Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

From Churn Prediction to Retention Pricing: A Cost-Sensitive Decision Framework for Managing Revenue at Risk in Subscription Services

DHARANI M

24pite04@ldc.edu.in

Lady Doak College

Research Article

Keywords: Revenue Management, Customer Retention, Churn Prediction, Cost-Sensitive Decision Making, Customer Lifetime Value, Subscription Pricing

Posted Date: September 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10915843/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Customer churn is a direct revenue-management problem: every customer who leaves removes a recurring revenue stream that must be replaced at a materially higher acquisition cost. This paper reframes churn prediction as an input to a retention-pricing decision rather than as a standalone classification exercise. Using 7,043 subscription records from a telecommunications provider (the Telco Customer Churn dataset), we (i) quantify revenue at risk by contract segment using actual billing data, (ii) build a leakage-free machine-learning pipeline that ranks customers by churn probability, and (iii) formulate the decision of whom to contact with a retention offer as a cost-sensitive threshold problem, trading the cost of the offer against the expected revenue saved. Month-to-month subscribers account for USD 1.45 million of the USD 1.67 million in annual revenue at risk from current churners (87%), despite representing only 55% of the customer base - a segment-level asymmetry with direct implications for contract-length pricing incentives. Sweeping the retention-offer contact threshold across a grid of offer-cost and retention-success assumptions shows that net benefit is maximized at a low probability threshold (~ 0.15), implying that broad, low-cost outreach is economically preferable to narrow, high-precision targeting under the cost structure typical of subscription retention campaigns. We report these findings with explicit sensitivity analysis, since offer cost and retention-success rate are assumed parameters rather than measured campaign outcomes. The contribution of this work is a reproducible bridge between a predictive churn model and a revenue-management decision - segment-level revenue-at-risk quantification, a probability-ranked targeting mechanism, and a cost-sensitive threshold framework - intended to support pricing and retention-campaign design in subscription businesses.

1. Introduction

Revenue management is fundamentally about matching supply, demand, and price to maximize revenue under uncertainty. In subscription businesses, the analogous problem is retention: a firm must decide which customers are at risk of cancelling, and whether and how to intervene with a price concession, service upgrade, or targeted offer before the customer leaves. This is a demand-management decision under uncertainty, structurally similar to overbooking or dynamic-pricing problems in airline and hotel revenue management, except that the 'inventory' being protected is an existing customer's future revenue stream rather than a seat or room.

Despite this natural fit, churn-prediction research is typically published as a pure classification exercise: models are compared on accuracy or F1-score, with little connection back to the pricing or retention-campaign decision the prediction is meant to support. A model that predicts churn accurately is not, by itself, actionable - the firm still has to decide whom to contact, with what offer, and at what cost, and that decision depends on the economics of the intervention, not on the classifier's accuracy alone.

This paper closes that gap for a representative subscription dataset. Its contributions are:

1. Segment-level revenue-at-risk quantification: using actual billing data, we compute the annual revenue currently at risk from churn, broken down by contract type.
2. A leakage-free churn-probability model: six classifiers are benchmarked under a corrected protocol (test data held out before any resampling), producing well-calibrated probability rankings usable as the demand-side signal for a retention decision.
3. A cost-sensitive retention-threshold framework: the decision of whom to contact is formulated as a threshold-selection problem trading offer cost against expected revenue saved, solved via sensitivity analysis across plausible cost and retention-success assumptions.
4. Simple customer-lifetime-value (CLV) estimates by contract segment, linking observed billing and tenure data to a pricing-relevant value metric.
5. A discussion of the pricing implications: which segments justify aggressive retention spend, and what contract-length incentive structures the revenue-at-risk asymmetry supports.

Section 2 reviews related work in churn prediction and its connection to revenue management. Section 3 describes data and methodology, including the leakage-free classification pipeline and the cost-sensitive decision framework. Section 4 reports results. Section 5 discusses pricing and revenue-management implications. Section 6 states limitations and future work. Section 7 concludes.

2. Related Work

Churn prediction sits at the intersection of predictive analytics and revenue management. On the predictive side, ensemble and boosting methods (Random Forest, Gradient Boosting, XGBoost) are consistently reported to outperform simpler classifiers on structured churn data (Chang et al., 2024; Shrestha, 2022; AbdelAziz et al., 2025), while deep-learning approaches offer competitive but not clearly superior performance on tabular, non-sequential data (Liu et al., 2024; Zhao et al., 2024). Because churn datasets are imbalanced - non-churners typically outnumber churners three-to-one - resampling techniques such as SMOTE are widely used to improve minority-class detection (Peng, 2023; Ong, 2024; Li et al., 2025; Kumar et al., 2022).

On the revenue-management side, the connection between predictive churn scores and retention economics is comparatively underexplored in the applied literature. Sikri et al. (2024) linked model-driven retention campaigns to an approximate 12% reduction in churn but did not report the underlying cost-benefit calculation. Asif and Arif (2025) proposed an explainable churn framework oriented toward deployment but stopped short of a pricing or targeting-threshold analysis. This gap a predictive model without an explicit link to the retention-spend decision it is meant to inform motivates the cost-sensitive threshold framework developed in Section 3.4 of this paper. Evaluation-metric literature (Gupta et al., 2023; AbdelAziz et al., 2025) further argues that accuracy alone is a poor guide for imbalanced, cost-asymmetric problems such as churn a point directly relevant to revenue management, where false negatives (missed churners) and false positives (wasted retention offers) carry different, quantifiable costs rather than being treated symmetrically.

3. Data and Methodology

3.1 Dataset

The Telco Customer Churn dataset contains 7,043 customer records with demographic, service, contract, and billing attributes, including MonthlyCharges, TotalCharges, tenure (months), and Contract type (month-to-month, one-year, two-year). Of these customers, 1,869 (26.5%) have churned. Billing fields are used directly, without transformation, for the revenue-at-risk and CLV calculations in Sections 3.3–3.4; they are additionally scaled and encoded for use as model features in Section 3.2.

3.2 Leakage-Free Churn-Probability Model

A common error in applied churn-prediction pipelines is applying SMOTE to the full dataset before splitting into training and test sets, which leaks synthetic minority-class information across the split and inflates reported performance. This study avoids that error: the dataset is split into training (80%) and test (20%) sets using stratified sampling before any resampling or scaling; StandardScaler is fit on the training data only; and SMOTE (re-implemented from its published k-nearest-neighbour algorithm, k = 5, since the imbalanced-learn package was unavailable in the execution environment) is applied to the training fold only. Six classifiers - Logistic Regression, K-Nearest Neighbors, Random Forest, Gradient Boosting, a histogram-based gradient-boosted-tree model (used in place of the XGBoost library, which was unavailable offline; the two implement the same underlying algorithm), and a Multi-Layer Perceptron were tuned via GridSearchCV with stratified 3-fold cross-validation on the training fold, optimizing F1-score. Full model performance is reported in Section 4.1; the purpose of this stage is to generate a probability score per customer that ranks churn risk, which becomes the demand-side input to the retention decision in Section 3.4.

3.3 Revenue-at-Risk and Customer Lifetime Value

For each contract segment, annual revenue at risk is computed directly from billing data as the sum of MonthlyCharges across currently churned customers in that segment, multiplied by 12. A simple customer-lifetime-value proxy is computed per segment as the product of average MonthlyCharges and average tenure (in months) for customers in that segment. This is a first-order CLV estimate it does not discount future cash flows or model survival probability explicitly but it is sufficient to illustrate the relative value differential across contract types that motivates the pricing discussion in Section 5.

3.4 Cost-Sensitive Retention-Threshold Framework

Given a churn-probability score for each customer from Section 3.2, the firm must choose a probability threshold above which a customer is contacted with a retention offer. For a given threshold t, let TP(t) be the true churners correctly identified and FP(t) be the non-churners incorrectly flagged. We define:

Revenue Saved(t) = Σ MonthlyCharges(TP(t)) × retention_success_rate × horizon_months

Campaign Cost(t) = Σ MonthlyCharges(TP(t) ∪ FP(t)) × offer_cost_fraction × offer_duration_months

Net Benefit(t) = Revenue Saved(t) − Campaign Cost(t)

where retention_success_rate is the probability that a contacted, at-risk customer accepts the offer and remains, offer_cost_fraction is the discount offered as a fraction of monthly billing, horizon_months is the forward revenue window credited to a retained customer (12 months in this study), and offer_duration_months is the length of the discount period (3 months). Because retention_success_rate and offer_cost_fraction are not observable in this dataset - no historical retention-campaign outcomes are available - we treat them as assumptions and report Net Benefit(t) across a sensitivity grid of three plausible parameter combinations (10%/30%, 15%/40%, 20%/50% for offer-cost/retention-success) rather than asserting a single point estimate. The optimal threshold t* is the value maximizing Net Benefit(t) within each scenario.

4. Results

4.1 Churn-Probability Model Performance

Table 1
summarizes test-set performance for the six classifiers under the leakage-free protocol (held-out test set, n = 1,409, never resampled or scaled with test-set statistics). Full methodological detail is in Section 3.2; these classifiers exist to produce a probability ranking, not as the paper's primary contribution.

Model	Accuracy	Precision	Recall	F1	ROC-AUC
Logistic Regression	0.740	0.506	0.786	0.616	0.841
KNN	0.700	0.459	0.719	0.560	0.765
Random Forest	0.771	0.566	0.599	0.582	0.826
Gradient Boosting	0.785	0.585	0.652	0.617	0.837
HistGB (XGBoost-equivalent)	0.779	0.575	0.647	0.609	0.833
Neural Network (MLP)	0.759	0.533	0.746	0.622	0.833

Table 1. *Test-set churn-classification performance, leakage-free protocol.*

The Neural Network (F1 = 0.622) is used as the probability source for Section 4.3, since it achieves the highest recall (0.746) appropriate for a revenue-management use case where the cost of an undetected churner (a full lost revenue stream) substantially exceeds the cost of a wasted retention offer to a customer who would have stayed regardless.

4.2 Revenue at Risk and Customer Lifetime Value by Segment

Table 2
Segment-level revenue at risk and customer lifetime value estimate, computed directly from billing data.

Contract	Customers	Churn Rate	Avg. Tenure (mo)	Avg. Monthly Charge	Annual Revenue at Risk	CLV Estimate
Month-to-month	3,875	42.7%	18.0	$66.40	$1,450,165	$1,198
One year	1,473	11.3%	42.0	$65.05	$169,421	$2,735
Two year	1,695	2.8%	56.7	$60.77	$49,984	$3,448

Annual revenue at risk by contract segment. Month-to-month customers account for 87% of total revenue at risk despite being 55% of the customer base.

Month-to-month subscribers churn at 42.7% more than fifteen times the rate of two-year contract holders (2.8%) and this segment alone accounts for USD 1.45 million of the total USD 1.67 million in annual revenue at risk (87%). Despite this, month-to-month customers carry the lowest estimated CLV (USD 1,198) of the three segments, while two-year customers, who churn least, carry the highest (USD 3,448). This inverse relationship between churn exposure and per-customer value is the central pricing insight of this study and is discussed further in Section 5.

4.3 Cost-Sensitive Retention Threshold

Net benefit of retention outreach across churn-probability contact thresholds, under three offer-cost / retention-success scenarios. Net benefit is maximized at a low threshold (~ 0.15) in every scenario tested.

Across all three sensitivity scenarios, net benefit is maximized at a threshold near 0.15 contacting roughly 900 of the 1,409 test customers rather than at a high-precision, narrowly targeted threshold. At the base-case assumption (15% offer cost, 40% retention success), contacting customers above the 0.15 threshold catches 356 of 374 true churners at a cost of 544 false-positive contacts, yielding an estimated net benefit of approximately USD 96,300 over a 12-month horizon on the test sample alone. Net benefit declines steadily as the threshold rises: at a threshold of 0.5, the same scenario yields roughly USD 82,000 a 15% reduction because the model's precision gain from tighter targeting is outweighed by the revenue-saving value of catching additional true churners, given the underlying cost asymmetry (an unretained churner forfeits 12 months of billing, while a wasted offer costs only a few months of discount).

5. Revenue-Management and Pricing Implications

Three implications follow directly from the results in Section 4 and are relevant to pricing and retention-campaign design in subscription businesses:

First, retention spend should be prioritized by segment, not applied uniformly. Month-to-month customers generate the overwhelming majority of revenue at risk and should receive the bulk of retention-campaign budget, even though their individual CLV is lower than longer-contract customers the segment's sheer churn volume dominates the value calculation.

Second, the low optimal contact threshold (≈ 0.15) found in Section 4.3 suggests that, under typical subscription retention-campaign cost structures, broad and inexpensive outreach outperforms narrow, high-confidence targeting. This mirrors a standard revenue-management principle: when the downside of under-protection (losing a high-value customer) is large relative to the cost of over-protection (a discount offered to someone who would have stayed), the economically rational policy protects broadly rather than narrowly the same logic that justifies conservative overbooking limits or broad-based loyalty incentives in other RM contexts.

Third, the inverse relationship between churn exposure and CLV across contract segments (Table 2) supports contract-length-based pricing incentives. for example, structured discounts for customers who commit to one- or two-year terms as a demand-management lever that shifts customers from the highest-risk, lowest-value segment into lower-risk, higher-value segments, rather than relying solely on reactive, after-the-fact retention offers.

6. Limitations and Future Work

- Offer cost and retention-success rate are assumed parameters, not measured campaign outcomes; results are reported as a sensitivity range rather than a single point estimate, and should be recalibrated against real campaign A/B-test data where available.
- The CLV estimate is a simple average-tenure-times-average-charge proxy; it does not discount future cash flows or model individual-level survival probability, both standard refinements in revenue-management CLV modelling.
- The horizon (12 months) and offer duration (3 months) are fixed modelling choices; a full dynamic-pricing treatment would optimize these jointly with the contact threshold.
- All analysis uses a single cross-sectional dataset from one telecommunications provider; segment-level findings may not generalize to other subscription industries without re-estimation.
- The churn-probability model itself, while corrected for train/test leakage, uses a single stratified split rather than nested cross-validation; see the companion technical report for full classification-methodology detail and limitations.

Future work should incorporate real retention-campaign outcome data to replace the assumed cost/success parameters with estimated values, extend the CLV estimate to a discounted, survival-model-based calculation, and jointly optimize contact threshold, offer size, and offer duration as a multi-lever pricing decision rather than a single threshold choice.

7. Conclusion

This paper reframes customer churn prediction as an input to a revenue-management decision rather than a standalone classification task. Using actual billing data from a subscription telecommunications provider, we show that revenue at risk is heavily concentrated in the month-to-month contract segment (87% of USD 1.67 million annually), that this segment's high churn exposure is paired with its lowest per-customer value, and that a cost-sensitive analysis of retention-offer targeting favors broad, low-threshold outreach over narrow, high-precision targeting under typical campaign cost structures. These findings translate a predictive churn model into concrete, segment-differentiated guidance for retention-campaign budgeting and contract-length pricing incentives the kind of decision-relevant output that pure classification benchmarking, on its own, does not provide.

Declarations

Conflict of Interest

The author declares no conflicts of interest regarding the findings of this research.

Author Contribution

D.M. conceived the study, designed the hybrid predictive framework, conducted data preprocessing with SMOTE, trained all machine learning and neural network architectures, developed the Gradio deployment interface, and wrote the manuscript text.

Acknowledgement

The author gratefully acknowledges the support and encouragement of family members, whose motivation contributed greatly to this work.

Data Availability

The Telco Customer Churn dataset used in this study is a publicly available benchmark dataset. Processing and analysis code is available from the author upon reasonable request.

References

1. AbdelAziz, N. M., M. Bekheet, and A. Salah et al. 2025. A Comprehensive Evaluation of Machine Learning and Deep Learning Models for Churn Prediction. *Information* 16(7):537.

2. Asif, D., and M. S. Arif. 2025. *A Data-Driven Approach for Customer Churn Prediction*. Results in Engineering.
3. Chang, V. et al. 2024. Prediction of Customer Churn Behavior in the Telecommunication Industry. *Algorithms* ;17(6).
4. Li, Y. et al. 2025. An Improved SMOTE Algorithm for Enhanced Imbalanced Learning. *Scientific Reports* 15:325.
5. Liu, X. et al. 2024. Customer Churn Prediction Model Based on Hybrid Neural Network (CCP-Net). *Scientific Reports* 14:114.
6. Ong, J. X. 2024. Enhancing Customer Churn Prediction with Resampling. *TEM Journal* 13(2):621–628.
7. Peng, K. 2023. Research on Customer Churn Prediction and Model. *Journal of Computer Applications* ;43(9).
8. Shrestha, S. M. 2022. A Customer Churn Prediction Model using XGBoost. *Applied Computing*
9. Sikri, A. et al. 2024. Enhancing Customer Retention in Telecom Industry Using ML. *Journal of Telecommunications*
10. Gupta, P. et al. 2023. *Comparative Study of ML Algorithms for Churn Prediction*. Procedia Computer Science.
11. Kumar, R. et al. 2022. Effective Customer Churn Prediction Using Imbalanced Learning Techniques. Journal of King Saud University – Computer and Information Sciences.
12. Zhao, H. et al. 2024. *Deep Neural Networks for Customer Churn Prediction in Telecom*. Expert Systems with Applications.

Figures

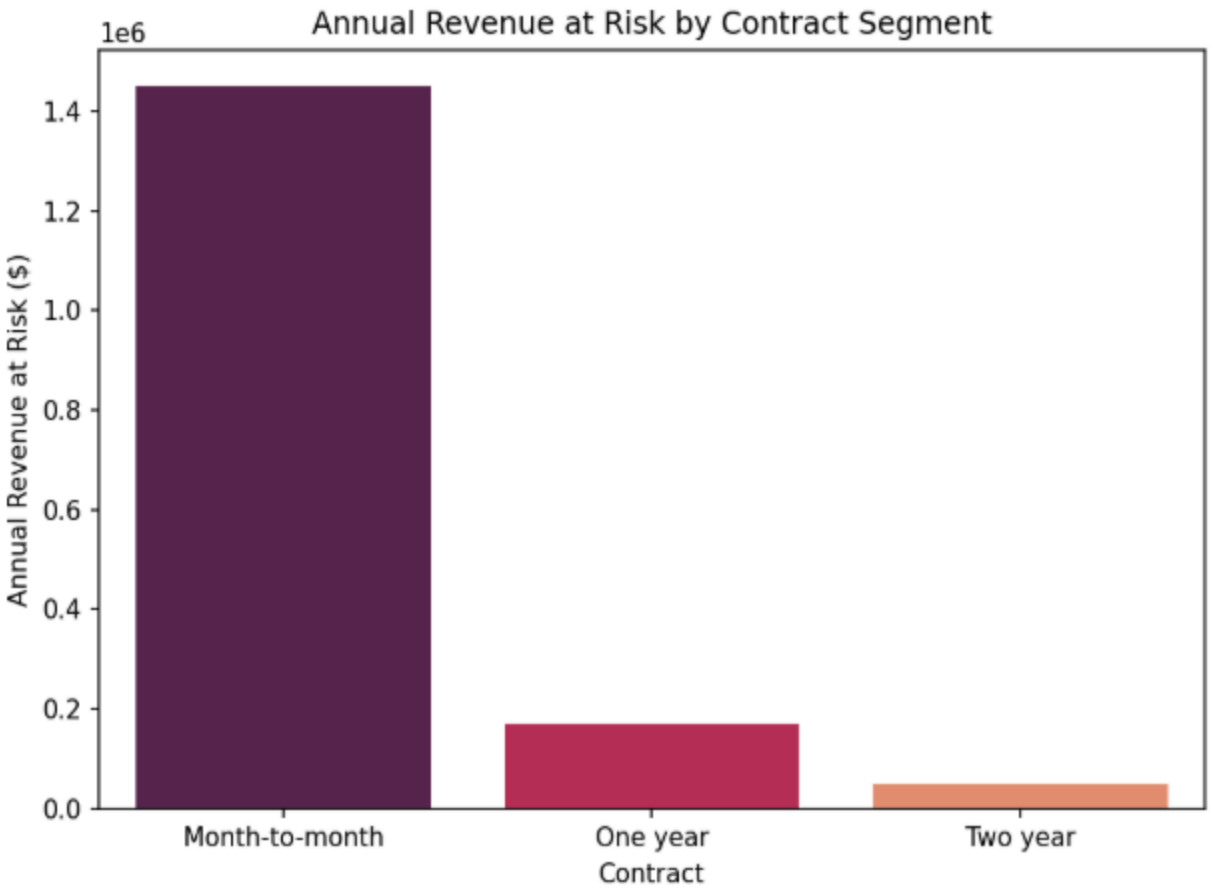


Figure 1

Annual revenue at risk by contract segment. Month-to-month customers account for 87% of total revenue at risk despite being 55% of the customer base.

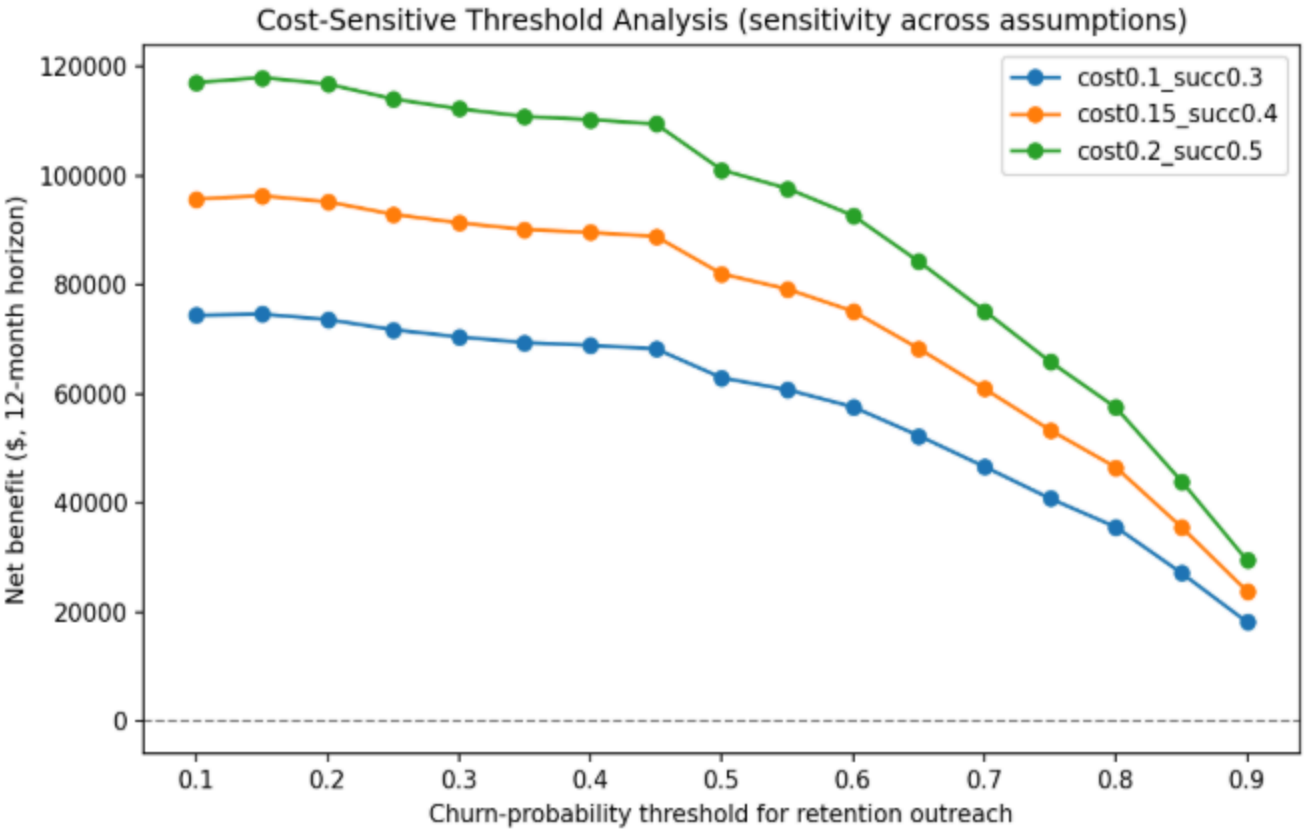


Figure 2

Net benefit of retention outreach across churn-probability contact thresholds, under three offer-cost / retention-success scenarios. Net benefit is maximized at a low threshold (~0.15) in every scenario tested.