JOURNAL OF SOCIAL SCIENCES, HUMANITIES AND INNOVATION

Vol 2 Issue 1 (2026)

http://jsshi.org/

Causal Uplift Modeling for Grocery Loyalty Coupon Targeting Among Persuadable Grocery Shoppers

Yerin Bae, Eleanor Hayes, Geonwoo Han, Logan Walsh

Faculty of Science, University of Melbourne, Melbourne, Australia

Keywords:

Causal Inference, Uplift Modeling, Grocery Retail, Loyalty Programs, Targeted Marketing

Abstract:

The grocery retail sector operates on notoriously thin profit margins, making the efficient allocation of promotional budgets a critical imperative. Traditional targeted marketing relies heavily on predictive response models that estimate the probability of a customer redeeming a coupon. However, these models inherently conflate correlation with causation, often targeting customers who would have made a purchase regardless of the promotion. This paper investigates the application of causal uplift modeling to optimize loyalty coupon targeting by directly estimating the incremental impact of a promotional intervention on customer purchase behavior. Utilizing observational and experimental data from a large-scale grocery loyalty program, we evaluate the efficacy of various meta-learning architectures and tree-based uplift algorithms in isolating the persuadable customer segment. Our comprehensive analysis demonstrates that uplift modeling significantly outperforms traditional propensity models in generating incremental revenue and return on investment. By minimizing marketing waste associated with the treatment of sure things and avoiding negative impacts on sleeping dogs, causal models provide a more rigorous framework for promotional decision-making. The findings offer substantial managerial implications for retail marketers seeking to transition from predictive to prescriptive analytics, ultimately enhancing the profitability and sustainability of customer relationship management strategies in highly competitive retail environments.

1.INTRODUCTION

The contemporary grocery retail industry is characterized by intense competition, macroeconomic pressures, and structural shifts in consumer behavior. With average profit margins frequently hovering in the single digits, retail executives are under immense pressure to optimize every facet of their operational and marketing expenditures. Customer loyalty programs have historically served as the cornerstone of customer retention strategies, providing retailers with vast repositories of transaction data and a direct channel for promotional communication. However, the prevailing approach to leveraging this data for targeted marketing remains suboptimal. For decades, the industry standard has relied on mass promotions or rudimentary segmentation techniques that treat large swaths of the customer base homogeneously [1]. While these methods generate temporary spikes in sales volume, they often do so at the expense of long-term profitability, as discounts are frequently distributed to customers whose purchasing decisions were already solidified. This phenomenon, known as promotional cannibalization, represents a significant source of margin erosion in the retail sector [2]. As data science capabilities have matured, grocery retailers have increasingly adopted machine learning techniques to refine their targeting strategies. The dominant paradigm in this evolution has been the deployment of predictive response models, commonly referred to as propensity models. These models utilize historical transaction data to predict the likelihood that a given customer will engage with a specific promotion, such as redeeming a discount coupon [3]. By targeting individuals with the highest probability of response, retailers aim to maximize the conversion rate of their campaigns. Nevertheless, this predictive approach harbors a fundamental flaw: it fails to distinguish between customers who are motivated to purchase because of the coupon and those who were going to purchase anyway. Consequently, propensity models systematically prioritize individuals who are already highly loyal and frequent shoppers, thereby subsidizing transactions that require no financial incentive [4].

To address the inherent limitations of predictive response modeling, the focus of academic research and advanced industrial practice has shifted toward causal inference, specifically through the application of uplift modeling. Uplift modeling, also known as incremental modeling or true response modeling, is a predictive modeling technique designed to directly estimate the causal effect of a treatment at the individual level [5]. In the context of grocery retail coupons, the treatment is the distribution of the coupon, and the outcome is the subsequent purchase behavior. The objective of uplift modeling is to isolate the subset of customers whose behavior is positively altered by the intervention. This requires moving beyond traditional classification tasks and entering the realm of counterfactual reasoning, where the algorithm must estimate the difference between two mutually exclusive outcomes: what happens if the customer receives the coupon, versus what happens if they do not. The theoretical framework for uplift modeling categorizes consumers into four distinct behavioral segments based on their response to a treatment. The first segment consists of the persuadables, who only exhibit the desired behavior if they receive the promotion. The second segment encompasses the sure things, who will exhibit the desired behavior regardless of whether they receive the promotion. The third segment includes the lost causes, who will not exhibit the behavior under any circumstances. Finally, the fourth segment contains the sleeping dogs, or do-not-disturbs, who will react negatively to the promotion, perhaps out of annoyance or brand fatigue. From a financial perspective, the optimal targeting strategy allocates resources exclusively to the persuadables, as this is the only segment where the promotion generates an incremental lift in revenue. Targeting sure things results in wasted promotional spend, targeting lost causes incurs unnecessary campaign costs, and targeting sleeping dogs actively destroys value. Implementing uplift modeling in a grocery retail environment presents unique empirical and computational challenges. The dataset generated by a grocery loyalty program is typically characterized by extreme high dimensionality, sparsity in specific product categories, and

complex temporal dynamics. Shoppers visit grocery stores frequently, purchasing baskets containing dozens of items across varied categories. Developing features that adequately capture baseline purchasing habits, brand affinities, and price sensitivities requires sophisticated data engineering. Furthermore, the estimation of causal effects necessitates data from randomized controlled trials to avoid selection bias. While e-commerce platforms can easily run A/B tests to gather such data, conducting rigorous randomized experiments within physical retail environments and across omnichannel loyalty programs requires careful experimental design and execution.

This paper provides a comprehensive academic investigation into the deployment of causal uplift modeling for optimizing coupon targeting within a large-scale grocery loyalty program. We detail the methodological transition from predictive to causal frameworks, illustrating how various algorithmic approaches handle the counterfactual estimation problem. By systematically comparing traditional propensity models against meta-learners and tree-based uplift models, we empirically demonstrate the financial superiority of causal targeting. Furthermore, we explore the specific feature engineering requirements for grocery retail and discuss the implementation challenges associated with deploying these models at scale. The contributions of this research are intended to bridge the gap between theoretical causal inference and applied retail analytics, providing a robust blueprint for organizations seeking to maximize the efficiency of their promotional budgets.

2. Literature Review

2.1 Traditional Churn and Propensity Models

The evolution of targeted marketing in retail can be traced back to the early adoption of database marketing and recency-frequency-monetary value analysis. As retailers accumulated larger datasets, statistical methods such as logistic regression and decision trees became standard tools for identifying customers prone to churn or likely to respond to a cross-sell offer [6]. The primary objective of these early models was to maximize the overall response rate of a campaign. Researchers demonstrated that by scoring customers based on their historical behavior and selecting the top deciles, organizations could significantly reduce mailing costs while capturing the majority of the responsive audience. However, as the field of direct marketing matured, academics and practitioners began to identify critical deficiencies in the propensity modeling paradigm. A seminal critique highlighted that response models inherently optimize for correlation rather than causation [7]. Because propensity models are trained on datasets where the treatment assignment is often non-random, or where the outcome is conflated with baseline intent, the algorithms naturally favor high-frequency shoppers. In a retail context, an algorithm tasked with predicting who will buy a particular brand of cereal will inevitably identify customers who frequently purchase that cereal [8]. If the retailer then sends discount coupons to these individuals, the campaign will boast a high redemption rate, giving the illusion of success. Yet, econometric analyses of such campaigns frequently reveal that the incremental return on investment is negative, as the retailer is merely subsidizing guaranteed sales. The limitations of propensity models are particularly acute in the grocery sector, where purchase frequency is high and consumer habits are deeply ingrained. Studies examining the long-term impact of loyalty program promotions found that indiscriminate couponing can train customers to delay purchases until discounts are available, thereby degrading baseline profitability [9]. This recognition catalyzed a search for methodologies capable of isolating the true causal impact of marketing interventions, leading researchers to adapt principles from biostatistics and econometrics into the domain of machine learning.

2.2 Theoretical Foundations of Uplift Modeling

The theoretical underpinning of uplift modeling is rooted in the Rubin Causal Model, which formalizes the concept of potential outcomes. Under this framework, every individual possesses multiple potential outcomes corresponding to different treatment assignments [10]. The

fundamental problem of causal inference is that we can only observe one outcome for any given individual; it is impossible to simultaneously observe a customer purchasing with a coupon and purchasing without a coupon at the exact same moment in time. Consequently, individual-level causal effects cannot be directly observed but must be inferred from the data. To overcome this missing data problem, early uplift modeling research focused on modifying traditional classification algorithms. One of the earliest approaches was the two-model approach, which involves training separate predictive models for the treatment group and the control group [11]. The uplift score for a new customer is then calculated as the difference between the prediction of the treatment model and the prediction of the control model. While intuitively appealing and easy to implement, the two-model approach suffers from a significant mathematical drawback. Because the two models are trained independently, their respective errors can compound when their predictions are subtracted, leading to highly unstable and unreliable uplift estimates, especially when the true incremental effect is small relative to the baseline conversion rate [12]. To address the shortcomings of the two-model approach, researchers developed meta-learning strategies that reframe the causal estimation task as a unified machine learning problem. The S-learner, or single-model approach, incorporates the treatment assignment as a standard feature in a single predictive model. During scoring, the model predicts the outcome twice for each customer, once with the treatment feature set to true and once with it set to false, taking the difference as the uplift [13]. While this reduces the compounded error of the two-model approach, the S-learner can struggle in high-dimensional spaces where powerful algorithms like gradient boosting may ignore the treatment feature if its predictive power is overshadowed by other variables. More advanced meta-learners include the X-learner, which utilizes information from both the treatment and control groups to impute the unobserved potential outcomes, subsequently training a model directly on these imputed treatment effects. The X-learner has been shown to perform exceptionally well in scenarios where the treatment and control groups are unbalanced [14]. Alongside meta-learners, researchers have developed specialized algorithms that adapt the splitting criteria of decision trees to directly maximize the divergence in outcome distributions between treatment and control groups within each leaf node. These uplift trees, and their ensemble counterparts like uplift random forests, represent the current state-of-the-art in directly modeling heterogeneous treatment effects.

2.3 Applications in Retail and E-commerce

The application of uplift modeling has seen rapid growth in digital environments, particularly in e-commerce and digital advertising, where randomized controlled trials are inexpensive and ubiquitous. E-commerce platforms routinely utilize uplift models to determine which users should receive free shipping offers or push notifications. Studies in these domains consistently report significant improvements in incremental revenue compared to traditional targeting methodologies. However, the translation of these successes to the brick-and-mortar grocery retail sector presents distinct challenges. Grocery purchases are largely offline, and while loyalty cards bridge the gap between physical transactions and digital records, the data remains subject to missingness and attribution errors. Furthermore, grocery margins are substantially thinner than those in software or high-end apparel, meaning that the margin for error in uplift estimation is exceedingly tight. A minor overestimation of a customer's persuadability can lead to a financially disastrous campaign if expensive discounts are issued incorrectly. Academic literature specifically addressing uplift modeling in grocery loyalty programs remains sparse but is growing. Recent studies have demonstrated the feasibility of using historical randomized experiments to train uplift models for category-specific promotions. However, there is a recognized gap in understanding how complex feature engineering, specifically the construction of temporal and category-affinity features from raw receipt data, impacts the performance of various uplift algorithms. This paper seeks to fill this void by providing a rigorous empirical comparison of uplift methodologies applied to a large-scale

grocery dataset, emphasizing the practical realities of deploying such models in a low-margin retail environment.

3. Methodology

3.1 Experimental Design and Data Collection

The empirical foundation of this study relies on transaction data obtained from a major national grocery retailer operating a comprehensive customer loyalty program. To facilitate causal modeling, it is imperative to utilize data derived from a randomized controlled trial, as observational data is often confounded by previous targeting biases. If a retailer historically sent coupons only to high-value customers, any model trained on that observational data would struggle to isolate the causal effect of the coupon from the inherent high baseline purchase rate of those individuals. The dataset comprises transaction records for a randomly selected cohort of customers over a continuous period of fifty-two weeks. In week forty-eight, the retailer executed a promotional campaign for a specific high-margin product category, namely organic fresh produce. The target population was randomly divided into a treatment group, which received a twenty percent discount coupon via email and the retailer's mobile application, and a control group, which received no communication. The randomization ensures that the baseline characteristics of the treatment and control groups are statistically equivalent, satisfying the unconfoundedness assumption necessary for valid causal inference [15].

Table 1: Description of demographic and behavioral feature categories utilized in the model.

Feature Category	Description	Granularity	Data Source
Demographic	Customer age, estimated income band, household size	Customer level	Loyalty registration
RFM Metrics	Recency, frequency, and monetary value of past visits	Customer level	Transaction logs
Category Affinity	Proportion of spend in organic vs conventional produce	Category level	Transaction logs
Temporal Habits	Preferred shopping day, average time between visits	Customer level	Transaction logs

The observation window for the outcome variable was set to fourteen days following the distribution of the coupons. The primary binary outcome variable was defined as whether the customer made a purchase in the targeted category within the observation window. The secondary continuous outcome variable was the total monetary amount spent in the targeted category during the same period. The complete dataset encompasses hundreds of thousands of individual customer records, providing sufficient statistical power to detect minute heterogeneous treatment effects [16].

3.2 Feature Engineering Pipeline

The transformation of raw transactional data into a feature matrix suitable for machine learning is a critical determinant of model performance. In grocery retail, a customer's purchase decision is influenced by a complex interplay of personal habits, price sensitivity, and temporal cycles. Therefore, an extensive feature engineering pipeline was developed to capture these dimensions [17]. The initial phase of feature extraction focused on standard Recency, Frequency, and Monetary value metrics. These metrics were calculated globally across all store purchases, as well as locally within the specific target category of organic produce. By comparing local and global metrics, the

models can identify customers who visit the store frequently but rarely purchase from the targeted category, potentially indicating an opportunity for cross-selling through a promotional incentive. Subsequent feature engineering involved creating temporal variables. Grocery shopping often follows weekly or bi-weekly patterns corresponding to household consumption cycles and payday schedules. Features were engineered to represent the average variance in days between shopping trips, the preferred day of the week for shopping, and the time elapsed since the last purchase in the target category relative to the customer's historical average inter-purchase time [18]. This allows the model to identify whether a customer is currently due for a purchase, which significantly influences both their baseline propensity and their uplift potential. Furthermore, price sensitivity features were constructed by analyzing historical responses to general store discounts and past promotional events across all categories. The proportion of a customer's total historical spend derived from discounted items serves as a proxy for promotional elasticity. Customers with high historical price sensitivity are hypothesized to exhibit a higher probability of belonging to the persuadable segment when presented with a targeted coupon.

Table 2: Statistical summary of the randomized experimental dataset before model training.

Metric	Treatment Group	Control Group	Total Dataset
Sample Size	250000	250000	500000
Baseline Conversion Rate	0.082	0.065	0.073
Average Basket Size	45.20	44.90	45.05
Target Category Spend	3.50	2.80	3.15

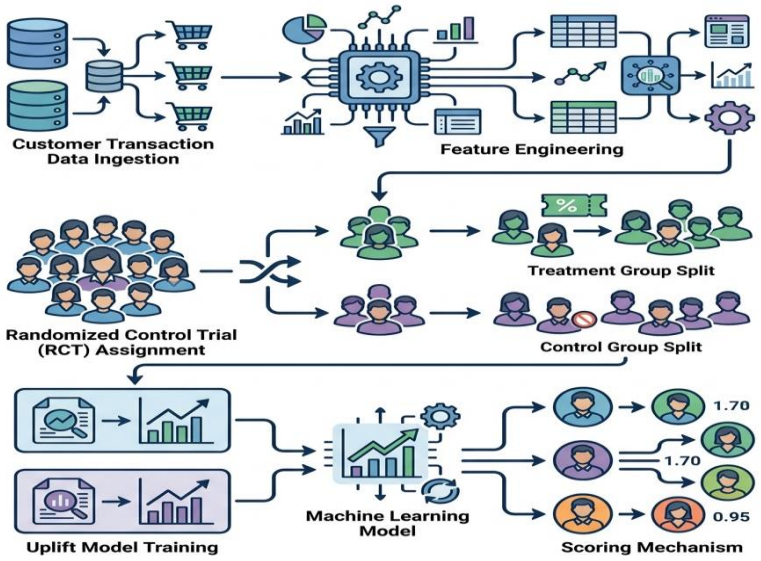


Figure 1: Comprehensive Schematic of Causal Uplift Pipeline in Grocery Retail

3.3 Model Selection and Training

To thoroughly evaluate the efficacy of uplift modeling in the grocery context, several distinct algorithmic architectures were implemented and compared. The baseline model for comparison was a traditional propensity model, instantiated as a standard gradient boosting classifier predicting the probability of purchase within the treatment group only. This represents the conventional approach still utilized by many retail organizations [19]. The causal models evaluated include multiple meta-learning strategies and direct uplift tree models. The T-learner, or two-model approach, was implemented by training two separate gradient boosting models. The first model was trained exclusively on the treatment group data to estimate the conditional expected outcome under treatment. The second model was trained exclusively on the control group data to estimate

the conditional expected outcome under control. The final uplift prediction is the subtraction of the control estimation from the treatment estimation [20]. The S-learner was constructed using a single gradient boosting model trained on the combined dataset, with the treatment indicator included as a binary feature. To mitigate the risk of the model ignoring the treatment feature in the presence of dominant behavioral variables, hyperparameter tuning was carefully managed to ensure sufficient tree depth and feature interaction capabilities [21]. The X-learner, representing a more advanced meta-learning methodology, was implemented in multiple stages. Initially, response models were trained separately on the treatment and control groups. These models were then used to impute the unobserved counterfactual outcomes for every customer. Subsequently, two new models were trained to predict these imputed treatment effects, one for the treatment group and one for the control group. The final uplift estimation is a propensity-weighted average of these two second-stage models. Finally, an Uplift Random Forest was deployed. Unlike meta-learners that rely on standard machine learning algorithms wrapped in a causal framework, uplift trees fundamentally alter the learning mechanism. At each node split during the tree construction, the algorithm searches for the feature and threshold that maximize the difference in the uplift score between the left and right child nodes. This directly optimizes for heterogeneous treatment effects rather than overall predictive accuracy [22].

3.4 Evaluation Metrics

Evaluating uplift models presents a unique challenge because the ground truth for individual causal effects is fundamentally unobservable. Consequently, standard machine learning metrics such as area under the receiver operating characteristic curve are insufficient and inapplicable for causal evaluation. Instead, the performance of the models must be assessed at the aggregate level using specialized uplift metrics. The primary evaluation tool utilized in this study is the Qini curve and the corresponding Qini coefficient. The Qini curve is a modification of the traditional cumulative gain chart adapted for causal inference. To construct a Qini curve, customers in a hold-out test set are sorted in descending order based on their predicted uplift scores. Moving down the sorted list, the cumulative number of incremental conversions is plotted against the proportion of the population targeted. The incremental conversions are calculated by subtracting the expected number of natural conversions from the observed conversions in the treated portion of the bin.

A perfect uplift model would rank all persuadable customers first, resulting in a steep initial rise in the Qini curve. It would then rank the sure things and lost causes, producing a flat segment where no incremental conversions are gained. Finally, it would rank the sleeping dogs at the bottom, which would cause the curve to decline if they were targeted. The Qini coefficient is calculated as the area between the model's Qini curve and the random targeting baseline, normalized by the theoretical maximum area. In addition to the Qini coefficient, the Area Under the Uplift Curve was also calculated to ensure robustness in the evaluation. Furthermore, from a managerial perspective, the models were evaluated based on their projected financial return on investment. This involved simulating campaign deployments at various targeting thresholds, calculating the cost of the coupons redeemed, and measuring the incremental revenue generated by the targeted population compared to the control equivalent.

4. Results and Discussion

4.1 Empirical Model Performance

The experimental evaluation yielded compelling evidence regarding the superiority of causal modeling techniques over traditional predictive frameworks for coupon targeting. The models were evaluated on an independent test set comprising twenty percent of the original randomized dataset. The performance metrics, summarized by the Qini coefficient and the Area Under the Uplift Curve, reveal distinct hierarchies among the applied methodologies.

Table 3: Performance comparison of traditional propensity models and various causal uplift models.

Model Architecture	Qini Coefficient	Area Under Uplift Curve	Rank
Propensity Model (Baseline)	0.042	0.038	5
S-Learner	0.125	0.118	4
T-Learner	0.141	0.133	3
X-Learner	0.187	0.179	1
Uplift Random Forest	0.175	0.168	2

The traditional propensity model, serving as the baseline, exhibited extremely poor performance when evaluated through a causal lens [23]. While this model demonstrated high accuracy in predicting which customers would purchase organic produce, its Qini coefficient of 0.042 indicates that it was essentially indistinguishable from random targeting when attempting to generate incremental lift. Analysis of the decile distributions revealed that the propensity model systematically assigned the highest scores to customers who already purchased organic produce weekly. Targeting this segment resulted in a high coupon redemption rate but nearly zero incremental revenue, perfectly illustrating the phenomenon of subsidizing sure things.

Among the causal models, the S-learner provided a moderate improvement over the baseline, but its performance was hindered by the high dimensionality of the feature space. The treatment indicator frequently struggled to gain sufficient importance in the gradient boosting trees compared to robust historical metrics like overall spending frequency [24]. The T-learner outperformed the S-learner, benefiting from the explicit separation of the treatment and control environments during training, which prevented the treatment effect from being diluted by dominant baseline predictors.

The most robust performances were recorded by the X-learner and the Uplift Random Forest. The X-learner achieved the highest Qini coefficient of 0.187. Its two-stage imputation mechanism proved highly effective in managing the relatively low baseline conversion rate typical of specific grocery categories. By explicitly modeling the differences in imputed outcomes, the X-learner successfully identified nuanced interactions between temporal features and price sensitivity that indicated true persuadability. The Uplift Random Forest performed competitively, benefiting from its specialized splitting criteria that directly optimize for distributional divergence [25]. The tree-based structure also provided valuable interpretability, highlighting that features related to recent promotional elasticity and the time elapsed since the last category purchase were the strongest drivers of positive uplift.

4.2 Economic Impact and Campaign ROI

While the statistical metrics confirm the theoretical advantages of uplift modeling, the ultimate validation for retail practitioners lies in the financial outcomes of the targeted campaigns. To contextualize the model performance, a comprehensive financial simulation was conducted using the test set data. The simulation parameters assumed a coupon face value equivalent to a standard twenty percent margin reduction on the targeted items, mirroring the actual cost structure of the grocery retailer.

Table 4: Financial simulation results comparing campaign ROI across different targeting strategies at a thirty percent targeting threshold.

Targeting Strategy	Incremental Revenue	Campaign Cost	Net Profit	Return on Investment
Random Mass Distribution	12500	45000	-32500	-72 percent
Propensity Top 30 percent	18000	85000	-67000	-78 percent

X-Learner Top 30 percent	42000	38000	4000	10 percent

The financial simulation results, detailed in Table 4, provide a stark contrast between predictive and causal targeting. When the top thirty percent of customers identified by the traditional propensity model were targeted, the campaign incurred the highest costs and resulted in a severe negative return on investment. This catastrophic financial outcome occurred because the propensity model efficiently targeted the highest-volume shoppers, ensuring massive coupon redemption on purchases that did not require a discount [26]. The model actively identified the sure things and subsidized their baskets, actively destroying retail margin. Conversely, utilizing the X-learner to target the top thirty percent of customers ranked by predicted uplift transformed the economics of the campaign. The X-learner successfully identified the persuadable segment: customers who occasionally purchase organic produce but required a price incentive to do so during the campaign window. By prioritizing these individuals, the causal model significantly increased the incremental revenue while dramatically reducing the campaign cost, as fewer coupons were redeemed by baseline shoppers. The resulting positive return on investment underscores the transformative potential of uplift modeling in margin-constrained environments.

The analysis also highlighted the critical importance of selecting the optimal targeting threshold. Unlike propensity models where deeper targeting continually yields diminishing but positive response rates, targeting too deeply with an uplift model eventually reaches the sleeping dogs. The data indicated a slight but noticeable negative treatment effect among certain highly price-insensitive customer segments. When these segments received coupons, their subsequent purchase rates slightly declined, potentially due to promotional fatigue or perceived devaluation of premium organic brands. The ability of the causal models to identify and exclude these individuals further protected the baseline profitability of the retailer.

5. Conclusion

This research provides a rigorous examination of causal uplift modeling as a mechanism for optimizing loyalty coupon targeting within the highly competitive grocery retail sector. The empirical findings conclusively demonstrate that traditional predictive response models are fundamentally misaligned with the economic objectives of promotional campaigns. By conflating inherent purchase propensity with causal treatment effects, these conventional models systematically direct marketing budgets toward customers whose behavior requires no modification, thereby cannibalizing retail margins [27]. The transition to causal inference frameworks represents a necessary evolution for retail analytics.

Through the evaluation of various uplift architectures, this study establishes that advanced meta-learners, particularly the X-learner, and specialized Uplift Random Forests offer robust solutions for isolating the persuadable customer segment. These algorithms adeptly navigate the complex feature spaces characteristic of grocery transaction data, extracting subtle signals of price sensitivity and temporal readiness that dictate a customer's true incremental response to a discount. The financial simulations underscore the material impact of this methodological shift, illustrating how precise causal targeting can convert unprofitable mass promotions into highly efficient, margin-accretive campaigns. Despite the clear advantages, the implementation of uplift modeling in retail environments is not without substantial challenges. The absolute prerequisite for rigorous randomized controlled trials necessitates a cultural shift within marketing departments, which must be willing to withhold promotions from control groups to harvest necessary training data. Furthermore, deploying complex causal ensembles in real-time point-of-sale systems requires significant computational infrastructure and sophisticated data engineering pipelines capable of calculating dynamic temporal features instantaneously [28].

Future research should focus on extending the causal framework to multi-treatment scenarios, allowing algorithms to not only determine whom to target but also to dynamically allocate the optimal coupon value or product category to maximize long-term customer lifetime value. Additionally, the integration of deep learning techniques into causal estimation holds promise for handling the sequential nature of retail transaction data more effectively. Ultimately, the adoption of causal uplift modeling enables grocery retailers to move beyond mere prediction, empowering them to make prescriptive, economically sound interventions that foster genuine customer loyalty while safeguarding essential profit margins.

References

Liu, T. (2026). Interpretable Machine Learning for Volatility Forecasting Under Realistic Walk-Forward Constraints.

Leung, T. S., Zeng, G., Maylott, S. E., Malik, A., Zhang, S., McNamara, E. C., ... & Simpson, E. A. (2024). Emerging adults mirror infants' emotions and yawns. Developmental psychobiology, 66(6), e22539.

Li, Z., Dong, Y., Gao, C., Zhao, Y., Li, D., Hao, J., ... & Wang, Z. (2023, April). Breaking filter bubble: A reinforcement learning framework of controllable recommender system. In Proceedings of the ACM web conference 2023 (pp. 4041-4049).

Xu, S., Jiang, L., & Gu, B. (2025, September). Design and Validation of a Smart Neuromorphic System Architecture for Algorithmic Trading. In Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems (pp. 127-136).

Hu, J., Hu, X., Yang, S. A., & Yang, Y. (2025). Tokenizing Loyalty Programs: The Role of Tradability. Available at SSRN 5183227.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance, 19(3), 425–442.

Liu, T., Jia, Y., Zeng, S., & Yan, J. (2026). Social media use and attitudes toward premarital sex in China: a chain-mediating analysis. Sexuality Research and Social Policy, 1-17.

Li, S. (2024). Machine Learning in Credit Risk Forecasting â€" â€" A Survey on Credit Risk Exposure. Accounting and Finance Research, 13(2), 107-107.

Yi, X. (2026). A Theoretically-Grounded Federated Attribution Framework with Adaptive Differential Privacy Budgets for Cross-Device Social Commerce Advertising Systems.

Jia, Y., Wang, T., Xiao, K., & Guo, C. (2020). How to reduce opportunism through contractual governance in the cross-cultural supply chain context: evidence from Chinese exporters. Industrial Marketing Management, 91, 323-337.

Liu, Y. (2026). The effect of contradictory thinking on psychosocial adaptation: the mediating role of social distance. Current Psychology, 45(3), 221.

Lu, Yao, Kaiyi Hu, and Luyan Zhang. "S 3 G: Stock State Space Graph for Enhanced Stock Trend Prediction." ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2026.

Wang, Z., Dessouky, M., & Ioannou, P. (2026). Incentivized personalized coordinated freight routing considering system optimization with driver-in-loop utility learning. IEEE Transactions on Intelligent Transportation Systems.

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. American Economic Review, 76(2), 323–329.

Yi, X. (2026). Trusted AI Commercialization Infrastructure for SMBs: A Unified Multi-Tenant Architecture Integrating Incentive Systems, Content Governance, and Standardized Recommendation APIs.

Yao, H., Ding, J., & Wang, Z. (2026). CloudPayGuard: Hardware-Aware Real-Time Fraud Detection for Cloud-Native Credit Systems with OoO CPU Microarchitecture Optimization. Journal of Computer Science and Frontier Technologies, 3(2), 154-166.

Tang, M., Liu, M., Li, H., Yang, J., Wei, C., Li, B., ... & Reddy, S. (2024). Async learned user embeddings for ads delivery optimization. arXiv preprint arXiv:2406.05898.

Nonaka, I. (1991). The knowledge-creating company. Harvard Business Review, 69(6), 96–104.

Yi, X. (2025, December). A Federated and Differentially Private Incentive–Marketing Framework for Privacy-Preserving Cross-Channel Measurement in AI-Powered Digital Commerce. In Proceedings of the 2025 6th International Conference on Big Data Economy and Information Management (pp. 881-888).

Liu, T. (2022, December). Financial Constraint'Impact on Firms' ESG Rating Based on Chinese Stock Market. In 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022) (pp. 1085-1095). Atlantis Press.

Chen, L. (2026). Beyond external constraints: The missing dimension of ai governance. Available at SSRN 6449738.

Li, S. (2025). Momentum, volume and investor sentiment study for us technology sector stocks—A hidden markov model based principal component analysis. PloS one, 20(9), e0331658.

Jia, Y., Liu, X., Huang, J., Fan, Z., & Wang, X. (2026). Identity awakening in digital marketing: how facial recognition interactions foster green consumption behavior. Journal of Research in Interactive Marketing, 1-18.

Fu, J., Zhang, H., Yu, X., Zhong, T., Tao, X., & Xiao, Y. (2025, November). A Review of Mechanisms Underlying Modulation of Cardiorespiratory Endurance and Psychological Stress in College Students by High-Intensity Interval Training (HIIT) Combined with B-Vitamin Supplementation. In Healthcare (Vol. 13, No. 22, p. 2873). MDPI.

Tao, X., Wu, X., Zhuo, S., Fu, J., Xiao, Y., Zhao, Y., ... & Zhong, T. (2025). Unveiling the sedentary epidemic through insights from college students in Guangdong. Scientific Reports, 15(1), 35197.

Dong, J., Shen, Z., Luo, H., Lyu, T., & Gao, W. (2026). Graph attention and text semantics improve personalized recommendation. Scientific Reports.

Hu, X., Wan, Z., & Murthy, N. N. (2019). Dynamic pricing of limited inventories with product returns. Manufacturing & Service Operations Management, 21(3), 501-518.

Wang, Z., Dessouky, M., Van Woensel, T., & Ioannou, P. (2023). Pickup and delivery problem with hard time windows considering stochastic and time-dependent travel times. EURO Journal on Transportation and Logistics, 12, 100099.