this combination that has allowed us to achieve results such as generating R\$19 million in additional revenue through automated campaigns that operate with a precision and scale unattainable by manual means.

5. Mathematical Modeling of LTV and Retention

Customer *Lifetime Value* (LTV) is the ultimate economic metric for assessing the long-term health of a customer base. At its core, LTV is the net present value of all future cash flows attributed to a customer throughout their relationship with the company.

However, in non-contractual environments such as retail banking (where the customer does not need to formally "cancel" to stop using the service, they simply stop transacting), calculating LTV becomes a complex stochastic challenge. Probabilistic models, such as the Pareto/NBD (*Negative Binomial Distribution*), are frequently used to estimate the probability of a customer still being active and to predict the number of future transactions.

Hyper-personalization directly impacts the two main variables in the LTV equation: transaction frequency (increasing daily utility and engagement) and customer lifetime value (reducing *churn*). Mathematically, small improvements in retention rate have an exponential effect on LTV, due to the compounding nature of cash flows over time. The *uplift*

The 19% increase in retention and conversion reported in the case study represents not only savings in acquisition costs, but a significant multiplication of the long-term equity value of the customer base.

Predictive behavioral segmentation allows the application of dynamic LTV models. Instead of calculating an average LTV for large segments, an individualized LTV is calculated, updated in real time with each interaction. This allows the bank to allocate incentive resources (such as *cashback*, annual fee waivers, or promotional rates) in a discriminatory and efficient manner. A customer with high potential LTV but a high risk of *churn* (e.g., a sudden decrease in app usage) justifies an aggressive investment in retention. Conversely, a customer with negative LTV should not receive incentives that only exacerbate the loss.

Modeling "critical moments of the journey" uses Markov chains or survival analysis models to identify the states in which the risk of abandonment is highest.



Intervention at these inflection points is what guarantees the "healing" of the journey. For example, the first credit card bill is a critical moment; if the customer doesn't understand the charges or has difficulty paying, the probability of *churn* skyrockets. Mathematical models identify these critical nodes and trigger GenAI to proactively intervene, explaining the bill or offering support.

Attributing revenue in multichannel and always *-on* campaigns also requires mathematical sophistication. Attribution models based on Game Theory (such as Shapley Values) or Markov chains are superior to heuristic "last-click" models because they can fairly allocate conversion credit among all touchpoints that influenced the customer's decision. This is essential to prove the ROI of data initiatives and justify continued investment in technology and analytics.

Beyond the monetary aspect, modern LTV should incorporate dimensions of social capital and influence. Customers who, in addition to being profitable, bring in new customers through referrals (network effect) possess a "viral LTV" that should be accounted for. Social graph analysis within the transaction database (e.g., PIX transfers between users) can reveal these influential nodes, enabling highly targeted and effective *Member-Get-Member* strategies .

Integrating financial data with satisfaction metrics (NPS, CSAT) into unified customer health models creates a holistic view that balances the pursuit of short-term profit with relationship sustainability. The product manager must be the guardian of this balance, using mathematics not to exploit the customer, but to maximize the exchange of mutual value, ensuring that revenue growth is a consequence of excellence in the service provided.

6. Algorithmic Governance and Ethics in Data Use

With the power of AI and hyper-personalization, ethical responsibility and the need for robust governance emerge. In a financial ecosystem, automated decisions about credit, interest rates, or product offerings have a direct impact on people's lives and economic well-being.

Algorithmic bias , where models learn and perpetuate historical biases contained in training data, is a real and unacceptable risk. AI governance must ensure that models are fair , explainable , and auditable.

The General Data Protection Law (LGPD) in Brazil and similar global regulations impose strict restrictions on how personal data can be used. Data architecture should be designed with *Privacy by Design*, ensuring anonymization and granular consent management. The customer should have transparency about why they are receiving a particular offer and the control to opt out of certain forms of data processing.



Trust is a bank's most valuable asset, and the misuse of data for excessive behavioral manipulation can irreversibly destroy that trust.

Explainability of models (XAI - *Explainable AI*) is particularly critical in complex models such as Deep Neural Networks or Decision Tree *Ensembles*. For regulatory and risk management purposes, it is unacceptable for the model to be a "black box". Techniques such as LIME (*Local Interpretable Model-agnostic Explanations*) or SHAP (*SHapley Additive exPlanations*) should be employed to translate the mathematical logic of the model into human-understandable explanations, justifying, for example, why a transaction was blocked or credit denied.

Governance also extends to the timeliness and quality of data. Outdated or incorrect data can lead to automated wrong decisions at scale, causing massive damage before they can be manually corrected. Continuous monitoring of data health (*Data Observability*) and *circuit breakers* (mechanisms that interrupt automations if they detect anomalies) are essential components of a secure and resilient AI infrastructure.

The role of the *Product Manager* in this context transcends technical aspects and delves into ethics. It is the responsibility of product leadership to define the ethical boundaries of persuasion. To what extent is it acceptable to use behavioral triggers to encourage consumption or credit acquisition? The "choice architecture" (*nudging*) designed by the algorithms should always aim for the client's financial well-being, aligning the bank's incentives with the user's financial health.

Cybersecurity is another pillar of governance. The centralization of intelligence and detailed behavioral data creates an attractive target for attacks. *Data Mesh* and microservices architecture helps segment risk, but security must permeate all layers, from event ingestion to model inference at the edge. Efficient integration between technology and business strategies presupposes that security is not a blocker, but an enabler of sustainable innovation.

Finally, algorithmic governance requires an organizational structure that supports shared responsibility. AI ethics committees, composed of technicians, lawyers, sociologists, and business executives, are needed to assess the social and individual impacts of new technologies before their large-scale *deployment*. Responsible innovation is the only path to longevity in the highly regulated and monitored financial sector.

7. CASE STUDY: QUANTITATIVE RESULTS AND EMPIRICAL ANALYSIS

The validation of the theses presented in this article is based on the analysis of results obtained through leading product *squads* focused on data and customer centricity within an institution.

A large financial institution. The implementation of a centralized platform for event consumption and 360° visibility was the driving force behind the transformation, reducing the time it takes for data to be made available to business areas from 7 days to real-time (*online*). This structural change was not just a technical improvement, but an unlocker of strategic value.

The first indicator of success was engagement with communications. Replacing "ruler" campaigns (based on fixed dates) with strategies based on behavioral segmentation and data events resulted in a 46% increase in engagement with transactional messages. Qualitative analysis indicated that contextual relevance was the determining factor; receiving a travel insurance offer at the exact moment the card is used to purchase an airline ticket is incomparably more effective than receiving the same offer in a generic monthly *newsletter*.

In financial terms, the orchestration of automated campaigns and personalization at scale generated R\$19 million in additional revenue in the first year of operation. This value was measured through rigorous control groups, ensuring that the attributed revenue was, in fact, incremental (causal) and not organic. The return on investment (ROI) of the data infrastructure proved to be highly positive, validating the allocation of capital in *Big Data* and AI technologies.

Customer retention, a critical metric for LTV, also experienced a significant positive impact.

The structuring of analytical models that enabled intelligent activation at critical moments in the customer journey increased conversion and retention by 19%. This demonstrates that the technology served as a safety net, capturing dissatisfied customers or those experiencing operational difficulties and offering proactive solutions before the desire to cancel became entrenched.

Backlog management and strategic prioritization were key to achieving these results.

The decision to focus first on high-volume, high-friction use cases ensured quick wins *that* helped fund and legitimize the expansion of the data program. The agile culture and the connection of multidisciplinary teams allowed for rapid iteration on models, adjusting parameters and creatives based on real market feedback.

The efficient integration of technology, analytics, and business strategies was the key to success. There were no silos; data engineers, data scientists, and business analysts worked with shared objectives and unified metrics. This organizational synergy is as important as the technological architecture; without it, the most sophisticated models remain academic exercises without impact on the "real world."

In short, the quantified results serve as robust empirical evidence that the application of causal inference and modern data architectures is not futurism, but a present and profitable necessity. The ability to transform terabytes of raw data into personalized experiences.

Generating millions in revenue is the contemporary definition of competitive advantage in the banking sector.

8. CONCLUSION

The evolutionary trajectory of digital product management in the financial sector unequivocally points to the primacy of data as a central strategic asset. This study demonstrated, through theoretical foundations and empirical evidence, that the transition from static demographic models to a predictive and causal behavioral approach is not merely an incremental improvement, but a complete redefinition of the banking value proposition. The integration of *Data Mesh*, Event-Driven Architectures, and Generative Artificial Intelligence forms the technological tripod that supports this new era of hyper-personalization.

The results presented—a drastic reduction in information latency, a significant *uplift* in engagement, and substantial incremental revenue generation—support the hypothesis that timeliness and contextual relevance are the determining variables for success in digital customer relationships. Technology, when properly orchestrated, ceases to be an operational cost center and becomes the primary engine for revenue generation and value retention.

However, technological sophistication brings with it challenges of equal magnitude in the field of governance and ethics. The ability to influence behavior through algorithms imposes on financial institutions the fiduciary duty to use this power responsibly, transparently, and fairly. Product leadership plays a crucial role as an ethical arbiter, ensuring that the pursuit of maximizing LTV does not compromise user trust and privacy.

The application of Causal Inference has proven essential in separating the signal from the noise, allowing for more efficient marketing investments and avoiding the waste of resources on spurious correlations. In the near future, the adoption of Autonomous AI Agents is expected to take personalization to a new level, where virtual assistants not only recommend products but also execute complex financial strategies on behalf of the client, autonomously and in an optimized way.

The professional experience analyzed in this study, combining solid technical training in Information Systems with cutting-edge certifications in AI and Analytics, exemplifies the profile of the professional needed to lead this transformation: a hybrid "translator" capable of navigating fluidly between the complexity of stochastic algorithms and the strategic imperatives of the business.

It can be concluded, therefore, that hyper-personalization at scale is not a final destination, but a continuous process of learning and adaptation. The institutions that master the art of listening to, interpreting, and responding to their customers' signals in real time, through a resilient and intelligent data infrastructure, will be the ones that survive and thrive in the digital economy. Customer lifetime value (LTV) will be maximized not by aggressive sales, but by unquestionable relevance and constant utility in every micro-moment of the financial journey.

The findings of this work suggest fertile avenues for future research, especially at the intersection of behavioral economics and reinforcement learning, exploring how AI systems can learn optimal incentive policies through continuous interaction with the market environment. The data revolution in the banking sector is only just beginning, and its economic and social impacts will continue to reverberate for decades to come.

9. BIBLIOGRAPHIC REFERENCES

1. AGARWAL, Ajay; GANS, Joshua; GOLDFARB, Avi. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, 2018.
2. ANDERSON, Chris. *The Long Tail: Why the Future of Business is Selling Less of More*. Hyperion, 2006.
3. DEGHANI, Zhamak. *Data Mesh: Delivering Data-Driven Value at Scale*. O'Reilly Media, 2022.
4. FADER, Peter S.; HARDIE, Bruce GS "Probability Models for Customer-Base Analysis". *Journal of Interactive Marketing*, v. 23, n. 1, p. 61-69, 2009.
5. GOODFELLOW, Ian; BENGIO, Yoshua; COURVILLE, Aaron. *Deep Learning*. MIT Press, 2016.
6. GUPTA, Sunil; LEHMANN, Donald R. *Managing Customers as Investments: The Strategic Value of Customers in the Long Run*. Wharton School Publishing, 2005.
7. HINTON, Geoffrey. "Deep Learning—A Technology With the Potential to Transform Health Care." *JAMA*, vol. 320, no. 11, p. 1101–1102, 2018.
8. IMBENS, Guido W.; RUBIN, Donald B. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015.
9. KLEPMANN, Martin. *Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems*. O'Reilly Media, 2017.
10. KOTLER, Philip; KARTAJAYA, Hermawan; SETIAWAN, Iwan. *Marketing 4.0: Moving from Traditional to Digital*. Wiley, 2016.
11. NEWMAN, Sam. *Building Microservices: Designing Fine-Grained Systems*. O'Reilly Media, 2015.
12. OSTERWALDER, Alexander; PIGNEUR, Yves. *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. Wiley, 2010.



13. PEARL, Judea. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge University Press, 2009.
14. PEARL, Judea; MACKENZIE, Dana. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.
15. PORTER, Michael E.; HEPPELMANN, James E. "How Smart, Connected Products Are Transforming Competition." *Harvard Business Review*, Nov. 2014.
16. RIES, Eric. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business, 2011.
17. RUSSELL, Stuart; NORVIG, Peter. *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson, 2020.
18. SCHWAB, Klaus. *The Fourth Industrial Revolution*. World Economic Forum, 2016.
19. THRUN, Sebastian; PRATT, Lorien (Eds.). *Learning to Learn*. Springer, 1998.
20. VASWANI, Ashish et al. "Attention Is All You Need". *Advances in Neural Information Processing Systems*, vol. 30, 2017.