

HEC MONTRÉAL
École affiliée à l'Université de Montréal

Essays on Pricing and Marketing Strategies in Digital Service Platforms

par
Melina Asadi

Thèse présentée en vue de l'obtention du grade de Ph. D. en administration
(spécialisation Ph.D. in Decision Science)

Janvier 2026

© Melina Asadi, 2026

HEC MONTRÉAL
École affiliée à l'Université de Montréal

Cette thèse intitulée :

Essays on Pricing and Marketing Strategies in Digital Service Platforms

Présentée par :

Melina Asadi

a été évaluée par un jury composé des personnes suivantes :

Maryam Daryalal
HEC Montréal
Présidente-rapportrice

Georges Zaccour
HEC Montréal
Directeur de recherche

FirstName LastName
HEC Montréal
Codirecteur de recherche

Emine Sarigollu
McGill University
Membre du jury

Bacel Maddah
American University of Beirut
Examineur externe

David Benatia
HEC Montréal
Représentant du directeur de HEC Montréal

Résumé

Nous étudions les stratégies de tarification optimales pour les plateformes numériques visant à maximiser leurs profits. Dans un premier temps, nous nous concentrons sur la tarification des menus pour les plateformes de services à la demande avec des paramètres de service fixes. Dans un second temps, nous étendons l'analyse en considérant les décisions endogènes relatives aux remises et à la qualité. Enfin, nous étudions une plateforme freemium optimisant ses efforts d'acquisition et de fidélisation, ainsi que ses décisions en matière de publicité et de tarification.

Dans le premier article, nous cherchons à déterminer la viabilité de la tarification à la carte pour les plateformes à la demande et la manière optimale de fixer les frais d'abonnement et de transaction. Nous prenons en compte des plateformes telles qu'Uber One et DoorDash DashPass, qui proposent des options de paiement à l'utilisation et d'abonnement premium. Les utilisateurs choisissent eux-mêmes leur option en fonction de leur fréquence d'utilisation et de l'évaluation du service, tandis que les taux de remise et la qualité du service restent fixes. Nous identifions un seuil de remise critique qui modifie fondamentalement les modèles de revenus des plateformes : en dessous de ce seuil, les plateformes s'appuient sur les revenus transactionnels avec des frais d'abonnement minimaux ; au-delà, elles privilégient les revenus issus des abonnements avec des frais de transaction réduits.

Dans le second article, nous analysons comment les stratégies de tarification optimales évoluent lorsque les plateformes peuvent choisir conjointement les frais d'abonnement, les prix des transactions, les taux de remise et la qualité du service. Nous examinons ce

phénomène à travers des scénarios d'optimisation de plus en plus complexes, depuis un modèle de référence avec des remises et une qualité fixes jusqu'à une optimisation complète des quatre variables. Notre étude vise à identifier les leviers stratégiques les plus importants lorsque les plateformes disposent d'un contrôle accru, à comprendre comment la structure du marché réagit à un champ de décision élargi et à déterminer si la segmentation à trois niveaux, couramment utilisée, reste optimale.

Dans le troisième essai, nous étudions l'optimisation des plateformes freemium afin de vérifier si la logique de seuil s'étend au-delà des contextes à la demande. Nous modélisons la base d'utilisateurs comme une chaîne de Markov à trois états : non-utilisateurs, utilisateurs gratuits et utilisateurs premium. La plateforme optimise simultanément les dépenses d'acquisition, les investissements de fidélisation, l'intensité publicitaire et le prix des abonnements. Les probabilités de transition sont des fonctions endogènes des décisions de la plateforme, reflétant l'impact des choix dans un domaine sur les résultats dans d'autres. Nous nous attachons à identifier les seuils critiques d'amélioration de la qualité, à comprendre les compromis entre acquisition et fidélisation et à déterminer comment l'efficacité de la fidélisation influence le pouvoir de fixation des prix.

Mots-clés

Tarification par les plateformes ; Conception des abonnements ; Qualité de service ; Services à la demande ; Maximisation du profit ; Hétérogénéité des clients

Méthodes de recherche

Recherche quantitative ; Analyse numérique

Abstract

We investigate optimal pricing strategies for profit-maximizing digital platforms. First, we focus on menu pricing for on-demand service platforms with fixed service parameters. Second, we extend the analysis by considering endogenous discount and quality decisions. Finally, we consider a freemium platform optimizing its acquisition and retention efforts, and its advertising and pricing decisions.

In the first essay, our aim is to determine when menu-based pricing is viable for on-demand platforms and how membership and transaction fees should be optimally set. We consider platforms such as Uber One and DoorDash DashPass that offer both pay-per-use and premium membership options. Users self-select between these options based on their usage frequency and service valuation, while discount rates and service quality remain fixed. We identify a critical discount threshold that fundamentally alters platform revenue models: below this threshold, platforms rely on transaction-based revenues with minimal membership fees; beyond it, platforms shift to membership-driven revenues with lower transaction charges.

In the second essay, we analyze how optimal pricing strategies change when platforms can jointly choose membership fees, transaction prices, discount rates, and service quality. We examine this through progressively richer optimization scenarios, from a benchmark with fixed discounts and quality to full optimization across all four variables. Our investigation focuses on understanding which strategic levers matter most when platforms have fuller control, how market structure responds to expanded decision scope, and whether the three-tier segmentation commonly assumed in practice remains optimal.

In the third essay, we investigate freemium platform optimization to examine whether threshold-driven logic extends beyond on-demand contexts. We model the user base as a Markov chain with three states: non-users, free users, and premium users. The platform optimizes acquisition expenditure, retention investment, advertising intensity, and subscription pricing simultaneously. Transition probabilities are endogenous functions of platform decisions, capturing how choices in one domain affect outcomes in others. We focus on identifying critical thresholds in quality enhancement, understanding acquisition-retention trade-offs, and determining how retention effectiveness reshapes pricing power.

Keywords

Platform pricing; Subscription design; Service quality; On-demand platforms; Profit maximization; Customer heterogeneity

Research Methods

Quantitative Research, Numerical Analysis

Contents

Résumé	iii
Abstract	v
List of Tables	xi
List of Figures	xiii
Preface	xvii
Introduction	1
References	9
1 Optimal Menu Pricing Strategies for On-Demand Service Platforms	11
1.1 Introduction	11
1.2 Literature Review	13
1.2.1 Research Gaps and Contributions	16
1.3 The Model	18
1.3.1 Model Framework and Scope	18
1.3.2 Consumer Behavior and Market Segmentation	18
1.3.3 Platform's Profit Maximization Problem	25
1.4 Numerical Analysis	28
1.4.1 Base Parameter Set	28
1.4.2 Sensitivity Analysis	29

1.4.3	Interaction Effects Between Discount Rate and Quality Enhancement	35
1.4.4	Numerical Validation and Solution Reliability	46
1.5	Customer Segmentation and Managerial Insights	48
1.6	Conclusion and Managerial Insights	52
1.6.1	Conclusion	52
1.6.2	Future Research Directions	53
1.6.3	Strategic Implications	54
Appendix A.		55
A.1	Impacts of Maximum Usage Frequency $x_{\max}$	55
A.2	Impacts of Maximum Customer Valuation $v_{\max}$	56
Appendix B.		58
B.1	Proof of Propositions 1 and 3.	58
B.2	Proof of Proposition 2.	61
References		62

2 Joint Optimization of Pricing and Service Quality in On-Demand Service

Platforms	67
2.1	Introduction 67
2.2	Model 70
2.2.1	Framework 70
2.2.2	Customer Choice 70
2.2.3	Platform Profit Maximization Problem 72
2.3	Numerical Results 76
2.3.1	Benchmark and Optimization Scenarios 76
2.3.2	Scenario II: Effect of endogenous d 78
2.3.3	Scenario III: Effect of Endogenous σ 82
2.3.4	Comparative Statics: Robustness of the Joint Optimization Strategy 89
2.4	Conclusion 98

References	99
3 Freemium Pricing and CRM Expenditures by a Digital Platform	101
3.1 Introduction	101
3.2 Literature Review	104
3.2.1 Customer Acquisition and Retention	104
3.2.2 Freemium Revenue Model	105
3.2.3 Our Contribution	107
3.3 The Model	108
3.3.1 Modeling Framework Overview	109
3.3.2 Customers' Behavior	110
3.3.3 Platform Optimization	115
3.4 Numerical Results	117
3.4.1 Benchmark Case	117
3.4.2 Sensitivity Analysis	120
3.4.3 Strategic Insights	135
3.4.4 Numerical Validation and Robustness Analysis	139
3.5 Conclusion	141
Appendix A.	142
A.1 Full Profit Model Derivation	142
References	142
General Conclusion	147
References	153

List of Tables

1.1	Notation for the General Model	20
1.2	Parameter specifications for sensitivity analysis	28
1.3	Parameter sensitivity analysis	30
2.1	Platform Optimization Scenarios	76
2.2	Comparative Statics: Scenario II (Endogenous Discount)	78
2.3	Comparative Statics: Scenario III (Endogenous Quality)	83
2.4	Comparative Statics: Scenario IV (Endogenous Quality and discount)	90
3.1	Notation for the General Model	114
3.2	Model Parameters and Ranges for Sensitivity Analysis	118
3.3	Parameter sensitivity analysis.	120

List of Figures

1.1	Customers' subscription preferences based on their usage frequency and valuation.	21
1.2	Customer segmentation and optimal pricing results for the base parameter set.	29
1.3	Effect of changing d on decision variables and platform profit	31
1.4	Effect d on customer fractions and usage frequencies	32
1.5	Effect of changing d on decision variables and platform profit	33
1.6	Effect σ on Customer fractions and Usage frequencies.	34
1.7	Effect of d on pricing variables and profit under low vs. high σ	37
1.8	Effect of discount rate d on customer fractions and usage frequencies under low vs. high σ	39
1.9	Effect of σ on pricing variables and profit under low vs. high d	43
1.10	Effect of σ on user fractions and usage frequencies under low vs. high d . . .	44
1.11	Comparison of results across different levels of σ and d	51
1.12	Comparison of results across different levels of $v_{\max}$ and $x_{\max}$	52
1.13	Impact of changing $x_{\max}$	57
1.14	Effect of changing $v_{\max}$	59
2.1	Scenario II: Sensitivity of platform outcomes to changes in $x_{\max}$	79
2.2	Scenario II: Sensitivity of platform outcomes to changes in σ	81
2.3	Scenario II: Sensitivity of platform outcomes to changes in k	82
2.4	Scenario III: Sensitivity of equilibrium outcomes to market usage scale $x_{\max}$. .	84
2.5	Scenario III: Sensitivity of platform outcomes to changes in $v_{\max}$	86

2.6	Scenario III: Sensitivity of platform outcomes to changes in d	87
2.7	Scenario III: Sensitivity of platform outcomes to changes in quality cost k . . .	89
2.8	Scenario IV: Sensitivity of platform outcomes to changes in $x_{\max}$	91
2.9	Scenario IV: Sensitivity of platform outcomes to changes in $v_{\max}$	94
2.10	Scenario IV: Sensitivity of platform outcomes to changes in k	97
3.1	Transition States	111
3.2	Impact of varying γ	121
3.3	Effect of γ on ρ_{ij}	122
3.4	Impact of varying k_H	124
3.5	Effect of k_H on ρ_{ij}	125
3.6	Impact of varying k_A	126
3.7	Effect of parameter k_A on ρ_{ij}	127
3.8	Impact of varying k_R	129
3.9	Effect of parameter k_R on ρ_{ij}	130
3.10	Impact of Varying ω_A	132
3.11	Effect of parameter ω_A on ρ_{ij}	133
3.12	Impact of varying ω_R	134
3.13	Effect of parameter ω_R on ρ_{ij}	135

Preface

This thesis is composed of three essays which are listed below:

- Asadi, M., Cetin, C. B., and Zaccour, G. Optimal Menu Pricing Strategies for On-Demand Service Platforms.
- Asadi, M. and Zaccour, G. Joint Optimization of Pricing and Service Quality in On-Demand Service Platforms.
- Asadi, M., Cetin, C. B., and Zaccour, G. Freemium Pricing and CRM Expenditures by a Digital Platform.

Introduction

Digital platforms have fundamentally transformed how consumers access services, creating complex new challenges in pricing strategy design. Whether in ride-hailing and food delivery or music streaming and video entertainment, these platforms now mediate billions of daily transactions worldwide. A defining characteristic of this landscape is the widespread adoption of tiered pricing structures which segment users based on their specific preferences and usage patterns. Whereas on-demand service platforms like Uber One and DoorDash DashPass offer membership programs combining upfront fees with discounted transaction charges (Chen et al., [2022](#)), content-centric freemium platforms such as Spotify and YouTube provide free ad-supported tiers alongside premium subscriptions (Wagner et al., [2014](#)). Despite these contextual differences, both models face a shared strategic imperative: designing pricing tiers that effectively segment heterogeneous customers while maximizing platform profitability.

Customer heterogeneity lies at the heart of platform pricing challenges. Users differ systematically in usage frequency, willingness to pay, and sensitivity to service quality and advertising. These differences create opportunities for price discrimination through menu-based offerings, yet they also introduce considerable complexity. A platform cannot optimize its membership fee in isolation, it interacts with the per-transaction price, discount rate, and quality level. These instruments jointly shape customer self-selection, and the boundaries between customer segments emerge endogenously from the platform's choices. Consequently, changes in one pricing instrument can shift market structure in ways that are difficult to anticipate.

The academic literature has addressed these challenges in a fragmented manner. Early research on digital platforms emphasized two-sided market structure and network effects (Armstrong, 2006; Rochet & Tirole, 2003). Subsequent work examined transaction-based and subscription-based pricing, showing that optimal structures depend on platform maturity and customer characteristics (Chen et al., 2022; Hu et al., 2023). However, most studies of menu pricing for on-demand platforms treat discount rates and quality levels as exogenous, restricting attention to price-setting alone. Research on freemium models has been similarly compartmentalized, with separate streams addressing conversion optimization and advertising strategy without examining their interaction (Alaei et al., 2023; DeValve & Pekeč, 2022). What remains missing is a unified framework that captures how these decisions jointly determine platform performance.

This thesis develops such a framework through three essays on platform pricing strategy. The central argument is that optimal platform pricing is governed by critical thresholds rather than smooth, incremental adjustments. At these thresholds, small changes in key parameters trigger discontinuous shifts in optimal strategy. This pattern emerges consistently across different platform types and strategic contexts, suggesting a fundamental principle of platform economics rather than an artifact of specific modeling choices. The research proceeds in three stages. The first essay examines menu pricing for on-demand platforms where the platform controls only its fee structure, while discount rates and quality levels remain fixed. The second essay expands the platform's strategic scope by endogenizing discount and quality decisions, revealing which levers matter most when platforms have fuller control. The third essay shifts to freemium platforms to test whether threshold-driven logic extends beyond on-demand contexts to subscription-based models with advertising. This progression, from constrained to expanded optimization, and from transaction-based to subscription-based platforms, builds a unified understanding of how customer heterogeneity, interdependent decisions, and critical thresholds shape platform strategy.

The first essay establishes the foundational framework by examining menu pricing under simple strategic conditions. The platform sets a membership fee and a per-transaction

price while discount rates and quality levels are given exogenously. This framework isolates the core tension in menu pricing: how to extract value from high-frequency users through membership while retaining lighter users who prefer pay-as-you-go pricing. Platforms such as Uber One, DoorDash DashPass, and Instacart+ navigate precisely this trade-off (Chen et al., 2022). Users self-select between pricing options based on their usage frequency and service valuation. Despite the widespread adoption of such programs, limited research addresses the joint optimization of membership and transaction fees when customers differ along both dimensions (Kung & Zhong, 2017).

The model considers a platform offering three options: non-participation, regular (pay-per-use) service, or premium membership with discounted transactions and enhanced quality. Customer valuations and usage frequencies are distributed across the population. The platform sets the membership fee and per-transaction price while discount rates and quality levels remain fixed parameters. Crucially, market segmentation is derived endogenously from utility-maximizing behavior rather than being assumed exogenously given. This approach reveals interactions between pricing decisions and customer composition that fixed-segment models would miss.

The central finding is a critical discount threshold that fundamentally alters how platforms should structure their revenue models. Below this threshold, profit maximization favors high per-transaction fees with negligible membership charges, the menu effectively collapses to a pay-per-use system, with transaction fees generating the vast majority of revenue. Once the discount crosses the threshold, optimal strategy reverses sharply: membership fees jump discontinuously, per-transaction fees drop substantially, and subscription revenue dominates. This discontinuity reflects a genuine economic phenomenon rather than a modeling artifact. Near the threshold, small changes in discount rates trigger fundamentally different optimal pricing structures.

A parallel threshold exists for service quality. When quality differentiation is weak, premium membership cannot be economically justified. At higher quality levels, however, a counterintuitive pattern emerges: platforms with superior service reach the critical discount threshold sooner than lower-quality platforms. High quality attracts engaged, high-

frequency users, but serving these users is costly. High-quality platforms must therefore shift to membership-based pricing earlier to avoid subsidizing heavy usage through discounts. Additional findings indicate that approximately half of potential customers remain non-users across a wide range of parameter values; that profit responds more strongly to usage frequency than to market size; and that pricing instruments exhibit clean separation, transaction fees depend primarily on customer valuations, while membership fees scale with usage capacity.

These results establish the baseline for the second essay. They demonstrate that even with limited strategic control, platforms face threshold-driven decisions that determine their fundamental business model.

Having established how platforms optimize pricing under fixed discount and quality parameters, the second essay asks a natural follow-up question: what changes when platforms can also choose their discount rates and quality levels? This extension reflects practical reality. Platforms actively determine not only their fee structures but also the magnitude of discounts offered to members and the level of quality improvement embedded in premium services. Treating these as exogenous overlooks decisions that managers routinely calibrate to influence market outcomes. By expanding from two decision variables to four, this essay captures the joint trade-offs between pricing, discounting, and service design, and highlights which strategic levers matter most.

To isolate distinct economic mechanisms, the analysis proceeds through progressively richer optimization scenarios. The benchmark replicates the first essay's setting, optimizing only membership and transaction fees. Subsequent scenarios endogenize discount optimization, then quality differentiation, and finally permit full optimization across all four variables. This stepwise relaxation preserves analytical tractability while quantifying the contribution of each lever to profitability and market structure.

The results reshape the interpretation of the first essay's findings. When platforms can freely choose discount rates, they optimally set discounts close to zero. The threshold identified in Essay 1 becomes less relevant because platforms simply avoid triggering it. More importantly, comparison across intermediate scenarios reveals a pronounced asym-

metry: endogenizing discounts yields only modest profit gains over the benchmark, while endogenizing quality yields substantial improvements. Quality differentiation dominates price discrimination. Price reductions cannibalize revenue from high-value customers without sufficiently expanding the market; quality investment, conversely, creates genuine value that supports higher prices.

A structural finding deepens the threshold logic from Essay 1. When discounts and quality become decision variables, the regular tier systematically vanishes, producing a purely binary market structure of premium subscribers and non-users. The three-segment market that might seem natural —non-users, regular users, premium members—collapses to two segments under full optimization. The platform achieves substantial profit improvements by combining near-zero discounts, strong quality differentiation, high transaction prices, and moderate membership fees. It serves roughly half the market as premium subscribers while deliberately excluding the rest. The elevated base price creates a cost differential that makes membership economically attractive even without transaction savings, as the quality enhancement alone justifies the membership fee for frequent users.

The analysis also uncovers winner-take-all dynamics driven by quality costs. When quality provision is nearly costless, platforms pursue extreme differentiation with correspondingly high membership fees. As quality costs rise, both optimal quality and profits decline rapidly. Together, the first two essays demonstrate that on-demand platforms face threshold-driven strategic decisions, with quality differentiation emerging as the dominant lever.

The third essay shifts focus to freemium platforms to examine whether threshold-driven logic extends beyond on-demand contexts to subscription-based models with advertising. This transition serves a methodological purpose: if critical thresholds appeared only in menu pricing settings, they might reflect specific features of that context rather than fundamental platform economics. Freemium platforms offer a compelling test case because they share the challenge of customer heterogeneity but differ substantially in structure. Revenue derives from both advertising to free users and subscriptions from premium users, and customer relationships evolve dynamically through acquisition, con-

version, and churn rather than through one-time self-selection.

This business model underpins platforms such as Spotify and YouTube, representing a substantial segment of the digital economy. Freemium platforms face distinctive optimization challenges. Free users generate advertising revenue but consume infrastructure resources; premium users provide higher-margin subscription revenue but require continuous value delivery to prevent churn. Unlike traditional businesses facing binary customer decisions, freemium platforms must simultaneously optimize multiple user segments with fundamentally different behaviors and revenue contributions (Wagner et al., 2014). The existing literature has examined these components separately: customer relationship management research focuses on acquisition-retention trade-offs, while two-sided market research analyzes advertising and pricing independently (Wei et al., 2023). No existing model jointly optimizes acquisition spending, retention investment, advertising intensity, and subscription pricing within a unified framework.

The model represents the platform’s user base as a Markov chain with three mutually exclusive states: non-users, free users, and premium users. The platform optimizes four decision variables: acquisition expenditure targeting non-users through marketing, retention expenditure targeting premium users through engagement, advertising intensity determining the ad burden on free users, and the subscription fee charged to premium users. A distinctive feature is that transition probabilities are endogenous functions of these choices, capturing how decisions in one domain affect outcomes in others. The platform maximizes long-run profit under steady-state conditions. The resulting problem is highly nonlinear that closed-form solutions are unavailable; numerical methods are employed to compute optimal strategies.

The analysis confirms that threshold-driven strategy extends to freemium platforms, though specific mechanisms differ from on-demand contexts. Premium quality enhancement exhibits a critical threshold at which churn from premium directly to non-user status effectively disappears. This threshold fundamentally reshapes platform economics and generates three phases. In a value-building phase, the platform emphasizes retention to reduce churn. At the critical threshold, retention reaches peak effectiveness. Beyond

this point, in a value-extraction phase, the platform can reduce retention spending and raise prices more aggressively, as ecosystem stickiness keeps users within the system even when they occasionally downgrade. This parallels the discount threshold in Essay 1 which both mark points where optimal strategy shifts discontinuously but the underlying mechanisms differ: on-demand platforms shift through price-segmentation mechanics, whereas freemium platforms shift through quality-churn dynamics.

Several additional results emerge. First, platforms allocate acquisition spending nearly four times higher than retention spending, even when cost and effectiveness parameters are symmetric. This contradicts conventional customer relationship management wisdom that retention should dominate acquisition. The structural bias arises from persistent premium churn: only a minority of premium users remain subscribed in the steady state, requiring continuous investment to refill the acquisition funnel. Second, the quality threshold generates non-monotonic relationships, challenging simple intuitions that widening the gap between free and premium tiers always increases premium adoption. Third, a retention paradox emerges: improved retention effectiveness enables higher pricing rather than expanding the premium user base, suggesting that retention investments function primarily as levers for margin improvement rather than user growth.

As in the earlier essays, the optimal market structure deliberately excludes a substantial portion of potential customers. The consistency of this finding across different platform types suggests that selective market coverage is a general principle of platform strategy rather than an artifact of specific models.

This thesis makes three primary contributions to the platform pricing literature.

First, it develops a unified perspective on how critical thresholds shape platform strategy. Across three essays examining different platform types and strategic scopes, optimal strategy is characterized by regime shifts requiring business model transformation rather than incremental adjustment. Discount rates trigger transitions from transaction-based to membership-based revenue in on-demand platforms. Quality parameters determine whether intermediate customer segments exist or vanish. Quality thresholds in freemium platforms enable shifts from retention-focused to pricing-focused strategies. The persis-

tence of threshold dynamics across diverse settings indicates that this is a fundamental feature of platform economics, arising from the interaction of customer heterogeneity with interdependent strategic instruments.

Second, the thesis reveals systematic patterns in how platforms should balance competing objectives. Quality-based differentiation consistently dominates price-based differentiation; in on-demand contexts, quality investment yields far greater profit improvement than discount strategies, while in freemium contexts, quality thresholds determine strategic regimes. Pricing instruments exhibit clean separation in on-demand platforms, with transaction fees depending mainly on valuations and membership fees scaling with usage capacity. Structural acquisition bias characterizes freemium platforms despite conventional emphasis on retention. All three essays find that optimal market structures deliberately exclude substantial portions of potential customers, challenging assumptions about universal coverage and demonstrating that serving low-value segments can diminish margins without proportional revenue gains.

Third, the thesis offers methodological contributions that provide foundations for future research: a menu pricing formulation with endogenous market segmentation, a step-wise optimization framework that isolates distinct economic mechanisms, and a Markov chain framework with endogenous transition probabilities for dynamic platform settings. The progressive structure of the thesis, moving from constrained to expanded optimization and from transaction-based to subscription-based platforms, offers a template for systematic investigation of complex strategic environments.

The remainder of this thesis is organized as follows. Chapter 2 presents the menu pricing model for on-demand service platforms with fixed discount and quality parameters, establishing the foundational framework and identifying the critical discount threshold. Chapter 3 develops the extended model with endogenous discounts and quality, demonstrating the dominance of quality-based differentiation and the emergence of binary market structures. Chapter 4 addresses freemium platform optimization through integrated analysis of acquisition, retention, advertising, and pricing decisions. Chapter 5 concludes with a synthesis of findings, theoretical implications, managerial guidelines, and direc-

tions for future research.

References

- Alaei, S., Makhdoumi, A., & Malekian, A. (2023). Optimal subscription planning for digital goods. *Operations Research*, 71(6), 2245–2266.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3), 668–691.
- Chen, D., Zhou, Y., Guan, X., & Lin, X. (2022). Transaction or membership? impact on on-demand delivery service platforms’ profits, consumer surplus, and labor welfare. *Journal of Systems Science and Systems Engineering*, 31(5), 563–593.
- DeValve, L., & Pekeč, S. (2022). Optimal price/advertising menus for two-sided media platforms. *Operations Research*, 70(3), 1629–1645.
- Hu, M., Qin, T., Wang, L., & Zhang, Z. J. (2023). *Subscription vs. spot pricing in on-demand economy* [Working paper].
- Kung, L.-C., & Zhong, G.-Y. (2017). The optimal pricing strategy for two-sided platform delivery in the sharing economy. *Transportation Research Part E: Logistics and Transportation Review*, 101, 1–12.
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029.
- Wagner, T. M., Benlian, A., & Hess, T. (2014). Converting freemium customers from free to premium—the role of the perceived premium fit in the case of music as a service. *Electronic Markets*, 24, 259–268.
- Wei, W., Liu, W., Tang, O., Dong, C., & Liang, Y. (2023). CSR investment for a two-sided platform: network externality and risk aversion. *European Journal of Operational Research*, 307(2), 694–712.

Chapter 1

Optimal Menu Pricing Strategies for On-Demand Service Platforms

1.1 Introduction

On-demand service platforms have increasingly adopted membership programs that combine discounted transaction fees and enhanced service quality with upfront payments (de Oliveira et al., [2023](#); Figliozzi & Unnikrishnan, [2021](#)). This pricing strategy creates a menu-based structure where users select between standard pay-per-use pricing and membership programs. For instance, Uber One offers a \$9.99 monthly plan that includes discounted ride and delivery fees along with perks such as priority driver matching (Uber Technologies Inc., [2024](#)). Similarly, DoorDash's DashPass and Instacart+ provide reduced service fees and enhanced features for premium users (DoorDash, [2018](#); Instacart, [2023](#)). These platforms segment users based on usage frequency and willingness to pay while building loyalty.

Platforms face complex optimization challenges in setting membership fees and per-transaction fees. The membership fee must be high enough to generate meaningful upfront revenue, while the per-transaction fee must balance revenue from regular customers with the attractiveness of membership plans for high-frequency users (Colby & Bell,

2016; Taylor, 2018). Despite widespread adoption, limited research addresses the optimal pricing structure, specifically how to jointly set membership and per-transaction fees and their impact on consumer behavior and platform profitability. This paper addresses this gap by developing an analytical framework for menu pricing optimization.

Motivated by these observations, our objective is to determine the optimal membership and transaction fees that maximize platform revenue. This requires understanding customer heterogeneity in usage frequency and valuation, and studying how these factors influence pricing decisions. Specifically, we address the following research questions:

1. How should platforms jointly optimize membership fees and per-transaction charges when customers differ in both usage frequency and service valuation? Under what conditions does one pricing strategy dominate the other?
2. What are the critical thresholds in discount rates and quality enhancements that drive customers' self-selection between pay-per-use and membership options? How do these thresholds interact with customer heterogeneity parameters?
3. How do different pricing regimes affect overall platform profitability, customer segmentation, and the viability of sustaining a three-tier market structure (premium, regular, and non-users)?

To answer these questions, we develop a novel mathematical framework for on-demand service platforms. Our results show that optimal pricing depends critically on discount intensity and quality differentiation, but in fundamentally different ways. For discount rates, we identify a critical threshold: below this level, profit maximization is achieved through high per-transaction fees and negligible membership charges, effectively reducing the platform to a pay-per-use system; above it, platforms must adopt membership-driven revenue models where upfront fees exceed per-transaction charges. This transition reflects the point at which discounts make membership sufficiently attractive to reshape pricing design and customer segmentation. For quality enhancements, the relationship is non-monotonic rather than threshold-based. Moderate quality improvements maximize

premium adoption, but excessive quality investment can screen out customers through high membership fees, reducing market coverage despite superior service. These findings reveal that menu pricing success depends on carefully calibrating both price discounts and service quality within specific viability zones.

In our model, on-demand service platforms offer menu pricing and customers can choose between pay-per-use at standard rates or premium membership with upfront fees and discounted transactions. While these platforms may operate as one-sided or two-sided service providers, we focus on the platform’s demand side to investigate pricing and consumer segmentation in detail.

The remainder of this paper is organized as follows. Section 1.2 reviews the relevant literature. Section 1.3 introduces the model. Section 1.4 presents sensitivity analyses and results, while Section 1.5 discusses customer segmentation. Section 1.6 provides concluding remarks and managerial insights. All proofs are collected in the Appendix.

1.2 Literature Review

The emergence of on-demand service platforms introduced pricing challenges absent from traditional markets. Early research emphasized their two-sided market structure. Rochet and Tirole, 2003 and Armstrong, 2006 developed foundational models showing how platforms balance supply and demand through pricing and network effects, while Caillaud and Jullien, 2003 analyzed the chicken-and-egg problem of platform launch. These insights were essential when firms like Uber and Lyft struggled to achieve critical mass on both sides.

As platforms matured and achieved supply-side liquidity, however, their *strategic focus shifted from growth to monetization*. Once network effects stabilized, competitive advantage increasingly depended on pricing sophistication. This shift is most visible in the *adoption of menu pricing* between 2018 and 2021, when DoorDash introduced DashPass, Instacart launched Instacart+, and Uber rolled out Uber One. These programs allow customers to self-select between pay-per-use and membership, combining upfront

fees with transaction discounts and enhanced services, a fundamental departure from the single-structure models that dominated early platform strategies.

Academic research has not kept pace with this evolution. Early studies examined pricing mechanisms in isolation. Transaction-based pricing maximizes flexibility but misses surplus from high-frequency users, while pure subscriptions generate predictable revenue but exclude occasional users. Two-part tariffs combine fixed and variable fees but typically impose both on all customers, limiting market coverage (Gangwar & Bhargava, 2023). Scholars also considered strategic customer behavior: Y. Chen and Hu (2020) showed static pricing can outperform dynamic pricing when customers anticipate surges, while W. Zhou et al. (2014) demonstrated how heterogeneity determines whether multiple classes of customers should be served. These contributions are valuable but rest on the assumption that platforms must commit to a single structure.

Later work compared alternative strategies. Cachon et al. (2017) analyzed static versus dynamic pricing for delivery services, while D. Chen et al. (2022) contrasted membership and transaction pricing for platforms with self-scheduling drivers, showing that emerging platforms benefit from memberships while mature ones prefer transactions. Hu et al. (2023) evaluated subscriptions, spot pricing, and two-part tariffs, finding that two-part tariffs and spot pricing dominate pure subscriptions, but again forcing platforms to adopt one strategy for all customers. W. Zhou et al. (2024) found that high-demand customers prefer membership while low-demand customers favor pay-per-use, directly suggesting menu benefits, yet their analysis still assumed single-model commitment. Other contributions point in the same direction: Taylor (2018) showed that valuation heterogeneity shapes preferences for membership versus per-transaction fees; Choi et al. (2020) linked heterogeneity to optimal pricing design; and Jacob and Roet-Green (2021) designed menus for ride-hailing without modeling driver dynamics. Together, these studies highlight the potential of menu pricing while stopping short of analyzing it directly, often citing mathematical complexity as a barrier.

Parallel research offers complementary insights. L. Chen (2019) showed that optimal profits often involve combining multiple fee types, while Jena and Ghadge (2023) and Nie

et al. (2022) demonstrated how commission design shapes profitability and supply chain outcomes. Quality differentiation has also been shown to enhance segmentation: Ma et al. (2025) found that heterogeneity in quality valuations justifies differentiated offerings, while Zhang and Tian (2024) showed that dynamic quality, price adjustment outperforms static approaches. Adjacent literatures confirm the value of menus in other digital settings: media platforms use ad-supported versus premium tiers (DeValve & Pekeč, 2022), digital goods rely on bundles (Alaei et al., 2023), and product, service systems employ contract menus (Jacob, 2023). Yet these contexts differ fundamentally from on-demand services, which are characterized by high-frequency, real-time usage and a distinctive trade-off between membership fees and per-transaction charges.

Empirical studies using proprietary platform data provide strong evidence on customer heterogeneity and pricing dynamics. Cohen et al. (2016) analyzed over 50 million Uber rides and documented demand elasticities between 0.4 and 0.6, while Hall et al. (2023) exploited Uber’s fare experiments to show how equilibrium constraints shape platform pricing. Despite these insights, analytical models have yet to explain the optimal thresholds for menu pricing transitions or to derive conditions under which platforms should shift between transaction-based and subscription-based regimes. Our framework addresses this gap by offering the first analytical derivation of critical discount and quality thresholds that determine optimal pricing structures.

In summary, practice has moved decisively toward menu pricing, but academic research remains fragmented. Existing studies either analyze single mechanisms, compare strategies without menus, or acknowledge menu pricing while avoiding its complexity. Parallel literatures demonstrate its potential, yet none fully capture the unique economics of on-demand platforms. This gap motivates our work: to develop a rigorous framework for optimal menu pricing when heterogeneous customers self-select between pay-per-use and membership options.

1.2.1 Research Gaps and Contributions

While the existing literature provides valuable insights into platform pricing strategies, it exhibits three critical limitations. First, most studies analyze platforms that employ a single pricing model, either pure transaction-based pricing (e.g., traditional ride-hailing), pure membership pricing (e.g., subscription services), or mandatory two-part tariffs where all customers must pay both fees. This approach fails to capture the reality of modern platforms like Uber One, DashPass, and Instacart+, which offer customers a choice between pricing structures. Second, when scholars acknowledge the existence of menu pricing (D. Chen et al., 2022; Kung & Zhong, 2017), they typically avoid analyzing it due to its mathematical complexity, leaving a significant gap in our understanding of optimal strategies. Third, prior work often treats service quality as homogeneous across tiers, ignoring how platforms use differentiated features to justify premium pricing and drive customer self-selection.

Our research makes three substantive contributions to address these gaps:

1. **Modeling Customer Choice in Menu Pricing.** We develop the first comprehensive analytical framework for platforms offering simultaneous access to both pay-per-use and membership options. Unlike two-part tariff models where all customers face identical fee structures, our model captures the self-selection process where customers optimally choose between: (i) regular service with standard per-transaction fees, (ii) premium membership with upfront fees plus discounted transactions and quality enhancements, or (iii) non-participation. This framework reflects the actual pricing architecture of leading platforms and reveals how customer heterogeneity in both usage frequency and value perception drives segmentation.
2. **Endogenous Demand Determination.** Rather than assuming exogenous demand functions, we derive market segmentation endogenously from utility-maximizing customer behavior. Our model shows how the platform’s pricing decisions (membership fee F and transaction fee c) jointly determine the boundaries between customer segments. This approach reveals non-trivial interactions, for instance, how

increasing membership fees can paradoxically increase regular usage by shifting the customer mix. These endogenous demand effects are crucial for understanding platform profitability but have been overlooked in models with fixed customer segments.

3. **Quality Differentiation and Strategic Thresholds.** We explicitly model how service quality enhancements (parameter σ) and discount rates interact with pricing to create distinct strategic outcomes. For discount rates, we identify a critical threshold at approximately 56%: below this level, platforms rely on transaction-based revenue with binary market structures; above it, they must fundamentally shift to membership-driven models with three-tier segmentation. For quality enhancements, the relationship is non-monotonic rather than threshold-based. We find that moderate quality levels (around $\sigma \approx 1.3$) maximize premium user adoption, but profit increases monotonically across the entire quality range. Quality investment enables membership pricing by creating sufficient value differentiation, but excessive quality combined with high membership fees can reduce premium adoption despite improving profitability. These findings provide managers with specific guidance on when to implement membership programs (discount threshold) and how to calibrate quality investment (balancing adoption and profit). Methodologically, our approach demonstrates that the perceived complexity of menu pricing models can be overcome through careful formulation and numerical analysis. By combining analytical derivations of customer segmentation with computational optimization of platform profits, we provide both theoretical insights and practical guidance for platform design.

The managerial relevance of our findings extends beyond pricing decisions. Our results show that successful menu pricing requires coordinated decisions about service quality, discount levels, and fee structures. Platforms cannot simply add membership options to existing transaction-based models; they must carefully calibrate the entire system

to achieve sustainable three-tier segmentation (premium, regular, and non-users) while avoiding market structures that leave significant profit on the table.

1.3 The Model

1.3.1 Model Framework and Scope

We consider an on-demand service platform that offers customers a menu of pricing options: they can either use the service on a pay-per-transaction basis or purchase a premium membership that provides discounted transaction fees and enhanced service quality.

The operational issues, like marketplace balance and incentive management, in two-sided markets have been widely studied (Armstrong, 2006; Belleflamme & Peitz, 2021; DeValve & Pekeč, 2022; Rochet & Tirole, 2003; Wei et al., 2023). We are complementing those by focusing on the pricing side of the problem. By doing so, we will be able to dive into the details of interconnections between customer characteristics, pricing and platform features. Our approach aligns with established literature that examines platform pricing from one side while treating the other as given (e.g., Rochet and Tirole, 2003 examine merchant fees taking consumer participation as given).

1.3.2 Consumer Behavior and Market Segmentation

Denote by m the potential number of users of the platform. Customers make their decisions on utility and pricing, and opt for the premium service, the regular service, or for not using the service at all.

Let p and r represent the number of customers who choose the premium and regular service, respectively, and let n denote the number of non-users, with $m = p + r + n$. Customers are differentiated in their valuation for the service provided by the platform and their usage frequency in a given period of time. The valuation of customers v is distributed uniformly in $[0, v_{\max}]$ where $v_{\max}$ corresponds to the highest valuation from a customer, (Bitran et al., 2008; D. Chen et al., 2022; Choi et al., 2020). Similarly, following B. Zhou

et al., 2009, the usage frequency x , i.e., the number of times a customer uses the service in a period, is distributed uniformly in $[0, x_{\max}]$, where $x_{\max}$ is the largest number of transactions by a customer. For example, a customer of a food delivery platform would be willing to pay v for a food delivery and uses the service x times in the specified period.

Let c be the fee charged by the platform per usage. The utility obtained by a regular customer from using the services of the platform once is given by the difference between their valuation and the regular transaction fee:

$$u^R = v - c, \quad (1.1)$$

and by a non-user:

$$u^N = 0. \quad (1.2)$$

If a customer subscribes to the premium service by paying the fixed fee of F , then she can enjoy a discount d on the transaction fee throughout the period, where $d \in (0, 1)$. The resulting per-transaction fee of a premium user would then be $(1 - d)c$ (Uber Technologies Inc., 2024). Further, the platform provides a better service experience to premium subscribers by enriching the service with additional features, like Uber One allows customers to match with top-rated drivers. This improved quality carries the customers' valuation of the platform's service upward (Wagner et al., 2014). Let $\sigma > 1$ represent the perceived value of the platform's service with additional features under the premium plan compared to the value under the regular plan (Uber Technologies Inc., 2024). The utility for premium users is given by:

$$u^P = \sigma v - (1 - d)c - \frac{F}{x}. \quad (1.3)$$

Clearly, the impact of the fixed cost on the utility is reduced as customers use the service more often. A customer will choose the option (non-subscriber, regular subscriber, or premium subscriber) that maximizes their utility.

Table 1.1: Notation for the General Model

Notation	Description
<i>Model Parameters</i>	
m	Market potential of a service.
x	Number of transactions that a consumer realizes, $x \in (0, x_{\max}]$.
v	Value that a consumer derives per transaction by using the platform, $v \in [0, v_{\max}]$.
d	Discount rate offered to premium service members, $d \in (0, 1]$.
σ	Premium service quality factor, $\sigma > 1$.
k	Quality cost coefficient, $k > 0$.
β^P	Fraction of customers choosing the premium service, $\beta^P \in [0, 1]$.
β^R	Fraction of customers choosing the regular service, $\beta^R \in [0, 1]$.
β^N	Fraction of non-users, $\beta^N \in [0, 1]$.
δ^P	Number of transactions from customers in the premium service, $\delta^P > 0$.
δ^R	Number of transactions from customers in the regular service, $\delta^R > 0$.
<i>Decision Variables</i>	
F	Membership price, $F > 0$.
c	Per-transaction cost, $c > 0$.

Figure 1.1 illustrates a conceptual market segmentation model for understanding customer behavior on the platform, based on the number of transactions they perform and the perceived value of each transaction. Customers are distributed according to two key factors: the perceived value of the service (v) and the frequency of use (x).

Customers select the option that maximizes their utility. Based on the utility they derive, customers are divided into three distinct groups: premium subscribers (P), regular subscribers (R), and non-subscribers (N). If the utility from the premium service exceeds that of the regular service, they choose the premium option (P). If the utility from the regular service is higher, they opt for the regular option (R). If they derive no positive utility from either option provided by the platform, they do not subscribe and are classified as Non-subscribers (N).




Figure 1.1: Customers' subscription preferences based on their usage frequency and valuation.

To determine the region for regular users, we examine the conditions $u^R > u^N$ and $u^R > u^P$. All customers having a valuation v greater than per-transaction fee c are going to use the platform, $c < v < v_{\max}$. However, their preference for whether to subscribe to the premium plan is also dependent on their platform usage frequency x . As their usage frequencies increase, they become more motivated to subscribe to the premium plan. The reason behind this trend is that the burden of fixed cost on a single transaction decreases with an increase in usage frequency. The condition $u^R > u^P$ holds when $x((\sigma - 1)v + dc) < F$. This inequality suggests that if the membership fee exceeds the total added value and discount on the transaction fee with the premium service, customers will prefer the regular service, which is illustrated in the R region of Figure 1.1. If the fixed fee of the premium plan can be compensated for with improved service and a discount on the transaction fee, the customers are better off with the premium plan and subscribe to it.

Even though the valuation should be higher than the transaction fee for a customer to use the service without a premium plan, it is not compulsory for the premium plan. With inflated valuations thanks to added features of the premium plan, the customers who did not use the platform could become premium plan subscribers. Hence, in addition to the ones with $v > c$ considered above, there is a customer group that switches from non-subscribers to premium plan users. Still, their valuation should be sufficiently high to

cover the discounted per-transaction fee with the inflated valuation, $(1-d)c < \sigma v$. The two groups of customers, the customers who opt-in thanks to the premium plan and the customers who prefer the premium plan over the regular plan, jointly create premium plan subscribers.

Finally, Figure 1.1 serves as a strategic tool for platforms. By understanding the distribution of customers across these regions, platforms can tailor pricing strategies, adjust service offerings, and implement targeted marketing campaigns to shift more customers into the premium service category or expand the base of regular users. Denote by β^P and β^R the fraction of customers joining the platform with the premium plan and with the regular plan, respectively, and by β^N the fraction of consumers not using the platform.

To save on notation, let

$$\Gamma = cd + (\sigma - 1)v_{\max} > 0, \quad (1.4)$$

$$\Omega = (d + \sigma - 1) > 0. \quad (1.5)$$

Proposition 1. *The fractions of customers in each service category are as follows:*

$$\beta^P = 1 - \frac{1}{x_{\max}v_{\max}} \left(\frac{x_{\max}c(1-d) + F}{\sigma} + \frac{F}{\sigma} \ln \left(\frac{x_{\max}c\Omega}{F} \right) + \frac{F}{(\sigma-1)} \ln \left(\frac{\Gamma}{c\Omega} \right) \right), \quad (1.6)$$

$$\beta^R = \frac{F}{(\sigma-1)x_{\max}v_{\max}} \ln \left(\frac{\Gamma}{c\Omega} \right), \quad (1.7)$$

$$\beta^N = \frac{1}{\sigma x_{\max}v_{\max}} \left(x_{\max}(1-d)c + F + F \ln \left(\frac{x_{\max}c\Omega}{F} \right) \right). \quad (1.8)$$

Proof. See Appendix B.1. □

Proposition 1 gives an interesting tool for determining the fraction of customers across different plans, enabling businesses to estimate shifts in customer behavior by adjusting problem parameters; transaction cost, discount rate, and service quality. The proposition suggests that platform factors, such as σ , and customer characteristics $x_{\max}$ and $v_{\max}$, play a crucial role in determining the distribution of customers across various service categories. In particular, increases in σ , $v_{\max}$, and $x_{\max}$ enhance customer inclination towards using the platform. An increase in $x_{\max}$ leads to decrease β^R and β^N , that is,

drives customers towards the premium plan, while a larger $v_{\max}$ reduces the fraction of non-subscribers. It is evident that when $x_{\max}$ and $v_{\max}$ values are small, then customers are less interested in using the platform.

Feasibility Conditions: The customer fractions β^P , β^R , and β^N must all be non-negative for the model to be valid. We first observe that $\frac{cd+(\sigma-1)v_{\max}}{c(d+\sigma-1)} \geq 1$, since $v_{\max} > c$. This ensures that $\beta^R \geq 0$ (equation (1.7)), with equality occurring only when the membership fee $F = 0$. To guarantee that β^N is also non-negative, we impose the following sufficient condition on the parameters:

$$\frac{F}{x_{\max}c} \leq d + \sigma - 1. \quad (1.9)$$

This condition states that if the premium service provides no additional features and no discount, then the fixed fee F must be zero, which is intuitive. It is a sufficient (but not necessary) condition to guarantee $\beta^N \geq 0$. As for β^P , the condition for its nonnegativity is more complex, involving all parameter values. We will later present a sufficient condition under which the premium demand (i.e., the number of transactions) is nonnegative.

Proposition 2. *The impact of the platform's decision variables on the distribution of consumers among the three groups is as follows:*

1. *If the platform increases the per-transaction fee c , then the fraction of non-users increases, while the fractions of regular and premium users decrease.*
2. *If the platform increases the subscription fee F , then the fractions of non-users and regular users increases, and the fraction of premium users decreases.*

Proof. See Appendix B.2. □

The results in the above proposition are intuitive. Increasing the transaction fee discourages users, whether they are regular or premium, to use the platform. A higher subscription fee shifts premium users to regular users and continue using the platform users or to abandon it.

In the next step, we determine the total usage of the platform from each customer group of individuals. Let δ^P and δ^R represent the number of transactions of premium and regular users, respectively.

Proposition 3. *The total number of transactions by the premium and regular users are given by*

$$\delta^P = \frac{1}{2x_{\max}v_{\max}} \left(x_{\max}^2 v_{\max} - \frac{x_{\max}^2 c(d-1) + 2Fx_{\max}}{\sigma} + \frac{F^2}{\sigma c \Omega} - \frac{F^2(v_{\max} - c)}{c \Omega \Gamma} \right) - \frac{x_{\max} c(1-d)}{\sigma v_{\max}}, \quad (1.10)$$

$$\delta^R = \frac{F^2}{2x_{\max}v_{\max}} \cdot \frac{v_{\max} - c}{c \Omega \Gamma}. \quad (1.11)$$

Proof. See Appendix B.1. □

Clearly, from equation (1.11), we have $\delta^R \geq 0$, with equality holding only if $v_{\max} = c$, which is an uninteresting case. For δ^P in equation (1.10) to be positive, the subscription fee F must satisfy the following condition:

$$F < \frac{x_{\max}}{c(\sigma + c) - v_{\max}} \left(c \Omega \Gamma + (v_{\max} - c) \sqrt{\sigma c \Omega \Gamma} \right). \quad (1.12)$$

Proposition 3 gives the total usage of the two groups of users as functions of the model's parameters, i.e., $x_{\max}$ (the maximum number of transactions), $v_{\max}$ (the maximum value derived per-transaction), d (the discount to premium users), and σ (the perceived quality of the premium service), and the decision variables c and F . These formulas are useful to service providers as they enable them to adjust (or at least account for) key parameters to optimize customer engagement across different service tiers.

For regular users, we have

$$\frac{\partial \delta^R}{\partial F} = \frac{F}{x_{\max} c v_{\max}} \frac{v_{\max} - c}{\Omega \Gamma} > 0, \quad (1.13)$$

$$\frac{\partial \delta^R}{\partial c} = \frac{1}{2} \frac{F^2 (cd(c - 2v_{\max}) + v_{\max}^2(1 - \sigma))}{x_{\max} c^2 v_{\max} \Omega \Gamma^2} < 0, \quad (1.14)$$

meaning that a higher membership fee boosts regular users' demand, while a higher per-transaction fee reduces their demand. The underlying reason for the increase in total usage

with a higher membership fee stems from switches from premium membership to regular membership by the high usage rate customers. As can be seen from Figure 1.1, higher usage customers are more inclined to prefer premium membership. Even the customers with a high usage rate cannot justify paying the inflated premium membership fee, which in turn increases the total usage by regular members.

Denote by p and r the number of customers in the premium and regular service, respectively, that is, $p = m\beta^P$ and $r = m\beta^R$. Moreover, let D^P and D^R represent the total number of transactions by customers in the premium and regular service, respectively, that is, $D^P = m\delta^P$ and $D^R = m\delta^R$.

Remark 1. *The parameter m is irrelevant to the outcome, as it merely scales the results without affecting their structure. Therefore, without loss of generality, we set $m = 1$ to simplify the presentation. Consequently, we have*

$$p = \beta^P, \quad r = \beta^R, \quad D^P = \delta^P, \quad D^R = \delta^R. \quad (1.15)$$

1.3.3 Platform's Profit Maximization Problem

The platform aims to maximize its total profit by selecting the membership fee F and the per-transaction fee c , while treating the quality enhancement level σ and the discount rate d as given. Profit is generated through two primary channels: membership fees collected from premium subscribers, represented by $F\beta^P$, and transaction fees earned from both premium and regular users.

Specifically, the platform receives $(1 - d)c\delta^P$ from premium users, who benefit from a discounted rate, and $c\delta^R$ from regular users. However, providing enhanced service quality to premium members incurs operational costs. The cost term $\frac{1}{2}k(\sigma - 1)^2\delta^P$ captures the expenses associated with quality improvements, where $k > 0$ represents the unit cost parameter, and $(\sigma - 1)^2$ reflects a convex cost structure in which marginal quality improvements become increasingly expensive.

This quadratic formulation follows established quality cost modeling in operations management, where convex cost structures capture the economic reality that achiev-

ing progressively higher quality levels requires disproportionately greater investments (Taguchi 1986; Cachon et al. 2017). In platform contexts, this captures realistic operational constraints where achieving higher quality levels (e.g., shorter wait times, premium vehicle availability, priority matching) requires disproportionately greater investments in driver incentives, fleet expansion, and technological infrastructure. The normalization $\sigma = 1$ ensures zero quality cost at baseline service levels, and δ^P reflects that these costs scale with premium customer usage. Accordingly, the platform's objective function can be expressed as:

$$\max_{F,c} \pi = \beta^P F + (1-d)c\delta^P + c\delta^R - \frac{1}{2}k(\sigma-1)^2\delta^P \quad (1.16)$$

This maximization is subject to the constraints of nonnegativity of the segment fractions and positivity of overall demand. Written in full, the profit becomes

$$\begin{aligned} \pi = & \frac{1}{2x_{\max}c\sigma v_{\max}(\sigma-1)\Gamma\Omega} \left(-F^2c\sigma(c-v_{\max})(\sigma-1) \right. \\ & + 2Fc\Gamma\Omega \left[-F\sigma \ln \left(\frac{\Gamma}{c\Omega} \right) + x_{\max}\sigma v_{\max}(\sigma-1) \right. \\ & \quad \left. \left. - (\sigma-1) \left(F \left[\ln \left(\frac{x_{\max}c\Omega}{F} \right) + 1 \right] - x_{\max}c(d-1) \right) \right] \right. \\ & \quad \left. - (\sigma-1) \left[c(d-1) + \frac{1}{2}k(\sigma-1)^2 \right] \left(F^2(\Gamma + \sigma(c-v_{\max})) \right. \right. \\ & \quad \left. \left. + x_{\max}c\Gamma\Omega \left[x_{\max}(c(d-1) + \sigma v_{\max}) - 2F \right] \right) \right) \end{aligned} \quad (1.17)$$

Assuming an interior solution, the first-order optimality conditions are as follows:

$$\begin{aligned}
\frac{\partial \pi}{\partial F} = & F \left[\frac{c(d-1) + \frac{1}{2}k(\sigma-1)^2}{x_{\max}c\Gamma\Omega} \left(1 - \frac{1}{v_{\max}}\right) \right. \\
& - \frac{c(d-1) + \frac{1}{2}k(\sigma-1)^2}{x_{\max}c v_{\max} \sigma \Omega} \left(1 - \frac{v_{\max}}{c}\right) \\
& + \frac{2}{x_{\max} \sigma v_{\max}} \ln \left(\frac{F}{x_{\max} c \Omega} \right) + \frac{2}{x_{\max} v_{\max} (\sigma-1)} \ln \left(\frac{c \Omega}{\Gamma} \right) \Bigg] \\
& + \left[\frac{2c(d-1) + \frac{1}{2}k(\sigma-1)^2}{\sigma v_{\max}} + 1 \right] = 0.
\end{aligned} \tag{1.18}$$

$$\begin{aligned}
\frac{\partial \pi}{\partial c} = & F^2 \left[\frac{2\Gamma\Omega - dc\Omega - d\Gamma}{2x_{\max}c\Gamma^2\Omega} - \frac{1}{x_{\max}c^2\sigma v_{\max}} - \frac{k(\sigma-1)^2}{4x_{\max}c\Gamma^2\Omega} (2d\Gamma + \sigma(\Gamma - cd)) \right] \\
& + F \left[\frac{2(d-1)}{\sigma v_{\max}} \right] \\
& + \frac{x_{\max}}{2} \left[1 - (d-1) + \frac{k}{2\sigma v_{\max}} ((\sigma-1)^2 - \sigma(d-1)) \right] = 0
\end{aligned} \tag{1.19}$$

Due to the high degree of nonlinearity, the two optimality conditions (2.13) and (2.14) do not have an apparent closed-form solution. To solve these conditions numerically, we use the interior point algorithm, which is efficient for tackling both linear and nonlinear programming problems. This algorithm iterates through the interior of the feasible region, rather than along its boundaries, making it especially effective for large-scale problems. Its polynomial-time complexity and ability to handle complex constraints make it well-suited for our task. We provide detailed implementation specifications and convergence analysis in Section 1.4.4. Remember we set $d = 0.6$ since "Our analysis reveals a critical threshold: menu pricing with meaningful membership fees only emerges when discount rates exceed approximately 56%. Below this threshold, platforms optimally set $F \approx 0$, effectively reverting to pure transaction-based pricing. This explains why successful membership programs (Uber One, DashPass) offer substantial discounts rather than token benefits."

1.4 Numerical Analysis

In this section, we present the results for a series of parameter constellations and conduct a sensitivity analysis to examine how key parameter values affect the outcomes.

1.4.1 Base Parameter Set

To investigate changes in the platform’s strategies, we begin with the following base parameter set:

Table 1.2: Parameter specifications for sensitivity analysis

Parameter	Description	Benchmark	Range
d	Discount rate	0.6	[0, 1]
σ	Quality enhancement	1.2	(1, 2]
k	Quality cost coefficient	0.5	[0, 1]
$x_{\max}$	Maximum transactions	10	[1, 10]
$v_{\max}$	Maximum value	1.0	[0, 1]

In the benchmark scenario, we set the maximum number of transactions per customer at $x_{\max} = 10$, capturing usage patterns from occasional to heavy users typical of on-demand platforms. The maximum valuation is normalized to $v_{\max} = 1$ for analytical convenience, as the qualitative insights remain unchanged across different valuation levels (see Section A.2). The quality uplift factor is set to $\sigma = 1.2$, reflecting a 20% improvement in service quality for premium subscribers.

To account for the operational expense of this enhancement, we introduce a quality cost coefficient of $k = 0.5$. Finally, we adopt a discount rate of $d = 0.6$, corresponding to a 60% reduction in per-transaction fees for premium users. As shown in the sensitivity analysis, a sufficiently high discount is necessary to incentivize customers to adopt the premium plan and for platforms to implement menu-based pricing effectively. These parameter choices yield the customer segmentation and pricing outcomes shown in Figure 1.2.




Figure 1.2: Customer segmentation and optimal pricing results for the base parameter set.

The membership fee for premium services is $F = 2.26$, which exceeds the per-transaction fee of $c = 0.61$. The fraction of non-users ($\beta^N = 0.53$) remains larger than both premium and regular users. As illustrated in Figure 1.2, premium customers account for a greater share of active users than regular customers, with $\beta^P = 0.30$ and $\beta^R = 0.17$, respectively.

Moreover, demand from premium users ($\delta^P = 2.18$) significantly exceeds that of regular users ($\delta^R = 0.37$), highlighting their stronger contribution to platform performance. This finding is consistent with expectations: premium customers, who use the platform more frequently, are less affected by the fixed membership fee. Consequently, even when premium and regular customers represent comparable fractions, premium users generate higher demand and greater overall value for the platform.

1.4.2 Sensitivity Analysis

We conduct a comprehensive sensitivity analysis to examine how key parameters affect the platform’s optimal pricing strategy, customer segmentation, and profitability. Table 3.3 summarizes the directional effects of parameter changes, while our detailed analysis highlights several critical threshold behaviors and strategic inflection points.

As variations in d affect the per-transaction fee for premium customers, we report separate figures for the fee paid by premium users, $c(1 - d)$, and the fee paid by regular users, c . However, changes in σ do not alter the fee difference between premium and

Table 1.3: Parameter sensitivity analysis

Parameter	F	c	β^P	β^R	β^N	δ^P	δ^R	π
d	$\nearrow$	$\nearrow \searrow$	$\searrow$	$\nearrow$	$\nearrow \searrow$	$\nearrow \searrow$	$\nearrow$	$\searrow$
σ	$\searrow \nearrow$	$\nearrow \searrow$	$\nearrow \searrow$	$\searrow \nearrow$	$\nearrow$	$\nearrow \searrow$	$\searrow \nearrow$	$\nearrow$
k	$\nearrow$	$\searrow$	$\searrow$	$\nearrow$	—	$\searrow$	$\nearrow$	$\searrow$
$x_{\max}$	$\nearrow$	—	—	—	—	$\nearrow$	$\nearrow$	$\nearrow$
$v_{\max}$	$\nearrow$	$\nearrow$	$\nearrow$	$\searrow$	—	$\nearrow$	$\searrow$	$\nearrow$

Note: $\nearrow$ indicates increase; $\searrow$ indicates decrease; $\nearrow \searrow$ indicates non-monotonic effect; — indicates no change.

regular users. Therefore, we plot c against σ to illustrate this relationship.

Effect of changing discount level d : Figure 1.3 reveals two distinct pricing regimes separated by a critical threshold. When discounts are modest ($d < 0.56$), platforms rely primarily on transaction-based revenue. Membership fees remain low ($F < 0.8$) while per-transaction fees rise substantially, eventually reaching the upper bound ($c \approx 1.0$). The platform charges high per-transaction fees to regular users while offering meaningful discounts to premium users, creating a price differential that makes the regular option unattractive. As Figure 1.4a shows, customers face essentially a binary choice: subscribe to premium ($\beta^P \approx 0.45$) or avoid the platform ($\beta^R \approx 0$). During this phase, transaction fees generate 80 – 90% of platform revenue.

There is a critical threshold ($d \approx 0.56$). At this point, the platform’s pricing strategy undergoes a fundamental shift. The membership fee jumps sharply from $F \approx 0.8$ to $F \approx 2.1$, while the per-transaction fee drops from $c \approx 1.0$ to $c \approx 0.66$ (Figure 1.3). This transition occurs because discount costs escalate while the platform cannot raise c beyond 1.0 (the maximum customer valuation). At $d = 0.55$ with $c = 1.0$, premium users pay only $c(1 - d) = 0.45$ per transaction, meaning the platform forgoes 0.55 per transaction from its highest-frequency customers. With c capped and discount costs growing, transaction-based revenue becomes unsustainable. The platform responds by restructuring to capture value upfront through membership fees. It deliberately lowers c to 0.66 to make the




Figure 1.3: Effect of changing d on decision variables and platform profit

regular option affordable, maintaining market coverage while shifting to membership-driven revenue.

After crossing the threshold, ($d > 0.56$), revenue composition reverses. Membership fees now contribute 60 – 70% of total revenue, compared to 10 – 20% before. Higher membership fees screen out customers with lower usage frequencies who cannot justify the upfront cost. Premium adoption falls from $\beta^P \approx 0.44$ to 0.33 (Figure 1.4a). To compensate for reduced premium volume, the platform attracts regular users through moderate transaction fees. The regular segment grows from near zero to $\beta^R \approx 0.12$. Despite this adjustment, overall profitability declines from $\pi \approx 1.57$ to 1.42, a reduction of 9.5%. The membership-intensive model generates more predictable revenue but serves fewer customers and operates less efficiently.

These findings explain real-world platform strategies. Services like *Uber One* and *DashPass* typically offer estimated effective discounts of 30 – 40%, positioning them safely below the critical threshold of 56%. This allows these platforms to maintain low membership fees while maximizing user base growth through transaction revenue. Platforms that offer aggressive discounts above 60% risk crossing into the high-membership




Figure 1.4: Effect d on customer fractions and usage frequencies

regime, which constrains market reach and reduces profit margins. The optimal strategy suggests maintaining discounts in the 0.40 to 0.55 range: high enough to attract frequent users, but conservative enough to preserve the more profitable transaction-based revenue model.

Effect of Changing Quality Enhancement σ : As quality increases from $\sigma = 1.01$ to 1.99, membership fees rise substantially from $F = 2.58$ to 5.85 (Figure 1.5). Per-transaction fees increase from $c = 0.51$ to a peak of 0.64 at $\sigma = 1.34$, then decline slightly to 0.59. Most importantly, profit increases monotonically by 35%, rising from $\pi = 1.25$ to 1.69. Higher quality raises customers' willingness to pay for premium service, and the platform captures this surplus primarily through membership fees rather than transaction fees, maintaining moderate c values that encourage usage while extracting value upfront. Unlike discount strategies that eventually erode profitability, quality investment expands both customer value and platform profits simultaneously.




Figure 1.5: Effect of changing d on decision variables and platform profit

At low quality ($\sigma = 1.01$), regular customers dominate with $\beta^R = 0.41$, while premium adoption remains minimal at $\beta^P = 0.09$ (Figure 1.6a). As quality increases to $\sigma = 1.09$, the segments reach parity with $\beta^P \approx \beta^R \approx 0.24$. Premium adoption then peaks at $\beta^P = 0.31$ at $\sigma = 1.34$, before gradually declining to 0.24 at $\sigma = 1.99$. Regular users fall to a minimum of $\beta^R = 0.15$ at $\sigma = 1.34$, then recover to 0.21. Non-users remain relatively stable between $\beta^N = 0.50$ and 0.56.

Below $\sigma = 1.1$, quality differentiation is too weak to justify premium membership, so most active users choose the regular tier. The crossover at $\sigma = 1.09$ marks the minimum viable quality threshold at which premium membership becomes economically rational. Beyond $\sigma = 1.34$, sharply rising membership fees (with F increasing up to 5.85) screen out moderate-frequency users, causing premium adoption to decline despite improving quality. The recovery of regular users at high σ reflects strategic segmentation: the platform preserves an affordable regular tier to capture price-sensitive users who benefit from higher quality but cannot justify the premium membership cost.

At $\sigma = 1.01$, regular users exhibit higher usage frequency ($\delta^R = 1.67$) than premium users ($\delta^P = 0.83$), a counterintuitive pattern indicating a market failure (Figure 1.6b). This relationship reverses as quality improves: premium usage rises sharply, peaks at 2.24 at $\sigma = 1.34$, and then declines to 1.83 at $\sigma = 1.99$. Regular usage follows the opposite




Figure 1.6: Effect σ on Customer fractions and Usage frequencies.

trajectory, falling from 1.67 to a minimum of 0.29 at $\sigma = 1.34$, before recovering to 0.55.

At low quality, membership fees screen out the wrong customers: high-frequency users who cannot justify paying upfront for negligible benefits remain in the regular tier. As quality improves beyond $\sigma = 1.1$, membership fees begin to function as intended, selecting high-frequency users who can amortize the fixed cost over many transactions. The inverted-U pattern in premium usage reflects the tension between quality-driven demand (increasing with σ) and price-driven exclusion (the rising membership fee). The post-peak decline after $\sigma = 1.34$ demonstrates that extremely high membership fees ($F > 4.0$) eventually price out even some high-frequency customers.

These findings explain the evolution of food delivery platforms. DoorDash, Uber Eats, and competitors initially competed through heavy promotional discounts to build market share. Over time, they shifted strategies, investing in faster delivery, broader restaurant selection, and reliability improvements before launching subscription programs like DashPass and Uber One. This sequence validates our finding that membership models fail when quality differentiation is insufficient.

When DoorDash introduced DashPass in 2018, the company had already spent years building out restaurant partnerships and delivery infrastructure. The subscription offered tangible benefits, free delivery, reduced fees, exclusive access that customers could actually experience. Reports indicated subscribers ordered substantially more frequently

than non-members, matching our model’s prediction that quality improvements attract high-frequency users who can justify membership costs through repeated usage.

Today’s subscription programs maintain accessible price points around \$10 monthly, targeting mass-market adoption rather than exclusive segments. Platforms have discovered that raising fees too aggressively reduces subscriber bases, even when adding premium features. This reflects the screening effect our model identifies: higher membership fees filter out moderate-frequency users, forcing platforms to balance revenue per customer against total market coverage.

The strategic lesson is clear, quality investment enables profitable subscriptions, but only after crossing a threshold where customers perceive meaningful differentiation. Platforms that attempted membership programs prematurely, before establishing service reliability, saw poor adoption among their most valuable users.

1.4.3 Interaction Effects Between Discount Rate and Quality Enhancement

This subsection examines parameter interactions while providing essential robustness checks. By studying how discount rate d and quality enhancement σ jointly affect platform outcomes, we verify that our critical threshold findings remain stable across different quality and discount environments. We conduct a two-way analysis: first varying d under extreme quality levels ($\sigma = 1.1$ vs. $\sigma = 1.8$), then varying σ under extreme discount levels ($d = 0.2$ vs. $d = 0.8$). This approach tests whether our main results, particularly the regime shifts at critical thresholds, are artifacts of specific parameter choices or represent fundamental platform dynamics.

Effect of Changing Discount d Across Quality Levels σ : Figures 1.7 and 1.8 illustrate the results for low-quality ($\sigma = 1.1$) and high-quality ($\sigma = 1.8$) scenarios. As we can see, quality shapes the discount threshold effects while confirming the robustness of the critical threshold phenomenon across different quality regimes.

Both quality scenarios show critical thresholds, but at different discount levels: high-

quality platforms ($\sigma = 1.8$) experience transition at $d \approx 0.4$, while low-quality platforms ($\sigma = 1.1$) sustain pay-per-use pricing until $d \approx 0.55$. This ordering reflects the impact of quality provision costs, the term $\frac{1}{2}k(\sigma - 1)\delta^P$ imposes greater cost pressure on high-quality platforms, eroding transaction-based revenue faster and necessitating earlier transition to membership-based pricing. At their respective thresholds, both scenarios share the same changing pattern, membership fees rising sharply (F jumping discontinuously) and per-transaction fees c dropping abruptly (Figures 1.7a and 1.7b). This shows that the platform has a breaking point where it must switch from charging mainly per-transaction to charging mainly membership fees.

The pricing structures differ markedly between quality regimes. High-quality platforms maintain positive membership fees ($F > 0$) even at low discount levels, starting around $F \approx 0.2$ and increasing gradually before jumping sharply at the threshold from approximately $F \approx 1.8$ to $F \approx 3.4$. In contrast, low-quality platforms keep membership fees near zero ($F \approx 0$) throughout the low-discount regime, relying entirely on transaction fees until the threshold at $d \approx 0.55$ forces a sudden shift to $F \approx 0.8$. Both continue increasing post-threshold, reaching $F \approx 6.6$ and $F \approx 4.6$ respectively at $d = 1.0$. This demonstrates that quality investment enables earlier adoption of membership-based revenue models, high-quality platforms can charge membership fees from the outset because customers value the enhanced service.

The per-transaction fee c also shows distinct patterns. Low-quality platforms show an initial increasing trend, rising from approximately $c \approx 0.73$ to the upper bound $c = 1.0$ before the threshold. High-quality platforms maintain $c = 1.0$ throughout their pre-threshold region. After their respective thresholds, both drop sharply and converge to similar declining trajectories, reaching $c \approx 0.56$ at $d = 1.0$. The net transaction revenue $c(1 - d)$ follows correspondingly different paths: relatively flat at approximately 0.66 for low quality until the threshold, versus steadily declining from approximately 0.9 for high quality (Figure 1.7d).

Customer segmentation patterns reveal important differences between quality regimes, Figures 1.8a. Before their respective thresholds, low-quality platforms actually achieve




Figure 1.7: Effect of d on pricing variables and profit under low vs. high σ .

higher premium adoption ($\beta^P \approx 0.49$) than high-quality platforms ($\beta^P \approx 0.44$). This occurs because high-quality platforms charge positive membership fees from the outset, creating a higher barrier to premium tier entry. At their thresholds, both scenarios experience a drop in premium adoption and the sudden emergence of the regular tier: β^R jumps from zero to approximately 0.11 in both cases, while β^P drops to around 0.33 for low quality and 0.31 for high quality. The non-user fraction β^N peaks at the threshold (reaching ≈ 0.55 for low quality and ≈ 0.61 for high quality) before declining as the regular tier absorbs customers in the post-threshold region. Both quality levels converge to similar segmentation at $d = 1.0$: $\beta^P \approx 0.22$, $\beta^R \approx 0.25$, and $\beta^N \approx 0.53$.

Usage intensity patterns also differ across quality regimes, Figures 1.8b. Premium usage δ^P shows contrasting threshold behavior: low-quality platforms see a temporary spike in premium usage at their threshold (δ^P jumping from ≈ 2.5 to ≈ 2.67) before declining,

while high-quality platforms experience an immediate drop (δ^P falling from ≈ 2.48 to ≈ 2.18). This difference reflects how the threshold transition restructures the premium tier, low-quality platforms shed their marginal premium customers first, temporarily increasing average usage among remaining members, while high-quality platforms experience broader customer rebalancing. Regular usage δ^R emerges only after the threshold in both cases, starting at approximately 0.18–0.20 and increasing steadily to ≈ 0.72 –0.76 at $d = 1.0$.

Despite facing earlier threshold pressure due to quality costs, high-quality platforms maintain substantially higher profits throughout the entire discount range. Platform profit for $\sigma = 1.8$ starts at approximately $\pi \approx 1.86$ and remains above $\pi \approx 1.62$ even at $d = 1.0$. Low-quality platforms achieve lower profits ($\pi \approx 1.57$ at low d) and decline to $\pi \approx 1.42$ at $d = 1.0$ (Figure 1.7c). This confirms that quality investment, despite imposing operational costs that trigger earlier regime transitions, generates sufficient value to maintain profit advantages across all discount environments. Investing in quality protects the platform’s profitability even when pricing strategies get challenging.

The discontinuity observed at the critical threshold represents a fundamental restructuring of optimal market segmentation. As the discount rate d intensifies pressure on variable margins, the quality provision cost structure, which scales with premium usage δ^P , renders the high-usage premium strategy unsustainable at low membership fees. The discontinuity in F (Figure 1.7a) and the simultaneous emergence of the regular tier ($\beta^R > 0$) indicates a transition from a volume-driven strategy to a surplus-extraction strategy. The platform effectively forces high-usage customers to self-select into a higher fixed-fee bracket, internalizing the escalating costs of their consumption through membership payments rather than per-transaction charges.

These findings reveal a fundamental strategic trade-off between pricing flexibility and profit resilience. Low-quality platforms sustain pay-per-use pricing across an extended discount range ($d \in [0, 0.55]$), whereas high-quality platforms encounter threshold constraints earlier ($d \approx 0.4$). However, this extended flexibility for low-quality platforms coincides with uniformly lower profit margins and more abrupt transitions when con-




Figure 1.8: Effect of discount rate d on customer fractions and usage frequencies under low vs. high σ .

constraints eventually bind. High-quality platforms, despite facing earlier threshold constraints, achieve three structural advantages: earlier adoption of membership-based revenue models, superior profit resilience across all discount environments, and smoother strategic transition paths. The consistent presence of regime shifts across both quality levels, albeit at different threshold locations, confirms that critical discount thresholds represent fundamental platform dynamics rather than artifacts of specific parameter configurations.

The introduction of quality provision costs reveals a critical strategic pivot point for platform managers, particularly those investing in high-quality service. Unlike cost-free quality models where pricing strategies evolve gradually, our results demonstrate that operational constraints, modeled here via usage-dependent quality costs $\frac{1}{2}k(\sigma - 1)\delta^P$, necessitate a discontinuous regime shift in revenue generation. Managers must recognize a specific *discount tolerance threshold*. Below this threshold, platforms can sustain revenue through transaction fees (c) while keeping membership barriers (F) low. However, as promotional discounts (d) deepen, the escalating costs of servicing high-frequency premium users rapidly erode margins. Once this threshold is breached, managers should not incrementally adjust fees; rather, they must radically restructure pricing by (1) sharply increasing membership fees to decouple revenue from variable usage costs, and (2) introducing a lower-tier regular option to retain price-sensitive customers who are priced out of the premium tier.

A counterintuitive implication emerges for premium platforms. Platforms with superior service quality ($\sigma = 1.8$) reach this threshold sooner, at lower discount levels ($d \approx 0.4$), than lower-quality platforms ($d \approx 0.55$). Managers of premium platforms must therefore transition to heavy reliance on membership fees earlier than their mass-market competitors to avoid subsidizing high-frequency, high-cost usage. This “quality paradox” arises because quality provision costs scale with both enhancement level ($\sigma - 1$) and premium usage intensity δ^P ; platforms that attract engaged, high-usage customers through quality improvements simultaneously incur greater per-transaction costs, accelerating the point at which transaction-based revenue becomes insufficient. Our numerical results quantify this effect: high-quality platforms maintain profit margins approximately 14% higher than low-quality counterparts at $d = 1.0$ ($\pi = 1.62$ versus $\pi = 1.42$), but only if they execute the membership-pricing transition at the appropriate threshold. This finding aligns with observed strategic sequencing in food delivery markets, where platforms such as DoorDash invested substantially in delivery infrastructure and restaurant partnerships prior to launching subscription programs, thereby establishing the service quality foundation necessary to sustain membership pricing amid promotional competition.

Effect of Changing Quality σ Across Discount Levels d : We now examine quality's impact under contrasting discount scenarios: low discount ($d = 0.2$) versus high discount ($d = 0.8$). This analysis reveals how discount intensity fundamentally alters both the mechanism and magnitude of quality's effect on platform outcomes, Figures 1.9 .

The most striking finding is that the two discount regimes operate under fundamentally different market structures. Under low discount ($d = 0.2$), the regular tier does not exist: $\beta^R = 0$ and $\delta^R = 0$ across the entire quality range (Figures 1.10a and 1.10b). The platform segments customers into only two groups, premium members and non-users, extracting value entirely through the premium tier. Under high discount ($d = 0.8$), all three customer segments coexist, with the regular tier absorbing a substantial fraction of the market ($\beta^R \approx 0.20$ – 0.29). This structural difference confirms that the three-tier segmentation emerges specifically as a response to discount-induced margin pressure, not as a universal feature of optimal platform design.

The pricing response to quality enhancement differs markedly across discount regimes. Under low discount, membership fees remain near zero ($F \approx 0$) for low quality levels ($\sigma < 1.4$), then increase gradually to $F \approx 1.9$ at $\sigma = 2.0$ (Figure 1.9a). Quality investment effectively “unlocks” the ability to charge membership fees, below a quality threshold, customers do not perceive sufficient value differentiation to justify fixed fees. Under high discount, membership fees start high ($F \approx 3.3$ at $\sigma = 1.1$) and increase steeply with quality to $F \approx 7.0$ at $\sigma = 2.0$. The platform has already transitioned to membership-based revenue extraction; quality investment amplifies this strategy rather than enabling it. This amplification is driven by the convexity of the cost function: under high discounts, the marginal revenue from transactions is negligible ($c(1 - d) \approx 0.1$). Consequently, the platform treats quality investment purely as a mechanism to increase consumer surplus, which is then extracted upfront via the membership fee F , effectively using F to internalize the escalating quality provision costs.

Per-transaction fees exhibit opposite patterns across regimes (Figure 1.9b). Under low discount, c increases from approximately 0.78 to the upper bound $c = 1.0$ as quality rises, reaching the constraint at $\sigma \approx 1.35$. Under high discount, c remains constrained to sub-

stantially lower values ($c \approx 0.54\text{--}0.58$), exhibiting slight non-monotonicity with a peak around $\sigma \approx 1.5$. The net transaction revenue $c(1 - d)$ reflects these differences starkly: under low discount, it rises from 0.55 to plateau at 0.70; under high discount, it remains essentially flat near 0.11–0.12 regardless of quality level (Figure ??). This confirms that high-discount platforms cannot rely on transaction fees for meaningful revenue, membership fees must bear the primary burden.

Platform profit increases with quality under both discount regimes, but with different dynamics (Figure 1.9c). Under low discount, profit rises from $\pi \approx 1.37$ to peak around $\pi \approx 1.79$ at $\sigma \approx 1.7$, then plateaus or slightly decreases, exhibiting diminishing returns to quality investment. Under high discount, profit increases more steadily from $\pi \approx 1.30$ to $\pi \approx 1.67$ at $\sigma = 2.0$, showing sustained returns throughout the quality range. Notably, the profit gap between discount regimes narrows substantially at high quality levels: from a difference of approximately 0.07 at $\sigma = 1.1$ to approximately 0.10 at $\sigma = 2.0$. Quality investment partially compensates for discount-induced profit erosion.

Customer segmentation responds differently to quality across regimes, Figures 1.10a. Under low discount, premium adoption *decreases* with quality (β^P falling from 0.50 to 0.37), while non-user fraction increases (β^N rising from 0.50 to 0.62). This counterintuitive pattern occurs because quality investment enables higher membership fees, which price out marginal customers even as remaining members generate greater value. Under high discount, premium adoption is already low ($\beta^P \approx 0.20\text{--}0.26$) and exhibits slight non-monotonicity, peaking around $\sigma \approx 1.4$ before declining. The regular tier fraction β^R shows a U-shaped pattern, minimizing around $\sigma \approx 1.5$, intermediate quality levels squeeze the regular tier most intensely.


Usage intensity patterns reveal quality's differential impact on customer engagement, Figures 1.10b. Premium usage δ^P is substantially higher under low discount ($\approx 2.4\text{--}2.5$) than high discount ($\approx 1.6\text{--}2.0$), reflecting the smaller, more engaged premium customer base in the low-discount regime. Under high discount, δ^P exhibits an inverted-U shape, peaking around $\sigma \approx 1.4$ before declining, quality costs eventually constrain usage intensity. Regular usage δ^R (which exists only under high discount) shows a similar U-



Figure 1.9: Effect of σ on pricing variables and profit under low vs. high d .

shaped pattern, minimizing around $\sigma \approx 1.5$ at $\delta^R \approx 0.52$ before increasing to $\delta^R \approx 0.70$ at $\sigma = 2.0$.

These findings yield distinct prescriptions for platforms operating under different competitive intensities. For platforms in low-discount environments, quality investment serves primarily as a *revenue model enabler*, it unlocks the ability to charge membership fees and shifts revenue composition from transactions to subscriptions. However, managers must recognize diminishing returns: profit plateaus beyond $\sigma \approx 1.7$, and further quality investment may not justify additional costs. For platforms in high-discount environments, quality investment serves as a *profit recovery mechanism*, it compensates for margin erosion and sustains profitability under competitive pressure. The sustained returns to quality investment under high discount suggest that platforms facing intense price competition should prioritize quality enhancement more aggressively than platforms in protected markets. The absence of the regular tier under low discount further implies that



(a) Customer fractions

(b) Usage frequencies

Figure 1.10: Effect of σ on user fractions and usage frequencies under low vs. high d

intermediate service offerings should be introduced contingent upon competitive pressure; in low-discount environments, a two-tier structure (premium plus non-users) outperforms three-tier segmentation.

The introduction of quality provision costs reveals a critical strategic pivot point for platform managers, particularly those investing in high-quality service. Unlike linear cost models where pricing strategies evolve gradually, our results demonstrate that convex operational constraints, modeled here via usage-dependent quality costs $\frac{1}{2}k(\sigma - 1)(\delta^P)^2$, necessitate a discontinuous regime shift in revenue generation. Managers must recognize a

specific *discount tolerance threshold*. Below this threshold, platforms can sustain revenue through transaction fees (c) while keeping membership barriers (F) low. However, as promotional discounts (d) deepen, the escalating marginal costs of servicing high-frequency premium users rapidly erode margins. Once this threshold is breached, managers should not incrementally adjust fees; rather, they must radically restructure pricing by (1) sharply increasing membership fees to decouple revenue from variable usage costs, and (2) introducing a lower-tier regular option to retain price-sensitive customers who are priced out of the premium tier.

A counterintuitive *Quality Paradox* emerges for premium platforms. Platforms with superior service quality ($\sigma = 1.8$) reach this threshold sooner, at lower discount levels ($d \approx 0.4$), than lower-quality platforms ($d \approx 0.55$). Managers of premium platforms must therefore transition to heavy reliance on membership fees earlier than their mass-market competitors to avoid the “profit trap” of subsidizing high-frequency, high-cost usage. This paradox arises because quality provision costs scale non-linearly with usage intensity; high-quality platforms induce greater consumption frequency δ^P , which pushes the platform into the steep portion of the convex cost curve faster. Our numerical results quantify this effect: high-quality platforms maintain profit margins approximately 14% higher than low-quality counterparts at $d = 1.0$ ($\pi = 1.62$ versus $\pi = 1.42$), but only if they execute the membership-pricing transition at the appropriate threshold. This finding aligns with observed strategic sequencing in food delivery markets, where platforms such as DoorDash invested substantially in delivery infrastructure and restaurant partnerships prior to launching subscription programs, thereby establishing the service quality foundation necessary to sustain membership pricing amid promotional competition.

This amplification is driven by the convexity of the cost function: under high discounts, the marginal revenue from transactions is negligible ($c(1 - d) \approx 0.1$). Consequently, the platform treats quality investment purely as a mechanism to increase consumer surplus, which is then extracted upfront via the membership fee F , effectively using F to internalize the escalating quality provision costs.

1.4.4 Numerical Validation and Solution Reliability

Given the nonlinear complexity of the menu pricing optimization problem, we implement a rigorous validation framework to ensure the reliability and economic validity of our numerical results. This includes establishing solution existence, demonstrating convergence, verifying optimality conditions, and confirming economic soundness.

The optimization problem seeks to maximize profit $\pi(F, c)$ over the feasible domain $\mathcal{D} = \{(F, c) : 0 \leq F \leq \bar{F}, 0 \leq c \leq v_{\max}, F \leq x_{\max}(d + \sigma - 1)c\}$. Although logarithmic terms in π introduce singularities at the boundaries $F = 0$ and $c = 0$, solution existence is guaranteed on the restricted domain $\mathcal{D}_\varepsilon = \{(F, c) \in \mathcal{D} : F \geq \varepsilon, c \geq \varepsilon\}$ for $\varepsilon > 0$. Since $\mathcal{D}_\varepsilon$ is compact and π is continuous, the Extreme Value Theorem ensures that an optimal solution exists. In practice, we set $\varepsilon = 10^{-6}$ to avoid numerical instability while preserving economic interpretability.

To solve this problem, we apply the trust-region interior-point algorithm (`trust-constr`), which is particularly effective for nonlinear programs with multiple constraints. This method uses logarithmic barrier functions to maintain strict feasibility while leveraging trust-region strategies for global convergence. Our formulation includes nine constraints: three bounds on market share fractions ($0 \leq \beta^i \leq 1$), one equality constraint ($\sum_i \beta^i = 1$), two non-negativity constraints on usage ($\delta^P, \delta^R \geq 0$), and three linear bounds. The algorithm converged successfully, achieving a gradient norm of $\|\nabla \mathcal{L}\| \approx 10^{-16}$, indicating first-order stationarity up to machine precision. The identified optimal solution is $F^* = 2.26$ (membership fee), $c^* = 0.61$ (per-transaction fee), and $\pi^* = 1.40$ (maximum profit). This solution is interior and satisfies $F^* > c^*$, which aligns with economic intuition that membership should command a premium over per-use transactions.

To verify first-order optimality, we check the Karush-Kuhn-Tucker (KKT) conditions. The Lagrangian gradient norm is below 10^{-15} , satisfying the stationarity condition. Primal feasibility holds with maximum constraint violations below 10^{-16} , and key solution values include: $\beta^P = 0.30$, $\beta^R = 0.17$, $\beta^N = 0.53$, with $\beta^P + \beta^R + \beta^N = 1.00$ (exact to machine precision), and $\delta^P = 2.18 > 0$, $\delta^R = 0.37 > 0$. Dual feasibility is satisfied, as

all Lagrange multipliers associated with inequality constraints are non-negative. Complementary slackness is enforced by the interior-point barrier formulation.

To confirm second-order sufficiency, we compute the eigenvalues of the Hessian of $-\pi(F, c)$ (since we minimize $-\pi$). The positive eigenvalues $\lambda_1 \approx 0.06$ and $\lambda_2 \approx 2.72$ confirm local concavity of π and thus a strict local maximum. The condition number indicates moderate conditioning and numerical stability.

To address the possibility of local traps due to non-convexity, we perform multi-start optimization from randomized initial points across the feasible region. The algorithm demonstrated exceptional stability, converging to the same interior solution in the vast majority of trials. This overwhelming convergence to a single solution from diverse starting points, combined with satisfied first- and second-order conditions, provides strong empirical evidence for global optimality and affirms the reliability of our solution.

Our numerical analysis encompasses a comprehensive parameter space to ensure robustness across diverse platform scenarios. Specifically, we systematically evaluated solutions across: discount rates from 0.1 to 1.0 (capturing minimal to complete fee elimination), quality factors from 1.01 to 2.0 (representing negligible to transformative service enhancements), maximum usage frequencies from 1 to 10 transactions per period, and maximum valuations from 0.5 to 1.0. This yields over 1,000 distinct parameter combinations examined in our sensitivity analysis. For each parameter set, we verified that: (i) the optimization converged within 1,000 iterations, (ii) KKT conditions were satisfied to numerical precision, and (iii) the resulting market segmentation remained economically interpretable. Convergence failures occurred only at extreme parameter boundaries (e.g., $d \rightarrow 1$ with $\sigma \rightarrow 1$) where the economic model becomes degenerate. The consistency of our numerical results across this extensive parameter space, combined with the systematic patterns identified in our sensitivity analysis, confirms that our findings are robust rather than artifacts of specific parameter choices.

Finally, the solution exhibits economically meaningful behavior. First, price discrimination is evident: $F^*/c^* \approx 3.73$, indicating the platform extracts greater surplus from heavy users. Second, market coverage is realistic: 47% of users adopt the platform

(30% premium, 17% regular), while 53% are priced out, typical in differentiated pricing schemes. Third, usage intensity is significantly higher for premium users: $\delta^P/\delta^R \approx 5.89$, justifying the two-part tariff structure. Together, these results validate both the mathematical soundness and the economic interpretability of the optimal strategy.

1.5 Customer Segmentation and Managerial Insights

We examine how variations in key platform parameters shape customer segmentation patterns and profitability.

Quality-Discount Interaction Effects on Market Structure: Figure 1.11 presents customer segmentation across nine scenarios, combining three quality levels ($\sigma \in \{1.1, 1.6, 2.0\}$) with three discount levels ($d \in \{0.1, 0.6, 1.0\}$), with quality cost fixed at $k = 0.5$.

Three patterns are evident. First, the non-user segment remains stable at 50% to 59% across all scenarios. Customers with $v < c$ cannot be profitably served regardless of quality or discount adjustments. Second, the regular tier (β^R) is highly sensitive to discounts: negligible at $d = 0.1$ (forcing a binary structure), but expanding to 24% to 32% at $d = 1.0$. Third, premium share contracts as quality increases. At $d = 0.1$, higher quality ($\sigma : 1.1 \rightarrow 2.0$) reduces β^P from 49.9% to 40.6% because optimal transaction fees rise ($c : 0.61 \rightarrow 1.00$), excluding marginal users despite higher utility.

These structural patterns translate into three actionable strategic insights for platform managers. 1. *The Discount-Profit Paradox.* Low discount scenarios theoretically yield higher profit than high discount scenarios. At $\sigma = 2.0$, profit falls from $\pi = 1.86$ ($d = 0.1$) to $\pi = 1.65$ ($d = 1.0$), an 11% decline. However, platforms often prefer the lower-profit, high-discount strategy because of revenue composition. At $d = 1.0$, the firm charges a substantial membership fee ($F = 7.99$), ensuring revenue stability through predictable recurring income. In contrast, $d = 0.1$ relies almost entirely on volatile transaction fees ($F \approx 0.79$). Real platforms face this tradeoff. DoorDash accepts lower profit margins from DashPass because subscription revenue is stable and predictable, making it easier to plan investments and operations (Natasha, 2025).

2. *The Regular Tier as a Conversion Pipeline.* At high discounts ($d = 1.0$), the regular tier retains 24% to 32% of the market. These customers are not lost, they are high-valuation users with insufficient usage frequency (x) to justify premium membership. This segment represents a prime target for usage-based incentives (e.g., “Complete 3 more orders to unlock free delivery”). By increasing x , platforms can push these users across the indifference boundary, the curved line $x = \frac{F}{(\sigma-1)v+dc}$ in Figure 1.11, converting them to the premium tier.

3. *Quality Creates Value; Discounting Redistributes It.* Comparing the two levers reveals a strategic asymmetry. Increasing quality ($\sigma : 1.1 \rightarrow 2.0$) at low discount boosts profit by 36% ($\pi : 1.37 \rightarrow 1.86$). Conversely, increasing the discount ($d : 0.1 \rightarrow 1.0$) at high quality reduces profit by 11% ($\pi : 1.86 \rightarrow 1.65$). Sustainable growth comes from service improvements, faster logistics, better matching algorithms, exclusive features, rather than aggressive price discounting, which cannibalizes the firm’s premium margins.

Demand and Value Interactions: Figure 1.12 examines nine scenarios combining usage frequency ($x_{\max} \in \{3, 6, 10\}$) and customer valuation ($v_{\max} \in \{0.5, 0.7, 1.0\}$).

Three patterns are evident. First, market segmentation remains stable across all scenarios, with $\beta^P \approx 27\%$ to 30% , $\beta^R \approx 17\%$ to 19% , and $\beta^N \approx 53\%$. This stability holds regardless of variations in $x_{\max}$ or $v_{\max}$, suggesting that the optimal market partition is a structural property of the distribution rather than the parameters.

Second, the results show variable separation. Transaction fees depend only on valuations (c scales with $v_{\max}$), independent of usage. Conversely, membership fees scale linearly with usage ($F \propto x_{\max}$). Economically, the two pricing instruments serve distinct roles: transaction fees screen users based on valuation (v), while membership fees extract surplus based on volume (x).

Finally, profit scales proportionally with $x_{\max}$, maintaining a constant profit-to-usage ratio ($\pi/x_{\max} \approx 0.140$). The platform’s value capture scales perfectly with customer engagement.

These results yield two actionable insights for platform managers. 1. *Scale Through Membership Fees Only.* Platforms entering larger markets (or markets with higher us-

age frequency) should adjust membership fees proportionally ($F \propto x_{\max}$) while keeping transaction prices fixed. The optimal transaction fee depends solely on customer valuations, not market size. This explains why DoorDash maintains consistent delivery fees across markets but varies DashPass pricing by region: transaction fees anchor to local willingness-to-pay, while membership fees scale with expected order frequency.

2. *Prioritize Engagement Over Acquisition.* Profit scales linearly with usage frequency, but market participation remains capped at approximately 47%. Doubling customer engagement ($x_{\max}$) doubles profit; doubling market reach (forcing $\beta^N \rightarrow 0$) does not. Platforms should invest in retention mechanisms (order reminders, loyalty rewards) rather than aggressive customer acquisition. The optimal market structure naturally excludes approximately 53% of the population because serving them would require lowering transaction fees to a level that cannibalizes margins from high-value users. Chasing universal coverage is profit-destructive.

Demand and Value Interactions: Figure 1.12 examines nine scenarios combining usage frequency ($x_{\max} \in \{3, 6, 10\}$) and customer valuation ($v_{\max} \in \{0.5, 0.7, 1.0\}$).

Three patterns are evident. First, market segmentation remains remarkably stable across all scenarios, with $\beta^P \approx 27\%$ to 30% , $\beta^R \approx 17\%$ to 19% , and $\beta^N \approx 53\%$. This stability holds regardless of variations in $x_{\max}$ or $v_{\max}$, suggesting that the optimal market partition is a structural property of the distribution rather than the parameters.

Second, the results show variable separation. Transaction fees depend only on valuations (c scales with $v_{\max}$), independent of usage. Conversely, membership fees scale linearly with usage ($F \propto x_{\max}$). The two pricing instruments serve distinct roles: transaction fees screen users based on valuation (v), while membership fees extract surplus based on volume (x).

Finally, profit scales proportionally with $x_{\max}$, maintaining a constant profit-to-usage ratio ($\pi/x_{\max} \approx 0.140$). The platform's value capture scales perfectly with customer engagement.

These results yield two actionable insights for platform managers. 1. *Scale Through Membership Fees Only.* Platforms entering larger markets (or markets with higher us-




Figure 1.11: Comparison of results across different levels of σ and d .

age frequency) should adjust membership fees proportionally ($F \propto x_{\max}$) while keeping transaction prices fixed. The optimal transaction fee depends solely on customer valuations, not market size. This explains why DoorDash maintains consistent delivery fees across markets but varies DashPass pricing by region: transaction fees anchor to local willingness-to-pay, while membership fees scale with expected order frequency.

2. *Prioritize Engagement Over Acquisition.* Profit scales linearly with usage frequency, but market participation remains capped at approximately 47%. Doubling customer engagement ($x_{\max}$) doubles profit; doubling market reach (forcing $\beta^N \rightarrow 0$) does not. Platforms should invest in retention mechanisms (order reminders, loyalty rewards) rather than aggressive customer acquisition. The optimal market structure naturally ex-

cludes approximately 53% of the population because serving them would require lowering transaction fees to a level that cannibalizes margins from high-value users. Chasing universal coverage is profit-destructive.




Figure 1.12: Comparison of results across different levels of $v_{\max}$ and $x_{\max}$.

1.6 Conclusion and Managerial Insights

1.6.1 Conclusion

This paper develops a profit-maximization model for on-demand service platforms employing menu pricing, where customers choose between transaction-based and membership-based plans. We incorporate heterogeneity in usage frequency and service valuation, and

explicitly model quality provision costs to study how platforms should set per-transaction and membership fees under varying discount and quality levels.

Optimal pricing depends on discount intensity and quality differentiation in fundamentally different ways. For discounts, a critical threshold exists around 56%. Below this level, platforms rely primarily on per-transaction fees with negligible membership charges. Above it, the optimal strategy shifts to membership-driven models where upfront fees dominate. For quality, the relationship is non-monotonic. Moderate quality improvements (around $\sigma \approx 1.3$) maximize premium user adoption, while profit increases monotonically across the entire quality range. Quality investment enables membership pricing by creating sufficient value differentiation, but higher quality combined with elevated membership fees reduces premium adoption despite improving profitability.

Three-tier segmentation (premium, regular, and non-users) emerges only when discounts exceed approximately 50%. At lower discount levels, platforms operate with binary market structures that limit strategic flexibility. Premium users generate approximately six times as many transactions as regular users ($\delta^P/\delta^R \approx 5.9$), justifying the economics of menu pricing. However, non-participation remains around 53% across parameter combinations, indicating that universal adoption is structurally unattainable.

Successful implementation of three-tier menu pricing requires careful analysis and disciplined execution. Platforms must precisely estimate customer heterogeneity, coordinate discount and quality calibration, conduct staged pilots near identified boundaries, and continuously monitor for drift into unstable regions. In practice, platforms should align pricing with operational capabilities and commit to coherent pricing architectures rather than incremental modifications to existing fee schedules.

1.6.2 Future Research Directions

This study opens several avenues for future investigation. First, while we examine optimal pricing under fixed discount and quality parameters, platforms in practice must determine these levels endogenously. Future work could extend the model to jointly optimize dis-

count rates and quality enhancements alongside membership and transaction fees. This would reveal how operational cost structures and competitive pressures shape the optimal balance between price-based and quality-based differentiation strategies.

Second, we employ uniform distributions for both usage frequency and customer valuations, following standard practice in platform pricing literature. However, real-world distributions may exhibit different characteristics such as right-skewed usage patterns (where most customers use the service infrequently but a small segment generates heavy usage) or bimodal valuation distributions (reflecting distinct customer types). Examining how alternative distributional assumptions affect the critical thresholds and viability zones identified in this study would strengthen the generalizability of our findings.

Third, empirical validation using transaction-level data from operational platforms would test the model’s predictions. Key empirical questions include: Do platforms exhibit the predicted threshold behavior around 56% discount rates? Does market segmentation collapse to binary structures at low discount levels as the model predicts? What is the actual relationship between quality investment and premium adoption? Partnerships with platforms like DoorDash, Uber, or Instacart could provide the granular data necessary to calibrate the model and validate its core insights.

Finally, this analysis abstracts from competitive dynamics by focusing on a monopolistic platform. In practice, multiple platforms often compete simultaneously, potentially affecting both pricing strategies and customer segmentation. Future research could examine how competition influences the critical thresholds, whether platforms converge to similar quality-discount positions or differentiate strategically, and how competitive intensity affects the profitability of menu pricing relative to simpler pricing structures.

1.6.3 Strategic Implications

Several fundamental principles emerge for platform strategy:

Market Structure Determinants Sustainable three-tier segmentation requires discounts exceeding approximately 50% ($d > 0.5$). Below this threshold, platforms operate

with binary markets that limit strategic flexibility. Quality enhancements increase profitability but do not guarantee three-tier structures.

Usage Frequency Criticality High-frequency use cases ($x_{\max} > 6$) are prerequisites for successful monetization. Low-frequency markets require fundamentally different approaches, either pure transaction models or alternative value capture mechanisms.

Quality Investment Returns Quality improvements yield sustained profit increases, with a 35% profit gain when increasing from low ($\sigma = 1.1$) to high quality ($\sigma = 2.0$). However, moderate quality levels ($\sigma \approx 1.3$) maximize premium user adoption, while higher quality creates exclusionary pricing that reduces market coverage despite higher profits.

Quality-Discount Trade-offs The optimal combination depends on platform priorities. For profit maximization, high quality with moderate discounts ($\sigma \approx 2.0$, $d \approx 0.6$) achieves $\pi = 1.69$. For market coverage, high discounts with lower quality maximize adoption. Platforms must align their positioning with strategic objectives rather than pursuing a universal optimal configuration.

Appendix A

A.1 Impacts of Maximum Usage Frequency $x_{\max}$

We examine how the platform's optimal pricing strategy and customer behavior respond to variations in $x_{\max}$, which represents the maximum expected number of customer transactions per period. We consider $x_{\max}$ as a variable ranging from 1 to 10, with increments of 0.2.

Figure 1.13a reveals that the membership fee F increases linearly with $x_{\max}$, while the per-transaction fee c remains constant at approximately 0.60 across all values of $x_{\max}$. This pricing pattern indicates that the platform adjusts total revenue by scaling the membership fee proportionally with expected customer engagement, rather than modifying

transaction fees. Such a strategy avoids discouraging higher transaction volumes, which could occur if service fees increased with usage expectations.

The customer segmentation shown in Figure 1.13c demonstrates remarkable stability across different usage frequency scenarios. The fractions of premium members (β^P), regular members (β^R), and non-members (β^N) remain essentially constant as $x_{\max}$ varies. This stability suggests that the platform's membership value proposition maintains consistent appeal across different usage intensity levels, with customer self-selection patterns remaining unchanged.


Figure 1.13d illustrates the usage behavior of different customer segments. Premium members' usage frequency (δ^P) increases substantially with $x_{\max}$, rising from approximately 0.1 to 2.1 transactions per period. In contrast, regular members' usage (δ^R) grows much more gradually, increasing from approximately 0.02 to 0.33. This divergence reflects the different incentive structures: premium members, having paid an upfront membership fee, increase usage more aggressively as the maximum feasible usage rises, while regular members remain constrained by per-transaction costs.

The combined effect of these dynamics is shown in Figure 1.13b, where platform profit increases linearly with $x_{\max}$. This growth stems from the platform's ability to extract higher membership fees from customers with greater expected usage, while maintaining stable customer segmentation and encouraging intensive usage by premium members.


A.2 Impacts of Maximum Customer Valuation $v_{\max}$

We examine how the upper bound of customer valuations affects platform outcomes by varying $v_{\max}$ from 0.2 to 1.0, while holding other parameters at their benchmark values ($d = 0.6$, $\sigma = 1.2$, $x_{\max} = 10$, $k = 0.5$).


Figure 1.14 reveals a fundamental property of the optimal pricing structure. Market segmentation remains remarkably stable across all valuation levels, with customer fractions holding constant at $\beta^P \approx 0.29$, $\beta^R \approx 0.17$, and $\beta^N \approx 0.53$ (Figure 1.14c). This stability demonstrates that the three-tier structure is a robust feature of the model, inde-




(a) Decision variables



(b) Platform's objective value



(c) Customer fractions



(d) Usage frequency

Figure 1.13: Impact of changing $x_{\max}$.

pendent of the absolute scale of customer valuations.

Both pricing instruments scale proportionally with $v_{\max}$ (Figure 1.14a). The membership fee increases from $F = 0.51$ at $v_{\max} = 0.2$ to $F = 2.26$ at $v_{\max} = 1.0$, while the per-transaction fee rises from $c = 0.10$ to $c = 0.61$. This linear relationship confirms that the transaction fee c depends exclusively on valuations, not usage frequency. Platform profit scales correspondingly, rising from $\pi = 0.23$ to $\pi = 1.40$ (Figure 1.14b).

Usage intensities show diminishing sensitivity to valuations. Premium usage δ^P increases from approximately 1.5 to 2.2, while regular usage δ^R declines from 1.0 to 0.4 (Figure 1.14d). This pattern reflects the interaction between higher valuations (which encourage usage) and higher transaction fees (which constrain it), with premium users better able to absorb the fee increases.


The strategic implication is clear: platforms can adjust both membership and transaction fees proportionally when entering markets with different valuation levels, without fundamentally altering their pricing architecture. The optimal market structure, customer segmentation, and strategic roles of the two pricing instruments remain invariant to the absolute scale of customer valuations.

Appendix B

B.1 Proof of Propositions 1 and 3.


The segmentation of customers is based on their utility profiles defined by their usage frequency x and the value they obtain from the usage of the service v . The areas based on the reasons explained below define the fractions of customers in each subscription segment, β^P and β^R . The integration of usage frequencies of customers in the defined regions provides us with the total demand coming from each customer segment, δ^P and δ^R .

If a customer has a lower valuation for the platform's service than the per-transaction fee, $v < c$, she only uses the platform thanks to the improved service quality of premium



(a) Decision variables

(b) Platform profit



(c) Customer fractions

(d) Usage frequency

Figure 1.14: Effect of changing v_{max} .

service and discount on the per-transaction fee, $x > \frac{F}{\sigma v - (1-d)c}$. This group of customers is represented in the first part of the integrations to derive β^P and δ^P .

If a customer has a valuation higher than the cost of service, $v > c$, that customer would definitely use the platform but choose between regular and premium subscriptions. If their usage frequency is high enough to compensate for bringing down the fixed cost's impact on a transaction or their valuation is high enough to justify paying the fixed cost for the improved service, $x > \frac{F}{(\sigma-1)v+dc}$, they subscribe to the premium service. This group of customers creates the second part of the integrations to derive β^P and δ^P . The customers with sufficiently high valuation to use the platform, $v > c$, but without high enough usage frequency or valuation, $x > \frac{F}{(\sigma-1)v+dc}$, they use the platform as regular subscribers.

$$\begin{aligned}
\beta^P &= \int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \int_{\frac{F}{\sigma v - (1-d)c}}^{x_{\max}} f(x, v) dx dv + \int_c^{v_{\max}} \int_{\frac{F}{(\sigma-1)v+dc}}^{x_{\max}} f(x, v) dx dv \\
&= \int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \int_{\frac{F}{\sigma v - (1-d)c}}^{x_{\max}} \frac{1}{x_{\max} v_{\max}} dx dv + \int_c^{v_{\max}} \int_{\frac{F}{(\sigma-1)v+dc}}^{x_{\max}} \frac{1}{x_{\max} v_{\max}} dx dv \\
&= \frac{1}{x_{\max} v_{\max}} \int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \left(x_{\max} - \frac{F}{\sigma v - (1-d)c} \right) dv + \frac{1}{x_{\max} v_{\max}} \int_c^{v_{\max}} \left(x_{\max} - \frac{F}{(\sigma-1)v+dc} \right) dv \\
&= 1 - \frac{x_{\max}(1-d)c+F}{\sigma x_{\max} v_{\max}} - \frac{F}{\sigma x_{\max} v_{\max}} \ln \left(\frac{c(d-1+\sigma)x_{\max}}{F} \right) \\
&\quad - \frac{F}{(\sigma-1)x_{\max} v_{\max}} \ln \left(\frac{cd + (\sigma-1)v_{\max}}{c(\sigma-1+d)} \right), \tag{1.20}
\end{aligned}$$

$$\begin{aligned}
\beta^R &= \int_c^{v_{\max}} \int_0^{\frac{F}{(\sigma-1)v+dc}} f(x, v) dx dv \\
&= \int_c^{v_{\max}} \int_0^{\frac{F}{(\sigma-1)v+dc}} \frac{1}{x_{\max} v_{\max}} dx dv \\
&= \frac{1}{x_{\max} v_{\max}} \int_c^{v_{\max}} \frac{F}{(\sigma-1)v+dc} dv \\
&= \frac{F}{(\sigma-1)x_{\max} v_{\max}} \ln \left(\frac{(\sigma-1)v_{\max} + dc}{(\sigma-1+d)c} \right). \tag{1.21}
\end{aligned}$$

$$\delta^P = \int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \int_{\frac{F}{\sigma v - (1-d)c}}^{x_{\max}} x f(x, v) dx dv + \int_c^{v_{\max}} \int_{\frac{F}{(\sigma-1)v+dc}}^{x_{\max}} x f(x, v) dx dv$$

$$\begin{aligned}
&= \int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \int_{\frac{F}{\sigma v-(1-d)c}}^{x_{\max}} x \frac{1}{x_{\max} v_{\max}} dx dv + \int_c^{v_{\max}} \int_{\frac{F}{(\sigma-1)v+dc}}^{x_{\max}} x \frac{1}{x_{\max} v_{\max}} dx dv \\
&= \frac{1}{2x_{\max} v_{\max}} \left[\int_{\frac{x_{\max}(1-d)c+F}{\sigma x_{\max}}}^c \left(x_{\max}^2 - \left(\frac{F}{\sigma v - (1-d)c} \right)^2 \right) dv + \int_c^{v_{\max}} \left(x_{\max}^2 - \left(\frac{F}{(\sigma-1)v+dc} \right)^2 \right) dv \right] \\
&= \frac{1}{2x_{\max} v_{\max}} \left[\frac{x_{\max}^2 c(\sigma-1+d) - 2Fx_{\max}}{\sigma} + \frac{F^2}{\sigma c(\sigma-1+d)} \left(\frac{(d+\sigma)c - v_{\max}}{dc + v_{\max}(\sigma-1)} \right) \right. \\
&\quad \left. + x_{\max}^2(v_{\max} - c) - \frac{F^2(v_{\max} - c)}{c(d+\sigma-1)(dc + v_{\max}(\sigma-1))} \right] \\
&= \frac{1}{2x_{\max} v_{\max}} \left(x_{\max}^2 v_{\max} - \frac{x_{\max}^2 c(d-1) + 2Fx_{\max}}{\sigma} + \frac{F^2}{\sigma(\sigma-1+d)c} \right. \\
&\quad \left. - \frac{F^2(v_{\max} - c)}{(\sigma-1+d)c(dc + (\sigma-1)v_{\max})} \right) - \frac{x_{\max}c(1-d)}{\sigma v_{\max}}, \tag{1.22}
\end{aligned}$$

$$\begin{aligned}
\delta^R &= \int_c^{v_{\max}} \int_0^{\frac{F}{(\sigma-1)v+dc}} x f(x, v) dx dv \\
&= \frac{F^2}{2x_{\max} v_{\max}} \int_c^{v_{\max}} \left(\frac{1}{(\sigma-1)v+dc} \right)^2 dv \\
&= \frac{F^2}{2x_{\max} v_{\max}} \cdot \frac{v_{\max} - c}{c(d+\sigma-1)(dc + v_{\max}(\sigma-1))}. \tag{1.23}
\end{aligned}$$

B.2 Proof of Proposition 2.

To prove the results, it suffices to compute the derivatives.

1. We have the following derivatives with respect to c :

$$\frac{\partial \beta^N}{\partial c} = \frac{1}{x_{\max} c v_{\max} \sigma} (F + x_{\max} c(1-d)) \geq 0, \tag{1.24}$$

$$\frac{\partial \beta^R}{\partial c} = -\frac{F}{x_{\max} c(v_{\max}(\sigma-1) + cd)} \leq 0, \tag{1.25}$$

$$\begin{aligned}
\frac{\partial \beta^P}{\partial c} &= \frac{F(v_{\max} - cd) + x_{\max} c(1-d)(v_{\max}(1-\sigma) - cd)}{x_{\max} c v_{\max} \sigma(v_{\max}(\sigma-1) + cd)} \leq 0 \\
&\Leftrightarrow \frac{F}{x_{\max} c} \leq \frac{(1-d)(v_{\max}(\sigma-1) + cd)}{v_{\max} - cd}. \tag{1.26}
\end{aligned}$$

Compute the difference

$$\frac{(1-d)(v_{\max}(\sigma-1) + cd)}{(v_{\max} - cd)} - (d + \sigma - 1) = d\sigma \frac{c - v_{\max}}{v_{\max} - cd} \leq 0, \tag{1.27}$$

therefore, $\frac{F}{x_{\max}c} \leq (d + \sigma - 1)$ and consequently $\frac{\partial \beta^P}{\partial c} \leq 0$.

2. We have

$$\frac{\partial \beta^N}{\partial F} = \frac{1}{x_{\max} \sigma v_{\max}} \ln \frac{x_{\max} c (d + \sigma - 1)}{F}, \quad (1.28)$$

and it is positive by Assumption 1. Further,

$$\frac{\partial \beta^R}{\partial F} = \frac{1}{(\sigma - 1) x_{\max} v_{\max}} \ln \left(\frac{cd + (\sigma - 1) v_{\max}}{c(d + \sigma - 1)} \right), \quad (1.29)$$

which is clearly positive. Finally,

$$\frac{\partial \beta^P}{\partial F} = -\frac{1}{x_{\max} \sigma (\sigma - 1) v_{\max}} \left(\sigma \left(\ln \frac{x_{\max} (v_{\max} (\sigma - 1) + cd)}{F} \right) - \ln \frac{x_{\max} c (d + \sigma - 1)}{F} \right), \quad (1.30)$$

and it is negative if

$$\sigma \ln \frac{x_{\max} (v_{\max} (\sigma - 1) + cd)}{F} - \ln \frac{x_{\max} c (d + \sigma - 1)}{F} \geq 0, \quad (1.31)$$

which holds true because $x_{\max} (v_{\max} (\sigma - 1) + cd) > x_{\max} c (d + \sigma - 1)$, $\frac{x_{\max} c (d + \sigma - 1)}{F} > 1$, and $\sigma > 1$.

References

- Alaei, S., Makhdoumi, A., & Malekian, A. (2023). Optimal subscription planning for digital goods. *Operations Research*, 71(6), 2245–2266.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3), 668–691.
- Belleflamme, P., & Peitz, M. (2021). *The economics of platforms*. Cambridge University Press.
- Bitran, G. R., Rocha e Oliveira, P., & Schilkrut, A. (2008). Managing customer relationships through price and service quality.
- Cachon, G. P., Daniels, K. M., & Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3), 368–384.

- Caillaud, B., & Jullien, B. (2003). Chicken & egg: competition among intermediation service providers. *RAND Journal of Economics*, 34(2), 309–328.
- Chen, D., Zhou, Y., Guan, X., & Lin, X. (2022). Transaction or membership? impact on on-demand delivery service platforms' profits, consumer surplus, and labor welfare. *Journal of Systems Science and Systems Engineering*, 31(5), 563–593.
- Chen, L. (2019). *The proportional fee, transaction fee and membership fee* [Working paper].
- Chen, Y., & Hu, M. (2020). Pricing and matching with forward-looking buyers and sellers. *Manufacturing & Service Operations Management*, 22(4), 717–734.
- Choi, T.-M., Guo, S., Liu, N., & Shi, X. (2020). Optimal pricing in on-demand-service-platform-operations with hired agents and risk-sensitive customers in the blockchain era. *European Journal of Operational Research*, 284(3), 1031–1042.
- Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). Using big data to estimate consumer surplus: the case of uber. *NBER Working Paper*, (22627).
- Colby, C., & Bell, K. (2016). The on-demand economy is growing, and not just for the young and wealthy. *Harvard Business Review*.
- de Oliveira, L. K., de Oliveira Carneiro, C. M., Maia, M. L. A., & França, J. G. d. C. B. (2023). Effects of the COVID-19 pandemic on the perception of on-demand delivery by the younger generation: an exploratory analysis for Brazil. *Research in Transportation Economics*, 100, 101334.
- DeValve, L., & Pekeč, S. (2022). Optimal price/advertising menus for two-sided media platforms. *Operations Research*, 70(3), 1629–1645.
- DoorDash. (2018). *DoorDash launches DashPass subscription and free customer pickup*. Retrieved October 1, 2024, from <https://www.prnewswire.com/news-releases/door-dash-launches-dashpass-subscription-and-free-customer-pickup-300693742.html>
- Figliozzi, M., & Unnikrishnan, A. (2021). Home-deliveries before-during COVID-19 lockdown: accessibility, environmental justice, equity, and policy implications. *Transportation Research Part D: Transport and Environment*, 93, 102760.

- Gangwar, M., & Bhargava, H. K. (2023). Pricing on-demand services: alternative ways of combining usage and access fees. *Production and Operations Management*, 32(1), 11–27.
- Hall, J., Kendrick, C., & Nosko, C. (2023). Equilibrium effects of surge pricing: evidence from uber. *The Quarterly Journal of Economics*.
- Hu, M., Qin, T., Wang, L., & Zhang, Z. J. (2023). *Subscription vs. spot pricing in on-demand economy* [Working paper].
- Instacart. (2023). *Instacart+ membership launch*. Retrieved October 1, 2024, from <https://www.instacart.com/company/pressreleases/instacartplus/>
- Jacob, J. (2023). Designing price-service menus in a product-service system. *International Journal of Systems Science: Operations & Logistics*, 10(1), 2235812.
- Jacob, J., & Roet-Green, R. (2021). Ride solo or pool: designing price-service menus for a ride-sharing platform. *European Journal of Operational Research*, 295(3), 1008–1024.
- Jena, S. K., & Ghadge, A. (2023). Price competition in ride-sharing platforms: a duopoly supply chain perspective. *Computers & Industrial Engineering*, 183, 109507.
- Kung, L.-C., & Zhong, G.-Y. (2017). The optimal pricing strategy for two-sided platform delivery in the sharing economy. *Transportation Research Part E: Logistics and Transportation Review*, 101, 1–12.
- Ma, H., Lu, D., & Fan, Z.-P. (2025). Product selection, optimal product launch quantity and service pricing of a B2C service sharing platform. *Computers & Industrial Engineering*, 199, 111212.
- Natasha. (2025, November). Beyond burritos: doordash’s quest for profitability [Accessed December 2025]. <https://42signals.com/blog/doordash-business-model/>
- Nie, T., Guan, S., Du, S., & Zhu, S. (2022). Pricing strategy selection for the sharing platform with heterogeneous laborers. *Computers & Industrial Engineering*, 169, 108241.
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029.

- Taylor, T. A. (2018). On-demand service platforms. *Manufacturing & Service Operations Management*, 20(4), 704–720.
- Uber Technologies Inc. (2024). *What is Uber One?* Retrieved November 27, 2024, from <https://help.uber.com/en/ubereats/restaurants/article/what-is-uber-one?nodeId=c038ce7c-bd44-43ba-ad15-35b92c845492>
- Wagner, T. M., Benlian, A., & Hess, T. (2014). Converting freemium customers from free to premium—the role of the perceived premium fit in the case of music as a service. *Electronic Markets*, 24, 259–268.
- Wei, W., Liu, W., Tang, O., Dong, C., & Liang, Y. (2023). CSR investment for a two-sided platform: network externality and risk aversion. *European Journal of Operational Research*, 307(2), 694–712.
- Zhang, C., & Tian, Y.-X. (2024). Joint optimization of dynamic pricing, quality, and production quantity in the presence of online reviews: a distributionally robust optimization approach. *Computers & Operations Research*, 169, 106720.
- Zhou, B., Katehakis, M. N., & Zhao, Y. (2009). Managing stochastic inventory systems with free shipping option. *European Journal of Operational Research*, 196(1), 186–197.
- Zhou, W., Chao, X., & Gong, X. (2014). Optimal uniform pricing strategy of a service firm when facing two classes of customers. *Production and Operations Management*, 23(4), 676–688.
- Zhou, W., Liu, H., Cao, P., & Wang, Y. (2024). Optimal pricing for on-demand service platforms with heterogeneous customer demand frequencies. *Transportation Research Part E: Logistics and Transportation Review*, 189, 103676.

Chapter 2

Joint Optimization of Pricing and Service Quality in On-Demand Service Platforms

2.1 Introduction

In the previous chapter, we modeled the platform's pricing problem using two decision variables: the membership fee (F) and the per-transaction fee (c). The discount rate (d) offered to premium subscribers and the quality factor (σ) representing enhanced service were treated as fixed parameters. This assumption simplified the analysis and allowed us to characterize critical thresholds in pricing strategies. However, it restricted the platform's role to adjusting prices for a pre-determined service product.

In practice, platforms such as Uber, DoorDash, and Instacart have full control over the design of their membership programs. They determine not only the fees but also the magnitude of the discounts provided to members and the level of quality improvement embedded in premium services. Treating d and σ as exogenously given overlooks the fact that managers routinely calibrate them to influence market structure and revenue outcomes.

This chapter extends the previous model by relaxing the assumption of fixed discount and quality parameters. We broaden the decision space so that the platform simultaneously chooses four decision variables: (F, c, d, σ) . This generalization provides a richer characterization of optimal menu pricing and captures the joint trade-offs between price, discounting, and service quality. It allows us to investigate how optimal strategies shift when platforms are free to calibrate all levers of their pricing and membership programs.

Based on this extended framework, the chapter addresses the following research questions:

1. How should platforms jointly optimize membership fees, per-transaction charges, discount levels, and service quality to maximize profitability?
2. How do endogenous service design choices alter the critical thresholds identified in earlier analyses?
3. Under what conditions do discount strategies and quality investments function as substitutes rather than complements?

To address these questions, we develop a model in which an on-demand service platform jointly optimizes membership fees, transaction prices, subscriber discounts, and service quality. Customers are heterogeneous in usage frequency and valuation, and the platform designs a menu to screen among them. By endogenizing discount depth and quality level—parameters typically fixed in prior work—the framework captures trade-offs that existing models cannot address.

Since this study extends the framework developed in the previous chapter, we shall refrain from repeating the whole literature review and but restrict our attention to studies specifically addressing nonlinear pricing and endogenous quality design to position our contribution relative to the baseline model.

The economic foundation of our work lies in the literature on menu pricing and price discrimination. While early studies on platform pricing focused on the choice between subscription and transaction fees in isolation (D. Chen et al., 2022; Hu et al., 2023), recent

research suggests that optimal profitability often requires combining these mechanisms. For instance, L. Chen, 2019 shows that combining multiple fee types allows platforms to better extract surplus. However, most existing studies typically assume that the service characteristics defining the premium tier are predetermined. Our work addresses the limitation noted by Gangwar and Bhargava, 2023, who argue that rigid tariff structures limit market coverage, by endogenizing the depth of the discount.

Our work also connects with the literature on vertical differentiation and quality choice. Research by Ma et al., 2025 highlights how firms optimize quality levels to mitigate cannibalization between products. In the platform economy, quality often translates to reduced waiting times or priority access. Zhang and Tian, 2024 demonstrate that dynamic adjustment of quality alongside price outperforms static approaches. While these studies explore quality differentiation, they rarely integrate it with the subscription membership structure common in modern delivery apps.

The primary contribution of this chapter is to bridge these streams by endogenizing the platform's service design within a menu-pricing framework. While the baseline model presented in the previous chapter optimized fees for a fixed service level, this extension aligns with the view that pricing and product design must be jointly optimized to achieve efficiency. By treating the discount rate (d) and quality factor (σ) as decision variables, we offer a general model that encompasses the previous chapter as a special case.

The analysis yields three main findings. First, quality investment dominates discounting as a differentiation strategy. When both instruments are available, the platform relies primarily on service quality and membership fees to extract surplus, while optimal discount levels remain low. Second, the three-tier market segmentation commonly observed in practice, with subscribers, regular users, and non-participants, loses its appeal when the levels of the four variables are optimized. Indeed, the regular segment tends to vanish, leaving a binary market structure with subscribers and non-users. Third, beyond the role of individual instruments, optimal pricing exhibits a scaling property: as markets grow, only the membership fee requires adjustment, while transaction prices and quality levels remain stable. The analysis also identifies cost and valuation thresholds that determine

when quality-led strategies are feasible and when membership programs become viable.

The remainder of the paper is organized as follows. Section 2.2 presents the model. Section 2.3 reports numerical results and sensitivity analysis. Section 2.4 concludes.

2.2 Model

2.2.1 Framework

We consider an on-demand service platform that offers customers a menu of pricing options. Customers may either use the service on a pay-per-transaction basis or purchase a premium membership that provides discounted transaction fees and enhanced service quality.

The market consists of a continuum of customers who are heterogeneous in terms of usage frequency and service valuation. Usage frequency $x \in [0, x_{\max}]$ captures how often a customer uses the platform, while service valuation $v \in [0, v_{\max}]$ reflects willingness to pay arising from differences in preferences or income. Both attributes are assumed to be independently and uniformly distributed.

The platform offers a two-tier pricing structure. Under the regular tier, customers pay a per-transaction fee $c > 0$ and receive baseline service quality. Under the premium tier, customers pay an upfront membership fee $F > 0$, obtain a discounted per-transaction fee $(1 - d)c$, where $d \in [0, 1)$ denotes the discount rate, and receive an enhanced service quality level $\sigma \geq 1$. The platform simultaneously chooses the four decision variables (F, c, d, σ) to maximize profit.

2.2.2 Customer Choice

Customer choice follows the same structure as in Chapter 1. Since the demand environment and utility specification are unchanged, the derivation of the indifference conditions and participation regions is not repeated here. The only difference is that the discount rate and service quality are now chosen by the platform rather than treated as fixed parameters.

Each customer is characterized by a pair (x, v) and selects the option that maximizes utility. Utility from non-participation is normalized to zero. A customer who uses the service without a membership obtains utility

$$u^R = v - c, \quad (2.1)$$

while a customer who purchases the premium membership obtains

$$u^P = \sigma v - (1 - d)c - \frac{F}{x}. \quad (2.2)$$

Premium membership is therefore more attractive to customers who value service quality, benefit from lower per-transaction prices, or use the service frequently, since the membership fee is spread over a larger number of transactions.

Utility maximization divides the market into three groups: non-users, regular users, and premium subscribers. As in Chapter 1, the boundaries between these groups are determined by indifference conditions in the (x, v) space. For notational convenience, define

$$\Gamma = cd + (\sigma - 1)v_{\max}, \quad (2.3)$$

$$\Omega = d + \sigma - 1. \quad (2.4)$$

Applying the same integration steps as in Chapter 1, the resulting market shares are

$$\beta^P = 1 - \frac{1}{x_{\max} v_{\max}} \left[\frac{x_{\max} c(1 - d) + F}{\sigma} + \frac{F}{\sigma} \ln \left(\frac{x_{\max} c \Omega}{F} \right) + \frac{F}{\sigma - 1} \ln \left(\frac{\Gamma}{c \Omega} \right) \right], \quad (2.5)$$

$$\beta^R = \frac{F}{(\sigma - 1)x_{\max} v_{\max}} \ln \left(\frac{\Gamma}{c \Omega} \right), \quad (2.6)$$

$$\beta^N = 1 - \beta^P - \beta^R. \quad (2.7)$$

Aggregate transaction volumes generated by premium and regular users are given by

$$\delta^P = \frac{1}{2x_{\max} v_{\max}} \left[x_{\max}^2 v_{\max} - \frac{x_{\max}^2 c(1 - d) + 2Fx_{\max}}{\sigma} + \frac{F^2}{\sigma c \Omega} - \frac{F^2(v_{\max} - c)}{c \Omega \Gamma} \right] - \frac{x_{\max} c(1 - d)}{\sigma v_{\max}}, \quad (2.8)$$

$$\delta^R = \frac{F^2}{2x_{\max} v_{\max}} \cdot \frac{v_{\max} - c}{c \Omega \Gamma}. \quad (2.9)$$

2.2.3 Platform Profit Maximization Problem

The platform maximizes total profit by simultaneously determining the membership fee F , the transaction fee c , the discount rate d , and the service quality level σ . This decision involves a trade-off between revenue generation and operational efficiency. While deeper discounts and higher service quality incentivize premium adoption and increase usage volume, they also raise the operational costs of quality provision, which scale with the quality differential and total premium usage.

The platform's profit function comprises three revenue streams, membership fees, premium transaction revenue, and regular transaction revenue, minus the cost of quality provision. Subject to non-negativity and market feasibility constraints, the maximization problem is:

$$\max_{F,c,d,\sigma} \Pi(F,c,d,\sigma) = \beta^P F + (1-d)c\delta^P + c\delta^R - \frac{1}{2}k(\sigma-1)^2\delta^P. \quad (2.10)$$

The explicit functional forms for market shares (β^i) and aggregate demands (δ^i) are defined in Section 2.2.2.

We expect the platform to evaluate the trade-off between discount-based and quality-based differentiation by comparing their per-transaction costs. A discount d forfeits revenue of $d \cdot c$ per premium transaction, while quality σ incurs a provision cost of $\frac{1}{2}k(\sigma-1)^2$ per premium transaction. Consequently, quality investment is the efficient choice when:

$$\frac{1}{2}k(\sigma-1)^2 < d \cdot c \implies k < \frac{2dc}{(\sigma-1)^2} \quad (2.11)$$

This inequality demonstrates that quality differentiation is viable only if provision costs k are low relative to the discount expense. Crucially, the quadratic structure implies that increasing the quality margin $(\sigma-1)$ rapidly constrains the feasible cost threshold; specifically, doubling differentiation reduces the acceptable k by a factor of four.

Optimality Conditions

By substituting the expressions for market shares β^i and aggregate usage levels δ^i derived in Section 2.2.2 into the objective function (2.10), we obtain the following profit function:

$$\begin{aligned} \Pi = & \frac{1}{2x_{\max}c\sigma v_{\max}(\sigma-1)\Gamma\Omega} \left[-F^2c\sigma(c-v_{\max})(\sigma-1) \right. \\ & + 2Fc\Gamma\Omega \left(-F\sigma \ln \left(\frac{\Gamma}{c\Omega} \right) + x_{\max}\sigma v_{\max}(\sigma-1) \right. \\ & \quad \left. \left. - (\sigma-1) \left[F \left(\ln \left(\frac{x_{\max}c\Omega}{F} \right) + 1 \right) - x_{\max}c(d-1) \right] \right) \right. \\ & \quad \left. - (\sigma-1) \left(c(d-1) + \frac{1}{2}k(\sigma-1)^2 \right) \left(F^2[\Gamma + \sigma(c-v_{\max})] \right. \right. \\ & \quad \left. \left. + x_{\max}c\Gamma\Omega \left[x_{\max}(c(d-1) + \sigma v_{\max}) - 2F \right] \right) \right], \end{aligned} \quad (2.12)$$

where $\Gamma = cd + (\sigma-1)v_{\max}$ and $\Omega = d + \sigma - 1$. Assuming an interior solution exists, the first-order optimality conditions with respect to each decision variable are derived below.

The partial derivatives with respect to the price instruments are:

$$\begin{aligned} \frac{\partial \Pi}{\partial F} = & F \left[\frac{c(d-1) + \frac{1}{2}k(\sigma-1)^2}{x_{\max}c\Gamma\Omega} \left(1 - \frac{1}{v_{\max}} \right) \right. \\ & - \frac{c(d-1) + \frac{1}{2}k(\sigma-1)^2}{x_{\max}c v_{\max} \sigma \Omega} \left(1 - \frac{v_{\max}}{c} \right) \\ & + \frac{2}{x_{\max}\sigma v_{\max}} \ln \left(\frac{F}{x_{\max}c\Omega} \right) + \frac{2}{x_{\max}v_{\max}(\sigma-1)} \ln \left(\frac{c\Omega}{\Gamma} \right) \Bigg] \\ & + \left[\frac{2c(d-1) + \frac{1}{2}k(\sigma-1)^2}{\sigma v_{\max}} + 1 \right] = 0. \end{aligned} \quad (2.13)$$

$$\begin{aligned}
\frac{\partial \Pi}{\partial c} = & F^2 \left[\frac{2\Gamma\Omega - dc\Omega - d\Gamma}{2x_{\max}c\Gamma^2\Omega} - \frac{1}{x_{\max}c^2\sigma v_{\max}} - \frac{k(\sigma - 1)^2}{4x_{\max}c\Gamma^2\Omega} (2d\Gamma + \sigma(\Gamma - cd)) \right] \\
& + F \left[\frac{2(d - 1)}{\sigma v_{\max}} \right] \\
& + \frac{x_{\max}}{2} \left[1 - (d - 1) + \frac{k}{2\sigma v_{\max}} ((\sigma - 1)^2 - \sigma(d - 1)) \right] = 0.
\end{aligned} \tag{2.14}$$

The first-order condition with respect to the discount rate captures the trade-off between margin reduction and volume expansion. Due to the interaction between pricing, demand, and quality provision, this condition is highly nonlinear:

$$\begin{aligned}
\frac{\partial \Pi}{\partial d} = & \frac{1}{2x_{\max}c\sigma v_{\max}(\sigma - 1)\Gamma^2\Omega^2} \left[\right. \\
& - F^2c\sigma\Omega(\sigma - 1)[c(c - v_{\max}) + \Gamma(c - v_{\max})] \\
& + 2Fc\Gamma\Omega \left\{ F\sigma[c\Omega - cd - v_{\max}(\sigma - 1)] + F(\sigma - 1)\Gamma - x_{\max}c(\sigma - 1)\Gamma\Omega \right\} \\
& + c(\sigma - 1) \left((d - 1) - \frac{1}{2}k(\sigma - 1)^2 \right) (-F^2\sigma(c - v_{\max})(c\Omega + \Gamma) - F^2\Gamma^2 + \underbrace{x_{\max}^2c^2\Gamma^2\Omega^2}_{(2.15)}) \\
& + c(\sigma - 1)\Gamma\Omega (F^2\sigma(c - v_{\max}) + F^2\Gamma + 2x_{\max}^2c^2(d - 1)\Gamma\Omega + x_{\max}^2c\sigma v_{\max}\Gamma\Omega) \\
& \left. - x_{\max}c^2(\sigma - 1)\Gamma^2\Omega^2(2F + x_{\max}c(d - 1)) \right] = 0.
\end{aligned}$$

The optimality condition for service quality reflects four forces: membership attraction, transaction volume changes, regular-tier cannibalization, and operational cost increases. The explicit derivative is:

$$\begin{aligned}
\frac{\partial \Pi}{\partial \sigma} = & \frac{1}{2x_{\max}c\sigma^2v_{\max}(\sigma-1)^2\Gamma^2\Omega^2} \left[\right. \\
& F^2c\sigma^2(\sigma-1)^2(v_{\max}-c)[v_{\max}\Omega+\Gamma] \\
& -2Fc\Gamma\Omega \left(-F\sigma^2(\sigma-1)[v_{\max}\Omega-\Gamma] + F\sigma^2\Gamma\Omega \ln\left(\frac{\Gamma}{c\Omega}\right) \right. \\
& \quad \left. -F\sigma(\sigma-1)^2\Gamma + (\sigma-1)^2\Gamma\Omega \left[F \ln\left(\frac{x_{\max}c\Omega}{F}\right) + F - x_{\max}c(d-1) \right] \right) \\
& -c(d-1)(\sigma-1)^2 \left(F^2\sigma^2(c-v_{\max})[v_{\max}\Omega+\Gamma] + F^2\Gamma[\sigma\Gamma+\Gamma\Omega] \right. \\
& \quad \left. + x_{\max}^2c\Gamma^2\Omega^2[2c(d-1)-c\sigma v_{\max}] - x_{\max}c\Gamma^2\Omega^2[2F+x_{\max}c(d-1)] \right) \quad (2.16) \\
& + k\sigma(\sigma-1)^3\Gamma\Omega \left(F^2\sigma(c-v_{\max}) + F^2\Gamma + 2x_{\max}^2c^2(d-1)\Gamma\Omega \right. \\
& \quad \left. + x_{\max}^2c\sigma v_{\max}\Gamma\Omega - x_{\max}c[2F+x_{\max}c(d-1)]\Gamma\Omega \right) \\
& - \frac{1}{2}k(\sigma-1)^4 \left(F^2\sigma^2(c-v_{\max})[v_{\max}\Omega+\Gamma] + F^2\Gamma[\sigma\Gamma+\Gamma\Omega] \right. \\
& \quad \left. + x_{\max}^2c\Gamma^2\Omega^2[2c(d-1)-c\sigma v_{\max}] - x_{\max}c\Gamma^2\Omega^2[2F+x_{\max}c(d-1)] \right) \left. \right] = 0.
\end{aligned}$$

Because of the high degree of nonlinearity evident in Equations (2.13)–(2.16), the system does not admit closed-form solutions. We therefore solve the platform's problem numerically using an interior-point algorithm. This approach efficiently handles the nonlinear objective and constraints while ensuring the feasibility of the solution. Implementation details and convergence diagnostics are provided in Section 1.4.4 in chapter 1.

2.3 Numerical Results

This section examines how optimal decision variables and platform outcomes respond to changes in discount rates and quality differentiation.

2.3.1 Benchmark and Optimization Scenarios

We examine how discount rates and quality differentiation affect optimal pricing strategies and platform outcomes. Table 2.1 presents four scenarios with progressively relaxed optimization constraints, using base parameters from Section 2.3 ($x_{\max}=10$, $v_{\max}=1.0$, $k=0.5$).

Table 2.1: Platform Optimization Scenarios

Scenario	Decision Variables	π	Optimal Values	Market Distribution	Usage Intensity	Fixed
Benchmark	F, c	1.40	$F^*=2.26$, $c^*=0.61$	$\beta^P=0.30$, $\beta^R=0.17$, $\beta^N=0.53$	$\delta^P=2.18$, $\delta^R=0.37$	$d=0.6$, $\sigma=1.2$
Scenario II	F, c, d	1.45	$F^*=0.33$, $c^*=1.0$, $d^*=0.5$	$\beta^P=0.47$, $\beta^R \approx 0$, $\beta^N=0.53$	$\delta^P=2.65$, $\delta^R \approx 0$	$\sigma=1.2$
Scenario III	F, c, σ	1.69	$F^*=6.16$, $c^*=0.59$, $\sigma^*=2.04$	$\beta^P=0.23$, $\beta^R=0.22$, $\beta^N=0.55$	$\delta^P=1.78$, $\delta^R=0.59$	$d=0.6$
Scenario IV	F, c, σ, d	1.90	$F^*=0.27$, $c^*=1.0$, $\sigma^*=1.96$, $d^* \approx 0$	$\beta^P=0.43$, $\beta^R \approx 0$, $\beta^N=0.57$	$\delta^P=2.32$, $\delta^R \approx 0$	—

In the benchmark specification, the platform optimizes membership fee F and per-transaction price c while holding discount rate ($d = 0.6$) and quality differential ($\sigma = 1.2$) fixed. This yields profit $\pi = 1.40$ and produces a three-tier market segmentation: 30% premium users, 17% regular users, and 53% non-participants. The optimal membership fee ($F^* = 2.26$) exceeds the transaction price ($c^* = 0.61$), establishing membership revenue as the primary income stream. This structure resembles standard ride-sharing models where casual users can utilize the service at standard rates without being forced into a subscription.

In Scenario II, introducing endogenous discount optimization produces a sharp structural shift. The platform selects a steep discount ($d^* = 0.5$) and raises the transaction price to its maximum ($c^* = 1.0$), collapsing the regular tier entirely ($\beta^R \approx 0$). Customers face a binary choice: subscribe or exit. Yet profit increases only 3.6% to $\pi = 1.45$. Steep discounts cannibalize revenue from customers who would pay full price without sufficiently expanding market coverage. Pure price discrimination, absent quality differentiation, generates minimal value.

Scenario III endogenizes quality differentiation while maintaining fixed discount ($d = 0.6$), demonstrating markedly superior profit potential. The optimal quality multiplier reaches $\sigma^* = 2.04$, constrained by quadratic quality cost $(1/2)k(\sigma - 1)^2\delta^P$ with $k = 0.5$. This quality investment enables the platform to extract $F^* = 6.16$ in membership fees, nearly triple the benchmark level, while maintaining lower transaction price ($c^* = 0.59$). The regular tier remains economically viable with $\beta^R = 0.22$ (22% market share), preserving three-tier segmentation. Profit rises 21% to $\pi = 1.69$, demonstrating that vertical differentiation through quality investment dominates horizontal differentiation through price discrimination when investment costs are moderate.

The fully endogenous specification (Scenario IV) optimizes all strategic variables simultaneously. The platform selects near-zero discount ($d^* \approx 0$), eliminating price-based discrimination entirely, while choosing moderate quality differentiation ($\sigma^* = 1.96$) and maximum transaction pricing ($c^* = 1.0$). This configuration yields the highest profitability ($\pi = 1.90$, a 36% improvement over benchmark) while maintaining purely binary market structure with $\beta^R \approx 0$. The optimal quality level falls slightly below Scenario III's $\sigma^* = 2.04$ because eliminating discounts reduces premium usage intensity ($\delta^P = 2.32$ versus $\delta^P = 2.65$ in Scenario II), lowering the marginal cost of quality provision. The high base price creates a cost differential that makes the modest membership fee ($F^* = 0.27$) attractive despite offering no transaction savings. Customers pay for quality access alone, representing pure vertical differentiation without horizontal price discrimination. To gain a deeper understanding of the mechanics driving these structural shifts, the following sections conduct a comprehensive sensitivity analysis across all scenarios, starting with the

Table 2.2: Comparative Statics: Scenario II (Endogenous Discount)

Parameter	Decision Variables			Market Distribution			Usage Intensity		π
	F	c	d	β^P	β^R	β^N	δ^P	δ^R	
$x_{\max}$	—	—	—	—	—	—	$\nearrow$	—	$\nearrow$
$v_{\max}$	—	$\nearrow$	—	—	—	—	—	—	$\nearrow$
σ	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$	$\searrow$	$\nearrow \searrow$	$\nearrow$	$\searrow$	$\nearrow$
k	—	$\nearrow$	—	$\searrow$	—	$\nearrow$	$\searrow$	—	$\searrow$

Notes: $\nearrow$ = increase; $\searrow$ = decrease; — = no change; $\nearrow \searrow$ = non-monotonic.

isolation of the discount rate in Scenario II.

This outcome is consistent with the cost asymmetry discussed in Section 2.2.3: under joint optimization, quality-based differentiation can be monetized through the menu, while discounting directly erodes premium transaction margins.

2.3.2 Scenario II: Effect of endogenous d

Table 2.2 summarizes the comparative statics for Scenario II. We examine how variations in usage capacity ($x_{\max}$), customer valuation ($v_{\max}$), fixed quality level (σ), and provision cost (k) impact the optimal decision variables and market structure.

Effect of Changing Usage Frequency ($x_{\max}$)

Figure 2.1 illustrates the effect of varying maximum usage frequency ($x_{\max}$) on the platform's equilibrium strategies. The analysis reveals that the maximum usage frequency acts as a pure scaling parameter. As $x_{\max}$ increases from 1.1 to 10, optimal pricing decisions remain unchanged (Figure 2.1a). The membership fee stays close to zero, the per-transaction price remains around 0.62, and the discount rate is approximately 0.03. Market shares are likewise stable (Figure 2.1c), with premium and non-user segments each accounting for about half of the market, while the regular-user segment remains negligible. In contrast, profit increases from 0.165 to 1.480 (Figure 2.1b), and aggregate premium usage rises from 0.28 to 2.5 (Figure 2.1d).

This pattern arises because $x_{\max}$ expands the support of the usage frequency distribu-

tion without altering pricing thresholds or subscription incentives. Higher values of $x_{\max}$ allow premium users to consume more frequently, increasing total transaction volume while leaving customer choices unchanged. As a result, profit scales proportionally with usage intensity.

Platforms operating in markets with high-frequency users achieve higher profitability through volume rather than pricing adjustments. Usage frequency determines the scale of profits, whereas pricing determines market segmentation.




Figure 2.1: Scenario II: Sensitivity of platform outcomes to changes in $x_{\max}$.

Effect of Changing Quality (σ)

The quality parameter σ triggers a sharp structural transition in platform strategy at a critical threshold of approximately $\sigma \approx 1.75$. Below this threshold, the platform maintains a binary market structure (Figure 2.2c) with $F \approx 0$ and steadily rising transaction fees (Figure 2.2a). The platform monetizes quality improvements by charging all users higher per-transaction prices, maintaining broad access to the enhanced service. Above the threshold, the platform abruptly switches to an exclusive gatekeeping model where membership fees spike to $F > 3$, transaction fees drop, and a three-tier market emerges with $\beta^R \approx 35\%$.

This transition occurs because quality provision becomes prohibitively expensive at high σ levels. When quality costs are moderate, the platform maximizes profit by providing enhanced service to everyone (Figure 2.2d). Once quality costs exceed a critical level, universal provision is no longer profitable. The platform must restrict access through high membership fees, effectively creating two service tiers: premium users who pay F for quality enhancement and regular users who pay only transaction fees for standard service. The sudden emergence of regular users reflects customers priced out of the premium tier who still value basic service above non-participation.

The profit pattern follows an inverted U-shape, peaking at $\sigma \approx 1.6$ before declining after the threshold (Figure 2.2b). This indicates a quality trap: maximum profitability occurs when quality is high enough to command premium transaction prices but not so high that gatekeeping becomes necessary. Exclusionary strategies reduce the active user base, lowering total profit despite higher per-customer revenue.

Platforms should avoid over-investing in quality beyond the level that can be profitably provided to the mass market. Exclusive, high-fee programs may be viable but are less profitable than scalable quality improvements accessible to all users.




Figure 2.2: Scenario II: Sensitivity of platform outcomes to changes in σ .

Effect of Changing Quality Cost Coefficient (k)

The quality cost coefficient, k , impacts total profit but does not alter the optimal strategy (Figure 2.3). As k increases from 0.1 to 1.0, all decision variables remain stable: the membership fee stays near zero, the transaction price remains around 0.7, and the discount rate is unchanged (Figure 2.3a). Market segmentation is similarly unresponsive, with premium users and non-users essentially splitting the market (50% each) while the regular tier remains empty (Figure 2.3c). Consequently, premium consumption remains constant at (Figure 2.3d). The only variable that shifts is profit, which drops from 1.615 to 1.515 as costs rise (Figure 2.3b).

These results are due to the fact that quality (σ) is exogenous in this scenario. Because the platform cannot adjust quality levels to offset rising costs, the cost term simply scales up, and the platform absorbs the expense rather than passing it on through strategic adjustments. For managers, this highlights the direct value of operational efficiency. Investments that lower, such as technology upgrades or logistical improvements, protect margins without requiring changes to pricing or market segmentation.




Figure 2.3: Scenario II: Sensitivity of platform outcomes to changes in k .

2.3.3 Scenario III: Effect of Endogenous σ

In Scenario III, the platform endogenizes quality differentiation (σ) while maintaining the discount rate fixed at $d = 0.6$. This configuration represents platforms that have flexibility

Table 2.3: Comparative Statics: Scenario III (Endogenous Quality)

Parameter	Decision Variables			Market Distribution			Usage Intensity		π
	F	c	σ	β^P	β^R	β^N	δ^P	δ^R	
$x_{\max}$	$\nearrow$	—	—	—	—	—	$\nearrow$	$\nearrow$	$\nearrow$
$v_{\max}$	$\nearrow$	$\nearrow$	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$
d	$\nearrow$	$\searrow$	—	$\searrow$	$\nearrow$	$\nearrow \searrow$	$\nearrow \searrow$	$\nearrow$	$\searrow$
k	$\searrow$	$\searrow$	$\searrow$	$\searrow$	$\nearrow$	$\searrow$	$\nearrow \searrow$	$\searrow \nearrow$	$\searrow$

Notes: $\nearrow$ = increase; $\searrow$ = decrease; — = no change; $\nearrow \searrow$ = non-monotonic.

to enhance service quality but retain a predetermined discount structure, whether due to strategic commitments, customer expectations, or internal policy constraints. With F , c , and σ as decision variables, this scenario examines how platforms optimize their pricing and membership strategies when quality becomes an active strategic lever.

Table 2.3 presents the comparative statics results for Scenario III, where the platform optimizes membership fee F , per-unit price c , and quality σ while discount rate d is exogenously imposed. The table captures the effect of varying the imposed discount rate (d), usage capacity ($x_{\max}$), maximum valuation ($v_{\max}$), and quality cost parameter (k) on optimal decision variables, market structure, and usage patterns.

Effect of Usage Capacity ($x_{\max}$)

Figure 2.4 illustrates the effect of varying usage capacity ($x_{\max}$) on the platform's equilibrium strategies. The analysis reveals complete structural stability. As $x_{\max}$ increases from 1 to 10, the optimal per-transaction price holds constant at $c \approx 0.59$ and the quality level remains fixed at $\sigma \approx 2.04$ (Figure 2.4a). Market segmentation is similarly invariant (Figure 2.4c), maintaining approximately 23% premium users, 22% regular users, and 55% non-users. The only strategic adjustment is a linear increase in the membership fee F , which drives profit growth from 0.17 to 1.68 (Figure 2.4b).

This pattern occurs because $x_{\max}$ expands usage potential without changing relative customer heterogeneity. Heavy users remain proportionally heavier than light users at any scale, so the platform does not need to alter its screening mechanisms (c and σ). Instead, the membership fee extracts the expanded surplus generated by higher usage

volumes (Figure 2.4d).

Tiered pricing scales effortlessly with usage growth. As markets mature and average usage intensity rises, platforms should adjust membership fees upward while keeping transaction prices and quality levels constant. This avoids costly redesigns of service tiers, allowing platforms to capture value through simple fee adjustments.




Figure 2.4: Scenario III: Sensitivity of equilibrium outcomes to market usage scale $x_{\max}$.

Effect of Maximum Valuation ($v_{\max}$)

Figure 2.5 illustrates the impact of increasing the maximum valuation ($v_{\max}$), which drives coordinated adjustments across all strategic levers. As $v_{\max}$ increases, the platform simul-

taneously raises quality, membership fees, and transaction prices (Figure 2.5a). Consequently, profit grows linearly (Figure 2.5b).

The market structure responds with a distinct shift: the premium segment expands while the regular segment shrinks (Figure 2.5c). Usage intensity follows a similar divergence, rising for premium users but falling for regular users (Figure 2.5d). This pattern reflects *quality-driven polarization*. Higher valuations justify greater quality investment, which necessitates higher fees. As the quality gap widens, the regular tier loses its appeal—it lacks the performance of the premium tier but is no longer cheap enough to justify casual use. The “middle option” is effectively squeezed out by the high-quality, high-price strategy.

Platforms in high-valuation markets should abandon middle-tier strategies. When willingness-to-pay is high, the optimal response is aggressive quality investment paired with premium pricing. Maintaining a robust regular tier in such markets is inefficient; profitability is maximized by moving toward high-margin, high-quality exclusivity.

Effect of Discount Rate (d)

Figure 2.6 illustrates the effect of the discount rate (d) on platform strategy. The discount rate triggers a sharp structural transition at $d \approx 0.38$. Below this threshold, the platform maintains low membership fees ($F < 1.5$), high transaction prices ($c \approx 1.0$), and quality at $\sigma \approx 1.9$ (Figure 2.6a). Market structure is binary: premium users comprise 43% and non-users 57%, with no regular tier (Figure 2.6c). Above the threshold, membership fees spike to $F > 4$, transaction prices drop to $c \approx 0.6$, and quality rises slightly to $\sigma \approx 2.0$. A three-tier market emerges with premium users falling to 19%, regular users appearing at 28%, and non-users at 53%. Premium usage drops from 2.45 to 2.0 at the threshold, while regular usage jumps from zero to 0.9 (Figure 2.6d). Profit declines monotonically from 1.55 to 1.35 as d increases (Figure 2.6b).

This threshold reflects a revenue structure constraint. At low discounts, minimal price differentiation between tiers allows accessible membership fees, driving mass premium adoption through high transaction volumes. At high discounts, revenue loss from dis-




Figure 2.5: Scenario III: Sensitivity of platform outcomes to changes in $v_{\max}$.

counted transactions forces substantially higher membership fees. High fees screen out casual users, who populate the regular tier instead. The platform lowers transaction prices to retain these users, shifting from transaction-based to membership-based revenue.

The finding is that aggressive discount programs reduce profitability. Deep discounts require prohibitive membership fees that exclude middle market users and lower overall profit. Platforms maximize profit through moderate discounts that enable accessible membership fees and broad premium adoption. Subscription platforms typically offer discounts around 30 to 40 percent rather than steep reductions that force exclusionary pricing models.




Figure 2.6: Scenario III: Sensitivity of platform outcomes to changes in d .

Effect of Quality Cost Coefficient (k)

Figure 2.7 reveals that the quality cost parameter k triggers hyperbolic sensitivity in strategic variables while leaving market structure invariant. At near-zero costs ($k < 0.05$), the platform pursues extreme differentiation (Figure 2.7a), with membership fees exceeding $F = 200$ and quality levels above $\sigma = 50$. As k increases from 0 to 0.1, membership fees collapse by over 90 percent to $F \approx 5$ and quality drops to $\sigma \approx 2$. Beyond $k > 0.4$, all variables stabilize with $F \approx 4$, $c \approx 0.6$, and $\sigma \approx 2$. This collapse drives a corresponding profit decline from extreme highs to modest levels around 1.5 (Figure 2.7b).

Despite this volatility in pricing, the market structure remains remarkably stable. Pre-

mum users hold at 23 percent, regular users at 22 percent, and non-users at 55 percent across all cost levels (Figure 2.7c). Premium usage intensity rises gradually from 1.63 to 1.78, while regular usage stays flat at 0.6 (Figure 2.7d). The transaction price c is invariant at approximately 0.6, indicating that the pricing structure depends on customer preferences rather than quality provision costs.

This pattern occurs because k affects what the platform can afford to offer without changing customer heterogeneity or segmentation thresholds. When quality is cheap to produce, the platform maximizes differentiation through extreme quality and extracts surplus via high membership fees. Rising costs force reductions in both quality and fees to maintain economic viability, but customer sorting between premium, regular, and non-participation remains unchanged.

Small operational efficiency advantages yield disproportionate returns. Platforms with marginally lower quality provision costs can sustain high-fee, high-quality equilibria that competitors cannot match. Beyond $k > 0.2$, platforms enter low-margin equilibria where significant differentiation becomes unaffordable, forcing convergence toward standard offerings.




Figure 2.7: Scenario III: Sensitivity of platform outcomes to changes in quality cost k .

2.3.4 Comparative Statics: Robustness of the Joint Optimization Strategy

Table 2.4 summarizes the sensitivity of the optimal strategy to changes in key market parameters: usage capacity ($x_{\max}$), customer valuation ($v_{\max}$), and the cost of quality (k). To conduct this analysis, we vary these parameters within the ranges specified in Table 1.2. These comparative statics reveal the stability of the platform's strategic levers under varying market conditions

Table 2.4: Comparative Statics: Scenario IV (Endogenous Quality and discount)

Parameter	Decision Variables				Market Distribution			Usage Intensity		π
	F	c	d	σ	β^P	β^R	β^N	δ^P	δ^R	
$x_{\max}$	$\nearrow$	—	—	—	—	—	—	$\nearrow$	—	$\nearrow$
$v_{\max}$	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$	$\searrow$	—	$\nearrow$	$\searrow \nearrow$	—	$\nearrow$
k	$\searrow$	—	—	$\searrow$	$\nearrow$	$\searrow$	$\searrow$	$\nearrow \searrow$	—	$\searrow$

Notes: $\nearrow$ = increase; $\searrow$ = decrease; — = no change; $\nearrow \searrow$ = non-monotonic.

Effect of Changing Market Usage Capacity ($x_{\max}$)

We examine the sensitivity of the optimal joint strategy, membership fee (F), per-use price (c), quality (σ), and discount rate (d), as the maximum market usage capacity ($x_{\max}$) scales from 1 to 10. The results reveal a fundamental structural property of the joint optimization model: *strategic invariance* with respect to market scale.


As illustrated in Figure 2.8, the platform’s optimal architecture remains structurally stable across all market sizes. The per-transaction price remains fixed at its upper bound ($c^* = 1.00$) and the discount rate remains zero ($d^* = 0$), while service quality stabilizes at $\sigma^* \approx 1.96$.

The only strategic variable that responds to market expansion is the membership fee (F^*), which scales linearly from \$0.03 to \$0.27. Consequently, customer segmentation remains effectively binary (approximately 43% Premium, 0% Regular, 57% Non-users). Driven by the linear increase in premium usage intensity (δ^P), total profit scales proportionally (a 9.5-fold increase) without requiring adjustments to the core pricing or quality structure.


The invariance of (c^*, σ^*, d^*) arises from the scaling properties of user demand. Increasing $x_{\max}$ extends the support of the customer distribution, implying that larger markets contain customers with proportionally higher usage needs. The platform captures this additional value not by redesigning its menu, but by allowing high-type customers to self-select into the existing premium tier, where they naturally consume more.

The optimal quality level (σ^*) remains constant because the marginal benefit of quality scales proportionally with the marginal cost of providing it. Similarly, a zero-discount


strategy ($d^* = 0$) persists because quality-based vertical differentiation strictly dominates price-based screening: offering discounts would cannibalize high-margin revenue from premium users with minimal gain in lower-value segments.




(a) Decision Variables: F and c




(b) Decision Variables: σ and d



(c) Customer Segment Fractions



(d) Demand Quantities



(e) Profit

Figure 2.8: Scenario IV: Sensitivity of platform outcomes to changes in $x_{\max}$.

Sensitivity Analysis: Impacts of Maximum Customer Valuation ($v_{\max}$)

We analyze the optimal platform strategy as the maximum customer valuation ($v_{\max}$) increases from 0.1 to 1.0, representing markets with varying willingness-to-pay ceilings. Figure 2.9 reveals that, unlike changes in market size ($x_{\max}$) where strategies remain structurally invariant, changes in customer valuation trigger a distinct regime shift in the revenue model. Specifically, we observe a structural transformation at a critical threshold of $v_{\max} \approx 0.75$. Below this threshold ($v_{\max} < 0.75$), the platform adheres to a strategy of transaction price dominance, where the per-use price (c^*) serves as the primary revenue driver, scaling linearly with market valuation from 0.07 to approximately 0.75 (Figure 2.9a). In this lower-valuation regime, the optimal membership fee is strictly zero ($F^* = 0$), indicating that the platform maximizes profit by removing entry barriers to capture value solely through usage fees. However, once valuation exceeds the critical threshold ($v_{\max} > 0.75$), a positive membership fee emerges, rising convexly to $F^* \approx 0.25$ at $v_{\max} = 1.0$ (Figure 2.9a). This signifies a transition to a two-part tariff structure. Throughout this transformation, the platform consistently prioritizes quality-led differentiation; service quality (σ^*) increases monotonically with valuation from 1.11 to 1.95, while the discount rate remains negligible ($d^* \approx 0$) (Figure 2.9b). This confirms that in high-valuation markets, the platform differentiates via vertical quality enhancement rather than horizontal price discounting.

The shift from a zero to a positive membership fee ($F^* = 0 \rightarrow F^* > 0$) highlights the economic limits of the membership model. In markets with lower maximum valuations, the potential consumer surplus is insufficient to support an upfront access fee without deterring participation. Consequently, the platform operates as a pure transaction-based service, capturing value solely through the per-use price (c), which tracks the users' willingness to pay. However, as $v_{\max}$ rises, the surplus generated by high-quality service (σ^*) expands beyond what can be captured by transaction pricing alone, given that c is tightly bounded by $v_{\max}$. Once this surplus exceeds the critical threshold, the platform introduces the membership fee (F) as a second instrument to extract the residual value from high-

type customers, effectively transitioning from a linear pricing model to a two-part tariff to capture the excess utility generated by premium service quality.

These findings suggest a tiered entry strategy based on regional economic maturity. In price-sensitive regions characterized by low willingness-to-pay, platforms should avoid paid memberships and focus on a single, high-quality transaction service with prices calibrated to the local ceiling; introducing an upfront fee in these markets would be counterproductive. Conversely, in affluent markets, platforms should unlock the “membership lever,” complementing high transaction prices with an upfront fee to monetize the premium service quality. This pattern validates the global segmentation strategies observed in the ride-hailing industry. In lower-income markets such as Southeast Asia and India, platforms like Grab and Uber historically operated on a pure pay-per-ride basis to minimize entry barriers. In contrast, in high-income markets like the United States and Canada, the same platforms aggressively market paid subscriptions (e.g., Uber One) to capture high-valuation surplus. This confirms that the monetization mechanism must evolve from pure transaction fees to a hybrid membership model only when the local willingness-to-pay threshold is crossed.




Figure 2.9: Scenario IV: Sensitivity of platform outcomes to changes in $v_{\max}$.

Sensitivity Analysis: Impacts of Quality Provision Cost (k)

We examine the platform's optimal strategy as the quality cost coefficient (k) varies from 0.0 to 1.0. This parameter captures the transition from low-marginal-cost industries (e.g., digital services) to high-marginal-cost industries (e.g., physical logistics). Figure 2.10

reveals that platform strategy is highly elastic to these costs.

As the cost of quality rises, the platform's value generation capability decays rapidly. Profit drops from $\pi \approx 3.7$ at $k = 0.05$ to $\pi \approx 1.6$ at $k = 1.0$ (Figure 2.10e). This decline is driven by a forced reduction in service differentiation: optimal quality (σ^*) falls from a high of ≈ 6.2 to ≈ 1.6 (Figure 2.10b). Correspondingly, the membership fee (F^*) collapses from an exclusionary price of $F \approx 19.2$ to near zero ($F \approx 0.0$) (Figure 2.10a). Throughout this range, the transaction fee c remains relatively stable, confirming that the platform adjusts quality levels rather than transaction prices in response to cost pressures (Figure 2.10a).

A critical structural break in market segmentation is observed at $k \approx 0.15$. Below this threshold ($k < 0.15$), the market supports a three-tier structure where a “Regular” user segment (β^R) exists alongside Premium and Non-users (Figure 2.10c). However, as costs exceed $k > 0.15$, the Regular tier vanishes entirely ($\beta^R \rightarrow 0$), and the market consolidates into a binary structure. Counterintuitively, as quality becomes more expensive and service levels degrade, the fraction of premium users (β^P) actually *increases*, jumping from $\approx 25\%$ at low costs to stabilize around $\approx 50\%$ at higher costs (Figure 2.10c).

This “Participation Paradox” reflects a strategic pivot from exclusionary segmentation to mass-market consolidation. When quality is inexpensive ($k < 0.15$), the platform creates a “luxury” tier with high differentiation and charges a prohibitive membership fee ($F \approx 19$). This effectively screens the market: only the highest-valuation users join the premium tier. However, as the marginal cost of quality rises, sustaining such a large value gap becomes impossible. To maintain revenue, the platform slashes F aggressively (from 19 to < 1), effectively commoditizing the premium tier. Consequently, “Regular” users are absorbed into the now-affordable “Premium” tier, explaining why β^P rises while β^R disappears. The platform effectively abandons the “VIP” approach in favor of a “volume” business model.

A second critical transition occurs when quality costs exceed $k = 0.72$. Throughout the range $k < 0.72$, the platform sets $d \approx 0$, relying exclusively on quality enhancement to attract premium users (Figure 2.10b). At $k = 0.72$, this strategy breaks down. The opti-

mizer introduces a positive discount for the first time ($d = 0.011\%$), which then increases systematically: $d = 0.13\%$ at $k = 0.73$, $d = 0.96\%$ at $k = 0.80$, and $d \approx 2.1\%$ at $k = 0.99$. This shift reflects a change in how the platform attracts customers. When quality investment is affordable ($k < 0.72$), premium users join for superior service ($\sigma = 6.16$). When quality becomes prohibitively expensive ($k > 0.72$), quality degrades to $\sigma \approx 1.55$ and the platform instead offers price savings to recruit members. However, this occurs under constrained conditions. The membership fee has collapsed from $F = 19.3$ to $F \approx 0$ (Figure 2.10a), the regular tier has vanished ($\beta^R \approx 0$) (Figure 2.10c), and the market structure is binary. The discount emerges not as a profitable strategy but as a substitute mechanism when rising operational costs eliminate quality investment as a viable option. This demonstrates that platforms resort to discounting only when cost structures prevent them from competing on service quality.

These dynamics suggest that a platform’s optimal business model is fundamentally shaped by its operational cost structure. For industries with low marginal costs of quality enhancement—such as music streaming services (e.g., Spotify)—the optimal strategy is “High Quality, Exclusive Access.” Because the cost of providing high-bitrate, ad-free streaming is negligible ($k \approx 0.05$), the platform can differentiate aggressively, effectively eliminating any “middle” tier (such as per-song purchasing) and converging toward a binary free vs. premium structure.

Conversely, for operationally intensive industries—such as last-mile delivery (e.g., DoorDash)—where improving service speeds incurs steep, convex costs ($k > 0.5$), the optimal strategy shifts to “Moderate Quality, Mass Adoption.” In this high-cost regime, offering a “super-premium” tier (e.g., guaranteed instant delivery) becomes prohibitively expensive. Instead, the platform provides a single broad-appeal membership (DashPass) with a low monthly fee targeted at maximizing volume, precisely mirroring our model’s prediction for the high- k environment.




Figure 2.10: Scenario IV: Sensitivity of platform outcomes to changes in k .

The prevalence of low optimal discount levels is driven by the relative cost of quality improvements compared to per-unit discounting. In the model, both discounting and quality enhancement increase the utility of premium users in proportion to their total usage. However, the associated costs differ substantially. Discounts reduce transaction revenue on every premium usage, whereas quality improvements impose a direct provision cost

that depends on the quality cost parameter k .

In the baseline specification, quality improvements are relatively inexpensive compared to discounting. As a result, increasing service quality generates additional customer utility at a much lower cost than offering per-unit discounts. This cost asymmetry leads the platform to rely primarily on quality-based differentiation while maintaining low discount levels. The optimizer therefore favors improving σ rather than increasing d , even when both instruments affect demand in similar ways.

This mechanism also explains why discounting becomes more prominent when the cost of quality provision increases. As k rises, the marginal cost of further quality improvements grows rapidly due to the convex structure of the quality cost function. Beyond a threshold, investing in quality is no longer cost-effective, and the platform substitutes toward discounting despite its revenue cannibalization effects. This pattern is consistent with the extension results, where positive discount levels emerge as quality provision becomes increasingly costly.

2.4 Conclusion

This chapter studies how an on-demand service platform should design its pricing and membership program when membership fees, transaction prices, discount rates, and service quality are all endogenous. The central question is whether subscription programs should rely primarily on price discounts or service differentiation when platforms can fully optimize their design.

The main result is that service quality dominates discounting as the optimal differentiation instrument across a wide range of conditions. Discounts reduce revenue on every subscriber transaction and therefore provide limited screening power. In contrast, service quality improvements can be targeted exclusively to subscribers and monetized through membership fees, allowing the platform to extract surplus more efficiently. As a result, optimal discount levels are driven close to zero when quality can be adjusted freely.

This shift toward quality-based differentiation has important implications for market

structure. Under joint optimization, the commonly assumed three-tier structure of subscribers, regular users, and non-users is not robust. Instead, the platform optimally serves a binary market in which users either subscribe to access enhanced service or exit the market. Intermediate pay-per-use segments become redundant once quality and pricing instruments are jointly optimized.

The analysis also clarifies the conditions under which this result holds. When the cost of improving service quality is sufficiently low, platforms optimally rely on quality-based membership programs with minimal discounting. As quality provision becomes more expensive, platforms substitute toward price-based incentives, although this shift is associated with lower profitability. These findings highlight that the design of subscription programs depends critically on the relative costs of quality provision and price incentives.

Overall, the results show that pricing structures and market segmentation are outcomes of optimization rather than fixed design choices. Platforms that can deliver service quality improvements efficiently should emphasize quality-based membership access rather than transaction-level discounts, while platforms facing high quality costs should expect weaker differentiation and greater reliance on price competition.

References

- Chen, D., Zhou, Y., Guan, X., & Lin, X. (2022). Transaction or membership? impact on on-demand delivery service platforms' profits, consumer surplus, and labor welfare. *Journal of Systems Science and Systems Engineering*, 31(5), 563–593.
- Chen, L. (2019). *The proportional fee, transaction fee and membership fee* [Working paper].
- Gangwar, M., & Bhargava, H. K. (2023). Pricing on-demand services: alternative ways of combining usage and access fees. *Production and Operations Management*, 32(1), 11–27.
- Hu, M., Qin, T., Wang, L., & Zhang, Z. J. (2023). *Subscription vs. spot pricing in on-demand economy* [Working paper].

- Ma, H., Lu, D., & Fan, Z.-P. (2025). Product selection, optimal product launch quantity and service pricing of a B2C service sharing platform. *Computers & Industrial Engineering*, 199, 111212.
- Zhang, C., & Tian, Y.-X. (2024). Joint optimization of dynamic pricing, quality, and production quantity in the presence of online reviews: a distributionally robust optimization approach. *Computers & Operations Research*, 169, 106720.

Chapter 3

Freemium Pricing and CRM Expenditures by a Digital Platform

3.1 Introduction

Many digital platforms, such as Spotify and YouTube, have implemented a freemium business model offering a basic service tier free of charge, with revenues generated from advertising, alongside a premium tier that provides enhanced features for a subscription fee and is advertising free. In a nutshell, such a model allows for building large user bases, generating dual revenue streams and creating upgrade paths for monetization (Kumar, 2019; Ross, 2018; Wagner et al., 2020)

The freemium business model represents a \$120+ billion global market growing at 15–20% annually (MoldStud, 2025), powering major platforms like Spotify (515 million users, \$13.2 billion revenue (Spotify, 2025)) and YouTube (2.7 billion users, \$29.2 billion advertising revenue (Business of Apps, 2024)).

Unlike traditional subscription or advertising-based businesses that face binary customer decisions, freemium platforms must simultaneously optimize multiple user segments with fundamentally different behaviors and revenue contributions. Free users generate advertising revenue but consume infrastructure resources, while premium users pro-

vide higher-margin subscription revenue but require continuous value delivery to prevent churn. This creates complex strategic interdependencies where actions that improve one metric often harm other metrics, increasing advertising intensity can boost short-term revenue but may accelerate premium conversions or trigger user churn, while enhancing premium features can improve retention but may reduce conversion rates if perceived value gaps become too wide. These optimization challenges are further complicated by competitive dynamics where platforms compete simultaneously on free service quality and premium value propositions, making traditional business optimization frameworks insufficient for freemium success (Kumar, 2019; Ross, 2018; Wagner et al., 2020).

The existing literature has typically examined the various components separately. Studies on customer relationship management focus on the acquisition-retention trade-off (Ben Rhouma & Zaccour, 2018; Blattberg & Deighton, 1996), while research on two-sided markets analyzes advertising and pricing decisions (Anderson & Gabszewicz, 2006; Sato, 2019). Recent work has explored tier differentiation (J. Li et al., 2024) and conversion dynamics (Niemand et al., 2019). However, this fragmented approach fails to capture the full system dynamics that determine platform performance. To the best of our knowledge, no existing model jointly optimizes acquisition effort, retention effort, advertising intensity, and subscription pricing within a unified framework that accounts for their interactions.

To fill this gap, we develop an integrated analytical model of freemium platform optimization. We model the platform's user base as a Markov chain with three states, namely, non-users, free users, and premium users, where transition probabilities depend endogenously on the platform's strategic choices. This approach captures how decisions in one area affect outcomes in others, enabling true system-wide optimization. Our research addresses three specific questions:

1. How should platforms jointly determine acquisition effort, retention effort, advertising intensity, and subscription pricing?
2. How do platform characteristics and user behavioral parameters shape the optimal

configuration of platform decisions?

3. Under what conditions does the platform prioritize acquisition versus retention, and how does this balance depend on the monetization strategy?

To address these questions, we develop a profit maximization model grounded in a steady-state Markov chain framework. The model characterizes the long-run distribution of users across three mutually exclusive states (non-users, free users, and premium users) and determines the platform’s optimal strategic choices under steady-state conditions. The platform’s decision variables include acquisition effort, retention effort, ad intensity in the free tier, and the subscription fee charged to premium users. These controls influence user transitions between states and, in turn, shape the steady-state composition of the user base and the platform’s overall profitability. Our results show that the optimal strategy depends on how users respond to advertising and pricing trade-offs. For example, higher advertising intensity can increase revenue from free users but may also reduce retention or trigger upgrades, depending on user sensitivity. Similarly, the subscription fee must be carefully set to balance conversion incentives with long-term revenue from premium users.

This paper contributes to the literature in three ways. First, we provide the first integrated model that jointly optimizes all four critical platform decisions, revealing interactions that single-dimension models miss. Second, we characterize how optimal strategies vary with user behavioral parameters, providing insights into platform differentiation. Third, we derive conditions under which platforms should emphasize acquisition versus retention, contributing to the broader customer relationship management literature.

This paper is structured as follows. Section 3.2 reviews the relevant literature. Section 3.3 presents the model, and Section 3.4 provides the numerical analysis, covering parameter calibration, sensitivity analysis, strategic insights, and numerical validation. Section 3.5 concludes. All technical derivations are provided in the Appendix.

3.2 Literature Review

Our work draws from, and contributes to, two streams of literature, namely, customer acquisition and retention and freemium revenue model.

3.2.1 Customer Acquisition and Retention

A substantial body of research has examined how firms should optimally allocate resources between customer acquisition and retention, particularly in the context of subscription-based services. Blattberg and Deighton (1996) were among the first to formally conceptualize this trade-off, suggesting that firms often overspend on acquisition despite the widely cited rule that acquiring a new customer can cost up to five times more than retaining one. Building on this, Pfeifer (2005) developed a mathematical model showing that firms should equalize the marginal costs of acquisition and retention, both tied to expected customer lifetime value (CLV). Subsequent research has developed sophisticated approaches to acquisition and retention optimization, including timing models (Schweidel et al., 2008), dynamic frameworks (Venkatesan & Kumar, 2004), and resource allocation models (Reinartz et al., 2005), though most work focuses on traditional binary customer relationships rather than the complex multi-state structures found in digital platforms.

Several factors influence the optimal balance between acquisition and retention efforts. Reinartz et al. (2005) directly address the customer profitability-maximizing balance between acquisition and retention resources, developing a modeling framework that determines optimal marketing spending allocation and distribution across communication channels. Kalyanam et al. (2025) examine advertising content's impact on both customer acquisition and retention. Customer base characteristics play a critical role, with Gregory (2016) demonstrating that firms with small customer bases should prioritize acquisition while larger firms should gradually shift emphasis toward retention as their base grows. Platform-mediated strategies introduce additional complexity, as Tang et al. (2021) show that sellers contracting with online platforms face trade-offs between price discrimination for acquisition versus quality differentiation for retention, with optimal strategy depend-

ing on consumer retention sensitivity to product quality and platform agency fees. Market dynamics introduce additional complexity, with Min et al. (2016) finding that acquisition costs are more sensitive to competitive pressure and market position than retention costs, suggesting asymmetric responses to environmental changes. Service quality considerations further complicate these decisions, as Afeche et al. (2017) show that when customers are sensitive to service quality, acquisition and retention efforts become interdependent, requiring integrated optimization approaches.

Mathematical modeling approaches have evolved to capture these complex relationships. Ben Rhouma and Zaccour (2018) developed a diffusion-based model incorporating CRM expenditures into customer equity maximization, using dynamic optimization to show that optimal acquisition and retention effort depends on market penetration effects and customer interaction patterns. Their approach employs differential equations to model customer base evolution, providing insights into how marketing investments influence long-term customer composition. Chang et al. (2020) extended this framework to competitive settings using differential game theory, demonstrating that optimal CRM allocations depend on firms' relative marginal customer equity and effectiveness in acquisition versus retention activities.

3.2.2 Freemium Revenue Model

Freemium business models create unique optimization challenges by combining dual revenue streams, advertising from free users and subscription fees from premium users, within a three-tier user structure. Unlike traditional subscription services, freemium platforms must simultaneously optimize user acquisition, conversion between tiers, and monetization across different user segments. Kumar (2014) established foundational principles for freemium success, emphasizing that platforms must balance free value provision with premium differentiation to drive sustainable conversion and retention. Mai and Hu (2023) extend this framework by developing dynamic optimization models for freemium services throughout their lifetime, comparing advertising-based monetization (Spotify-

inspired) versus feature-limitation strategies (Dropbox-inspired), and demonstrating that platforms should prioritize growth in early stages before shifting focus to monetization after building critical user mass.

The dual revenue structure creates complex trade-offs between advertising intensity and subscription pricing. Sato (2019) models freemium as optimal menu pricing in two-sided markets, demonstrating that offering an ad-supported free tier alongside an ad-free premium tier maximizes profits under heterogeneous user preferences for ad tolerance. Anderson and Gabszewicz (2006) show that in two-sided markets, advertising revenue and subscription revenue interact non-linearly, as advertising intensity affects both current revenue and future conversion rates. X. Li et al. (2024) demonstrate that these revenue streams are often conflicting, with subscription divisions underpricing relative to platform-optimal levels due to failure to account for positive effects of subscription fees on advertising revenue.

Platform transitions from free to fee-based models reveal critical insights for freemium optimization. Pauwels and Weiss (2008) analyze how firms successfully move from free to “free and fee” models, finding that this transition involves fundamental trade-offs between advertising revenue losses and subscription revenue gains. Their empirical analysis shows that different marketing actions have asymmetric effects on free versus paid subscriptions, while free-to-fee conversion campaigns create a double-edged effect, increasing subscription revenue while reducing advertising revenue, demonstrating the complex interactions between advertising intensity and membership fee decisions.

Platform-specific factors complicate these optimization decisions. J. Li et al. (2024) explore optimal tier differentiation strategies but treat advertising intensity and premium pricing as fixed parameters, missing critical interactions between these decisions. Lin (2020) addresses this gap by modeling subscription pricing and advertising exposure as jointly optimizable levers, revealing that personalized approaches enable platforms to extract surplus from diverse user segments. Prasad et al. (2003) extend this framework by proposing hybrid models with varying advertising loads across tiers.

Conversion dynamics represent a critical component of freemium optimization. Nie-

mand et al. (2019) identify the “zero-price effect,” in which users’ attachment to free offerings creates conversion barriers unless premium value is clearly differentiated. Wagner et al. (2014) show that successful upgrades depend on perceived fit between user needs and premium features, suggesting that feature design and pricing must be jointly optimized with marketing efforts. Runge et al. (2022) investigate the impact of pricing promotions on freemium conversion rates and revenue via field experiments, but focus solely on pricing strategies without considering interactions with acquisition effort, retention efforts, or advertising intensity.

3.2.3 Our Contribution

Despite significant advances in both customer relationship management and freemium business model research, existing models face a fundamental limitation: they cannot handle the joint optimization of acquisition spending, retention effort, advertising intensity, and membership fees that freemium platforms require. Traditional CRM optimization models assume binary customer states and single revenue streams (Gregory, 2016; Reinartz et al., 2005), while freemium platforms operate with three-state user structures and dual revenue sources that create complex interdependencies between all strategic decisions. Freemium pricing models optimize tier differentiation and subscription fees but treat user transitions as exogenous (Kumar, 2014; Sato, 2019), missing how acquisition and retention efforts shape user composition over time. Even advanced work addressing revenue stream interactions focuses on advertising-subscription trade-offs without considering marketing investments (X. Li et al., 2024; Pauwels & Weiss, 2008). Empirical studies emphasize user retention or acquisition strategies in isolation (Lee et al., 2013; Sciglimpaglia & Raafat, 2022), failing to capture the systematic interactions that determine platform performance. The result is a fragmented literature that provides optimization insights for individual decisions but cannot guide platforms facing integrated strategic choices where advertising intensity affects conversion rates, membership fees influence churn patterns, and marketing investments determine user composition. Even

advanced work like Mai and Hu (2023), which addresses dynamic freemium optimization across growth and monetization phases, focuses on temporal transitions rather than the simultaneous optimization of acquisition, retention, advertising, and pricing decisions that platforms face at any given time.

This paper addresses this gap by developing the first integrated optimization model for freemium platform management that simultaneously determines acquisition effort, retention effort, advertising intensity, and membership fees. We model user behavior through a three-state Markov chain with endogenous transition probabilities that depend on all platform decisions, capturing how strategic choices influence user movements between non-user, free-user, and premium-user states. The model is solved numerically for steady-state profit maximization, yielding optimal strategies under varying behavioral and market conditions. Our approach reveals systematic interactions that single-dimension models miss, demonstrates how optimal acquisition-retention trade-offs depend on advertising and pricing decisions, and characterizes conditions under which platforms should prioritize different strategic levers. By integrating all four core decisions into a unified framework, this research contributes to multiple literature streams while providing actionable insights for platform strategy in settings where monetization and user growth must be balanced over the long term.

3.3 The Model

We consider a freemium platform, e.g., a music streaming service, that generates revenue through advertising shown to free-tier users and membership fees paid by premium users. To grow and sustain its user base, the platform invests in marketing activities directed at both user acquisition and retention. In this study, we develop a profit-maximization model in which the platform endogenously and jointly optimizes four key managerial levers: acquisition effort, retention effort, advertising intensity for free users, and subscription pricing for premium users. For simplicity, we refer to the freemium platform simply as “the platform” throughout the remainder of this paper.

To build the model, we begin by specifying the platform’s key decision variables and their functional roles. Acquisition effort, denoted by A , targets non-users through broad outreach activities, such as digital advertising, social media campaigns, and promotional offers, to increase platform visibility and attract new users. Retention effort, denoted by R , is exclusively targeted at premium users, aiming to keep them engaged and subscribed to the platform over time.

In addition to marketing allocation, the platform controls two income-related levers: the advertising intensity (H) experienced by free-tier users, and the subscription fee (C) charged to premium users. These four decision variables, i.e., A , R , H , and C , form the basis of the platform’s strategic controls. In the following section, we formalize how these variables influence user behavior, platform engagement, and profitability within a steady-state Markov framework.

3.3.1 Modeling Framework Overview

Having established the platform’s four key decision variables, we now develop a formal model of how these decisions influence user behavior and platform profitability. Our approach makes three key modeling choices that distinguish this work from prior literature. First, we use a Markov-chain framework to model user state transitions. Consistent with Pfeifer and Carraway (2000), Markov models capture multi-state customer relationships and migration patterns better than static segmentation, recognizing that users move dynamically between engagement levels as their platform experience evolves. Second, we endogenize transition probabilities as functions of platform decisions. This allows us to analyze how changes in acquisition effort, retention effort, advertising intensity, and pricing jointly affect the steady-state user distribution, a capability absent in models that treat these decisions independently. Third, we focus on steady-state equilibrium rather than dynamic optimization. While this simplifies the analysis, it captures the long-run strategic trade-offs that are central to sustainable freemium platform management.

We develop this framework in two steps: customer behavior modeling (Section 3.3.2)

and platform optimization (Section 3.3.3).

3.3.2 Customers' Behavior

At any time, each customer is in one of three mutually exclusive and collectively exhaustive states: non-user (N), free user (F), or premium user (P), with the total population normalized to one. Non-users are individuals who have not adopted the platform. Free users are active users who access the platform without payment but experience advertising. Premium users are subscribers who pay a fee for enhanced, ad-free service.

Customer movement between states follows a Markov process where transition probabilities depend on the platform's strategic decisions. The key insight is that the platform can influence these probabilities through its choice of A, R, H , and C . Figure 3.1 illustrates the possible transitions between the three user states and their dependence on the platform's decision variables.

The transition probabilities ρ_{ij} denote the probability that a user moves from state i to j , where $i, j \in \{N, F, P\}$. As illustrated in Figure 3.1, each transition depends on specific combinations of the platform's decision variables. The transition matrix is defined by $\mathcal{P}$.

In addition to the platform's strategic decisions, transition probabilities are influenced by exogenous parameters. Two more influential of them are the baseline quality (b) and the premium quality multiplier (γ). The baseline quality (b) captures the appeal of the standard service available to all users, including core features such as content accessibility, platform usability, and overall user experience. A higher b enhances user satisfaction and increases the likelihood of adoption and retention. The premium quality multiplier (γ) represents the added value exclusive to premium users, such as offline access, improved media quality, or curated content. While retention effort (R) promotes engagement through external means like targeted campaigns, γ reflects built-in enhancements that raise the intrinsic appeal of the premium tier. Together, b and γ influence users' perceived value of the platform and their transitions across user states.




Figure 3.1: Transition States

Following Carroni and Paolini (2020), we assume that a higher level of advertising intensity (H) increases advertising revenue but decreases satisfaction among free users, which in turn decreases the transitions into free usage, $\frac{\partial \rho_{iF}}{\partial H} < 0$ for $i \in \{N, F, P\}$. Similarly, a higher subscription fee (C) increases revenue from premium users but lowers satisfaction, reducing the likelihood of transitions into the premium-user state, where $\frac{\partial \rho_{iP}}{\partial C} < 0$ for $i \in \{N, F, P\}$.

Building on the work of Blattberg and Deighton (1996), Berger and Bechwati (2001), and Ben Rhouma and Zaccour (2018), we use exponential functions to describe how marketing expenditures affect customer acquisition and retention rates.

When customers are in the non-user state (N), their transition to the free-user state is influenced by acquisition effort (A) and advertising intensity (H). Increasing A raises the probability of becoming a free user ($\frac{\partial \rho_{NF}}{\partial A} > 0$), while higher advertising intensity reduces it ($\frac{\partial \rho_{NF}}{\partial H} < 0$). This relationship is represented by the following model:

$$\rho_{NF} = 1 - e^{-(k_b b + k_A A - k_H H)},$$

where k_A and k_b ¹ are coefficients associated with acquisition effort and platform quality, respectively. The parameter k_H reflects users' sensitivity to advertising, with higher values indicating lower tolerance, while lower values implying greater tolerance.

We assume that the probability of a non-user instantly converting into a premium subscriber is effectively zero; consequently, we set

$$\rho_{NP} = 0,$$

in the transition matrix. This assumption is both theoretically grounded and empirically supported. Theoretically, the freemium model is designed around experience-based value discovery, i.e., users must experience the service to understand premium benefits (Kumar, 2014). Empirically, industry data consistently show that over 95% of premium conversions originate from active free users (Wagner et al., 2014; Yoganasimhan et al., 2023). While some platforms may have enterprise sales or targeted premium trials, these represent separate market segments outside our consumer-focused model. This assumption is standard in the freemium literature (J. Li et al., 2024; Sato, 2019) and is essential for capturing the sequential nature of freemium conversion funnels.

For customers in the free-user state (F), transitions are influenced by advertising intensity (H) and baseline service quality (b). A higher level of ad intensity reduces the probability of remaining a free user ($\frac{\partial \rho_{FF}}{\partial H} < 0$), as excessive ads may decrease user satisfaction. We represent this relationship as:

$$\rho_{FF} = 1 - e^{-(k_b b - k_H H)}.$$

The probability of transitioning from the free-user state to the premium-user state (ρ_{FP}) depends on advertising intensity (H , subscription fee (C), and the relative quality of the premium tier. This probability increases with higher advertising intensity ($\frac{\partial \rho_{FP}}{\partial H} > 0$), as users may upgrade to avoid ads, and decreases with a higher subscription fee ($\frac{\partial \rho_{FP}}{\partial C} < 0$).

¹Note that k_b and b always appear as a $(k_b b)$, implying that only their combined effect influences the model outcomes. However, we retain k_b separately for consistency across the formulation, so that every parameter or decision variable (e.g., H , R , C) is associated with a corresponding effectiveness coefficient (k_i). The impact of k_b is thus fully aligned with the impact of b .

The transition probability is given by

$$\rho_{FP} = 1 - e^{-(\gamma k_b b + k_H H - C)},$$

where $\gamma k_b b$ reflects the quality of the premium service relative to the baseline, $\gamma > 1$.

Finally, for customers in the premium-user state (P), transitions are influenced by retention effort (R), advertising intensity (H), and the subscription fee (C). The probability of downgrading to the free-user state (ρ_{PF}) decreases with retention effort ($\frac{\partial \rho_{PF}}{\partial R} < 0$), decreases with greater advertising intensity ($\frac{\partial \rho_{PF}}{\partial H} < 0$), and increases with higher membership fees ($\frac{\partial \rho_{PF}}{\partial C} > 0$). Premium users may downgrade if they perceive the service as less valuable relative to the fee paid or if the user experience becomes less satisfactory. This relationship is captured by the following expression:

$$\rho_{PF} = 1 - e^{-(k_b b + C - k_R R - k_H H)},$$

where k_H is the user's sensitivity to ad intensity.

The probability of remaining in the premium state (ρ_{PP}) increases with retention efforts ($\frac{\partial \rho_{PP}}{\partial R} > 0$), decreases with higher subscription fees ($\frac{\partial \rho_{PP}}{\partial C} < 0$) and is defined by

$$\rho_{PP} = 1 - e^{-(\gamma k_b b + k_R R - C)},$$

where $\gamma k_b b$ captures the enhanced quality of premium service relative to the baseline.

To wrap up, the transition matrix is given by

$$\mathcal{P} = \begin{bmatrix} e^{-(k_A A + k_b b - k_H H)} & 1 - e^{-(k_A A + k_b b - k_H H)} & 0 \\ e^{-(k_b b - k_H H)} + e^{-(\gamma k_b b + k_H H - C)} - 1 & 1 - e^{-(k_b b - k_H H)} & 1 - e^{-(\gamma k_b b + k_H H - C)} \\ e^{-(k_b b + C - k_R R - k_H H)} + e^{-(k_R R + \gamma k_b b - C)} - 1 & 1 - e^{-(k_b b + C - k_R R - k_H H)} & 1 - e^{-(k_R R + \gamma k_b b - C)} \end{bmatrix}$$

For all transition probabilities to lie within the interval $[0, 1]$, the following conditions must be satisfied:

$$k_b b \geq k_H H, \quad (3.1)$$

$$k_R R + k_H H - k_b b \leq C \leq \gamma k_b b + \min\{k_H H, k_R R\}, \quad (3.2)$$

$$1 \leq e^{-(k_b b - k_H H)} + e^{-(\gamma k_b b + k_H H - C)}, \quad (3.3)$$

$$1 \leq e^{-(k_b b + C - k_R R - k_H H)} + e^{-(k_R R + \gamma k_b b - C)}. \quad (3.4)$$

Table 3.1: Notation for the General Model

Notation	Description
<i>Model Parameters</i>	
b	Basic service offered by the platform, where $b > 0$
γ	Quality factor of the premium service, where $\gamma > 1$
k_H	Customers' advertising sensitivity, where $0 < k_H < 1$
k_i	Effectiveness coefficient for $i \in \{A, R, b\}$, where $0 < k_i < 1$
ω_i	Cost coefficient for marketing activity $i \in \{A, R\}$, where $0 < \omega_i < 1$
<i>State Variables</i>	
q_i	Fraction of customer type i , where $i \in \{N, F, P\}$ and $0 < q_i < 1$
<i>Decision Variables</i>	
A	Acquisition effort directed at non-users, where $A > 0$
R	Retention effort directed at premium users, where $R > 0$
H	Advertising intensity for free users, where $H > 0$
C	Membership fee for premium users, where $C > 0$

The interpretation of these inequalities is intuitive. Equations 3.1 and 3.2 ensure that the platform's feature quality b must exceed the deterrent effects of advertising intensity ($k_H H$) and/or subscription fee (C) to have an inflow to that state. In other words, to attract and retain customers effectively, the platform's value proposition must outweigh the frictions introduced by advertisements and the required retention effort. Equations 3.3 and 3.4 guarantee that $\rho_{FN} \geq 0$ and $\rho_{PN} \geq 0$, respectively.

At steady-state, the distribution of customers stops varying over time, meaning the inflow and outflow of each state are balanced. This equilibrium condition is expressed as $Q = Q\mathcal{P}$, where $Q = (q_N, q_F, q_P)$ represents the steady-state probability vector, with each element q_j denoting the long-run probability of a customer being in state j . Expanding this condition yields the system of equations

$$q_N = q_N(1 - \rho_{NF} - \rho_{NP}) + q_F(1 - \rho_{FF} - \rho_{FP}) + q_P(1 - \rho_{PF} - \rho_{PP}),$$

$$q_F = q_N\rho_{NF} + q_F\rho_{FF} + q_P\rho_{PF},$$

$$q_P = q_N\rho_{NP} + q_F\rho_{FP} + q_P\rho_{PP}.$$

These equations characterize the long-run behavior of customer transitions and ensure the

stability of each group's population.

In equilibrium, the system reaches the state where customer inflows and outflows are balanced across all categories. By solving the system of equations for the steady-state probabilities (q_N, q_F, q_P) , we can analyze how key market variables such as advertising revenue (H) , subscription fees (C) , and marketing efforts (A, R) affect customer distribution. Understanding these dynamics provides valuable insights for optimizing acquisition and retention strategies, ultimately improving market performance and long-term sustainability.

3.3.3 Platform Optimization

The platform's optimization problem is given as follows:

$$\max_{A, R, H, C} \pi = Hq_F + Cq_P - \left(\frac{1}{2} \omega_A A^2 q_N + \frac{1}{2} \omega_R R^2 q_P \right)$$

The platform generates revenue from two sources and incurs costs from marketing activities. Advertising revenue Hq_F is the product of ad intensity (H) and the fraction of free users (q_F) , while subscription revenue Cq_P is the product of subscription fees (C) and the fraction of premium users (q_P) . In terms of marketing costs, we assume that the cost of acquisition effort is given by $\frac{1}{2} \omega_A A^2 q_N$, and the cost of retention effort is $\frac{1}{2} \omega_R R^2 q_P$. The convexity of these cost functions reflects the idea that marketing activities are subject to marginal decreasing returns in scale.

Noting that $q_N = 1 - q_F - q_P$, we rewrite the profit function as follows, which highlights the contribution of each user segment:

$$\max_{A, R, H, C} \pi = q_F \left(H + \frac{1}{2} \omega_A A^2 \right) + q_P \left(C + \frac{1}{2} \omega_A A^2 - \frac{1}{2} \omega_R R^2 \right) - \frac{1}{2} \omega_A A^2. \quad (3.5)$$

In this formulation, free users contribute advertising revenue minus their share of acquisition costs, while premium users contribute subscription revenue plus acquisition cost savings minus retention costs. The platform also bears baseline acquisition costs regardless of user distribution.

Before deriving the optimal strategy, we characterize how the user distribution evolves to a steady-state. The key insight is that the fraction of users in each state depends on the complex interplay of all transition probabilities in the system. In equilibrium, the flow into each state equals the flow out:

$$q_F = \frac{\rho_{NF} + q_P(\rho_{PF} - \rho_{NF})}{1 - \rho_{FF} + \rho_{NF}}, \quad q_P = \frac{q_F \rho_{FP}}{1 - \rho_{PP}}, \quad q_N + q_F + q_P = 1.$$

Setting inflows equal to outflows and using the constraint $q_N + q_F + q_P = 1$, we can solve for the steady-state distribution, which are as follows:

$$q_F = \frac{\rho_{NF}(1 - \rho_{PP})}{\rho_{FF}\rho_{PP} - \rho_{FF} + \rho_{FP}\rho_{NF} - \rho_{FP}\rho_{PF} - \rho_{NF}\rho_{PP} + \rho_{NF} - \rho_{PP} + 1}, \quad (3.6)$$

$$q_P = \frac{\rho_{FP}\rho_{NF}}{\rho_{FF}\rho_{PP} - \rho_{FF} + \rho_{FP}\rho_{NF} - \rho_{FP}\rho_{PF} - \rho_{NF}\rho_{PP} + \rho_{NF} - \rho_{PP} + 1}. \quad (3.7)$$

The resulting closed-form expressions show that the steady-state fraction of free users (q_F) is driven by both the acquisition rate (the probability of non-users joining as free users) and the churn rate from the premium tier, while the fraction of premium users (q_P) depends jointly on the ability to acquire new users and convert them to premium status. The complex denominator in these expressions captures the dynamic interactions among retention, conversion, and churn across all tiers. Economically, these steady-state formulas reveal that optimizing one tier impacts the others: aggressive acquisition can increase both free and premium user bases, strong premium retention can shrink the free tier, and conversion and churn rates introduce feedback loops that shape the long-run user composition. Therefore, understanding these dynamics is essential for designing joint strategies that maximize platform performance.

Substituting for the transition probabilities in (3.6)-(3.7) leads to the following steady-state values for the fractions of free and premium users, respectively:

$$q_F = \frac{e^{Ak_A - Hk_H + bk_b} - 1}{e^{Ak_A} + e^{-Hk_H + Rk_R} - e^{-C + Rk_R + b\gamma k_b} + e^{Ak_A - Hk_H + bk_b} + e^{Ak_A - 2C + 2Rk_R + b\gamma k_b} - e^{Ak_A - C - Hk_H + 2Rk_R} - 1} \quad (3.8)$$

$$q_P = \frac{(-e^{-C + b\gamma k_b} - e^{Ak_A - 2Hk_H + bk_b} + e^{Ak_A - C - Hk_H + b\gamma k_b + bk_b} + e^{-Hk_H}) e^{Rk_R}}{e^{Ak_A} + e^{-Hk_H + Rk_R} - e^{-C + Rk_R + b\gamma k_b} + e^{Ak_A - Hk_H + bk_b} + e^{Ak_A - 2C + 2Rk_R + b\gamma k_b} - e^{Ak_A - C - Hk_H + 2Rk_R} - 1} \quad (3.9)$$

Now, substituting for q_F and q_P from (3.8)-(3.9) into (3.5) leads to a very long expression for the profit (see A.1). Schematically, we have a highly nonlinear profit function of the

form

$$\pi = f(A, R, H, C, q_F(A, R, H, C), q_P(A, R, H, C)),$$

which involves exponential transition probabilities, fractional user distributions, four coupled decision variables with cross-effects, and multiple threshold effects. The highly non-linear structure of the optimization problem precludes analytical solutions, making a numerical approach necessary. We employ the interior point algorithm for several reasons: it provides polynomial-time convergence guarantees for this class of problems; it naturally handles the inequalities in (3.1)-(3.4) without requiring constraint transformation; and it searches within the feasible region's interior, thereby avoiding the boundary-related numerical issues that can affect other methods.

We verify solution quality through several approaches. The algorithm consistently converges to the same solution from diverse starting points, first-order optimality conditions are satisfied (with gradient norms below 10^{-8}), and constraint violations are negligible (less than 10^{-10}). Sensitivity analysis further confirms smooth, continuous solution paths without discontinuous jumps, indicating stable and robust optimization behavior.

3.4 Numerical Results

In this section, we numerically solve the model for a large number of parameter values to assess parameter impacts comprehensively. First, we illustrate the kind of insight that we can get from the model in a benchmark case, and next run a sensitivity analysis to assess the impact of varying the parameter values on customer acquisition, churn, and retention across different states, and consequently on the platform long-term customer basis and profitability.

3.4.1 Benchmark Case

As a benchmark (or base case), we set the parameter values as follows: The impact and cost parameters have been set equal to 0.5 in the benchmark to simplify the presenta-

Table 3.2: Model Parameters and Ranges for Sensitivity Analysis

Parameter	Benchmark Value	Sensitivity Range
Basic service quality (b)	0.7	—
Premium quality multiplier (γ)	1.3	[1.0, 2.0]
Advertising impact (k_H)	0.5	[0.1, 1.0]
Quality impact (k_b)	0.5	[0.1, 1.0]
Acquisition impact (k_A)	0.5	[0.1, 1.0]
Retention impact (k_R)	0.5	[0.1, 1.0]
Acquisition effort cost coefficient (ω_A)	0.5	[0.1, 1.0]
Retention effort cost coefficient (ω_R)	0.5	[0.1, 1.0]

tion and discussion of the results. Each of these parameters will be evaluated for their varying values between 0.1 and 2.0 in the sensitivity analysis. We set the basic service quality b available for free users at 0.70 and the premium multiplier γ at 1.30. These values reflect platforms where the free tier provides substantial value (like Spotify’s full music library) while premium features add meaningful enhancement (ad removal, offline access) worth approximately 30% additional value. This specification aligns with industry analysis showing premium tiers typically provide 25-40% incremental value (Kumar, 2014).

Remark 2. *The sensitivity ranges have zero as their lower bound, allowing us to capture the case in which the relevant parameter has no impact, and an upper bound equal to twice the benchmark value. With intervals of this magnitude, we believe that the analysis covers a very rich set of parameter constellations.*

Solving the optimization problem with the benchmark parameter constellation, we obtain the following solution:

$$(A, R, H, C) = (0.83, 0.21, 1.04, 0.66), \quad \pi = 0.50,$$

with the corresponding transition matrix given by²

$$\mathcal{P} = \begin{bmatrix} 0.56 & 0.45 & 0.00 \\ 0.30 & 0.16 & 0.54 \\ 0.18 & 0.52 & 0.30 \end{bmatrix}$$

Substituting for the probability transitions in (3.6)-(3.7), we get $q_N = 0.355$, $q_F = 0.365$, and $q_P = 0.280$. That is, 36.5% of the population are free users, and 28.0% are premium users.

In this optimal solution, we note that the ratio of acquisition to retention effort is nearly 4 (3.95 to be precise), even though the cost and effectiveness parameters are fully symmetric. This challenges the conventional wisdom that retention should dominate acquisition in mature markets and suggests that, for freemium platforms, the underlying economic structure creates a persistent acquisition bias, even in steady-state. Rather than stemming from parameter asymmetries, this allocation reflects the inherent growth orientation of such platforms, which continue to prioritize expanding the user base over maximizing value from existing customers.

The transition matrix exposes difficulties in retention across all user segments. First, 55% of individuals exposed to acquisition effort remain in the non-engaged state, signaling a waste in acquisition effort. For premium users, which is the platform's most valuable segment, only 30% remain subscribed in the next period, with 52% downgrading to the free tier and 18% churning entirely. This 70% premium attrition rate is particularly alarming given these users have already demonstrated a willingness to pay. The free tier shows better conversion dynamics with 54% upgrading to premium, but this success is overshadowed by the subsequent premium retention failure, creating a "leaky bucket" scenario where the platform must continuously refill both the acquisition funnel and the premium tier.

²Due to rounding, the probabilities in a given row may sum to slightly more than one. In addition, values reported as zero are not necessarily equal to zero but are instead negligibly small.

Table 3.3: Parameter sensitivity analysis.

Parameter	Decision Variables				Customer Fractions			
	A	R	H	C	q_N	q_F	q_P	π
b	$\nearrow$	$\nearrow$	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$	$\nearrow$	$\nearrow$
γ	$\nearrow$	$\nearrow\searrow$	$\searrow\nearrow$	$\nearrow$	$\searrow\nearrow$	$\nearrow\searrow$	$\nearrow\searrow$	$\nearrow$
k_H	$\searrow$	$\nearrow$	$\searrow$	$\searrow$	$\nearrow\searrow$	$\searrow\nearrow$	$\nearrow$	$\searrow$
ω_A	$\searrow$	$\searrow$	$\searrow$	$\searrow$	$\searrow$	—	$\nearrow$	$\searrow$
k_A	$\nearrow\searrow$	$\nearrow$	$\nearrow$	$\nearrow$	$\nearrow\searrow$	$\nearrow$	$\searrow$	$\nearrow$
ω_R	—	$\searrow$	—	$\nearrow$	$\nearrow$	$\nearrow$	$\searrow$	$\searrow$
k_R	—	$\nearrow\searrow$	$\searrow$	$\nearrow$	$\nearrow$	$\searrow$	$\nearrow$	$\nearrow$

Note: An incline up arrow ($\nearrow$) indicates an increase, incline down ($\searrow$) a decrease, and bidirectional ($\nearrow\searrow$ or $\searrow\nearrow$) a non-monotonic relationship. A dash (—) indicates no observed change.

3.4.2 Sensitivity Analysis

We vary in turn the value of a parameter and assess its impact on decision variables, profit, and customer basis. In each case, the other parameters are kept at their benchmark values. Table 3.3 summarizes the results of the sensitivity analysis.

Varying γ

Our results show how premium quality enhancement (γ) reshapes platform economics through three distinct phases. Figure 3.2a shows that acquisition effort (A) increases monotonically from 0.76 to 1.25 as γ rises from 1.0 to 2.0, reflecting the increasing value of each acquired customer due to improved retention. In contrast, retention effort (R) follows an inverted U-shape: it rises from 0.14 to a peak of 0.286 at $\gamma = 1.71$, then declines to 0.13 by $\gamma = 2.0$. This non-monotonic trend arises from a substitution effect between premium quality and marketing effort. Initially, platforms must increase R to mitigate poor retention, but beyond $\gamma = 1.71$, premium-driven retention reduces the need for retention marketing. The profit π increases steadily from 0.35 to 0.78 as γ rises (Figure 3.2b), indicating that quality improvements consistently enhance profitability, despite requiring

greater acquisition effort, by enabling premium pricing and higher user retention.

As Figure 3.2c shows, the platform’s monetization levers evolve strategically with rising γ , though through distinct mechanisms. Advertising intensity (H) remains stable near 1.05 for $\gamma < 1.7$, but rises sharply to 1.30 at $\gamma = 2.0$, implying that higher premium quality permits heavier ad monetization without significant user attrition. Membership fee (C), representing premium subscription pricing, increases with quality level of the premium tier.




Figure 3.2: Impact of varying γ .

A close examination of Figures 3.2d and 3.3 reveals a key inflection point at $\gamma = 1.7$, marking a regime shift in which both advertisement load and subscription price can increase concurrently, supported by the enhanced value buffer provided by improved premium quality. As γ increases, the steady-state population distribution shifts significantly. The non-engaged population (q_N) declines from 45% to 18%, the free tier (q_F) remains stable around 35 – 40%, and the premium tier (q_P) more than doubles from 20% to 44%. These changes reflect improvements in user engagement and monetization, driven by evolving transition probabilities in the underlying Markov matrix.

The non-to-free switching rate improves modestly from 42% to 49%, suggesting that early engagement depends more on onboarding design than premium features. In contrast, the free-to-premium conversion rate rises substantially from 48% to 66%, demonstrating that premium quality directly drives monetization intent. Premium retention improves from a low 18% to 44%, while the premium-to-free downgrade rate remains high (52–56%), underscoring the ongoing challenge of sustaining premium subscriptions.

At $\gamma = 1.74$, direct premium-to-non churn drops to zero, indicating peak retention capability (Figure 3.2a). However, this also marks a turning point: the premium-to-free downgrade rate begins rising. Recognizing diminishing returns from retention effort, the platform pivots, cutting retention effort, increasing acquisition effort, and raising ad loads for free users (Figure 3.2c). As a result, q_N begins to rise again, while q_F and q_P decline (Figure 3.2d).

Importantly, we highlight the non-monotonic relationship between premium feature quality (γ) and the platform’s strategic decisions, offering new insights beyond existing theories. Prior studies, such as Koch and Benlian (2017) and Wagner et al. (2014), argue that increasing the gap between free and premium features should consistently increase the share of premium users. While our findings initially support this view, we observe that the relationship holds only up to a certain threshold.



Figure 3.3: Effect of γ on ρ_{ij} .

As γ increases moderately, the platform raises both acquisition effort and membership fees, capitalizing on improved free-to-premium conversions. However, when γ becomes sufficiently large, a strategic shift emerges: premium users nearly stop churning, prompting the platform to raise subscription prices more aggressively, reduce retention effort, and invest more in acquisition and advertising. Simultaneously, some premium users begin downgrading to the free tier instead of leaving the platform altogether.

This behavior leads to a new paradigm, in which users remain active but increasingly resist paying for premium access. These findings challenge the assumption that “more differentiation is always better.” Instead, they suggest the existence of an optimal range of premium quality that maximizes profit without triggering downgrade dynamics. Thus, our results extend current freemium theories by showing that excessively enhancing premium features can backfire, not by driving users away, but by incentivizing them to step down rather than exit.

Varying k_H

As ad sensitivity (k_H) increases from 0.1 to 1.0, the platform undergoes a fundamental strategic transformation, evident in its spending patterns and profitability. Figure 3.4a shows that acquisition effort (A) plummets from 2.0 to 0.5, demonstrating that when users are highly sensitive to ads, aggressive acquisition becomes economically inviable due to rapid churn. Retention effort (R) exhibits a threshold effect, remaining at zero for $k_H < 0.3$ when users tolerate ads, then rising sharply to 0.3 as sensitivity increases, indicating that retention efforts only become necessary when ad intolerance begins to threaten the user base. Profit (π) declines monotonically from 2.5 to 0.3, with the steepest drop occurring between $k_H = 0.1$ and $k_H = 0.4$, indicating a critical trade-off: while ad-tolerant users allow for extreme short-term profitability, such exploitation undermines long-term sustainability (Figure 3.4b).




Figure 3.4: Impact of varying k_H .

Figure 3.4c shows how the platform's monetization strategy adjusts in response to rising ad sensitivity. Advertising intensity (H) shows the most extreme reaction, falling from an exploitative 7.0 (roughly $10\times$ the base service value) at $k_H = 0.1$ to just 0.5 at high sensitivity levels. This follows an inverse power relationship, reflecting the platform's diminishing capacity to monetize through ads as users become less tolerant. Subscription cost (C) declines more gradually from 0.9 to 0.6, suggesting that subscription prices are less elastic to ad sensitivity than ad load. Still, the platform must reduce pricing to preserve attractiveness as ads become a weaker conversion lever. This shift marks a transformation from “push” monetization, coercing upgrades through overwhelming ad loads, to “pull” monetization, where users upgrade based on perceived value as k_H rises.

Figures 3.4d and 3.5 exhibit how customers fractions and transition matrices evolve in response to ad sensitivity. The non-engaged population (q_N) follows a U-shaped curve: starting high at 42% when extreme ad loads drive away even tolerant users, reaching a low of 32% at moderate sensitivity ($k_H \approx 0.6$), then rising again as high ad sensitivity constrains monetization levers. The free tier (q_F) exhibits an inverse U-shape, peaking at

40% around $k_H = 0.5$, while the premium population (q_P) increases monotonically from 20% to 31%. This counterintuitive result indicates that greater ad sensitivity can actually expand the premium base, as the platform must shift from pressure to value in order to convert and retain users.



Figure 3.5: Effect of k_H on ρ_{ij} .

The transition matrices show clear shifts in user behavior as ad-sensitivity (k_H) increases. At low k_H , the platform achieves strong acquisition and conversion rates, i.e., 63% of non-users join the free tier, and 51% of them upgrade to premium. However, this setting also results in extremely high premium churn, with 98% of subscribers leaving the paid tier. As k_H rises, fewer users are willing to join the free tier due to discomfort with ads, leading to a marked decline in acquisition. In contrast, retention within the premium tier improves substantially, with churn dropping to 64%, as ad-sensitive users choose to maintain their ad-free access. The conversion rate from free to premium improves only slightly and then levels off, since the platform lowers ad intensity to retain free users; this softer ad environment reduces churn from the platform but also lessens the urgency to upgrade. Overall, the matrices show how increasing ad aversion reduces platform entry, strengthens premium retention, stabilizes conversion, and limits downgrade behavior from premium to free. Around $k_H = 0.3$, the platform begins to adjust its strategy by introducing retention efforts and reducing ad pressure to maintain user engagement across

tiers.

Varying k_A

Figure 3.6a shows that when acquisition effectiveness (k_A) is low, acquisition effort (A) initially increases. However, beyond its highest point (say $k_A \approx 0.8$ – 1.0), acquisition effort begins to go down. This decline reflects rising acquisition effectiveness: as acquisition becomes more effective, less spending is required to achieve user growth. Nevertheless, until this effectiveness is realized, platforms must invest heavily in acquisition. Retention effort (R) increases modestly from 0.20 to 0.27, suggesting that improved acquisition effectiveness allows for slightly greater retention effort, but it does not fundamentally alter the retention strategy. Profit exhibits a nuanced pattern, initially rising as k_A improves from very low levels, peaking in the optimal effectiveness zone, and stabilizing as the platform reaches operational efficiency. This confirms that effectiveness, not just spending volume, drives long-term profitability.




Figure 3.6: Impact of varying k_A .

Figure 3.6c illustrates how the platform's pricing strategy adapts to acquisition ef-

fectiveness. Advertising intensity (H) undergoes a sharp transition, from 0.75 at low k_A to 1.4 at high k_A , and then plateaus once k_A exceeds about 0.9. This pattern indicates that effective acquisition supports aggressive ad monetization because churned users can be quickly replaced. When activation is nearly guaranteed, platforms can afford to push users harder. In contrast, premium pricing (C) continues to increase, from 0.55 to 0.97, as k_A rises. This reflects both higher user quality (better-activated users with higher willingness to pay) and increased pricing power (platforms can be selective in whom to retain). The simultaneous increase in both H and C shows how effective acquisition transforms the platform’s approach from cautious and retention-oriented to confident and growth-focused.



Figure 3.7: Effect of parameter k_A on ρ_{ij}

Figure 3.6c shows that around $k_A \approx 0.8$ – 1.0 , H reaches a ceiling, while C continues rising. This is because at this threshold, we have that $\rho_{FF} \approx 0$ (Figure 3.7): users who are free-tier tend not to remain stuck there. Additionally, ρ_{NF} increases sharply, signaling near-guaranteed activation. Importantly, C does not appear in early-funnel transitions like $N \rightarrow F$, so higher pricing does not deter user acquisition. This insight is critical: users do not encounter the premium price until they are already engaged. When k_A is high, activation is nearly guaranteed (with $\rho_{NF} > 90\%$), user quality improves, which means that engaged users convert more reliably, and the platform gains pricing flexibility,

as it can afford to lose price-sensitive users without compromising overall performance. Thus, H plateaus because it negatively affects early transitions ($N \rightarrow F$, $F \rightarrow F$), creating a natural upper bound where further increases would reduce engagement. In contrast, C continues rising, as it only influences late-funnel monetization decisions where the platform can afford to be selective. This dynamic highlights a core insight: platforms cannot monetize users who churn early (the H constraint), but they can maximize revenue from those who reach the premium conversion point (the C opportunity).

Figure 3.6d shows that as k_A increases, the non-engaged user segment (q_N) decreases, while the free user fraction (q_F) grows. Surprisingly, the premium user fraction (q_P) declines substantially. This pattern emerges from saturation effects in user quality: while activation improves, fewer users convert to premium, as the platform's strategy focuses more on activation and less on aggressive premium pushing.

Figure 3.7 illustrates the mechanism behind these changes. At low k_A , a large fraction of users remain non-engaged despite acquisition efforts. As k_A increases toward 1.0, we observe steady improvement in activation rates, with the most dramatic changes occurring in the 0.6–1.0 range.

Varying k_R

Increasing the effectiveness of retention (k_R) leads to an increase, at a decreasing rate, in retention effort (R). This concave relationship between k_R and R is intuitive as one expects to see marginal decreasing returns in any marketing activity (Figure 3.8a). Interestingly, even though retention becomes more effective as k_R increases, the number of premium users (q_P) stays almost the same, hovering around 28 – 29%, Figure 3.8d. The reason lies in how the platform adjusts its strategy. As retention tools become more effective, the platform increases its retention effort, but at a decreasing rate. In other words, the platform relies on the improved efficiency of the tools rather than ramping up spending. While users stay longer, the platform doesn't need to invest substantially more to maintain this group. This leads to a counterintuitive but logical outcome: better retention tools do not necessarily result in more premium users, they simply reduce the cost of keeping the

same number.

The platform also adjusts its monetization strategy. It slightly lowers the ad intensity (H) drops from 1.05 to 1.02), reducing pressure on free users. At the same time, it raises the premium price (C increases from 0.62 to 0.74), Figure 3.8c. This is a smart move. Indeed, when retention is strong, users are less likely to churn, so the platform can afford to charge more. This price increase does not hurt premium user numbers because those who stay are more committed. It also explains why conversion from free to premium users ($F \rightarrow P$) drops slightly; the higher price discourages some new upgrades, but the platform makes up for it with better margins on existing users.

The transition matrices tell a clear story (Figure 3.9). Premium user retention ($P \rightarrow P$) improves from 26% to 41% as k_R increases, and fewer users downgrade ($P \rightarrow F$ drops). However, fewer free users convert to premium because the rising prices make it harder to justify the upgrade. This reveals a key insight: retention effectiveness increases pricing power, not user growth. When tools work well, platforms raise prices, reduce pressure, and prioritize revenue per user over total user numbers.




Figure 3.8: Impact of varying k_R .




Figure 3.9: Effect of parameter k_R on ρ_{ij}

In the real-world, this explains why platforms with strong retention tools behave differently. Early-stage products with poor tools use low prices and high ad pressure to attract and retain users, e.g., Youtube in its early days. In contrast, mature platforms like Salesforce use strong retention mechanisms (e.g., integrations, switching costs) to raise prices and deepen profit margins without chasing user growth.

Even as retention tools become stronger, acquisition effort (A) and overall profit (π) remain steady. This tells us that retention effectiveness is a *local optimization*, it does not change the big picture of platform economics. Instead, it gives the platform flexibility to adjust internal levers like pricing and ad load.

This shows that pricing power grows almost linearly with k_R . In other words, the better the retention tools, the more the platform can charge. The user base stays the same, but its *value* increases.

The most important insight is this: strong retention efforts do not make platforms bigger, they make them more profitable. This is why some of the most successful companies have small user bases but high prices and strong loyalty. They are not trying to win on size, they are winning on value.

But there is a warning in the results too. As prices rise and conversion drops, platforms risk slowing down their growth. If they rely too much on retention tools, they may become

overconfident and create barriers for new users. The best strategy balances the power of retention with the need for ongoing user acquisition.

Varying ω_A

Figure 3.10a shows that when the cost coefficient of acquiring new users (ω_A) increases from 0.1 to 0.99, acquisition effort (A) sharply decreases from 2.197 to 0.472, a 79% drop. Higher costs make acquiring new customers expensive, pushing platforms away from aggressive growth strategies toward more balanced approaches. Retention effort (R) drops 15% from 0.237 to 0.201, which is very modest given that ω_A is ten times higher. This signals that retention decision is (almost) independent of acquisition costs. Platform profit (π) falls by 27% from 0.639 to 0.466, demonstrating platforms' ability to adapt by shifting from high-growth to efficiency-based strategies. Figure 3.10b confirms profitability can be maintained across a broad range of acquisition costs.

Figure 3.10c illustrates how revenue sources adapt to changes in acquisition costs. Advertising intensity (H) decreases by 35%, from 1.398 to 0.911, as acquisition costs rise. This indicates that platforms rely heavily on advertisements when acquiring users is inexpensive, but significantly reduce ads to improve the user experience as customer acquisition becomes costly. Similarly, the premium fee (C) decreases by 29% from 0.845 to 0.601, showing platforms must offer more attractive premium pricing when acquiring users becomes expensive. This coordinated reduction in advertising and premium pricing reflects a strategic shift: costly user acquisition demands better user experiences to retain and convert customers effectively.

Figure 3.10d exhibits how user segments change with acquisition costs. Premium users (q_P) steadily increase from 24.3% to 29.3% as acquisition costs rise, as the platform becomes more focused on monetizing a smaller, higher-quality user base. The share of non-engaged users (q_N) initially slightly increases from 37.9% to 38.5%, then declines to 34.3%. This pattern indicates that moderate acquisition costs improve average user quality, while very high costs limit total new user acquisition. Free-tier users (q_F) remain relatively unchanged, fluctuating modestly around 36.4% to 37.9%, reflecting a stable

segment.

Transition matrices further clarify these shifts, Figure 3.11. At low acquisition cost ($\omega_A = 0.10$), platforms achieve high activation (67% from non-user to free) and strong conversion from free to premium (53%), but have poor premium retention (only 17%), creating a high-churn environment. In contrast, at high acquisition cost ($\omega_A = 0.99$), initial activation falls to 38%, but premium retention significantly improves to 34%, with stable conversion rates around 53%. Retention rates improve because higher costs naturally select more committed users, reducing the need for aggressive monetization through ads or high fees.




Figure 3.10: Impact of Varying ω_A .

Overall, these results highlight a fundamental strategic shift: when acquiring new users becomes expensive, platforms transition from volume-driven growth to quality-focused retention. High acquisition costs select for users with stronger intrinsic motivation, enhancing long-term retention and platform sustainability. Consequently, premium conversion remains stable even as retention dramatically improves, because users opting

for premium plans are primarily motivated by genuine service value rather than aggressive promotions or discounts.



Figure 3.11: Effect of parameter ω_A on ρ_{ij}

Varying ω_R

As expected, retention effort (R) decreases when ω_R increases (Figure 3.12a). From a monetization perspective, Figure 3.12c indicates that the platform keeps the ad intensity (H) constant, while gradually reducing the premium price (C) as retention costs rise. This strategy prioritizes ad revenue stability while adjusting premium pricing to encourage ongoing subscriptions. As retention effort only targets premium users, acquisition (new users entering), activation (free users engaging), and ad tolerance remain unchanged.




Figure 3.12: Impact of varying ω_R .

The interpretation of our findings are as follows: when retention cost goes up, the platform adapts by enhancing premium quality, adjusting prices, and relying on natural user transitions rather than forcing retention. Profits remain remarkably stable, dropping only slightly from 0.51 to 0.50 as ω_R rises. This highlights that retention has a minor effect on profit; most of the value stems from acquisition and conversion. In practice, this is why firms like B2B platforms prioritize content and premium development over aggressive churn management (Figure 3.12a).



Figure 3.13: Effect of parameter ω_R on ρ_{ij}

The most striking result is that the overall user distribution (almost) does not change when we vary ω_R , even when it is increased tenfold. Non-users (q_N), free users (q_F), and premium users (q_P) stay stable at 35.5%, 36.5%, and 28%, respectively (Figure 3.12d). One explanation is that the platform maintains equilibrium by adjusting retention effort inversely with cost. When retention becomes more expensive, the platform simply spends less, keeping the system balanced.

3.4.3 Strategic Insights

Our sensitivity analysis highlights how core platform parameters shape optimal strategy and market outcomes.

Quality Factor of the Premium Service (γ):

Our analysis identifies a critical strategic inflection point at the threshold $\gamma = 1.7$, where the platform's optimal strategy fundamentally transforms due to a precise economic mechanism. At this threshold, premium-to-non-user transitions vanish ($\rho_{PN} = 0$), meaning dissatisfied premium users now downgrade to the free tier instead of leaving the ecosystem entirely, creating strong *ecosystem stickiness* that markedly enhances the platform's pricing power.

When quality differentiation is lower than the threshold ($\gamma < 1.7$), the gap between free and premium tiers is too narrow to create meaningful user retention within the platform ecosystem. The platform must focus on *value-building strategies*, devoting significant resources to retention because modest increases in service quality generate high returns by reducing complete churn. In this regime, premium users frequently exit the ecosystem entirely when dissatisfied, making user base preservation the strategic imperative before attempting revenue optimization.

When quality differentiation exceeds the threshold ($\gamma > 1.7$), the premium tier becomes sufficiently attractive that even dissatisfied users prefer to downgrade rather than abandon the platform completely. This ecosystem stickiness enables the platform to prioritize *value-extraction strategies*, aggressively increasing subscription pricing while reducing retention effort, since users are effectively “retained” within the platform ecosystem even when they downgrade.

This phase transition reflects the economic relationship between premium value (γk_{bb}) and retention probability (ρ_{PP}): when γ reaches the threshold ($\gamma = 1.7$), premium value suffices to prevent total user exit, regardless of pricing pressure. Strategically, this implies that moderate differentiation around the threshold maximizes the premium user base, while higher differentiation enables greater revenue extraction per premium subscriber.

Notably, our model’s identification of a non-monotonic relationship between premium differentiation and platform performance challenges the conventional wisdom that greater differentiation always improves outcomes (J. Li et al., 2024; Wagner et al., 2014). Empirically, several industry trends reflect the three-phase evolution predicted by our framework. For example, content platforms such as Netflix appear to mirror our predicted early stage (when quality differentiation is lower than the threshold) transition by investing heavily in original content and maintaining modest premium pricing during market expansion. Similarly, software platforms like Adobe Creative Cloud illustrate aspects of when quality differentiation is higher than the threshold dynamics by leveraging established user bases for aggressive premium pricing; nonetheless, such pricing power may reflect market dominance, switching costs, or ecosystem integration rather than quality differentiation alone.

These patterns suggest our quality differentiation results are subject to significant boundary conditions: the $\gamma = 1.7$ inflection assumes competitive markets and may shift under strong network effects, regulatory intervention, or asymmetric competition. Moreover, user heterogeneity by segment, geography, or demographic, and dynamic factors such as learning, adaptation, and technological change, may all justify departures from the aggregate predictions of our static optimization framework.

Advertisement Sensitivity (k_H):

The relationship between advertisement sensitivity (k_H) and platform strategy shows a fundamental shift in monetization philosophy, driven by the exponential nature of user response. The mathematical foundation lies in the acquisition function, $\rho_{NF} = 1 - e^{-(k_A A + k_b b - k_H H)}$, which illustrates that advertising harm scales exponentially with user sensitivity: doubling k_H causes the same ad intensity (H) to inflict dramatically more acquisition damage. This dynamic generates three distinct regimes. In the *exploitation regime* ($k_H < 0.3$), users are highly tolerant of ads, so the platform maximizes short-term revenue by aggressively increasing ad pressure (e.g., $H = 7.0$), with little or no investment in retention, since users are willing to tolerate poor experiences. As k_H enters the *transition zone* ($0.3 \leq k_H \leq 0.6$), user sensitivity becomes moderate, prompting a more balanced approach: the platform reduces ad load and increases retention effort to preserve engagement as tolerance declines. In the *value creation regime* ($k_H > 0.6$), users are highly averse to ads, and the platform must shift to “pull” monetization, focusing on premium quality and differentiation (e.g., $H = 0.5$) to retain and convert users. The exponential sensitivity explains why the optimal ad intensity collapses from $H = 7.0$ to $H = 0.5$ as k_H rises; platforms cannot pursue both aggressive growth and monetization in markets with ad-averse users. These insights have practical implications: conducting market research to assess ad tolerance is crucial for effective strategy, platform maturity may naturally increase user sophistication (and thus k_H), and competitive dynamics can force a shift from exploitation toward value-based models.

Acquisition Sensitivity (k_A):

Acquisition effectiveness k_A creates a dual constraint structure. While increasing k_A

enhances user activation, it also imposes a natural ceiling on sustainable advertising: once activation rates are high (90%+), any further increase in ad intensity (H) starts to damage acquisition more than it benefits monetization, creating a hard upper bound on advertising leverage. In contrast, premium pricing (C) is far less constrained, as high k_A effectively filters out low-value users, allowing the platform to raise prices aggressively and rely on continued acquisition to replenish the customer base. This dual constraint mechanism explains why successful platforms often emphasize acquisition quality over advertising volume, and why premium pricing can rise with market maturity.

Retention Sensitivity (k_R):

Our analysis of retention parameter reveals a counterintuitive but robust pattern: improvements in retention effectiveness primarily empower higher pricing rather than expanding the premium user base. Specifically, when the marginal cost of retention (ω_R) increases by an order of magnitude, both the user distribution (with $q_N \approx 35.5\%$, $q_F \approx 36.5\%$, and $q_P \approx 28\%$) and platform profit (remaining stable around 0.50) are largely unaffected, while optimal retention effort (R) declines as expected. Similarly, as the effectiveness of retention effort (k_R) increases, retention effort (R) exhibits an inverted U-shape, peaking at intermediate values of k_R and declining thereafter, yet the share of premium users remains essentially unchanged. Notably, improvements in retention effectiveness allow for higher subscription pricing (C rising from 0.62 to 0.76) without eroding the premium customer base. This “retention paradox” challenges the standard CRM perspective (Reinartz et al., 2005) that improved retention should expand the customer base. Instead, our findings suggest that, in the freemium context, better retention tools translate into pricing power, platforms can increase revenue per user through higher prices without affecting overall user allocation. For managers, this underscores that retention efforts should be viewed primarily as a lever for boosting margins rather than as a tool for user base growth, and that the optimal user mix will remain stable even with enhanced retention programs.

3.4.4 Numerical Validation and Robustness Analysis

To ensure the reliability of our optimization results, we implemented a comprehensive validation framework. This included convergence diagnostics, robustness checks across multiple starting points, sensitivity analysis, Monte Carlo simulation, and constraint verification. Together, these steps demonstrate the stability of our numerical solution and its consistency with the underlying economic assumptions.

The core optimization was conducted using the `trust-constr` interior point method, which is well-suited for constrained nonlinear problems. The algorithm successfully converged, achieving a gradient norm of 6.86×10^{-9} , well below the common tolerance threshold of 10^{-6} , while satisfying all inequality and equality constraints exactly. The optimizer terminated based on the `gtol` condition, indicating that first-order optimality conditions were met. Furthermore, convergence was efficient, requiring only 15 function evaluations. We also verified the structure of the optimized transition matrix, confirming that each row sums to one and all probability values lie strictly within the $[0, 1]$ interval.

To ensure the solution was not a result of favorable initialization, we performed multi-start validation by solving the optimization problem from five different random starting points. Remarkably, all trials converged to the same optimal solution, with values $A = 0.8276$, $R = 0.2063$, $H = 1.0401$, and $C = 0.6570$. The profit outcome was also consistent across all runs at $\pi = 0.499834$, and the maximum deviation across all decision variables was less than 6.4×10^{-6} . This consistency strongly suggests that the solution is not only stable but represents a global optimum. This multi-start validation exceeds standards commonly reported in computational optimization literature, where single-algorithm results are often accepted without convergence verification from multiple initial conditions.

We further evaluated the model's robustness through sensitivity analysis. By perturbing key input parameters by $\pm 5\%$ and $\pm 10\%$, we assessed the impact of parameter uncertainty on profit outcomes. The model exhibited high stability with respect to several parameters, including premium quality (γ), acquisition cost (ω_A), and retention cost

(ω_R), all of which led to less than 5% variation in profit. Some changes in advertisement sensitivity (k_H) even led to modest improvements in profitability. However, the model showed higher sensitivity to the baseline service quality (b), with profit impacts ranging from -5.1% to $+23.4\%$. These findings align with intuitive expectations, as core service quality is fundamental to user satisfaction and engagement.

To test the model’s behavior under broader uncertainty, we conducted a Monte Carlo simulation using ten randomized parameter sets and starting points. Seven of these scenarios resulted in successful convergence, with profit outcomes ranging from 0.260 to 0.945. The average profit across these scenarios was 0.622 ± 0.238 . This result reinforces the model’s applicability across a wide range of market conditions, while also identifying edge cases where economic feasibility constraints are violated and convergence is not achieved.

Constraint satisfaction was also carefully verified at the optimal solution. All transition probabilities were within bounds, and eigenvalue analysis confirmed that the largest eigenvalue equals unity with all others having modulus less than one, ensuring steady-state existence and uniqueness. Furthermore, economic feasibility conditions, such as the requirement that platform quality ($k_b b$) outweighs advertising harm ($k_H H$) were satisfied. Pricing constraints for the membership fee also held under all tested conditions.

These validation results provide high confidence that our numerical solution represents a genuine economic optimum rather than a computational artifact, meeting high validation standards. While alternative optimization algorithms (SLSQP, COBYLA) encountered numerical difficulties due to the high density of nonlinear constraints (4 in total) and complex exponential transition functions, the trust-constr interior point method’s superior handling of such problems is well-established in the optimization literature (Nocedal & Wright, 2006; Vanderbei, 1999). Specifically, LOQO and related methods demonstrate the efficacy of interior-point approaches for structured, constrained quadratic problems, reinforcing the trust in our solver’s performance. This algorithmic choice was validated by the consistent convergence across all robustness tests.

To summarize, the validation framework provides strong evidence of the robustness

and reliability of our numerical solution. The model met or exceeded standard thresholds in all key areas. Gradient convergence was well within the tolerance limit, the solution was consistent across all random starts, sensitivity analysis showed high parameter stability, and Monte Carlo testing demonstrated solid performance across diverse scenarios. All constraints were fully satisfied, confirming economic coherence.

3.5 Conclusion

This study presents a unified framework for optimizing acquisition, retention, advertising, and pricing decisions in freemium platforms. By modeling user transitions as a Markov process and solving for long-run profit maximization, we uncover how strategic levers interact in non-obvious ways. Our results highlight that platforms often favor acquisition over retention due to structural churn, that stronger retention tools enhance pricing power more than user growth, and that excessive premium differentiation can unintentionally trigger downgrades rather than exits.

From a managerial perspective, our findings emphasize four key principles: (1) target moderate premium differentiation rather than extremes, (2) align monetization strategies with user engagement levels, (3) focus on acquisition quality over volume, ensuring that new users are likely to engage and convert, and (4) leverage retention to support pricing power rather than simply expanding the user base. These insights reveal the inherent trade-offs in freemium strategy, where maximizing user count and maximizing revenue per user often require opposing approaches.

Our model's limitations suggest some potential directions for future research. First, the current static optimization framework could be extended to dynamic settings that capture learning effects over time, evolving technologies, and competitive market responses. Incorporating user heterogeneity and analyzing strategic interactions across multiple platforms would also offer deeper insights. Moreover, empirical validation through field experiments and natural experiments would enhance the real-world applicability of our theoretical predictions.

As the freemium model continues to expand across digital industries, understanding the underlying optimization dynamics becomes increasingly crucial. This study provides an integrated framework that unifies acquisition, retention, advertising, and pricing decisions. In doing so, it lays a foundation for future research while underscoring the complex strategic challenges that distinguish freemium platforms from traditional business models.

Appendix A

A.1 Full Profit Model Derivation

Substituting q_F and q_P from equations (3.8)–(3.9) into the objective function (3.5), we obtain:

$$\pi = f(A, R, H, C, q_F(A, R, H, C), q_P(A, R, H, C)), \quad (3.10)$$

The complete expression is given below:

$$\pi = \frac{\left(\begin{aligned} &a0.5A^2\omega_A e^{3Ak_A+4Rk_R+2k_HH+6b\gamma+5b} + 0.5A^2\omega_A e^{3Ak_A+2Rk_R+2k_HH+5b\gamma+5b+2C} \\ &- 0.5A^2\omega_A e^{3Ak_A+4Rk_R+k_HH+5b\gamma+5b+C} + H e^{2Ak_A+2Rk_R+2k_HH+5b\gamma+5b+2C} \\ &- H e^{3Ak_A+2Rk_R+k_HH+5b\gamma+6b+2C} + (-0.5R^2\omega_R + C) e^{2Ak_A+3Rk_R+2k_HH+6b\gamma+5b+C} \\ &+ (0.5R^2\omega_R - C) e^{2Ak_A+3Rk_R+k_HH+5b\gamma+5b+2C} \\ &+ (-0.5A^2\omega_A + 0.5R^2\omega_R - C) e^{3Ak_A+3Rk_R+k_HH+6b\gamma+6b+C} \\ &+ (0.5A^2\omega_A - 0.5R^2\omega_R + C) e^{3Ak_A+3Rk_R+5b\gamma+6b+2C} \end{aligned} \right) e^{-2Ak_A-2Rk_R-k_HH-5b\gamma-5b}}{e^{Rk_R+2C} - e^{k_HH+2C} + e^{Ak_A+b+2C} - e^{Ak_A+2Rk_R+C} + e^{Ak_A+k_HH+2C} + e^{Ak_A+2Rk_R+k_HH+b\gamma} - e^{Rk_R+k_HH+b\gamma+C}}$$

References

Afeche, P., Araghi, M., & Baron, O. (2017). Customer acquisition, retention, and service access quality: optimal advertising, capacity level, and capacity allocation. *Manufacturing & Service Operations Management*, 19(4), 674–691.

- Anderson, S. P., & Gabszewicz, J. J. (2006). The media and advertising: a tale of two-sided markets. *Handbook of the Economics of Art and Culture*, 1, 567–614.
- Ben Rhouma, T., & Zaccour, G. (2018). Optimal marketing strategies for the acquisition and retention of service subscribers. *Management Science*, 64(6), 2609–2627.
- Berger, P. D., & Bechwati, N. N. (2001). The allocation of promotion budget to maximize customer equity. *Omega*, 29(1), 49–61.
- Blattberg, R. C., & Deighton, J. (1996). Manage marketing by the customer equity test. *Harvard business review*, 74(4), 136–144.
- Business of Apps. (2024). Youtube statistics (2024) [Accessed: 2025-07-22]. <https://www.businessofapps.com/data/youtube-statistics/>
- Carroni, E., & Paolini, D. (2020). Business models for streaming platforms: content acquisition, advertising and users. *Information Economics and Policy*, 52, 100877.
- Chang, S., Zhang, Z., Wang, X., & Dong, Y. (2020). Optimal acquisition and retention strategies in a duopoly model of competition. *European Journal of Operational Research*, 282(2), 677–695.
- Gregory, J. (2016). Dynamic customer acquisition and retention management. *Production and Operations Management*, 25.
- Kalyanam, K., Thomadsen, R., & Zhao, N. (2025). The impact of advertising content on customer acquisition and retention for subscriptions of physical goods: insights from a field experiment. *Available at SSRN 5338643*.
- Koch, O. F., & Benlian, A. (2017). The effect of free sampling strategies on freemium conversion rates. *Electronic Markets*, 27(1), 67–76.
- Kumar, V. (2014). Making freemium work. *Harvard Business Review*, 92(5), 27–29.
- Kumar, V. (2019). Designing freemium: understanding and optimizing the customer conversion journey [Available at <https://vineetkumars.github.io/Papers/DesigningFreemiumS2019.pdf>]. *Harvard Business School Working Paper*.
- Lee, C., Kumar, V., & Gupta, S. (2013). Designing freemium: a model of consumer usage, upgrade, and referral dynamics. *MimeoLeung, Michael (2013), Two-step estimation of network formation models with incomplete information*.

- Li, J., Wu, J., Zhang, J., & Ji, X. (2024). Freemium design: optimal tier differentiation models for content platforms. *Transportation Research Part E: Logistics and Transportation Review*, 188, 103612.
- Li, X., Balasubramanian, H., Chen, Y., & Pang, C. (2024). Managing conflicting revenue streams from advertisers and subscribers for online platforms. *European Journal of Operational Research*, 314(1), 241–254.
- Lin, S. (2020). Two-sided price discrimination by media platforms. *Marketing Science*, 39(2), 317–338.
- Mai, Y., & Hu, B. (2023). From growth to monetization: managing freemium services. Available at SSRN 4349210.
- Min, S., Zhang, X., Kim, N., & Srivastava, R. K. (2016). Customer acquisition and retention spending: an analytical model and empirical investigation in wireless telecommunications markets. *Journal of marketing research*, 53(5), 728–744.
- MoldStud. (2025). Freemium vs premium: analyzing revenue potential for cross-platform apps [Accessed: 2025-07-22]. <https://moldstud.com/articles/p-freemium-vs-premium-analyzing-revenue-potential-for-cross-platform-apps>
- Niemand, T., Mai, R., & Kraus, S. (2019). The zero-price effect in freemium business models: the moderating effects of free mentality and price–quality inference. *Psychology & Marketing*, 36(8), 773–790.
- Nocedal, J., & Wright, S. J. (2006). *Numerical optimization* (2nd). Springer.
- Pauwels, K., & Weiss, A. (2008). Moving from free to fee: how online firms market to change their business model successfully. *Journal of Marketing*, 72(3), 14–31.
- Pfeifer, P. E. (2005). The optimal ratio of acquisition and retention costs. *Journal of targeting, measurement and analysis for marketing*, 13, 179–188.
- Pfeifer, P. E., & Carraway, R. L. (2000). Modeling customer relationships as markov chains. *Journal of Interactive Marketing*, 14(2), 43–55.
- Prasad, A., Mahajan, V., & Bronnenberg, B. (2003). Advertising versus pay-per-view in electronic media. *International Journal of Research in Marketing*, 20(1), 13–30.

- Reinartz, W., Thomas, J. S., & Kumar, V. (2005). Balancing acquisition and retention resources to maximize customer profitability. *Journal of marketing*, 69(1), 63–79.
- Ross, N. (2018). Customer retention in freemium applications. *Journal of Marketing Analytics*, 6(4), 127–137.
- Runge, J., Levav, J., & Nair, H. S. (2022). Price promotions and “freemium” app monetization. *Quantitative marketing and economics*, 20(2), 101–139.
- Sato, S. (2019). Freemium as optimal menu pricing. *International Journal of Industrial Organization*, 63, 480–510.
- Schweidel, D. A., Fader, P. S., & Bradlow, E. T. (2008). A bivariate timing model of customer acquisition and retention. *Marketing Science*, 27(5), 829–843.
- Sciglimpaglia, D., & Raafat, F. (2022). Freemium marketing: use of demand-side research in market segmentation strategy. *Journal of Strategic Marketing*, 30(7), 667–690.
- Spotify. (2025, February). Spotify reports fourth quarter 2024 earnings [Accessed: 2025-07-22]. <https://newsroom.spotify.com/2025-02-04/spotify-reports-fourth-quarter-2024-earnings/>
- Tang, Y., Chen, R. R., & Guan, X. (2021). Daily-deal market with consumer retention: price discrimination or quality differentiation. *Omega*, 102, 102330.
- Vanderbei, R. J. (1999). Loqo: an interior point code for quadratic programming. *Optimization Methods and Software*, 11(1-4), 451–484.
- Venkatesan, R., & Kumar, V. (2004). A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing*, 68(4), 106–125.
- Wagner, T. M., Benlian, A., & Hess, T. (2020). Designing for freemium: understanding the effects of platforms and premium conversion strategies in business models. *Information Systems Journal*, 30(2), 229–260. <https://doi.org/10.1111/isj.12262>
- Wagner, T. M., Benlian, A., & Hess, T. (2014). Converting freemium customers from free to premium—the role of the perceived premium fit in the case of music as a service. *Electronic Markets*, 24, 259–268.
- Yoganarasimhan, H., et al. (2023). Personalization and user engagement: evidence from freemium platforms. *Marketing Science*.

Conclusion

This thesis investigates how digital platforms should design pricing structures when customers differ in their usage patterns and willingness to pay. Programs like Uber One and DoorDash DashPass combine membership fees with discounted transactions, while Spotify and YouTube offer free ad-supported tiers alongside premium subscriptions (Chen et al., 2022; Wagner et al., 2014). Despite their prevalence, these pricing strategies involve complex interdependencies that existing research has examined only in pieces. This thesis develops a unified framework for understanding these interactions, showing that optimal platform pricing is governed by critical thresholds where small parameter changes trigger fundamental shifts in strategy and market structure.

The research progresses through three essays that systematically expand the platform's strategic scope. The first essay examines on-demand platforms that optimize membership fees and per-transaction prices while discount rates and service quality remain fixed. The second essay expands this to full optimization where platforms jointly choose membership fees, transaction prices, discount rates, and service quality, providing a complete map of on-demand platform pricing. The third essay shifts to freemium platforms where the focus becomes advertising intensity, subscription pricing, and customer relationship management through acquisition and retention expenditures. This progression shows that quality differentiation dominates price-based strategies across platform types, that market structures deliberately exclude substantial customer segments, and that critical thresholds characterize optimal platform design.

The first essay, entitled *Optimal Menu Pricing Strategies for On-Demand Service*

Platforms, addresses a practical problem: platforms like Uber One and DoorDash Dash-Pass have widely adopted menu pricing, yet it remains unclear when these programs work and how platforms should set their fees (Kung & Zhong, 2017). The platform faces competing objectives. It must extract value from high-frequency users through membership while retaining lighter users who prefer pay-per-use pricing. The essay develops a profit-maximization model where customers self-select between membership and pay-per-use based on their usage frequency and service valuation. The platform optimizes membership fees and per-transaction prices while discount rates and service quality are treated as given parameters.

The analysis uncovers a striking pattern: discount rates and quality levels affect market structure in fundamentally different ways. For discounts, a critical threshold emerges. Below this level, platforms cannot justify substantial membership fees and rely primarily on per-transaction revenues, creating a binary market where users either pay premium prices per transaction or exit entirely. Once discounts cross this threshold, the economics shift dramatically. Membership fees jump, transaction prices fall, and a true three-tier structure emerges with premium subscribers, regular pay-per-use customers, and non-users. Near the threshold, small changes in discount policy trigger qualitatively different optimal pricing architectures. Quality differentiation works differently. Moderate quality improvements maximize premium adoption, creating a sweet spot where membership appeal is strongest. Yet profitability increases monotonically with quality across the entire range, so platforms face a strategic choice: optimize for market coverage with moderate quality, or maximize profit with high quality that creates exclusionary pricing. Regardless of this choice, approximately half of potential customers remain non-users across parameter combinations. Menu pricing success depends on identifying narrow viability zones defined by specific discount and quality configurations rather than making incremental pricing adjustments. Certain limitations in the first essay could be addressed in future research. The model employs uniform distributions for usage frequency and customer valuations; examining alternative distributions such as heavy-tailed or bimodal patterns would test whether the critical discount threshold remains robust when a small fraction

of very heavy users coexist with many low-activity users. Additionally, empirical validation using transaction-level data from platforms like Uber One or DoorDash would test whether observed pricing regimes exhibit the predicted threshold dynamics and market structure transitions.

The second essay, entitled *Joint Optimization of Pricing and Service Quality in On-Demand Service Platforms*, provides a complete map for on-demand platform pricing by endogenizing discount and quality decisions alongside membership and transaction fees. The analysis proceeds through carefully designed scenarios that progressively relax constraints: first adding discount optimization while holding quality fixed, then adding quality optimization while fixing discounts, and finally permitting simultaneous optimization across all four variables. This stepwise approach isolates how much value each instrument contributes and how market structure responds to expanded strategic control.

The results reshape the interpretation of the first essay's findings. When platforms can freely choose discount rates, they optimally set discounts close to zero. The critical threshold identified in Essay 1 becomes less relevant because platforms simply avoid triggering it. More importantly, comparing intermediate scenarios reveals a pronounced asymmetry: endogenizing discounts yields only modest profit gains over the benchmark, while endogenizing quality generates substantial improvements. Quality differentiation dominates price discrimination because discounts reduce revenue on every subscriber transaction, whereas quality improvements can be monetized through membership fees without sacrificing transaction revenues.

A deeper structural insight emerges under full optimization. When both discounts and quality become decision variables, the three-tier market structure commonly assumed in practice typically collapses to a binary outcome. The regular user segment, those who pay full transaction prices without membership, systematically vanishes. Instead, platforms optimally pursue a stark strategy: combine near-zero discounts with strong quality differentiation, high transaction prices, and moderate membership fees. Roughly half the market subscribes to access enhanced service, while the rest is deliberately excluded. The elevated base price creates a sufficient cost differential to make membership economi-

cally attractive even without transaction savings (Armstrong, 2006). This binary structure contradicts intuitions from the first essay, where three-tier segmentation appeared viable within specific parameter ranges. The resolution is straightforward: when platforms control quality and discounts endogenously, they face different trade-offs than when these parameters are fixed exogenously. Fixed parameters might accidentally place a platform in a three-tier regime, but optimization actively steers away from this configuration toward simpler, more profitable binary structures. The analysis also reveals winner-take-all dynamics driven by quality costs: when quality provision is nearly costless, platforms pursue extreme differentiation with correspondingly high fees, but as quality costs rise, both optimal quality and profits decline rapidly. The analysis abstracts from competitive dynamics by focusing on a monopolistic platform. Extending the framework to examine strategic interactions when multiple platforms compete for users who can switch or multi-home would reveal whether quality differentiation remains dominant under rivalry and how the identified thresholds shift in competitive equilibrium. Additionally, incorporating supply-side considerations such as driver participation and capacity constraints would yield more comprehensive platform models that link pricing strategy with operational realities.

Having established how quality dominance and threshold-driven strategy operate in on-demand contexts, a natural question arises: do these patterns reflect fundamental platform economics, or are they specific to menu pricing structures? This question motivates the third essay's shift to freemium platforms, which face similar challenges of customer heterogeneity but differ substantially in structure. Revenue comes from advertising to free users and subscriptions from premium users, and customer relationships evolve dynamically through acquisition, conversion, and churn rather than through one-time self-selection into pricing tiers.

The third essay, entitled *Freemium Pricing and CRM Expenditures by a Digital Platform*, examines platforms like Spotify and YouTube that must simultaneously optimize advertising intensity, subscription pricing, acquisition expenditures, and retention expenditures. The user base evolves as a Markov chain with three states: non-users, free users,

and premium users. Transition probabilities depend endogenously on platform decisions. This creates complex interdependencies: acquisition spending affects how many users enter the free tier, advertising intensity influences conversion to premium and churn back to free status, retention investment reduces premium churn, and subscription pricing affects both conversion and attrition. The platform maximizes long-run profit under steady-state conditions.

The analysis confirms that threshold-driven strategy extends to freemium platforms, though specific mechanisms differ from on-demand contexts. Premium quality enhancement exhibits a critical threshold at which churn from premium directly to non-user status effectively disappears. Below this threshold, platforms operate in a value-building phase, investing heavily in retention to prevent users from abandoning the ecosystem entirely. At the threshold, retention reaches peak effectiveness. Beyond it, platforms enter a value-extraction phase where they can reduce retention spending and raise prices more aggressively, confident that ecosystem stickiness keeps users within the system even when they occasionally downgrade to the free tier. This parallels the discount threshold in Essay 1. Both mark points where optimal strategy shifts discontinuously, but the underlying mechanisms differ: on-demand platforms shift through price-segmentation mechanics, whereas freemium platforms shift through quality-churn dynamics. Additional findings challenge conventional customer relationship management wisdom. Platforms allocate acquisition spending substantially higher than retention spending even when cost and effectiveness parameters are symmetric. This structural bias arises from persistent premium churn: only a minority of premium users remain subscribed in the steady state, requiring continuous investment to refill the acquisition funnel. A retention paradox also emerges: improved retention effectiveness enables higher pricing rather than expanding the premium user base, so retention investments function primarily as levers for margin improvement rather than user growth. As in the earlier essays, optimal market structures deliberately exclude substantial portions of potential customers, reinforcing that selective coverage is a general principle of platform strategy rather than a context-specific finding. This essay relies on a steady-state Markov chain framework in which free and premium

tiers are treated primarily as static substitutes, with users choosing among options in each period. Future research could extend the model to fully dynamic settings that allow for complementarity between tiers, where free-tier usage acts as a discovery process that gradually increases willingness to pay for premium features. Incorporating learning and evolving preferences would allow platforms to optimize intertemporal strategies rather than steady-state allocations alone. Further work could also derive the transition probabilities from explicit utility-maximizing consumer behavior, strengthening the behavioral foundations of the model and linking the functional forms used here to underlying choice primitives. Finally, empirical validation using panel data from freemium platforms would make it possible to test whether the predicted threshold effects and acquisition-retention tradeoffs arise in practice.

The methodological progression across the three essays, from constrained to expanded optimization and from transaction-based to subscription-based platforms, demonstrates that certain patterns are fundamental rather than accidental. Quality-based differentiation dominates price-based approaches whether platforms employ menu pricing or freemium models. Critical thresholds govern optimal strategy whether in discount rates for on-demand platforms or quality parameters for freemium platforms. Binary market structures emerge under optimization whether platforms serve customers through one-time transactions or ongoing subscriptions.

This thesis makes three primary contributions to the platform pricing literature. First, it develops a unified perspective on how critical thresholds shape platform strategy. Optimal platform design is characterized by regime shifts requiring business model transformation rather than incremental adjustment. Second, it reveals systematic patterns in how platforms should balance competing objectives: quality differentiation consistently dominates price discrimination, pricing instruments exhibit clean separation, and optimal market structures deliberately exclude substantial customer segments. Third, it offers methodological contributions through a menu pricing formulation with endogenous market segmentation, a stepwise optimization framework that isolates distinct economic mechanisms, and a Markov chain framework with endogenous transition probabilities for

dynamic settings.

Together, these essays establish that successful platform pricing requires identifying and navigating critical thresholds, investing primarily in quality differentiation rather than price-based incentives, and accepting that profit maximization typically leads to selective market coverage rather than universal adoption. These insights apply across diverse platform contexts from ride-hailing and food delivery to music streaming and video entertainment. As digital platforms continue to mediate an expanding share of economic activity, understanding these optimization dynamics becomes increasingly important for both platform strategy and regulatory policy.

References

- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3), 668–691.
- Chen, D., Zhou, Y., Guan, X., & Lin, X. (2022). Transaction or membership? impact on on-demand delivery service platforms' profits, consumer surplus, and labor welfare. *Journal of Systems Science and Systems Engineering*, 31(5), 563–593.
- Kung, L.-C., & Zhong, G.-Y. (2017). The optimal pricing strategy for two-sided platform delivery in the sharing economy. *Transportation Research Part E: Logistics and Transportation Review*, 101, 1–12.
- Wagner, T. M., Benlian, A., & Hess, T. (2014). Converting freemium customers from free to premium—the role of the perceived premium fit in the case of music as a service. *Electronic Markets*, 24, 259–268.

