

RESEARCH ARTICLE

Intelligent Customer Propensity Modeling and Causal Decision Engines: Integrating Machine Learning, Explainable Artificial Intelligence, And Financial Technology for Predictive Decision-Making

Aoi Suzuki

School of Computational Finance and Artificial Intelligence, Waseda University, Tokyo, Japan

VOLUME: Vol.06 Issue01 2026

PAGE: 207-214

Copyright © 2026 European International Journal of Multidisciplinary Research and Management Studies, this is an open-access article distributed under the terms of the Creative Commons Attribution-Noncommercial-Share Alike 4.0 International License. Licensed under Creative Commons License a Creative Commons Attribution 4.0 International License.

Abstract

The increasing digitization of financial services, marketing systems, and customer relationship platforms has produced unprecedented volumes of behavioral and transactional data. Organizations across industries now seek advanced analytical methods capable of transforming these data streams into actionable intelligence for predicting customer behavior, estimating response probabilities, and optimizing decision-making processes. Machine learning and artificial intelligence have emerged as critical technologies enabling the development of predictive decision engines capable of modeling customer propensity, churn, risk, and engagement patterns. Despite significant progress, several methodological challenges remain, including issues related to causal inference, model transparency, explainability, and operational deployment in financial ecosystems. This research investigates the integration of machine learning-based propensity prediction with causal inference methodologies and explainable artificial intelligence to develop robust decision engines for financial and digital service environments.

The study synthesizes theoretical insights from machine learning frameworks, causal inference literature, uplift modeling approaches, and financial technology research to propose a comprehensive predictive architecture capable of supporting decision-making under uncertainty. Drawing upon contemporary developments in predictive analytics, propensity scoring, uplift modeling, explainable AI, and data-driven decision science, the article explores how organizations can construct intelligent systems that not only predict customer behavior but also estimate the causal effects of interventions such as marketing campaigns, credit offers, and financial recommendations. The research further discusses how modern machine learning platforms, including open-source ecosystems such as Python-based data science libraries, enable scalable deployment of predictive models in real-world financial infrastructures.

The findings suggest that combining causal modeling frameworks with machine learning architectures significantly enhances the interpretability, robustness, and strategic value of predictive systems. In addition, explainable AI techniques can improve regulatory compliance and transparency within financial institutions by providing insights into model behavior and decision logic. The study concludes by outlining a future research agenda focusing on the integration of agentic artificial intelligence, retrieval-augmented knowledge systems, and blockchain-enabled financial infrastructures for next-generation predictive decision engines.

KEY WORDS

Medical simulation, patient safety, anesthesia education, deliberate practice, crisis resource management, artificial intelligence.

INTRODUCTION

The global economy has entered an era characterized by the widespread digitization of business operations, financial services, and consumer interactions. Organizations across sectors increasingly rely on data-driven decision-making to optimize customer engagement, manage financial risk, and enhance operational efficiency. This transformation has been driven by advances in computational infrastructure, big data technologies, and artificial intelligence algorithms capable of analyzing complex behavioral patterns in large-scale datasets. Machine learning has become one of the most influential technologies within this digital transformation landscape, enabling predictive modeling capabilities that were previously unattainable using traditional statistical techniques (Ahmed et al., 2022).

Financial institutions, marketing organizations, and digital platforms increasingly employ predictive analytics to understand customer behavior, estimate the likelihood of future events, and design personalized interventions aimed at improving engagement and profitability. One of the most prominent applications of predictive analytics is customer propensity modeling, which involves estimating the probability that an individual will perform a specific action, such as purchasing a product, responding to a marketing campaign, defaulting on a loan, or discontinuing a service subscription. These predictive capabilities allow organizations to allocate resources more efficiently, tailor their offerings to individual customer needs, and design targeted strategies for retention and acquisition (Provost & Fawcett, 2013).

The concept of customer propensity prediction has evolved significantly over the past several decades. Early predictive models relied primarily on classical statistical methods such as logistic regression and discriminant analysis. These techniques provided valuable insights into customer behavior but were limited in their ability to capture nonlinear relationships, complex feature interactions, and high-dimensional data structures. As machine learning technologies matured, more sophisticated predictive models emerged, including ensemble learning methods, neural networks, and deep learning architectures capable of modeling complex behavioral patterns across diverse data sources (Ahmed et al., 2022).

In modern financial and marketing environments, predictive decision systems often operate within large-scale digital infrastructures that integrate transactional data, behavioral logs, demographic information, and contextual signals. The availability of such rich datasets enables organizations to build advanced decision engines capable of estimating customer responses to various interventions, including targeted promotions, financial product recommendations, and personalized service offerings. These decision engines have become central components of contemporary customer relationship management systems and digital marketing platforms (Ngai & Wu, 2022).

Despite these advancements, several challenges remain in the design and deployment of predictive decision systems. One of the most significant challenges relates to the distinction between correlation-based prediction and causal inference. Traditional machine learning models are primarily designed to identify statistical patterns within historical data, enabling them to predict future outcomes based on observed correlations. However, such models do not necessarily capture causal relationships between variables, which limits their ability to estimate the effects of interventions such as marketing campaigns or credit policy changes (Robins et al., 2000).

Causal inference methodologies have therefore gained increasing attention in recent years as organizations seek to develop predictive systems capable of estimating not only what is likely to happen but also what would happen under different decision scenarios. Propensity score modeling represents one of the most widely used techniques for addressing causal inference challenges in observational data environments. By estimating the probability that an individual receives a particular treatment or intervention, propensity scores enable researchers to construct balanced comparison groups and estimate causal effects more reliably (Rosenbaum & Rubin, 1983).

In addition to causal inference challenges, another critical issue facing predictive decision systems is model transparency and interpretability. Many advanced machine learning

algorithms, particularly deep learning models, operate as complex computational structures whose internal decision logic can be difficult for human analysts to interpret. This lack of transparency raises significant concerns in regulated industries such as finance and insurance, where organizations must demonstrate accountability and fairness in their decision-making processes (Alapati & Valleru, 2023).

Explainable artificial intelligence has emerged as a promising approach for addressing these challenges by providing tools and methodologies that enable analysts to understand how machine learning models generate predictions. By revealing the relative importance of input features, identifying decision pathways, and detecting potential biases, explainable AI techniques contribute to the development of transparent and trustworthy predictive systems (AI Shiam et al., 2024).

Upadhyay (2025) emphasized that conversational AI, personalization engines, and voice AI have emerged as critical components of modern customer experience strategies. These technologies enable businesses to provide real-time assistance, personalized recommendations, and seamless interactions, thereby improving customer satisfaction and operational performance. The findings suggest that AI-based customer engagement tools contribute significantly to competitive advantage in digital markets.

This distinction between prediction and causal impact is particularly important in marketing and financial decision-making contexts. For example, a customer who is already likely to purchase a product may respond to a promotional campaign even without receiving the promotion. Targeting such customers may therefore generate limited incremental value. Uplift modeling helps organizations identify customers whose behavior can be influenced by targeted interventions, thereby improving the efficiency and effectiveness of marketing strategies (Radcliffe, 2007).

Recent advances in artificial intelligence have further expanded the capabilities of predictive decision systems. Developments in deep learning, agentic AI architectures, retrieval-augmented knowledge systems, and blockchain-enabled financial infrastructures have created new opportunities for integrating predictive analytics with automated decision-making platforms. These technologies are reshaping the financial services industry by enabling real-time analytics, personalized customer interactions, and automated

operational processes (Parker, 2023).

Furthermore, financial technology innovations have introduced new forms of digital infrastructure that facilitate data sharing, transaction automation, and algorithmic decision-making. Blockchain technology, for example, provides decentralized systems for recording financial transactions, while artificial intelligence algorithms enable automated risk assessment and fraud detection. The convergence of these technologies is transforming the structure of modern financial ecosystems and creating new opportunities for data-driven decision-making (Ahluwalia et al., 2020).

Given these developments, there is a growing need for integrated research frameworks that combine machine learning prediction, causal inference methodologies, explainable AI techniques, and financial technology infrastructures. Such frameworks would enable organizations to develop predictive decision engines capable of supporting strategic decision-making while maintaining transparency, reliability, and regulatory compliance.

This research aims to address this need by exploring the theoretical foundations and practical implications of integrating machine learning-based propensity prediction with causal inference models and explainable artificial intelligence. By synthesizing insights from diverse research domains, the study proposes a comprehensive architecture for predictive decision systems capable of supporting intelligent decision-making in complex financial and digital environments.

METHODOLOGY

The methodological framework of this research is designed to integrate theoretical insights from machine learning, causal inference, financial analytics, and explainable artificial intelligence into a unified predictive decision engine architecture. The objective is not to test a specific empirical dataset but rather to develop a conceptual research framework that explains how modern predictive systems can be designed to estimate customer propensity, identify causal effects of interventions, and support decision-making processes within financial and digital service environments.

The methodological approach draws upon established principles of data science, predictive analytics, and decision theory. The foundation of the framework is the data-driven decision-making paradigm described in contemporary data

science literature, which emphasizes the integration of statistical analysis, computational modeling, and domain knowledge to extract actionable insights from complex datasets (Provost & Fawcett, 2013). Within this paradigm, predictive modeling is viewed as a systematic process involving data collection, feature engineering, model training, evaluation, and deployment.

A widely recognized process model for structuring predictive analytics projects is the Cross-Industry Standard Process for Data Mining, commonly known as CRISP-DM. This framework organizes the development of predictive models into several iterative phases that include business understanding, data understanding, data preparation, modeling, evaluation, and deployment. The CRISP-DM methodology provides a structured workflow that ensures alignment between analytical objectives and business requirements while facilitating systematic experimentation and model refinement (Wirth & Hipp, 2000).

In the context of customer propensity modeling, the methodological process begins with the identification of relevant data sources that capture customer interactions, demographic characteristics, transactional histories, and behavioral patterns. These data sources often include structured datasets such as financial transaction records and customer account information, as well as unstructured data such as digital interaction logs, text communications, and online browsing behaviors. The integration of diverse data sources is essential for capturing the multidimensional nature of customer behavior in modern digital environments (Dumais et al., 2014).

Once the data sources have been identified, the next step involves data preprocessing and feature engineering. Feature engineering is a critical component of predictive modeling because the quality and relevance of input features significantly influence the performance of machine learning algorithms. In customer propensity modeling, features may include demographic attributes, historical purchase patterns, frequency of interactions with digital platforms, and responses to previous marketing campaigns. The process of feature engineering often requires domain expertise to identify variables that capture meaningful behavioral signals (Nimmagadda, 2022).

Machine learning algorithms are then applied to the prepared dataset to estimate the probability that each customer will

exhibit a specific behavior. Various machine learning models can be employed for this purpose, including decision trees, ensemble methods, neural networks, and support vector machines. The selection of appropriate algorithms depends on factors such as dataset size, feature dimensionality, computational constraints, and interpretability requirements.

Open-source machine learning libraries such as Python-based data science platforms provide powerful tools for implementing predictive models at scale. These platforms support a wide range of machine learning algorithms and enable researchers to conduct experiments, evaluate model performance, and deploy predictive systems in production environments (Pedregosa et al., 2011).

However, predictive modeling alone is insufficient for understanding the causal impact of interventions on customer behavior. To address this limitation, the methodological framework incorporates causal inference techniques that enable analysts to estimate treatment effects using observational data. One of the most widely used approaches in causal inference is the propensity score method, which involves estimating the probability that an individual receives a particular treatment based on observed characteristics (Rosenbaum & Rubin, 1983).

Propensity score modeling allows researchers to construct balanced comparison groups that approximate the conditions of randomized experiments. By matching individuals with similar propensity scores, analysts can compare outcomes between treated and untreated groups while reducing the influence of confounding variables. This approach enables more accurate estimation of causal effects in situations where controlled experiments are not feasible (Rosenbaum & Rubin, 1985).

Another important methodological component is uplift modeling, which focuses on estimating the incremental effect of an intervention on customer behavior. Unlike traditional response models, uplift models aim to identify individuals whose behavior changes as a result of a specific action. This approach allows organizations to target interventions more effectively by focusing on individuals who are most likely to respond positively to a given treatment (Radcliffe & Surry, 1999).

Recent research has also explored the use of ensemble learning techniques to improve the stability and accuracy of

uplift models. Ensemble methods combine predictions from multiple models to produce more robust estimates of treatment effects. Weighted ensemble approaches have been proposed to reduce volatility in uplift predictions and enhance the reliability of decision-making systems (Rößler et al., 2021).

The methodological framework further incorporates explainable artificial intelligence techniques designed to improve model transparency and interpretability. Explainable AI methods analyze the internal structure of machine learning models to identify the features and interactions that contribute to specific predictions. These techniques enable analysts to understand how models generate outputs and detect potential biases or inconsistencies in the decision process (Alapati & Valleru, 2023).

The growing emphasis on omnichannel retailing has encouraged many DTC brands to reconsider wholesale partnerships. According to Shounik (2025), wholesale re-expansion serves as an effective risk-management strategy by reducing operational uncertainties and enhancing business stability. These findings indicate that integrating multiple sales channels can strengthen overall brand performance and customer accessibility.

RESULTS

The conceptual implementation of the proposed predictive decision engine demonstrates how integrating machine learning, causal inference, and explainable artificial intelligence can significantly enhance the quality and reliability of customer propensity predictions. Although the present research does not rely on a single empirical dataset, the theoretical framework allows us to examine the expected performance characteristics and operational implications of such systems within financial and digital service environments.

One of the most significant outcomes of the proposed framework is the improved ability to distinguish between correlation-based predictions and causal treatment effects. Traditional predictive models typically estimate the probability that a customer will perform a specific action based on historical patterns observed in data. While these predictions can be useful for forecasting behavior, they do not necessarily reveal whether an intervention would change that behavior. The integration of propensity score modeling addresses this

limitation by enabling analysts to construct balanced observational comparisons that approximate experimental conditions (Rosenbaum & Rubin, 1983).

When predictive machine learning models are combined with propensity score methodologies, organizations gain the ability to estimate not only the likelihood of customer responses but also the causal impact of targeted actions. For example, a financial institution offering a credit product can analyze whether a promotional offer genuinely increases the probability of customer adoption or whether the observed responses are simply driven by preexisting customer preferences. This distinction is crucial for optimizing marketing strategies and ensuring that promotional resources are allocated efficiently.

Another important result of the framework is the enhancement of decision-making precision through uplift modeling. Traditional response models identify customers who are most likely to respond to marketing campaigns, but they do not differentiate between customers who would respond regardless of the campaign and those whose behavior is actually influenced by the intervention. Uplift modeling addresses this challenge by estimating the incremental effect of marketing actions on individual customers, allowing organizations to identify those whose behavior is most likely to change as a result of targeted efforts (Radcliffe, 2007).

The application of uplift modeling within the predictive decision engine significantly improves the efficiency of marketing campaigns and customer engagement strategies. By focusing on customers with the highest predicted uplift, organizations can reduce unnecessary expenditures on individuals who are unlikely to change their behavior. This targeted approach leads to more effective resource allocation and higher overall return on investment for marketing initiatives.

Another key finding relates to the role of explainable artificial intelligence in enhancing transparency and trust within predictive systems. As machine learning models become more complex, their internal decision processes can become difficult for human analysts to interpret. This lack of transparency poses significant challenges in regulated industries such as finance and insurance, where decision-making processes must be auditable and compliant with regulatory standards.

Explainable AI techniques address this challenge by providing

mechanisms for interpreting model predictions and identifying the factors that influence decision outcomes. For example, feature attribution methods can reveal which variables contribute most strongly to a particular prediction, enabling analysts to verify that the model is relying on meaningful signals rather than spurious correlations. Such insights are essential for ensuring that predictive systems operate in a fair and responsible manner (Al Shiam et al., 2024).

The results also highlight the importance of scalable computational infrastructure for implementing predictive decision engines. Modern machine learning frameworks, including open-source platforms for data analysis and modeling, enable organizations to train and deploy predictive models on large datasets containing millions of customer records. These platforms support advanced algorithms and provide tools for model evaluation, cross-validation, and performance optimization (Pedregosa et al., 2011).

Within financial services environments, predictive decision engines can be integrated into various operational workflows, including credit risk assessment, fraud detection, and customer relationship management. For example, financial institutions can use predictive models to estimate the likelihood that a borrower will default on a loan, allowing them to adjust credit terms accordingly. Similarly, predictive analytics can identify patterns indicative of fraudulent transactions, enabling organizations to detect and prevent financial crimes in real time (Ali et al., 2022).

Another important result concerns the integration of predictive analytics with emerging financial technologies such as blockchain and decentralized finance platforms. Blockchain technology provides a secure and transparent infrastructure for recording financial transactions, while artificial intelligence algorithms enable automated analysis of transactional data. The combination of these technologies creates opportunities for developing intelligent financial systems capable of managing risk, verifying identities, and facilitating secure digital transactions (Ahluwalia et al., 2020).

The research framework also demonstrates the potential benefits of integrating agentic artificial intelligence systems into predictive decision engines. Agentic AI refers to autonomous software systems capable of performing complex tasks, making decisions, and interacting with digital environments without continuous human supervision. When combined with predictive analytics, such systems can

dynamically adapt to changing market conditions and customer behaviors, enabling organizations to respond more effectively to emerging opportunities and risks (World Economic Forum, 2025).

DISCUSSION

The integration of machine learning, causal inference, and explainable artificial intelligence represents a significant advancement in the development of predictive decision systems for financial and digital service environments. The findings of this research suggest that combining these methodological approaches can substantially enhance the accuracy, interpretability, and strategic value of predictive analytics.

One of the most important implications of the proposed framework is the recognition that predictive accuracy alone is insufficient for effective decision-making. Traditional machine learning models often focus exclusively on maximizing predictive performance, which can lead to models that identify statistical patterns without capturing causal relationships. While such models may generate accurate forecasts, they may also produce misleading insights when used to evaluate the impact of interventions.

Causal inference methodologies provide a critical complement to predictive modeling by enabling analysts to estimate treatment effects and understand how specific actions influence outcomes. The integration of propensity score methods and uplift modeling into predictive decision engines allows organizations to move beyond simple prediction and toward causal reasoning. This shift represents a fundamental transformation in the way organizations use data to guide strategic decisions (Robins et al., 2000).

Another significant implication concerns the growing importance of transparency and accountability in algorithmic decision-making. As artificial intelligence systems become more deeply embedded in financial and social infrastructures, concerns about algorithmic bias, fairness, and ethical responsibility have become increasingly prominent. Explainable AI techniques provide valuable tools for addressing these concerns by enabling analysts to understand and evaluate the behavior of complex machine learning models (Alapati & Valleru, 2023).

However, the implementation of explainable AI also raises several challenges. Some interpretability techniques provide

only approximate explanations of model behavior, which may not fully capture the underlying decision logic. Additionally, there may be trade-offs between model complexity and interpretability, as highly accurate models often rely on intricate computational structures that are difficult to explain in simple terms. Future research should therefore focus on developing new methods for balancing predictive performance with interpretability.

Another limitation of predictive decision engines relates to the quality and representativeness of the underlying data. Machine learning models are highly sensitive to biases present in training datasets. If the data used to train a model contain systematic biases or incomplete information, the resulting predictions may perpetuate or amplify those biases. This issue is particularly relevant in financial services, where biased models could lead to discriminatory lending practices or unequal access to financial opportunities (Ahmed et al., 2022).

Addressing these challenges requires rigorous data governance practices, including data auditing, bias detection, and continuous monitoring of model performance. Organizations must also implement ethical guidelines and regulatory compliance frameworks to ensure that predictive systems operate in a fair and responsible manner.

Looking forward, several emerging technological developments are likely to influence the evolution of predictive decision engines. Retrieval-augmented generation systems represent one such innovation, combining machine learning models with external knowledge retrieval mechanisms to improve the accuracy and contextual relevance of predictions. These systems enable predictive models to access and integrate information from large knowledge repositories, thereby enhancing their ability to generate informed recommendations (Merritt, 2023).

Agentic artificial intelligence systems represent another promising direction for future research. These systems are designed to perform complex tasks autonomously by interacting with digital environments and learning from feedback. In the context of predictive analytics, agentic AI could enable decision engines to continuously refine their models based on real-time data and adapt their strategies in response to changing conditions (OpenAPIHub, 2024).

Blockchain technology may also play a significant role in the future of predictive analytics by providing secure and

transparent mechanisms for managing data and verifying transactions. Decentralized data infrastructures could facilitate collaboration between organizations while preserving data privacy and security.

CONCLUSION

The rapid expansion of digital technologies and data-driven business practices has created unprecedented opportunities for organizations to leverage predictive analytics in decision-making processes. Machine learning has emerged as a powerful tool for analyzing complex datasets and identifying patterns that can inform strategic actions. However, the effectiveness of predictive systems depends not only on their ability to generate accurate predictions but also on their capacity to provide interpretable insights and estimate the causal impact of interventions.

This research has presented a comprehensive conceptual framework for integrating machine learning, causal inference methodologies, and explainable artificial intelligence into predictive decision engines designed for financial and digital service environments. By combining predictive modeling techniques with propensity score methods, uplift modeling, and interpretability tools, organizations can develop more robust and transparent systems capable of supporting intelligent decision-making.

The findings highlight several key benefits of this integrated approach, including improved marketing efficiency, enhanced risk management capabilities, greater transparency in algorithmic decision-making, and stronger alignment with regulatory requirements. At the same time, the research identifies important challenges related to data quality, model bias, interpretability limitations, and the complexity of deploying predictive systems at scale.

Future research should explore the integration of emerging technologies such as agentic artificial intelligence, retrieval-augmented knowledge systems, and blockchain-enabled financial infrastructures. These innovations have the potential to further enhance the capabilities of predictive decision engines and enable new forms of intelligent automation within financial ecosystems.

Ultimately, the continued evolution of predictive analytics will depend on the ability of researchers and practitioners to develop methodologies that combine computational power with rigorous scientific reasoning and ethical responsibility. By

integrating machine learning with causal inference and explainable AI, organizations can move toward a new generation of intelligent decision systems that not only predict the future but also understand the mechanisms that shape it.

REFERENCES

1. Accenture. (2024). Pulse of Change survey, May 2024.
2. Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. *Research in International Business and Finance*, 61.
3. Ahluwalia, S., Mahto, R. V., & Guerrero, M. (2020). Blockchain technology and startup financing: A transaction cost economics perspective. *Technological Forecasting and Social Change*, 151.
4. Al Shiam, S. A., Hasan, M. M., Pantho, M. J., Shochona, S. A., Nayeem, M. B., Choudhury, M. T. H., & Nguyen, T. N. (2024). Credit risk prediction using explainable AI. *Journal of Business Management Studies*, 6(2), 61–66.
5. Alapati, N. K., & Valleru, V. (2023). The impact of explainable AI on transparent decision-making in financial systems. *Journal of Innovation and Technology*, 6(1).
6. Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences*, 12(19).
7. Dumais, S., Jeffries, R., Russell, D., Tang, D., & Teevan, J. (2014). Understanding user behavior through log data and analysis. Springer.
8. Krishnan, G., Bhat, A. K., & Shah, J. (2025). Decision engine: Propensity prediction in the financial industry based on customer data features. In *Artificial Intelligence and Sustainable Innovation* (pp. 107-112). CRC Press.
9. Upadhyay, H. (2025). Consumer Experience Trends Based on AI Features: A Comprehensive Analysis of Conversational AI, Personalization Engines, and Voice AI. *Frontiers in Emerging Artificial Intelligence and Machine Learning*, 2(11), 6–15.
<https://doi.org/10.64917/feaiml/Volume02Issue11-02>
10. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12, 2825–2830.
11. Pondel, M., Wuczyński, M., Gryncewicz, W., Łysik, Ł., Hernes, M., Rot, A., & Kozina, A. (2021). Deep learning for customer churn prediction in e-commerce decision support. *Business Information Systems*.
12. Provost, F., & Fawcett, T. (2013). Data science for business: What you need to know about data mining and data-analytic thinking.
13. Radcliffe, N. (2007). Using control groups to target on predicted lift: Building and assessing uplift model. *Direct Marketing Analytics Journal*.
14. Radcliffe, N., & Surry, P. (1999). Differential response analysis: Modeling true responses by isolating the effect of a single action.
15. Robins, J. M., Hernán, M. A., & Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5), 550–560.
16. Rosenbaum, P. R. (1991). A characterization of optimal designs for observational studies. *Journal of the Royal Statistical Society: Series B*, 53(3), 597–610.
17. Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.
18. Rosenbaum, P. R., & Rubin, D. B. (1985). Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39(1), 33–38.
19. Rossi, P. E., & Allenby, G. M. (2003). Bayesian statistics and marketing. *Marketing Science*, 22(3), 304–328.
20. Röbber, J., Roman, T., & Detlef, S. (2021). To treat, or not to treat: Reducing volatility in uplift modeling through weighted ensembles. *Proceedings of the Hawaii International Conference on System Sciences*.
21. Shounik, S. (2025). The Great DTC Reset as Stress Management: Evidence that Wholesale Re-Expansion Reduces "Operating Tail Risk" in Consumer Brands.. *Advances in Consumer Research*, 2(6), 1221-1231.