Research Article

A Decision Support Framework for Customer Loyalty Program Managers: Reward Mix Optimization

Ayşe Salı[1*], Şahika Koyun Yılmaz[2], Meryem Ezgi Aslan[3], Gizem Temelcan Ergenecoşar[4]

[1] Metric Software and Consultancy, Orcid ID: https://orcid.org/ 0009-0005-3512-5430, E-mail: ayse.sali@metric.net
[2] Metric Software and Consultancy, Orcid ID: https://orcid.org/ 0000-0002-2589-3568, E-mail: sahika.koyun@metric.net
[3] Yildiz Technical University, Orcid ID: https://orcid.org/0000-0003-2925-2968, E-mail: muluhan@yildiz.edu.tr
[4] Beykoz University, Orcid ID: https://orcid.org/ 0000-0002-1885-0674, E-mail: gizemtemelcan@beykoz.edu.tr

[*] Correspondence: ayse.sali@metric.net

Received: 17 May 2025
Revised: 15 August 2025
2nd Revised: 18 November 2025
Accepted: 25 November 2025
Published: 27 November 2025

Reference: Salı, A., Koyun Yılmaz, Ş., Aslan, M. E., & Temelcan Ergenecoşar, G. (2025). A decision support framework for customer loyalty program managers: Reward mix optimization. *The European Journal of Research and Development, 5*(1), 228–245.

Abstract

Customer Loyalty Programs are a proven methodology for establishing and maintaining customer relationships. With the development of mobile technologies and the power of digitalization, what was once a simple punch card has now evolved into a full-fledged mobile application. The paradigm shift has opened up research areas on an individual customer level, especially in non-contractual traditional commerce, which was previously impossible due to a lack of loyalty data. The cost and budget of Customer Loyalty Programs increase with their strategic value. Balancing the attractiveness of a reward to the customer with the unit cost to the organization is essential for designing effective programs. In this study, we propose a framework that combines the attractiveness and unit cost of rewards to provide an optimized reward mix, thereby aiding Customer Loyalty Program managers in their decision-making processes.

Keywords: Convex Optimization, Customer Loyalty Programs, Multi-Layer Perceptron, Perceived Value, Unit Reward Cost

1. Introduction

The customer is long considered an asset of a firm [1]. The relationship between firms and customers is reciprocal. Customer related metrics such as customer lifetime value, customer engagement are critical with Customer Loyalty *Programs (CLPs)* as well considering the saturation of programs [2].

While attracting new customers is essential, it has been repeatedly shown that retaining existing customers is more cost-effective and holds greater potential for profitability [3]. Especially in non-contractual markets such as Fast Moving Consumer Goods (FMCG), monitoring or predicting retention, customer churn rates, and wallet-share on an individual basis is very costly, if not impossible. E-commerce has the data to monitor and predict such metrics; however, traditional commerce generates such data in the presence of *CLPs*. CLPs are a proven methodology for establishing and maintaining customer relationships. Customer engagement is critical driver of brand loyalty, and higher customer engagement rates lead to higher customer retention rates [4].

A loyalty program is a structured reward system designed to foster loyalty among repeat customers by cultivating emotional connections and promoting brand advocacy [3], [5-7]. Earlier examples include simple punch cards, where the number of purchases was tracked on a physical card, and a free product or gift was offered upon reaching a certain number of purchases. With the development of mobile technologies, these physical programs have been transferred to the digital environment, revealing a potential beyond ensuring repeat purchases. In their study, Andia-Reyna and Malasques-Villanueva show that "*… the emerging technologies significantly impact customer loyalty*", and data science is the key tool for loyalty programs to reach their full potential [8-9]. CLPs, by nature, are programs designed to incentivize repeat purchases, and with the power of digitalization, they have become spaces for gathering data on an individual level. Traditional markets, where most commerce is still face-to-face, can track trends, satisfaction levels, purchase periods, and other key metrics. This rich and valuable data has the potential to expand research and implementation of predictive analysis [7,10].

The research around CLPs is clustered around organizational and customer perspectives [3,5,7,8,11]. The common ground of these studies is that both organizations and customers benefit from loyalty programs by strengthening their relationships. Although there are two perspectives, studies show that they are complementary, each influencing the other. In their literature review [11] categorize research elements as psychological, design, and operational elements, illustrating the complementary nature of organization and customer. For example, the *perceived value* (customer value) of a reward on the customer side is reflected as the *unit reward cost* to the organization. With the enriched data available through digitalized programs, the research is clustered around topics such as customer experience, individualized reward targeting, and pattern and trend identification [8].

The Loyalty Program Trends 2025 reports that the top three trends concerning loyalty programs are Customer Lifetime Value, lowering customer churn, and increasing purchasing frequency [2]. These trends are closely related to customer engagement, therefore it is critical to design and manage CLPs that will drive more engagement.

In this study, we propose a framework that combines perceived value (reward attractiveness) and unit reward cost to provide an optimized reward mix, aiding Customer Loyalty Program managers in budgeting and operational processes. Masked Multi-Layer Perceptron models, and convex optimization models are developed for generating probability distribution functions for reward subsets.

This paper is organized as follows: In Section 2, we provide definitions for the problem and relevant terminology. In Section 3, we outline the methodology for the proposed framework. In Section 4, we summarize the results from a case study. Finally, in Section 5, we discuss the results and explore further research topics.

2. The Case Study and Terminology

In this study, we examine a case of a customer loyalty program campaign in the FMCG sector. The program is a mobile app that allows customers to claim rewards upon purchasing campaign-affiliated products. A campaign has a predefined reward set, and each reward is a subset of rewards that vary in value. A customer can claim a single reward for each purchase, but is not limited to a maximum number of rewards. The exact amount of rewards or a sub-reward is not predetermined. The campaign budget constraint limits the total amount of rewards that can be allocated. This amount depends on the distribution of rewards. The program manager plans the distribution ratio for each reward subset based on previous campaigns and experience. The most significant issues regarding the current planning process are as follows1:

- The planning process is person-dependent and manager-oriented.

The lack of organizational memory and experience-oriented decision-making leads to inefficient campaign management. When management changes, the knowledge disappears, and the experience becomes vague at best. The learning curve for managing a campaign effectively varies from individual to individual, contributing to the inefficiency of campaign management.

- Reward preference trends are not tracked, and do not influence the planning process.

[1] Not in importance order.

The driving force of the current decision-making is to maximize the total number of claims by the customer, therefore increasing sales. The perceived value of a reward affects the reward preference trends. Consequently, the preference trends influence which rewards are highly attractive, thus customers try to claim more. However, preference trends are not taken into account, leading to inefficient use of campaign budget.

- The cost of reward distribution cannot be predicted due to a lack of reward preference tracking.

The reward distribution ratio produces an expected value for the daily cost, ideally. From a temporal perspective, the rewards that the customer will highly prefer depend on the previous day's reward distribution. The current planning process assumes independence between rewards and no temporal effect. As a result, even though an expected value for cost can be calculated, it is not a reliable prediction.

- The budget distribution is not uniform throughout the year. The daily cost of total rewards distributed decreases unpredictably throughout the year.

As mentioned earlier, the driving force behind the current process is to maximize the total number of claims by customers, and decision-making is person-dependent. This leads to subjective allocation, based on personal views and experience. While one manager chooses to allocate a larger budget at the beginning of the year, another chooses to allocate it more uniformly throughout the year.

- The reward distribution is not uniform throughout the year.

The personal preferences of managers also affect the distribution of rewards. Highly preferred rewards are typically used up earlier in the year and cannot be distributed uniformly throughout the year. As a result, customers' interest in campaigns tends to decline towards the end of the year. The distribution ratio within reward subsets is crucial to maintaining customer interest throughout the year. However, the current planning process handles the reward as independent.

2.1.Terminology

A customer loyalty program typically has three essential components: rewards, customers, and the delivery system. The case we are examining also comprises these three components, along with metrics such as unit reward cost, total cost, and total rewards claimed (TRC). In addition, we propose a new metric, the attractiveness of a reward, to measure the relative attractiveness of rewards within the active rewards set. In this section, we provide definitions we use in the case and in the framework we developed for its solution.

2.1.1. Customer

A customer is an individual who purchases items affiliated with the campaigns and uses the mobile app to claim a reward. Customers are not limited by a maximum number of rewards they can claim within a day or throughout the campaign. Each transaction is eligible for a reward, provided the customer is a program member and has accepted the terms and conditions of the program.

2.1.2. The Delivery System

The customer loyalty program is managed with a mobile application. Customers download the application from application stores and sign up with a valid email address or phone number. Then they can claim rewards proving their purchase through the application. Upon approval of the purchase, the application lists the active rewards for the customer to select from. Finally, a voucher is rewarded according to the reward's distribution probability function, as given in Eq 1.

$$\boldsymbol{P(x_i)} = \begin{cases} \boldsymbol{p(x_{i1})} \\ \boldsymbol{p(x_{i2})} \\ \vdots \\ \boldsymbol{p(x_{ij})} \end{cases}, \sum_{j} \boldsymbol{p(x_{ij})} = \mathbf{1} \; \boldsymbol{and} \; \forall \; \mathbf{0} \leq \boldsymbol{p(x_{ij})} \leq \mathbf{1} \qquad \textit{Eq 1}$$

2.1.3. Rewards

A reward is a voucher for partner organizations with variable values. The campaign rewards set ($I = (1, \ldots, i)$) is constituted of partner organizations. The campaign reward set is the union of the active (I_{active}) and inactive ($I_{inactive}$) reward sets. Each reward is a set of vouchers ($J = (1, \ldots, j)$) with a range of values corresponding to a unit reward cost (c_{ij}). The value of vouchers and the total budget are contractually agreed upon at the beginning of the year. Managers' preferences shape how the vouchers are distributed throughout the year.
For each reward i, the daily expected cost ($E(c_i)$) is calculated as given in Eq 2. Then the daily average unit reward cost ($\tilde{c}_i$) is calculated as given in Eq 3.

$$\boldsymbol{E(c_i)} = \sum_{j} \boldsymbol{c_{ij} x_{ij}}, \qquad \forall \boldsymbol{i} \qquad \textit{Eq 2}$$

$$\boldsymbol{\tilde{c}_i} = \frac{\sum_j \boldsymbol{c_{ij} x_{ij}}}{\sum_j \boldsymbol{x_{ij}}}, \qquad \forall \boldsymbol{i} \qquad \textit{Eq 3}$$

The daily total expected cost ($E(c)$) is calculated as given in Eq 4

$$\boldsymbol{E(c)} = \sum_{i}\sum_{j} \boldsymbol{c_{ij}x_{ij}} \qquad \text{Eq 4}$$

The total number of rewards claimed daily is defined as the *Total Reward Claim* (TRC) and calculated as given in Eq 5. TRC is the metric that measures the amount of customer engagement with the campaign.

$$\boldsymbol{TRC} = \sum_{i}\sum_{j} \boldsymbol{x_{ij}} \qquad \text{Eq 5}$$

where $\boldsymbol{x_{ij}}$ is the number of claimed voucher j for reward i.

The attractiveness of a reward is the daily rate at which customers select it, given the active set of rewards. It is calculated for the historical data as the share of total claimed vouchers for each reward $\boldsymbol{i}$ in $\boldsymbol{TRC}$**,** as given in Eq 6.

$$\boldsymbol{a_i} = \frac{\sum_j \boldsymbol{x_{ij}}}{\boldsymbol{TRC}}, \qquad \forall \boldsymbol{i} \qquad \text{Eq 6}$$

2.2.Data

The historical data provided within the scope of the case study covers a period of 18 months and consists of the number of reward vouchers claimed daily. The active/inactive status of each reward varies throughout the data period. No reward is active one day and inactive the next; active/inactive statuses remain stable for at least a month for each reward.

There are 17 rewards, each with a varying number of vouchers, defined throughout 18 months. The size of a reward subset varies between 1 and 6, and there are a total of 59 distinct vouchers at maximum throughout the data period. Similar to rewards, the active/inactive statuses of vouchers change throughout the data. Therefore, the size of a reward subset is capped at its maximum value found in the historical data.

In this study, we used the first 12 months of data for training strictly. The latter 6 months are incrementally added weekly to simulate a real planning process. Although a forecasting phase is designed within the framework, it has not been implemented at this stage in order to increase the accuracy of the performance evaluation by reducing data variability.

The proposed framework consists of five main parts: (i) data pre-processing, (ii) forecasting, (iii) prediction, (iv) optimization, and (v) reporting. The forecasting and prediction phases are parallel, as they do not feed information to each other. The optimization phase follows the forecasting and attractiveness prediction phases.

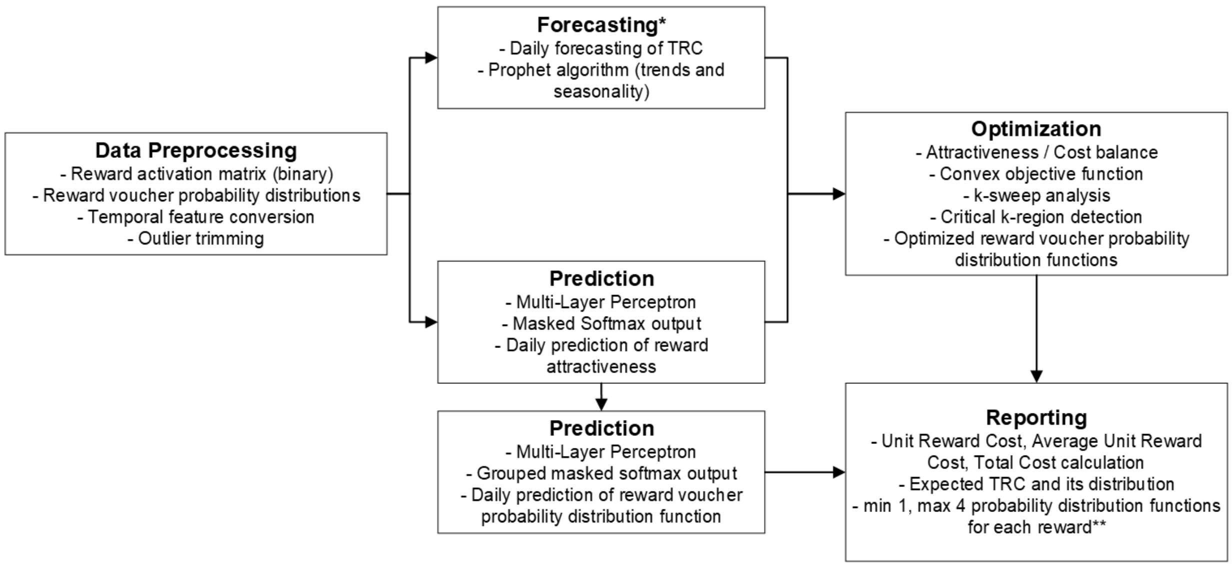

Figure 1 Proposed framework workflow[2]

3.1.Data Pre-Processing

The historical data provided consists of the number of daily reward claims for each reward's vouchers (x_{ij}). The pre-processing steps are as follows:

- Data transformation

The historical data did not include the active/inactive status of rewards; therefore, we transformed the data under the assumption that if at least one voucher is claimed in a reward subset, *then that reward was active*. If a reward's status suddenly changes for a single day, it is considered an outlier case and is replaced with the opposite value to ensure consistency.

- Temporal transformations

To capture the seasonal characteristics of the campaign, temporal features (month, day of the month, week, day of the week, and day of the year) were obtained using sine and cosine transformations.

[2] * Forecasting phase is designed as a part of the framework, however it is not implemented as part of this study. Details are in Section 3.2.
** For each day at least one probability distribution function for each reward is generated with the prediction phase. If optimization phase yielded feasible solutions three critical k values are selected. As a result a maximum of four probability distribution functions are presented for any given day.

- Outlier trimming

To improve data quality, outlier analysis is performed, and outlier values are labeled. Since outliers do not appear in all reward subsets within a single day, their values are replaced with the weighted average of the previous 3 days and the following 3 days.

- Performance metric calculations

Average unit reward costs, total reward costs, total daily cost, and total number of rewards claimed (TRC) with simple calculations.

- Attractiveness calculation

Since the attractiveness of a reward (a_i) is the share of total claimed vouchers of the reward subset in the TRC, a_i is calculated as given in Eq 6. The daily calculated a_i ratios form a vector whose elements sum to 1.

- Reward probability distribution function calculation

A reward probability distribution function is a discrete probability function, as defined in Eq 1. Each probability is the share of the actual vouchers claimed over the actual total vouchers claimed for that reward subset (Eq 7).

$$p(x_{ij}) = \frac{x_{ij}}{\sum_j x_{ij}}, \qquad \forall i, j \tag{Eq 7}$$

3.2.Forecasting

Total reward claim (TRC) time series obtained in the data pre-processing phase is used to forecast TRC_{t+1} value with the Prophet model [12]. The Prophet model is chosen due to its efficacy in modeling trends, seasonality, and special periods (such as limited-time vouchers and product launches) [12-13]. Forecasted TRC_{t+1} value is then used to calculate the performance metrics of predicted and optimized solutions.

Within the scope of this case study, production results are still being accumulated for evaluation. Therefore, we do not use forecasted TRC_{t+1} values to reduce data variability in performance evaluation for the testing period.

3.3.Prediction

The prediction phase is two-fold: In the first phase, the active/inactive reward matrix serves as the input matrix, and the attractiveness matrix represents the target values. Using a multi-layer perceptron model (Figure 2) with a masked softmax output setup [14-16], we predict an attractiveness vector. This vector is then used to distribute the actual TRC, and fed as a parameter to the optimization model. This model excludes inactive input characteristics using a binary mask, producing a single integrated probability distribution based on the remaining valid characteristics.

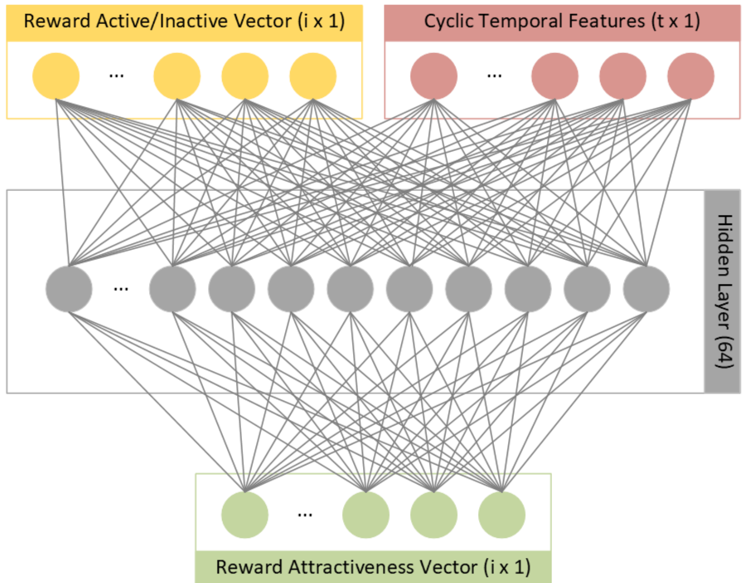

Figure 2 Masked Softmax Output MLP

The second phase produces probability distribution function values for rewards with a multi-layer perceptron model (Figure 3). The output is set up as masked grouped softmax [14-16] to ensure each reward's probability function sums to 1. The input data is the historical attractiveness of rewards, and the output is the corresponding probability distribution function values. The purpose of this model is to process semantically related feature groups separately, producing independent probability distributions for each reward. In the output, there are n active rewards[3], there are n vectors where the sum of the elements of each vector equals 1. Consequently, the second model ensures that probability values are only produced for the active rewards.

[3] The corresponding value in the input vector is greater than 0.

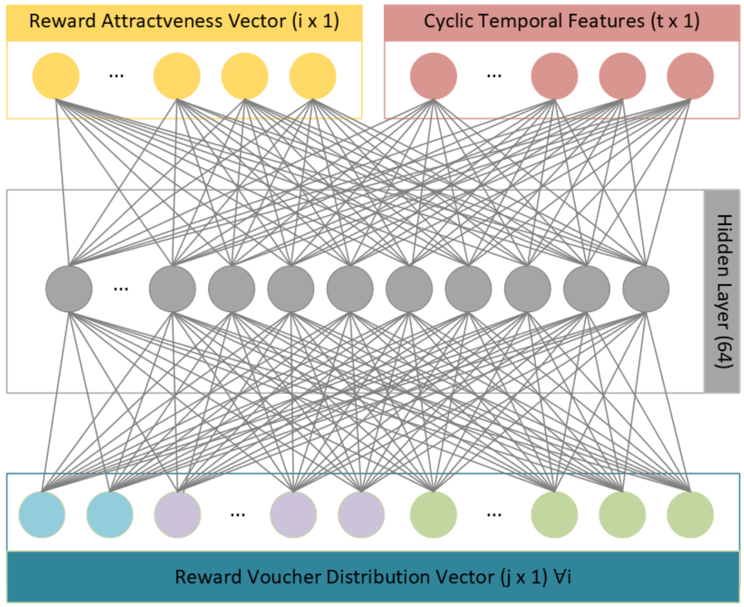

Figure 3 Masked and Grouped Softmax Output MLP

Both model takes a two-part input vector ($\vec{I}_{(t+i)\times 1}$). Here t is the number of temporal features of the corresponding date as sine and cosine transformations, while i is the number of rewards. Reward active/inactive vector's elements are binary, while reward attractiveness vector's elements belong to the half-open interval [0,1). The models are designed with a single Rectified Linear Unit (ReLU) hidden layer. ReLU is preferred because it enables the model to learn nonlinear relationships and offers computational efficiency.

To ensure the models learn from the active reward set, a masked loss strategy was implemented. Instead of using a standard loss function that averages over all outputs, we first computed the per-element loss (squared error) to produce a loss tensor. This tensor was then subjected to an element-wise multiplication with the input mask, effectively nullifying the error for any inactive rewards. The final loss was calculated by summing the remaining errors and normalizing by the number of active rewards. This approach guarantees that the inactive rewards do not contribute to the gradient updates during training.

In summary, the first phase aims to predict customer behavior, while the second phase aims to determine which probability distributions are most suitable for creating predicted attractiveness ratios.

3.4.Optimization

The classical linear utility functions assume that marginal utility is constant; therefore, when the cost component increases, the solution favors low-cost rewards, creating *corner*

point clusters in the solution space. This leads to a decrease in reward variety and high-cost rewards being excluded from solution sets [17]. Actual user behavior, however, exhibits diminishing marginal utility. Therefore, the utility term in the model is defined in logarithmic form. Thus, (i) overloading with low-cost rewards is prevented, (ii) high-cost rewards are represented at a reasonable level in solutions, and (iii) the cost versus reward attractiveness diversity balance is maintained.
From mathematical perspective $log(1 + x_{ij})$ is strictly concave over the entire domain, and the cost component $\sum_{ij} c_{ij}x_{ij}$ is linear. The difference between these two components (Eq 8) results in a concave objective function.

$$\max z = \sum_i a_i \log(1 + x_{ij}) - k \sum_i \sum_j c_{ij}x_{ij} \qquad \textit{Eq 8}$$

where the k coefficient is a value between 0 and 1 that determines the weight of the attraction over the cost ratio. Since all constraints in the model are linear, the solution space is convex. Therefore, the model corresponds to the optimization class defined in the literature as *concave maximization over a convex feasible set*. According to convex optimization theory, such problems have a single global optimum, and CVXPY's Discipline Convex Programming (DSP) framework directly supports this structure [19]. As a result, the model can produce stable, repeatable, and computationally reliable solutions under large-scale campaign conditions [20].
As a result, the model balances cost and customer preferences (reward attractiveness) according to the k coefficient. As the value of k increases, the model begins to behave more frugally, meaning it prefers low-cost rewards over high-cost ones. As the value of k decreases, the opposite occurs: the model prefers high-cost rewards over low-cost rewards (Figure 4, Eq 9).

$$\begin{gathered} 0 \leq k \leq 0.25 \Rightarrow \text{Generous distribution} \\ 0.25 < k < 0.75 \Rightarrow \text{Normal distribution} \\ 0.75 \leq k \Rightarrow \text{Frugal distribution} \end{gathered} \qquad \textit{Eq 9}$$

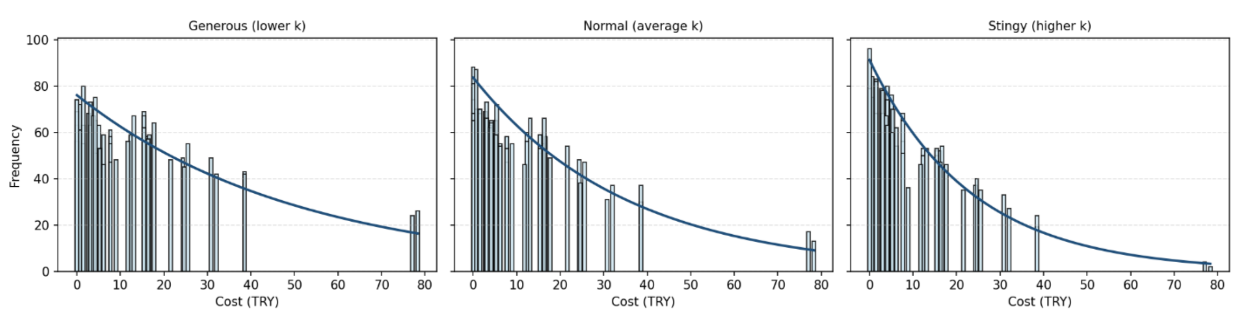

Figure 4: Optimal reward distribution scenarios based on attractiveness/cost ratio (k)

3.4.1. Parameter Sweeping and Critical k Selection

The model is solved in the interval $k = (0,1]$, increasing k by 0.01 each iteration, and the total cost ($M(k)$) of the solution is recorded. Points of sudden change[4] are identified on the resulting cost curve. These points represent critical regions where the system undergoes structural transformation in its cost equilibrium. Consequently, it is possible to track at which k values the system changes its behavior.

Additionally, sweeping k uniformly in the (0,1] range allows for a systematic examination of how system behavior changes across different regimes. This parameter sensitivity analysis-based approach is widely used in multi-scenario generation and decision support models [21]. Each k scenario is evaluated comparatively using total utility, total cost, reward diversity, and constraint compliance metrics. This structure allows for the identification of structural breakpoints (critical k intervals) in model behavior, allowing the campaign manager to analytically evaluate the most efficient equilibrium level.

3.4.2. Final Solution

The model is re-solved for the specified critical k values, and the optimal distribution $[x]^*_{k_{critical}}$ is obtained for each solution. These distributions represent the system's behavior at different budget sensitivities (generous, frugal). The model reveals not only a single optimal point but also multiple equilibrium points valid at different levels of k. Thus, the campaign manager can select the equilibrium point that aligns with their strategic priorities among total utility, budget utilization, and reward diversity.

In the final stage, feasible solutions are generated by integerization, min-max quantities, stock and lot limits, and operational constraints. When necessary, soft constraints or penalty functions are used to select the best feasible solutions, which are the closest to the theoretical optimum. This step ensures the integration of the theoretical model with field conditions, in accordance with the literature on multi-dimensional assignment problems [18].

3.4.3. Constraints and The Complete Model

- Demand band constraints (Allowed deviation in TRC), (Eq 10)

This constraint ensures that the total demand (TRC) remains within deviation range ($\pm r$). Thus, the model can produce a solution that is sensitive to forecasting errors but stable.

- Budget constraint (Eq 11)

This constraint ensures that the total cost does not exceed the specified budget (B_i). In scenario-based studies, it is activated or deactivated to test optimal distributions under different budget levels.

[4] The regions with the maximum $\Delta M(k)$

- Reward variety constraint (Eq 12)

This constraint requires that at least one voucher is given for each active reward. It prevents the complete exclusion of high-cost rewards, thereby maintaining the balance of diversity and fairness in the distribution.

- Decision variable domain (Eq 13)

The mathematical model for the convex optimization is as follows:

$$\max z = \sum_i a_i \log(1 + x_{ij}) - k \sum_i \sum_j c_{ij} x_{ij}$$

st:

$$\begin{aligned} &\sum_i x_{ij} \geq TCS(1-r) \\ &\sum_i x_{ij} \leq TCS(1+r) \end{aligned} \qquad \text{Eq 10}$$

$$\sum_j c_{ij} x_{ij} \leq B_i, \quad \forall i \qquad \text{Eq 11}$$

$$x_{ij} \geq 1, \quad \forall i \in I_{active} \qquad \text{Eq 12}$$

$$x_{ij} \geq 0, \quad \forall i \in I \qquad \text{Eq 13}$$

where x_{ij} is the daily amount of voucher j of reward i (decision variable), a_i is the attractiveness ratio of reward i, c_{ij} is the unit reward cost of voucher j of reward i, $k \in (0,1]$ is the ratio of $^{Attractiveness}/_{Cost}$, and $r \in [0,0.05]$ is the reward claim deviation ratio.

3.5. Reporting

The performance metrics are calculated for predicted and optimized solutions, and presented to campaign managers for evaluation and decision-making. The results are presented in comparison to aid in the decision-making process; however, the final decision remains with the manager.

4. Results

The test was conducted over a 5-month (146-day) period. During the optimization phase, the 3 most critical k-value scenarios were selected for each day, resulting in a total of 432 scenarios. No feasible solution was found for 2 days during the 146-day test period. The scenarios suggest that a trade-off between the attractiveness of a reward and the unit reward cost is often the most optimal solution. When the scenarios were examined daily, there were only three days in which all three scenarios (one generous, one normal, and one frugal) were present. The number of days where only one type of scenario was presented was 30 days for the generous, 26 days for the normal, and 10 days for the frugal scenario. 39% of all the scenarios suggested a generous distribution, while 18% suggested a frugal distribution.

The predicted distributions resulting from the prediction phase show a less steep trend, and decrease towards the end of the testing period, compared to the actual costs (Figure 5).

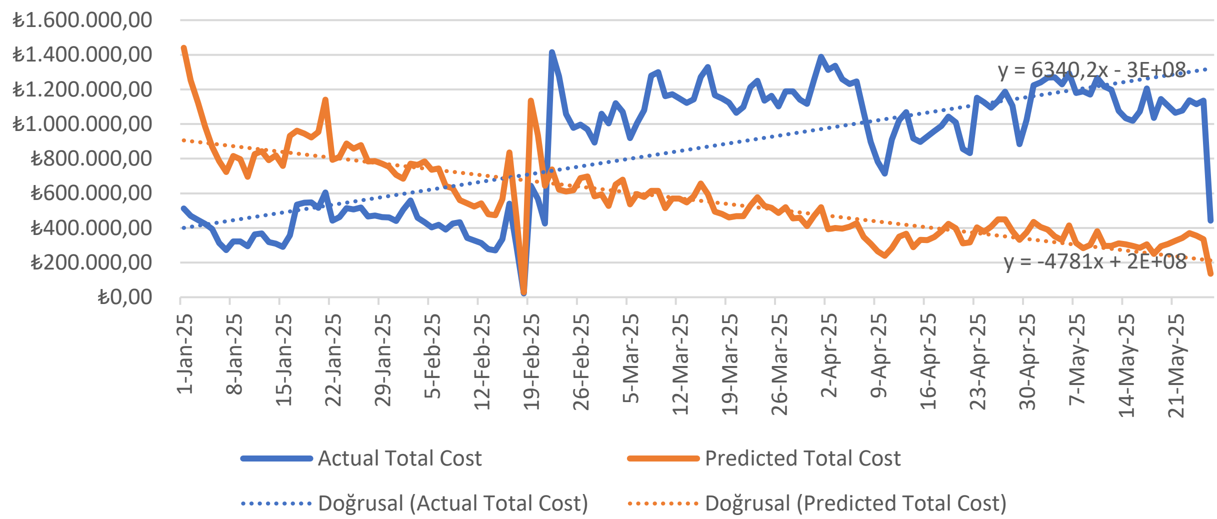

Figure 5 Daily Total Reward Claimed Cost (TRY)

Since the distributions obtained from the optimization phase are for different k-values for each day, they cannot be directly compared with the actual costs. Table 1 shows the 3 days where there were three types of scenarios present.

Table 1 Daily Total Reward Claim Cost Comparison[5]

			Optimized		
Date	Actual	Predicted	Generous	Normal	Frugal
Sample Day 1	₺507.752,47	₺858.411,58	₺1.595.723,32 $k = 0.05$	₺1.551.922,15 $k = 0.27$	₺1.414.467,95 $k = 0.82$
Sample Day 2	₺968.265,02	₺696.541,91	₺260.605,07 $k = 0.03$	₺172.133,76 $k = 0.63$	₺242.672,53 $k = 0.93$
Sample Day 3	₺893.907,91	₺583.602,55	₺228.060,34 $k = 0.02$	₺224.900,44 $k = 0.55$	₺214.472,06 $k = 0.90$

Four rewards with voucher subset sizes greater than 1 are evaluated in terms of total cost. In all four cases, the predicted cost trend is less steep than the actual cost trend, similar to the total cost comparison (Figure 6-9).

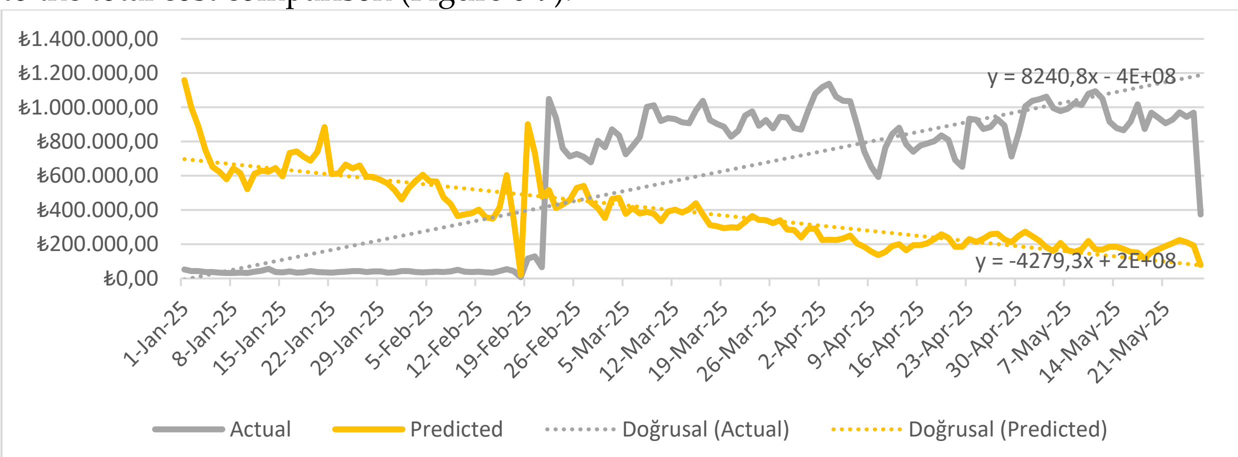

Figure 6 Reward 1 Daily Total Claim Cost

[5] A fixed rate of 40TRY/USD

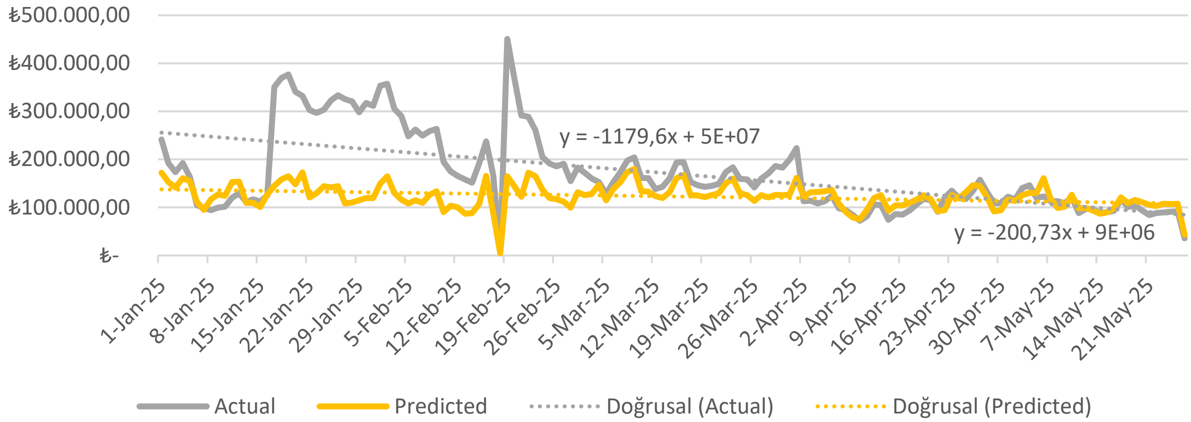

Figure 7 Reward 2 Daily Claim Cost

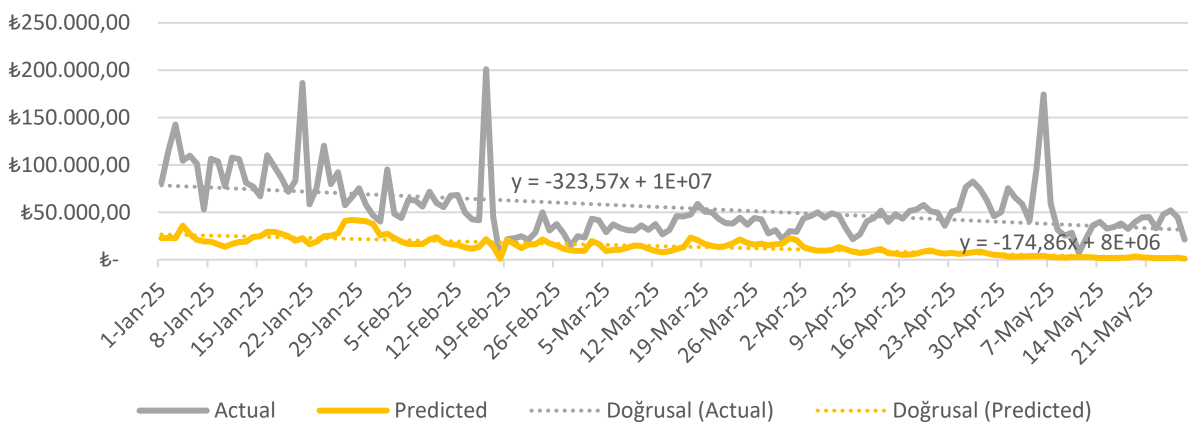

Figure 8 Reward 3 Daily Claim Cost

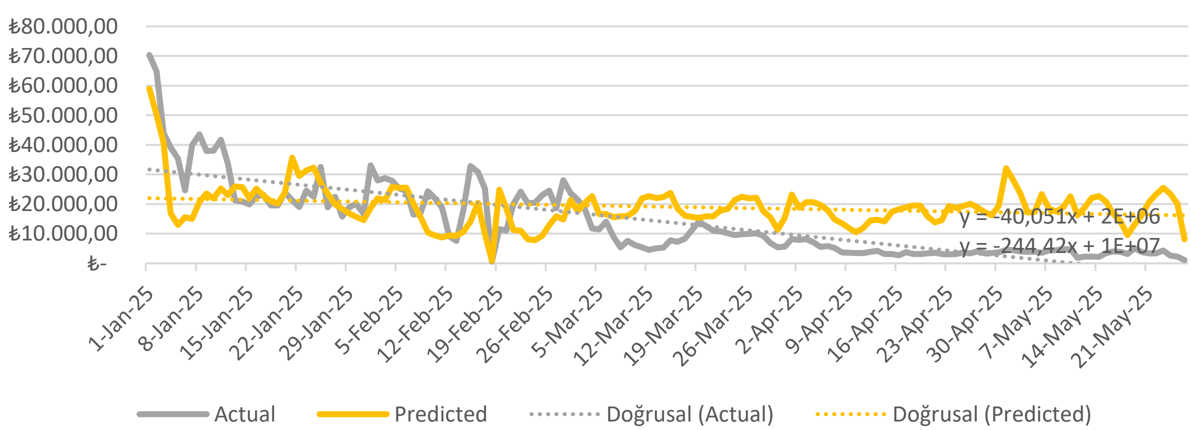

Figure 9 Reward 4 Daily Claim Cost

Figure 10 illustrates the voucher variety for four rewards during the testing period.

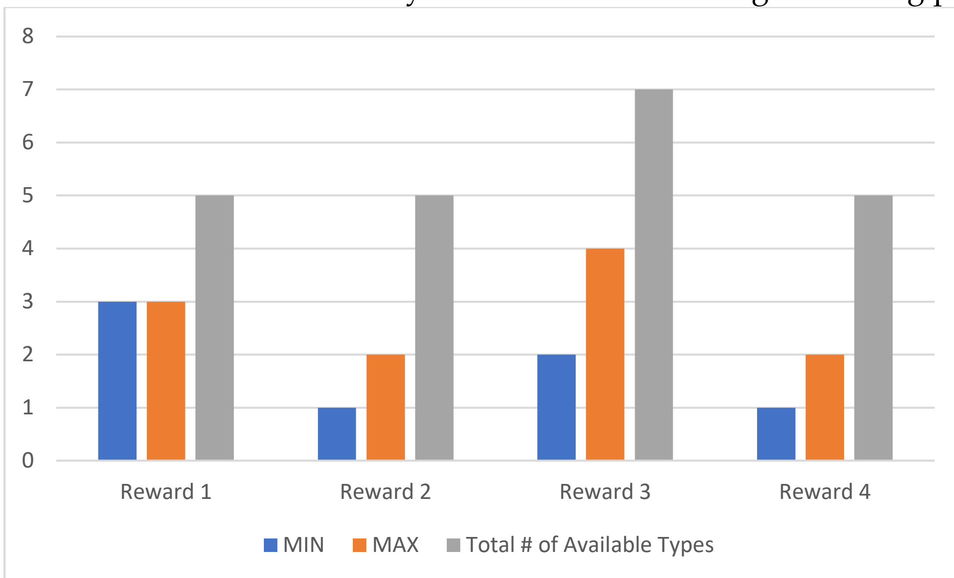

Figure 10 The Number of Active Voucher Variability for Rewards (Prediction Results)

5. Conclusions

In this study, we examined a case study in managing daily reward and voucher distribution for a customer loyalty program. We propose a framework to (i) forecast daily customer engagement, (ii) predict customer behavior in regard to reward attractiveness, then predict a voucher probability distribution, and finally (iii) search for alternative optimal solutions. Results show that a less steep trend for daily Total Reward Claimed cost with the prediction model. The voucher variety within rewards can be improved with increasing actual samples with high variety to train predictive models. There are alternative solutions with different levels of weight for attractiveness over cost ratio. In conclusion, the results show that an analytical framework for managing customer loyalty program can provide a budget efficient solution set than experience-oriented approach.

References

[1] Gupta, S., & Lehmann, D. R. (2003). Customers as assets. *Journal of Interactive Marketing, 17*(1), 9-24.

[2] Open Loyalty. (2025). *Loyalty Program Trends 2025.* Open Loyalty. 11 7, 2025 tarihinde https://www.openloyalty.io/resources/loyalty-program-trends

[3] Mary, B. (2025, May 26). *Customer loyalty programs and their long-term impact.* doi:http://dx.doi.org/10.2139/ssrn.5269142

[4] Zaware, N., & Deokate, R. (2025). Customer engagement impact on brand loyalty: a comprehensive analysis. *Asian Journal of Management and Commerce, 6*(1), 590-591.

[5] Bariha, P. P. (2020). Customer loyalty program and retention relationship. *Psychology and Education, 57*(9), 5069-5074.

[6] Kim, B.-D., Shi, M., & Srinivasan, K. (2001). Reward programs and tacit collusion. *Marketing Science, 20*(2), 99-120.

[7] Chen, Y., Mandler, T., & Meyer-Waarden, L. (2021). Three decades of research on loyalty programs: A literature review and future research agenda. *Journal of Business Research*, 179-197. doi:https://doi.org/10.1016/j.jbusres.2020.11.057

[8] Andia-Reyna, J., & Malasquez-Villanueva, Y. (2025). Impact of emerging technologies on customer loyalty: a systematic review. *International Journal of Advanced Computer Science and Applications, 16*(3), 204-212.

[9] Danaher, P. J., Sajtos, L., & Danaher, T. S. (2020). Tactical use of rewards to enhance loyalty program effectiveness. *International Journal of Research in Marketing, 37*(3), 505-520. doi:https://doi.org/10.1016/j.ijresmar.2020.02.005

[10] Son, Y., Yim, D., & Oh, W. (2017). Loyalty analytics: predicting customer behavior using reward redemption patterns under mobile-app reward scheme. *Thirty Eighth International Conference on Information Systems.*

[11] Kim, J. J., Steinhoff, L., & Palmatier, R. W. (2021). An emerging theory of loyalty program dynamics. *Journal of the Academy of Marketing Science, 49*, 71-95. doi:https://doi.org/10.1007/s11747-020-00719-1

[12] Zunic, E., Korjenic, K., Hodzic, K., & Donko, D. (2020). Application of Facebook's prophet algorithm for successful sales forecasting based on real-world data. *International Journal of Computer Science & Information Technology, 12*(2), 23-36.

[13] Nadar, A. T., Chandane, S., Raj, G. N., Pasi, N. M., & Patil, Y. A. (2023). Automated energy billing with blockchain and the prophet forecasting model: a holistic approach. *IEEE International Conference on Multidisciplinary Research in Technology and Management MRTM 2023.*

[14] Guo, H., Tang, R., Ye, Y., Li, Z., & He, X. (2017). DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. https://arxiv.org/abs/1703.04247

[15] Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., . . . Liu, X. (2016). Wide & Deep Learning for Recommender Systems. *CoRR.* https://arxiv.org/abs/1606.07792

[16] PyTorch Foundation. (2025). *PyTorch documentation*. PyTorch: https://docs.pytorch.org/docs/stable/index.html

[17] Ammar, H. A., Adve, R., Shahbazpanahi, S., Boudreau, G., & Srinivas, K. V. (2022). Downlink Resource Allocation in Multiuser Cell-Free MIMO Networks With User-Centric Clustering. *IEEE Transactions on Wireless Communications, 21*(3), 1482-1497.

[18] Bertsekas, D. (2020). Constrained Multiagent Rollout and Multidimensional Assignment with the Auction Algorithm. https://arxiv.org/abs/2002.07407 adresinden alındı

[19] CVXPY. (2025). *Open source Python-embedded modelling language for convex optimization*. https://www.cvxpy.org/

[20] Khan, A. A., Adve, R. S., & Yu, W. (2020). Optimizing Downlink Resource Allocation in Multiuser MIMO Networks via Fractional Programming and the Hungarian Algorithm. *IEEE Transactions on Wireless Communications, 19*(8), 5162-5175.

[21] Gutin, G., & Karapetyan, D. (2009). A Memetic Algorithm for the Multidimensional Assignment Problem. *Engineering Stochastic Local Search Algorithms. Designing, Implementing and Analyzing Effective Heuristics*, (s. 125-129).