Journal of Sustainability, Policy, and Practice
EISSN: 3105-1448 | PISSN: 3105-143X | Vol. 2, No. 1 (2026)

Article

Causal Impact of Cashback Campaigns on Post-Marketing Default Behavior in Consumer Lending

Anna Sokolov [1], Dmitry Kuznetsov [1] and Ivan Petrov [1,*]

[1] Faculty of Economic Sciences, HSE University, Moscow, 109028, Russia
* Correspondence: Ivan Petrov, Faculty of Economic Sciences, HSE University, Moscow, 109028, Russia

Abstract: This study examines how cashback marketing campaigns influence subsequent credit-risk outcomes. Using an AB testing framework implemented on a dataset of 640,000 loan applicants, the treatment group received a one-time cashback incentive, while the control group did not. Propensity-score weighting and double-machine-learning estimators were applied to isolate causal effects. Results show that campaign exposure increases short-term activation rates by 18.4% but also raises 90-day delinquency risk by 6.7 percentage points. Heterogeneity analysis reveals significantly larger default spillovers among new-to-bank customers (increase of 11.3 pp). A cost-risk decomposition suggests that while the campaign boosts revenue in the first month, cumulative losses surpass gains after 5.6 months. The findings highlight the risk implications of customer-acquisition incentives.

Keywords: causal inference; marketing spillover; AB test; consumer lending; default risk

Published: 21 February 2026

1. Introduction

Consumer lending has expanded rapidly in recent years as digital distribution channels and automated underwriting have substantially reduced customer-acquisition costs. To accelerate portfolio growth, many lenders rely on cashback offers, reward points, and similar monetary incentives to increase application, approval, and activation rates. In credit-card markets, reward programs have long been shown to influence spending behaviour and portfolio expansion [1,2]. Comparable incentive schemes have now become common in personal-loan and buy-now-pay-later products. At the same time, regulatory and industry assessments caution that aggressive promotions may disproportionately attract borrowers with weaker financial positions, raising concerns that short-term increases in utilisation could be followed by higher delinquency and default rates [3].

Most existing research in consumer credit focuses on improving default prediction at loan origination. A large empirical literature shows that gradient boosting, ensemble learning, and neural-network models often outperform traditional scorecards in terms of predictive accuracy and risk discrimination [4]. Recent studies show that embedding forward-looking macroeconomic scenarios into consumer-finance credit-risk models can substantially improve portfolio-level risk assessment and stress-testing outcomes compared with borrower-only approaches [5]. However, these frameworks typically impose a stable relationship between macroeconomic conditions and default risk and do not consider how borrower behaviour may respond differently across product designs or acquisition incentives. Despite these advances, the dominant emphasis remains predictive rather than causal: models are designed to forecast default risk more precisely, but they rarely quantify how changes in product design or acquisition incentives directly alter subsequent repayment behaviour. A separate stream of research examines behavioural determinants of loan default. Reviews in this literature highlight the roles of financial literacy, social environment, and psychological factors in shaping repayment outcomes [6,7]. Evidence from mobile and online lending platforms suggests that soft information, such as social-network signals or digital footprints, can significantly influence delinquency risk [8]. Experimental studies further show that relatively small operational interventions-such as disbursement timing, reminder messages, or personalised nudges-can reduce default rates [9]. While these findings underscore that borrower behaviour responds to institutional and design features, they rarely focus on monetary acquisition incentives such as cashback and seldom link such incentives to medium-term credit-risk outcomes. Insights from marketing and household finance provide a related but incomplete perspective. Theoretical models indicate that cashback and reward programs can affect consumption and saving choices even without invoking behavioural biases [10]. Empirical studies document that rewards and discounts influence credit-card usage and household spending patterns [11]. Similar effects are observed in e-commerce settings, where cashback schemes alter purchase and return behaviour [12]. Research on platform pricing further shows that lenders adjust credit limits, interest rates, and promotional intensity to balance growth and risk objectives [12,13]. However, this literature primarily examines demand responses and utilisation, with limited attention to detailed loan-level delinquency dynamics following promotional campaigns. Recent advances in causal-inference methods offer new tools to move beyond correlational analysis. Double-machine-learning and related approaches have been applied to estimate the effects of credit expansions and credit-limit changes on household spending and repayment behaviour [14]. In the context of consumer lending, recent evidence shows that promotional incentives can raise loan take-up while also affecting subsequent default risk, highlighting the need to evaluate acquisition strategies jointly in terms of growth and risk [15]. Nevertheless, robust causal evidence on how one-time cashback incentives influence post-marketing delinquency in consumer-loan portfolios remains scarce. Much of the existing work focuses on credit-card usage or short-term activation outcomes, relies on observational data subject to unobserved heterogeneity, or examines only very short horizons following treatment [16]. Moreover, heterogeneity across borrower groups-such as new-to-bank versus existing customers-has received limited attention, despite indications that their responses to incentives may differ substantially [17].

In this study, we examine the causal impact of cashback incentives on consumer-loan performance using large-scale experimental data. The analysis draws on a randomised A/B test covering approximately 640,000 loan applicants, in which a one-time cashback offer was assigned to the treatment group. To estimate campaign effects in the presence of high-dimensional borrower characteristics, we combine propensity-score weighting with double-machine-learning estimators. This approach allows us to quantify the effects of the incentive on short-term activation and subsequent 90-day delinquency, to assess heterogeneity between new-to-bank and existing customers, and to trace the dynamic trade-off between acquisition gains and downstream credit losses through a cost-risk decomposition. By linking randomised treatment assignment with detailed loan-level outcomes, the framework provides a systematic way to evaluate how monetary acquisition incentives affect medium-term credit performance in consumer lending.

2. Materials and Methods

2.1. Sample and Study Setting

The study uses loan-application data from a large consumer-lending platform. The dataset contains 640,000 applicants who submitted requests for unsecured installment loans during a continuous business period. All applicants completed basic eligibility checks, including identity verification, minimum income confirmation, and fraud screening. The analysis focuses on first-time or recently inactive applicants because they were the target users of the cashback campaign. Each record includes demographic features, credit-bureau information, internal behavior variables, device details, and repayment outcomes. Loan performance was tracked for at least 120 days after disbursement to measure early-stage delinquency. The study environment reflects routine lending operations and consistent credit policies across regions.

2.2. Experimental Design and Control Assignment

A randomized AB test was used to measure the effect of the cashback incentive. Applicants were randomly assigned to a treatment group, which received a one-time cashback reward after activation, or to a control group, which did not receive any incentive. Both groups followed the same underwriting rules, loan limits, and pricing conditions. Random assignment helps ensure that the two groups are similar before the intervention. The cashback amount was fixed for all treatment users to avoid variation caused by different reward levels. Group size and assignment consistency were checked throughout the experiment to confirm that the planned design was followed.

2.3. Measurement Procedures and Quality Control

Loan outcomes were recorded using standard operational rules. Activation occurred when an approved applicant completed signing and received the funds. Delinquency was measured with a 90-day past-due status, which is widely used as an indicator of early default. Data quality checks included completeness review, cross-checking with internal payment records, and verification against credit-bureau updates. Device logs, timestamps, and payment histories were screened to remove duplicated or inconsistent entries. Summary statistics from the cleaned data were compared with the raw data to confirm that the cleaning process did not distort sample patterns. All quality-control steps were completed before treatment-effect estimation.

2.4. Data Processing and Model Specification

Data preprocessing included simple scaling of numerical features, encoding of categorical variables, and removal of variables with very low variation. Propensity-score weights were calculated using a gradient-boosted model to balance the treatment and control groups. The main treatment effect was estimated with a double-machine-learning (DML) method, which controls for many covariates while keeping unbiased estimates for the treatment variable.

The baseline model takes the form:

$$Y_i=\alpha+\tau D_i+X_i'\beta+\varepsilon_i,$$

Where Y_i is the delinquency outcome, D_i is the treatment indicator, and X_i is the covariate vector.

A simple revenue-loss ratio was also computed to track the campaign's financial path [18]:

$$R(t)=\frac{\sum_{k=1}^{t}\mathrm{Revenue}_k}{\sum_{k=1}^{t}\mathrm{Loss}_k}$$

This indicator identifies the month when accumulated losses exceed early-stage revenue gains.

2.5. Ethical Considerations and Data-Use Compliance

All data were de-identified before analysis, and no personal identifiers were retained. The study followed the lending institution's data-governance rules and applied the same promotional rules used in daily operations. The experiment was limited to eligible users who met standard credit-policy requirements. All procedures were consistent with regional data-protection laws and internal guidelines for responsible use of financial data.

3. Results and Discussion

3.1. Effect of the Cashback Offer on Activation and Delinquency

The cashback offer produced a clear rise in short-term activation. After covariate balancing and treatment-effect estimation, the activation rate in the treatment group was 18.4% higher than that of the control group. This confirms that a direct monetary reward is an effective stimulus for completing the loan-signing process. However, the same group showed a 6.7-percentage-point increase in 90-day delinquency. This indicates that a portion of the additional activated borrowers had weaker repayment ability or lower repayment priority. Figure 1 presents these two outcomes and highlights the shift between higher activation and higher repayment risk. Similar patterns have been noted in recent research showing that even well-calibrated credit models cannot fully offset behavioral responses to product incentives [19]. Compared with studies that examine only prediction accuracy, our results show that campaign design itself can alter the underlying risk mix of the activated population [20].

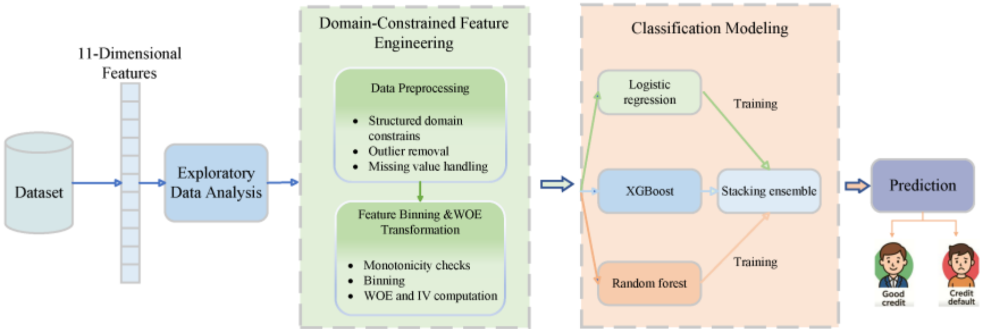

Figure 1. Comparison of activation rates and 90-day delinquency between the treatment and control groups under the cashback campaign.

3.2. Differences across Customer Groups

The effect of the cashback incentive varied strongly across borrower types. Among new-to-bank users, the 90-day delinquency rate increased by 11.3 percentage points. This is nearly twice the average effect. Existing customers showed a smaller and statistically weaker shift. These results suggest that applicants without a prior lending history may be more sensitive to short-term monetary rewards and may activate loans despite limited repayment capacity. This finding agrees with earlier work in digital-credit settings, where some user profiles show higher sensitivity to promotional messages and higher default tendencies [21]. Evidence from reward-based purchase studies also shows that monetary incentives can lead to inconsistent post-transaction behavior for certain consumer groups. In the current context, this inconsistency appears as poorer repayment performance rather than weaker satisfaction, which has direct consequences for portfolio-level risk control.

3.3. Cost-Risk Balance and Time to Loss Crossover

Short-term financial returns improved after the campaign, but the gains did not last long. Using observed interest income, fees, and realized losses, we calculated the cumulative revenue-loss ratio. In the first month, the ratio was above one, meaning that added loan volume covered both the incentive cost and expected early losses. As time passed, more treatment-group borrowers moved into delinquency, causing the ratio to fall. The ratio dropped below one after about 5.6 months, which marks the point where losses began to exceed the campaign's early gains. Figure 2 illustrates this crossover point. Similar warnings appear in recent studies emphasizing that early-period metrics may hide later-period credit losses, even when institutions apply advanced credit-risk models (Ding et al., 2025; Noriega et al., 2023). Work on cashback incentives in digital commerce also shows that monetary rewards can change later behavior in ways that reduce long-term value (Guo, 2024). For lenders, these results indicate that campaign performance should be assessed across a time window that reflects the full repayment cycle.

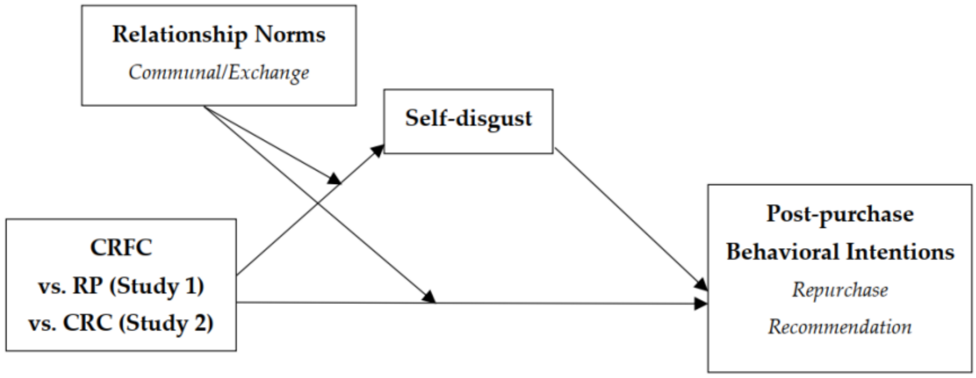

Figure 2. Monthly revenue-loss ratio showing the break-even point at which cumulative losses exceed initial campaign gains.

3.4. Comparison with Prior Studies and Implications

The results add new evidence to both the credit-risk and marketing-incentive literature. Previous work on machine-learning credit models has stressed improvements in prediction and segmentation [22]. The present findings show that-even with stable underwriting rules and model-based risk controls-a simple acquisition tool such as a one-time cashback can reshape the profile of activated borrowers and raise delinquency risk. Existing research on cashback programs mainly examines purchase behavior, review behavior, or customer satisfaction. Here, the evidence extends to loan repayment, where changes in customer composition directly affect portfolio losses. These findings suggest that marketing and risk teams should jointly evaluate the acceptable lift in default risk and define the maximum time window for break-even. The results come from one institution and one campaign design, so they may not generalize to all credit products. Future work could examine different incentive levels, timing of reward delivery, or communication strategies to determine which combinations support growth without raising default risk beyond operational tolerance.

4. Conclusion

This study evaluated how a one-time cashback offer influences loan activation and early repayment outcomes. The results show that the incentive raises activation, but it also increases 90-day delinquency. The rise in delinquency is especially large among new-to-bank applicants, which suggests that the campaign attracts borrowers who have weaker repayment ability. The cost-risk results further show that the campaign brings net gains only in the first few months. As more borrowers fall behind on repayment, cumulative losses become higher than the early revenue lift. These findings show that short-term indicators do not fully reflect the long-term effect of acquisition incentives. The results provide useful guidance for lenders that aim to expand their customer base while keeping portfolio risk at a manageable level. The analysis is based on a single institution and one campaign design, which limits generalization. Future studies can test other incentive amounts, reward timing, and communication methods to identify campaign designs that support growth without adding extra credit risk.

References

1. R. Agrawal, A. Khanna, and S. Hamdare, "Analyzing and rewarding credit card spending habits in India: A machine learning approach," *International Journal of Computational Intelligence Systems*, vol. 18, no. 1, p. 165, 2025. doi: 10.1007/s44196-025-00899-0
2. Y. Ming, J. Chen, and C. Li, "The impacts of acquisition modes on achieving customer behavioral loyalty: An empirical analysis of the credit card industry from China," *International Journal of Bank Marketing*, vol. 39, no. 1, pp. 147-166, 2021.
3. G. Lapierre-Berger, M. C. Cohen, and J. Puyana-Bastin, "When platforms go public, standards drop," *Production and Operations Management*, 2025. doi: 10.1177/10591478241310217

4. D. Tran, and A. W. Tham, "Accuracy comparison between feedforward neural network, support vector machine and boosting ensembles for financial risk evaluation," *Journal of Risk and Financial Management*, vol. 18, no. 4, p. 215, 2025. doi: 10.3390/jrfm18040215
5. J. Wang, and Y. Xiao, "Assessing the spillover effects of marketing promotions on credit risk in consumer finance: An empirical study based on AB testing and causal inference," 2025.
6. G. N. Kyeyune, and J. M. Ntayi, "Empowering rural communities: The role of financial literacy and management in sustainable development," *Frontiers in Human Dynamics*, vol. 6, p. 1424126, 2025. doi: 10.3389/fhumd.2024.1424126
7. S. J. S. Krishna, M. Aarif, N. K. Bhasin, S. Kadyan, and B. K. Bala, "Predictive analytics in credit scoring: Integrating XGBoost and neural networks for enhanced financial decision making," In *Proceedings of the 2024 International Conference on Data Science and Network Security*, 2024, pp. 1-6.
8. T. Li, J. Xia, S. Liu, and E. Hong, "Strategic human resource leadership in global biopharmaceutical enterprises: Integrating HR analytics and cross-cultural," 2025. doi: 10.20944/preprints202511.0572.v1
9. J. Nahar, M. A. Rahaman, M. Alauddin, and F. Z. Rozony, "Big data in credit risk management: A systematic review of transformative practices and future directions," *International Journal of Management Information Systems and Data Science*, vol. 1, no. 4, pp. 68-79, 2024. doi: 10.62304/ijmisds.v1i04.196
10. X. Gu, J. Yang, and M. Liu, "Optimization of anti-money laundering detection models based on causal reasoning and interpretable artificial intelligence and its empirical study on financial system stability," *Optimization*, vol. 21, p. 1, 2025. doi: 10.22541/essoar.176126742.22902505/v1
11. G. Santilli, "Climate risks in the banking industry," *Springer*, 2025.
12. L. Tan, Z. Peng, Y. Song, X. Liu, H. Jiang, S. Liu, and Z. ... Xiang, "Unsupervised domain adaptation method based on relative entropy regularization and measure propagation," *Entropy*, vol. 27, no. 4, p. 426, 2025.
13. J. P. Allegret, and R. Cergibozan, "Determinants of the European sovereign debt crisis: Application of logit, panel Markov regime switching model and self-organizing maps," *Entropy*, vol. 25, no. 7, p. 1032, 2023.
14. W. Zhu, Y. Yao, and J. Yang, "Real-time risk control effects of digital compliance dashboards: An empirical study across multiple enterprises using process mining, anomaly detection, and interrupted time series," 2025.
15. C. R. Brewin, L. Atwoli, J. I. Bisson, S. Galea, K. Koenen, and R. Lewis-Fernández, "Post-traumatic stress disorder: Evolving conceptualization and evidence, and future research directions," *World Psychiatry*, vol. 24, no. 1, pp. 52-80, 2025.
16. J. Wang, and Y. Xiao, "Research on credit risk forecasting and stress testing for consumer finance portfolios based on macroeconomic scenarios," 2025. doi: 10.1145/3785706.3785751
17. M. Fleischer, D. Das, P. Bose, W. Bai, K. Lu, M. Payer, and G. ... Vigna, "ACTOR: Action-guided kernel fuzzing," In *Proceedings of the 32nd USENIX Security Symposium*, 2023, pp. 5003-5020.
18. E. A. Morrison, D. A. Adu, and D. Kimani, "Propelling the transition to circular economy: Exploring the role of corporate circular economy performance-based incentive policy," *Business Strategy and the Environment*, 2025.
19. W. Zhu, J. Yang, and Y. Yao, "How cross-departmental collaboration structures mitigate cross-border compliance risks: Network causal inference based on ManpowerGroup's staffing projects," 2025. doi: 10.20944/preprints202510.1339.v1
20. G. Benigno, A. Foerster, C. Otrok, and A. Rebucci, "Estimating macroeconomic models of financial crises: An endogenous regime-switching approach," *Quantitative Economics*, vol. 16, no. 1, pp. 1-47, 2025.
21. X. Gu, J. Yang, and M. Liu, "Research on a green money laundering identification framework and risk monitoring mechanism integrating artificial intelligence and environmental governance data," 2025. doi: 10.21203/rs.3.rs-7927453/v1
22. Q. G. Nguyen, L. H. Nguyen, M. M. Hosen, M. Rasel, J. F. Shorna, M. S. Mia, and S. I. Khan, "Enhancing credit risk management with machine learning: A comparative study of predictive models for credit default prediction," *The American Journal of Applied Sciences*, vol. 7, no. 1, pp. 21-30, 2025.