

Retail Platforms, Information Provision, and Consumer Returns

Reza Nazari Khanmiri

A Thesis

in

The John Molson School of Business

Presented in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy (Business Administration) at

Concordia University

Montréal, Québec, Canada

September, 2025

© Reza Nazari Khanmiri, 2025

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis prepared

By: **Reza Nazari Khanmiri**

Entitled: **Retail Platforms, Information Provision, and Consumer Returns**

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Business Administration)

complies with the regulations of this University and meets the accepted standards with respect to originality and quality.

Signed by the Final Examining Committee:

Dr. Raymond Paquin Chair

Dr. Benny Mantin External Examiner

Dr. Tieshan Li Arms-Length Examiner

Dr. Xiaodan Pan Examiner

Dr. Satyaveer Chauhan Examiner

Dr. Xiao Huang Supervisor

Dr. Georges Zaccour Co-supervisor

Approved by _____
Dr. Tracy Hecht, Graduate Program Director

2026 _____
Dr. Anne-Marie Croteau, Dean of Faculty

Abstract

Retail Platforms, Information Provision, and Consumer Returns

Reza Nazari Khanmiri, Ph.D.

Concordia University, 2025

Online retail platforms broaden the customer base but introduce uncertainty about product fit and increase product returns. As a result, information provision and return policy design are pivotal to retailers' profitability and customer welfare. This thesis studies how retailers can coordinate four managerial levers – return policy, influence on customers, information provision and price – to achieve retail objectives in the era of social media. Three essays build a staircase of complexity, employing a game-theoretical framework to explore issues from monopoly to competition.

The first essay considers a monopoly in which the firm determines both the refund and the influence level on customers. Closed-form results reveal that the two can act as substitutes: refunds may be more suitable for customers with low probability of product fit, while influence works better when it is moderate. Moreover, partial refunds consistently outperform influence, effectively serving as their substitute if permitted. Restricting the firm to a binary choice between no refund and a full refund may be conducive to simultaneous gains in profitability and consumer welfare.

The second essay discusses how the retailer should properly endogenize information provision by fine-tuning signal quality. The key finding is that when a high probability of product fit eliminates the strategic value of customer influence and refund, the retailer not only lacks incentive to enhance signal accuracy but may even find excessive precision to be detrimental for conversion. In contrast, with low fit probability, both information tools – through information precision and persuasive influence – and refunds become valuable alternatives for inducing sales.

The third essay analyzes competitive pricing and influence on customers in a duopoly. The equilibrium characterizes how the sellers would craft price–influence decisions to gain competitive advantage. In

particular, the seller with lower probability of fit relies on influence to attract marginal consumers, whereas the higher-fit seller gains no benefit from employing influence. Nonetheless, equilibrium dictates symmetry in pricing. Interestingly, influence on customers can strengthen the set of customers who are loyal to the rival. Furthermore, higher signal quality or external price constraints narrow the set of loyal customers, thereby increasing the strategic value of influence.

Acknowledgments

My deepest gratitude goes to my supervisors, Professor Xiao Huang and Professor Georges Zaccour. Professor Huang's exceptional guidance shaped every stage of this thesis. Her patience, steady availability, and genuine care created an environment in which rigorous work could flourish, teaching me not only methods and results but also the sustaining habits of clarity, curiosity, and persistence. I am equally indebted to Professor Georges Zaccour for his generous mentorship. His counsel was reliably available, his support unfailing, and his humility in discussion exemplary. His guidance provided both direction and confidence, encouraged clearer arguments, and strengthened the substance and presentation of this work.

My thanks extend as well to Professor Xiaodan Pan, my committee member, for her exceptional kindness and sustained support. She went above and beyond – making time, offering help, and providing careful readings and thoughtful suggestions. I am further indebted to my colleagues in the Supply Chain and Business Technology Management Department at the John Molson School of Business and at HEC Montréal (GERAD), whose helpful discussions and constructive atmosphere contributed to many aspects of this research. I also wish to recognize the department and program staff, whose dedication made this journey smoother and more manageable.

I am profoundly grateful to my family and friends for their encouragement, patience, and belief in me. In particular, I wish to thank my mother, whose love and sacrifices have sustained me throughout my studies, and my partner, whose care and active involvement were invaluable in helping me complete this thesis.

Finally, I would like to acknowledge my examiners, Professor Satyaveer Chauhan, Professor Tieshan Li, and Professor Benny Mantin, for graciously agreeing to participate in the evaluation of this thesis.

Contents

List of Figures	x
List of Tables	xiii
1 Introduction	1
1.1 Two Sides of the Digital Retail Coin	1
1.2 Chapter Overview	3
1.3 Motivation and Objectives	3
1.4 Thesis–Level Research Questions	6
1.5 Analytical Lenses and Cross–Study Linkages	7
1.5.1 Information Design as Bayesian Persuasion	7
1.5.2 Return Policy and Money–Back Guarantees	7
1.5.3 Competitive Pricing and Spatial Search	8
1.5.4 Cross–Study Synthesis	9
1.6 Theoretical Foundations and Modelling Framework	9
1.6.1 Unified Product Fit–Uncertainty Setting	10
1.7 Solution Concept and Methodological Blueprint	11
1.7.1 Notation Summary	11
1.8 Summary of Analytical Contributions and Key Results	11
1.9 Thesis Structure	13
2 Literature Review	15

2.1	Introduction	15
2.2	Return Policy: Behavioural and Cost Implications	16
2.2.1	Lenient Versus Strict Return Policies	16
2.2.2	Signalling and Return Policies	17
2.2.3	Competition and Return Policy Design	18
2.2.4	Behavioural and Long-Term Considerations	19
2.3	Electronic Word-of-Mouth and Influence Strategies	20
2.3.1	Firms' Role in Managing e-WOM	20
2.3.2	Influencer Marketing and Credibility Management	21
2.3.3	Firm Strategy and the Limits of Influence	23
2.4	Signal Informativeness and Information Disclosure Design	24
2.4.1	Bayesian Persuasion and Strategic Signal Design	25
2.5	Spatial Competition and Consumer Search	26
2.6	Integrated Modelling of Pricing, Information, and Returns	28
2.7	Research Gaps	30
2.7.1	Positioning	31
3	Influencing Customers and Product Returns	33
3.1	Overview	33
3.2	Analytical Setup	35
3.2.1	Product, Preferences, and Uncertainty	35
3.2.2	Sequence of Events	36
3.2.3	Modelling the Influence Mechanism	37
3.3	Benchmark: No Refund ($r = 0$)	39
3.4	Influencing Customers under an Exogenous Refund Rate	44
3.5	Joint Refund Design and Influence Configuration	54
3.5.1	Binary Refund	54
3.5.2	Continuous Refund	60
3.6	Robustness Analysis with Performance-Linked Influence Cost	62

3.7	Summary	64
4	The Impact of Baseline Informativeness in Customer Influence and Returns	66
4.1	Overview	66
4.2	Analytical Setup	67
4.3	Exogenous Influence Level	68
4.4	Joint Configuration of Baseline Informativeness and Influence	73
4.4.1	No Refund	73
4.4.2	Exogenous Refund Rate	75
4.5	Joint Design of Return Policy, Baseline Signal Informativeness, and Influence Configuration	77
4.5.1	Binary Refund Case	79
4.6	Summary	80
5	Strategic Influence and Pricing Decisions with Post-Purchase Considerations in Duopoly Markets	82
5.1	Overview	82
5.2	Analytical Setup	84
5.2.1	Sellers	85
5.2.2	Customers	85
5.2.3	Customer Influence	85
5.3	Benchmark Case: Exogenous Prices	88
5.3.1	Timing of the Game	88
5.3.2	Equilibrium Outcomes	89
5.4	Customer Influence and Pricing	101
5.4.1	Equilibrium Outcomes	101
5.5	The Impact of Product Returns	106
5.5.1	Equilibrium Outcomes	106
5.6	Summary	113

6	Discussion and Implications	114
6.1	Chapter Overview and Purpose	114
6.2	Coalescing Findings Across Studies	115
6.2.1	Revisiting the Core Questions	115
6.3	Practical Insights and Strategic Relevance	120
6.3.1	Practical Relevance and Translational Considerations	120
6.4	From Model to Managerial Implications	123
7	Conclusion and Future Research	129
7.1	Key Insights and Contributions to the Knowledge	129
7.1.1	Integrated Retail Strategy	129
7.1.2	Signal Design Before Persuasion	129
7.1.3	Practical Implications	130
7.2	Limitations and Boundary Conditions	131
7.3	Directions for Future Research	132
	Bibliography	135
A	Chapter 3	154
B	Chapter 4	164
C	Chapter 5	170
C.1	Explicit Analytical Forms	170
C.2	Analytical Proofs and Numerical Continuations	173

List of Figures

Figure 3.1	Optimal Influence Level varying with the prior belief	42
Figure 3.2	Optimal Influence Level varying with the signal quality	43
Figure 3.3	Optimal Influence Level varying with price/valuation	43
Figure 3.4	Minimum Return-affected Purchase Threshold given a Positive Signal varying with Refund r ($p = 1, \gamma = 0.7$)	46
Figure 3.5	Return-affected price/valuation for Purchase varying with Refund r ($p = 1, \gamma = 0.7$)	46
Figure 3.6	Threshold for the Cost of the Influence Tool $\tilde{C}$ as a Function of the prior belief ρ_0 ($p = 1, v = 2, \gamma = 0.7$)	50
Figure 3.7	Optimal Level of Influence a_r^* as a Function of Prior Belief ρ_0 ($p = 1, v = 2, \gamma = 0.7$)	51
Figure 3.8	Optimal Level of Influence a_r^* as a Function of the Refund r ($p = 1, \gamma = 0.7, \rho_0 = 0.5$)	51
Figure 3.9	Optimal Level of Influence a_r^* as a Function of the Baseline Informativeness γ for low price/valuation ($p = 1, v = 2, \rho_0 = 0.3$)	52
Figure 3.10	Optimal Level of Influence a_r^* as a Function of the Baseline Informativeness γ for high price/valuation ($p = 1.8, v = 2, \rho_0 = 0.6$)	53
Figure 3.11	Joint Binary Refund and Influence Configuration ($k = 0.75, \gamma = 0.7, c = 0$): I/II – Full Refund, No Influence; III – No Refund, Moderate to High Influence; IV – No Refund, No Influence.	55
Figure 3.12	Threshold for the Cost of the Influence Tool under the Optimal Binary Refund ($p = 1, \gamma = 0.7$)	57
Figure 3.13	Threshold for the Level of Informativeness of the Influence Tool under the Optimal Binary Refund ($p = 1, \gamma = 0.7$)	57

Figure 3.14 Expected Profit (Π) and Customer Utility (U) under the Optimal Binary Refund ($p = 1, v = 2, \gamma = 0.7, c = 0$)	58
Figure 3.15 Expected Profit (Π) and Customer Utility (U) under an Optimal Binary Refund ($p = 1.8, v = 2, \gamma = 0.7, c = 0$)	59
Figure 3.16 Expected Profit (Π) and Customer Utility (U): Exogenous Refund vs. Optimal Continuous Refund ($p = 1, v = 2, \gamma = 0.7, c = 0$)	61
Figure 4.1 Optimal Baseline Signal Quality varying with customer prior ($k = 0.7, a = 0.3$) . . .	70
Figure 4.2 Optimal Baseline Signal Quality varying with the Influence Level ($k = 0.7, \rho_0 = 0.4$)	71
Figure 4.3 Optimal Baseline Signal Quality varying with price/valuation ($a = 0.3, \rho_0 = 0.4$) . .	71
Figure 4.4 Optimal Baseline Signal Quality, γ^* , and Influence Level, a^* varying with ρ_0 ($k = 0.7$)	75
Figure 5.1 Seller 1's secure market boundary ρ_1 vs. Seller 2's Influence Level a_2 ($p_1 = 0.7, p_2 =$ $0.6, v = 1, \gamma = 0.7$)	92
Figure 5.2 Seller 2's secure market boundary ρ_2 vs. Seller 1's Influence Level a_1 ($p_1 = 0.7, p_2 =$ $0.6, v = 1, \gamma = 0.7$)	93
Figure 5.3 Seller 1's secure market boundary vs Baseline Signal Quality ($p_1 = 0.7, p_2 = 0.6, v =$ $1, a_2 = 0.5$)	93
Figure 5.4 Seller 2's secure market boundary vs Baseline Signal Quality ($p_1 = 0.7, p_2 = 0.6, v =$ $1, a_2 = 0.5$)	94
Figure 5.5 Purchase thresholds ρ_1, ρ_2 , and k as a function of a_1 ($a_2 = 0.5, \Delta p = 0.1$)	95
Figure 5.6 Purchase thresholds ρ_1, ρ_2 , and k as a function of a_2 ($a_1 = 0.4, \Delta p = 0.1$)	96
Figure 5.7 Purchase thresholds ρ_1, ρ_2 , and k as a function of $\Delta p = p_1 - p_2$ ($a_1 = 0.4, a_2 = 0.5$)	96
Figure 5.8 Optimal Influence Level of Seller 2, a_2^* , and expected profits vs. misfit level, ρ_0 ($\gamma = 0.7, v = 1, p_1 = 0.7$, and $p_2 = 0.6$).	99
Figure 5.9 Optimal Influence Level of Seller 1, a_1^* , and expected profits vs. misfit level, ρ_0 ($\gamma = 0.7, v = 1, p_1 = 0.7$, and $p_2 = 0.6$)	100
Figure 5.10 Contour Plot of the competitive region size as a function of Price Difference and Influence Level ($\gamma = 0.7$).	104
Figure 5.11 Effect of Seller 1's refund rate r_1 on choice thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ ($p_1 = 0.7, p_2 =$ $0.6, \gamma = 0.7, v = 1, a_i = 0.4, r_2 = 0.9$)	108

Figure 5.12	Effect of Seller 2's refund rate r_2 on choice thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ ($p_1 = 0.7, p_2 = 0.6, \gamma = 0.7, v = 1, a_i = 0.4, r_1 = 0.9$)	108
Figure C.1	Function $\rho_1(a_2)$ over the interval $a_2 \in [0, 1]$	174
Figure C.2	First derivative of function $\rho_1(a_2)$ with respect to a_2 over the interval $a_2 \in [0, 1]$. . .	174
Figure C.3	Function $\rho_2(a_1)$ over the interval $a_1 \in [0, 1]$	174
Figure C.4	First derivative of function $\rho_2(a_1)$ with respect to a_1 over the interval $a_1 \in [0, 1]$. . .	175
Figure C.5	4D Grid Analysis of $\rho_1(a_2)$ for parameters $\Delta p, \gamma$, and a_2	176
Figure C.6	4D Grid Analysis of $\frac{\partial \rho_1(a_2)}{\partial a_2}$ for parameters $\Delta p, \gamma$, and a_2	176
Figure C.7	4D Grid Analysis of $\rho_2(a_1)$ for parameters $\Delta p, \gamma$, and a_1	177
Figure C.8	4D Grid Analysis of $\frac{\partial \rho_2(a_1)}{\partial a_1}$ for parameters $\Delta p, \gamma$, and a_1	177
Figure C.9	4D Grid Analysis of $\rho_1(\Delta p)$ for parameters $\Delta p, \gamma$, and a_2	178
Figure C.10	4D Grid Analysis of $\rho_2(\Delta p)$ for parameters $\Delta p, \gamma$, and a_1	178
Figure C.11	4D Grid Analysis of ρ_1 for parameters $\Delta p, \gamma$, and a_2	179
Figure C.12	4D Grid Analysis of ρ_2 for parameters $\Delta p, \gamma$, and a_1	180

List of Tables

Table 1.1	Treatment of Model Elements Across Phases	9
Table 1.2	Glossary of Notations	12
Table 2.1	Selected Studies on e–WOM, Refund Design, and Pricing	32
Table 5.1	Numerical Equilibrium Results with Different Refund Values	111

Chapter 1

Introduction

“Retail is no longer a place to go, but a continuum of touch points we inhabit; we don’t go *to* retail anymore, we *live* it.” — Doug Stephens, *Re-engineering Retail* (2017)

1.1. Two Sides of the Digital Retail Coin

The last decade has redrawn the boundaries of retail. A shopper might watch an unboxing video on TikTok during breakfast, add the item to an online cart at lunch, test it physically in a pop-up after work, complete the purchase from a mobile app on the commute home, and – if dissatisfied – drop the return at an automated locker the next day. The industry labels this seamless integration *omnichannel commerce*, but for the firms orchestrating it, the journey is anything but seamless. In particular, two emerging frictions have quietly intensified:

- **Pre-purchase information asymmetry** Consumers rely on electronic Word-of-Mouth (e-WOM), which includes reviews, influencer endorsements, social media discourse, and other forms of user-generated content, to infer quality attributes that cannot be verified *ex ante*. Yet, the credibility of those cues is increasingly blurred by paid sponsorships or algorithmic amplification. This dynamic underpins the concept of customer influence, which refers to the capacity of firms to shape consumer purchasing behaviour through accessible and often decentralized, information provision mechanisms.

- **Post-purchase mismatch costs** As part of the broader post-purchase experience, return policies have become a central mechanism for reducing consumer risk and enhancing purchase confidence. Generous refund policies stimulate demand but fuel what analysts now call the *return tsunami*: 900 billion US\$ were sent back in 2024 in the United States alone, with clothing and electronics return rates north of 25% ([National Retail Federation and Happy Returns, 2024](#)).

Retail decision-makers therefore face a multi-instrument optimization problem that spans strategic pricing, channel integration, information design (e.g., product information, consumer reviews, and influencer-generated content), customer-assurance mechanisms (e.g., returns and warranties), and broader considerations such as the design of product-and-service bundles – while simultaneously managing fulfillment capacity, sustainability goals, and regulatory scrutiny.

Retail platforms are intermediated commerce environments that facilitate transactions between consumers and sellers while actively shaping the informational and contractual terms of exchange. Unlike traditional retailers, which primarily control inventory and pricing, retail platforms centralize and govern key market design elements such as product information display, consumer reviews, influencer content, and return and refund policies. By aggregating demand and standardizing post-purchase processes, platforms reduce search and transaction costs but simultaneously mediate how consumers form beliefs about product fit and quality. In this sense, retail platforms function not only as marketplaces but also as information designers and risk-allocation mechanisms, making the joint calibration of influence and refund policies a first-order strategic concern rather than an auxiliary operational choice.

Academic research, however, has largely approached these levers in silos: marketing scholars examine e-WOM credibility and demand elasticity; operations scientists minimize fulfillment and reverse-logistics costs; and industrial-organization economists focus on price competition and spatial search. Developing integrated models that weave together even a few of these strands can offer a more holistic lens on the broader optimization challenge. This thesis focuses on a selected subset – namely, the informativeness of signals in information provision, the design of return policies, and pricing in the context of competitive dynamics.

1.2. Chapter Overview

To address these intertwined strategic challenges, the dissertation adopts a staircase structure that layers complexity in measured steps. Each step represents a deliberate progression, with Chapter 1 establishing the conceptual foundation upon which the subsequent analyses build. It outlines the motivation and objective of this work, and presents its positioning within the broader landscape of digital retail strategy. It articulates the central research questions, previews the theoretical tools employed, and introduces the structure of the chapters that follow. The main purpose of this chapter is to set the stage for a review of relevant literature and to clarify the anticipated contribution of this work.

1.3. Motivation and Objectives

The digital commerce landscape is undergoing a transformation where operational burdens increasingly accompany the rise in influence-driven marketing. Global estimates place the aggregate cost of product returns at 1.8% of total retail revenue and 15 million tonnes of carbon emissions annually ([KPMG International, 2024](#)). As consumers grow more reliant on mediated sources of product information, the signals they encounter – endorsements, reviews, influencer content – play a central role in shaping their expectations and decisions. Surveys show that 85% of people are highly likely to act on an influencer’s recommendation, and 93% say online reviews are “very” or “somewhat important” when purchasing a general item ([Avery and Israeli, 2020](#); [Skeldon, 2023](#)). In 2023, the shopping platform LTK alone facilitated 5 billion US\$ in sales via influencer channels ([Gutfreund, 2024](#)), and the global influencer marketing industry was projected to reach 24.2 billion US\$ in 2024 ([Statista, 2024](#)). Despite lingering disruptions in global supply chains and ongoing geopolitical uncertainty, the industry reached 24 billion US\$, demonstrating the robustness of this shift ([Influencer Marketing Hub, 2024](#)).

Influencer marketing, a subset of digital marketing, extends beyond surface-level promotion to function as a strategic signalling system. Grounded in signalling theory, influencers serve as intermediaries whose perceived expertise and curated personas bridge the gap between product

representation and consumer perception. The effectiveness of this mechanism hinges on the quality of the signal: how accurately the influencer's content aligns with product attributes and customer preferences. To bolster signal credibility, firms invest in various tactics: certified reviewer programs, rich media content, and selective disclosure of information (Hao and Tan, 2019). Yet audits highlight a growing trust deficit in digital marketplaces. For instance, studies still find that 20–30% of five-star reviews exhibit suspicious or deceptive patterns (Ott et al., 2012; UK Competition and Markets Authority, 2023).

At the same time, firms face increasing difficulties managing the consistency and credibility of influence as the online ecosystem becomes more fragmented. Today's consumers interact with brands across multiple platforms and devices, and their decisions are influenced by an evolving mix of social content, algorithmic recommendations, and peer-generated reviews (Yadav and Pavlou, 2014). This complexity is particularly pronounced in the context of experience goods, where pre-purchase inspection is impossible and consumers must rely entirely on mediated signals. The stakes of misinformation are thus substantial.

In parallel with these shifts in influence strategies, operational challenges tied to product returns are escalating. By 2024, the average return rate for U.S. retail was 13.2%, down from 14.5% in 2023 and 16.5% in 2022, but still well above the 8.1% seen in 2019 (Capital One Shopping, 2025). Online channels experience even higher rates: e-commerce returns averaged 24.5% in 2024, with apparel at 26% and footwear at 19% (Capital One Shopping, 2025; Radial, 2024). Fraudulent and abusive returns accounted for 15.1% of total returns, costing retailers 104 billion US\$ in 2024 (Capital One Shopping, 2025). The most common drivers remain fit and description mismatches: 53% of apparel returns stem from sizing issues and 28% from products not matching online representations (PowerReviews, 2023). McKinsey & Company also reports that nearly 70% of apparel returns stem from misfit or stylistic issues (McKinsey & Company, 2021). These dynamics have prompted retailers to adopt virtual try-on tools and enforce stricter return windows to mitigate losses (Financial Times, 2024; Kapner, 2023). In response, many other retailers and platforms are retreating from unconditional return policies. Brands such as Zara, H&M, and Abercrombie & Fitch now impose return fees to manage reverse

logistics costs (Kapner, 2023; Meyersohn, 2023).¹

In this evolving landscape, retailers must coordinate two sets of levers. Monetary tools such as price and refund, interact with informational tools such as signal quality and influence intensity. While price remains the most immediate input, its optimal calibration becomes interdependent with the other levers once consumers are modelled as forward-looking Bayesian learners.

Thus, this thesis zeroes in on the triad of influence, information signal quality, and return policy:

- **Influence**, resources allocated to shaping or amplifying reviews, endorsements, or other informational signals. Its effectiveness depends on signal quality that governs consumer belief updates.
- **Signal quality**, the intrinsic credibility of a message source, influenced by verification, richness of content, and alignment with consumer expectations. An accurate signal enhances the informativeness and can reduce mismatch-driven returns.
- **Return-policy (proxied by refund amount (r))**, operationalized through the fraction refunded, the return window, inspection stringency, and restocking fees (Oghazi et al., 2018; Janakiraman et al., 2016). This thesis focuses on the refund fraction as the key decision variable, given its direct interpretability. Full refunds ($r = 1$) offer maximum insurance; partial refunds ($0 < r < 1$) are more efficient under high uncertainty (Su, 2009; Altug, 2023).

Therefore, the objective of this thesis can be framed as twofold: first, positive (descriptive and predictive) to characterize how profit-maximizing retailers control the triad across various competitive and informational environments; and second, normative (prescriptive) to derive guidance for the responsible deployment of influence, disclosure, and refund policies, minimizing avoidable returns and misinformation while encouraging sales and preserving consumer trust in omnichannel retail.

¹Zara began charging a fee for online returns in the U.S. (while maintaining free in-store returns), H&M implemented per-return fees for mailed returns for non-members, and Abercrombie & Fitch introduced shipping or handling fees for returns unless customers opt for store credit or in-store returns.

1.4. Thesis–Level Research Questions

Building on the foundational insights gleaned from the broader literature, the thesis addresses three focal questions:

- (i) **Monopoly trade-off** – How should a single retailer co-optimize customer influence and refund when consumers adjust their purchase decisions based on the information they observe?
- (ii) **Signal-quality optimization** – When a retailer can shape the inherent informativeness of the information medium it controls, what is the profit-maximizing combination?
- (iii) **Competitive equilibrium** – In a duopoly setting, how does a retailer’s pricing decision intertwine with its behaviour-shaping levers (influence, baseline informativeness, return-policy leniency) to shape competitive outcomes?

These overarching questions establish the analytical agenda for the dissertation. Each will be branched, refined, and scrutinized within the context of the three main studies presented in Chapters 3–5:

- **Chapter 3**, driven by the fact that influence stimulates demand yet amplifies future refunds, fixes baseline informativeness and shows that the optimal influence level is non-monotonic in the customer prior fit belief.
- **Chapter 4**, motivated by the idea that perfect accuracy maximizes welfare but leaves little scope for profitable persuasion and an overly rosy signal risks disbelief, makes baseline informativeness a decision variable, capturing marketing investments in richer, more credible content. The firm now looks into the benefits of *informativeness* and *persuasiveness* simultaneously.
- **Chapter 5**, embeds two competing sellers in a Hotelling duopoly and shows that the weaker seller uses greater influence while the stronger rival relies on its natural fit advantage.

1.5. Analytical Lenses and Cross-Study Linkages

The dissertation’s three essays share a common conceptual core and a shared toolkit of analytical ideas. The studies treat *marketing influence*, *return policy*, and *pricing* as jointly chosen instruments in an optimization problem, yet each activates only the subset required for its particular focus. This section therefore weaves the relevant analytical lenses into one narrative, making clear how insights cascade from the monopoly baseline of Chapter 3, through the signal-quality extension of Chapter 4, and into the competitive duopoly of Chapter 5.

1.5.1 Information Design as Bayesian Persuasion

Optimal disclosure is modelled with the strategic–information framework first formalized by [Rayo and Segal \(2010\)](#) and later popularized as *Bayesian persuasion* by [Kamenica and Gentzkow \(2011\)](#). A *sender* (the firm) commits to a signalling device such as an influencer campaign, curated reviews, or virtual-fit content that maps the unobservable state of product–consumer *fit* into a noisy public signal. By choosing the device’s *intensity level* (a), the seller raises the probability that the signal recommends purchase, but at the cost of lower posterior precision: more mismatched buyers show up at the returns stage.

1.5.2 Return Policy and Money–Back Guarantees

A retailer’s return policy bundles several design choices of refund fraction (r), eligibility window, inspection rules, and restocking fees, that together determine how lenient or restrictive the post-purchase experience feels to shoppers ([Janakiraman et al., 2016](#)). The refund fraction, $r \in [0, 1]$, is used as the focal leniency parameter; $r = 1$ is a full *money-back guarantee* (MBG). The literature identifies five main functions and implications of return policies:

- Behavioural Spill-overs. Excessive leniency encourages *bracketing* (ordering multiple sizes with the intent to return) and *wardrobing* (buy-use-return). Survey data from 5,400 U.S. shoppers show that extending the return window beyond 60 days raises bracketing incidence

by 18%, whereas an abrupt policy tightening cuts repeat-purchase intent by 12% ([Abdulla et al., 2024](#)).

- **Cost and Carbon.** Returns impose additional operational burdens such as reverse logistics, repackaging, and processing, which amplify emissions and costs beyond the original transaction.
- **Operational Trade-off.** Returns create a trade off between encouraging demand and incurring processing costs; when these costs are high, more selective refund policies become optimal. Analytical models show that once the per-unit processing cost of a return exceeds roughly 7% of the sales price, a well-calibrated partial refund dominates an MBG by screening out low-certainty buyers while preserving most of the risk-reduction benefit ([Davis et al., 1995](#); [Che, 1996](#)).
- **Strategic Role.** An MBG signals quality but may ignite a “leniency race” that erodes industry profit if rivals match it ([Shulman et al., 2011](#)).

Hence, the refund choice must be coordinated with pricing and influence – the core theme of this thesis.

1.5.3 Competitive Pricing and Spatial Search

Drawing inspiration from [Li et al. \(2019\)](#)’s idea of treating Hotelling “location” as a proxy for *fit probability*, Chapter 5 places two symmetric retailers on that probability line. Heterogeneous consumers – indexed by their prior relative probability of fit across sellers – are assumed to engage in spatial search and choice. The dynamics play out on the intertwined margins that determine the shape of demand:

- (i) **Price** shifts surplus directly, so that the market split point is altered and demand is reallocated;
- (ii) **Influence and signal quality** tilt posterior beliefs, effectively reordering consumers along the fit-probability line and modifying competitive intensity near the centre;

- (iii) **Refund leniency** insures *ex post* mismatches, thereby altering price elasticity and the option value of search under match uncertainty.

Competition is thus structured along the intertwined margins that jointly locate the indifferent consumer; the levers are chosen in equilibrium and interact through payoff, belief-updating, and risk-insurance mechanisms to characterize how these margins map into the fit-probability line under competition.

1.5.4 Cross-Study Synthesis

Thus, taken together, the essays trace a conceptual arc from monetary insurance to informational influence and finally, to their interplay under competitive dynamics. Table 1.1 summarizes how each lens is activated across studies. The modular architecture ensures tractability by allowing additional layers – such as signal quality in Chapter 4 or competition in Chapter 5 – to be introduced independently, while ensuring that any resulting policy shifts can be attributed to that specific complexity rather than to hidden model adjustments.

Table 1.1: Treatment of Model Elements Across Phases

Scope	Chapter 3	Chapter 4	Chapter 5
Market Structure	Monopoly	Monopoly	Duopoly
Customer Influence	Decision Variable	Both	Decision Variable
Signal Informativeness	Input Parameter	Both	Input Parameter
Return Policy	Both	Both	Input Parameter
Price	Input Parameter	Input Parameter	Decision Variable

* Both = The study models each scope sequentially, treating it as an input parameter at certain stages and as an endogenous decision variable at others.

This blueprint preserves parsimony while letting each essay add exactly one substantive layer, so comparative statics remain transparent across chapters.

1.6. Theoretical Foundations and Modelling Framework

To determine how a firm’s full set of strategic levers can work in concert rather than in isolation, the thesis builds on a single analytical framework that – aside from study-specific refinements –

serves as the backbone for all three phases of the research.

Throughout, *expected-utility theory* maps beliefs to choices, *return-screening models* handle the refund–cost trade-off, and the *Bayesian-persuasion framework* guides information design. Spatial competition appears only in Chapter 5, requiring no new consumer primitives – merely heterogeneous priors across the Hotelling line.

1.6.1 Unified Product Fit–Uncertainty Setting

A (or two) risk-neutral firm(s) that sell an experience good to a continuum of consumers facing product fit-uncertainty are considered.² Neither the firm(s) nor the consumers can observe the realized fit prior to purchase. Consumers share a prior belief $\rho_0 \in (0, 1)$ that the product will fit; they update this belief after receiving a *noisy binary signal* (good or bad fit) generated by an influence medium.

Influence tool and tuning. The seller(s) can (i) *acquire* an influence tool and (ii) *tune* its level $a \in [0, 1]$ by, e.g., curating reviews, financing sponsorships, or boosting positive posts. Tuning increases the signal’s recall – more “good-fit” consumers are persuaded to buy – but reduces its precision, so more mismatched buyers eventually surface. When the baseline signal quality I is itself a choice variable (Chapter 4), the firm can also invest to increase baseline informativeness, for example, by adopting verified reviews, richer sizing tools, etc.

Return policy. Where applicable, and prior to sales, each firm announces a refund fraction $r \in [0, 1]$, the most tractable proxy for overall return, policy leniency, complementing other levers such as eligibility window and inspection stringency. Higher r invites more hesitant consumers by offering greater insurance, but also increases the expected reverse-logistics cost.

Price and competition. In monopoly (Chapters 3–4), price p is treated as exogenous to isolate the influence–return trade-off. In duopoly (Chapter 5), two competing sellers $i \in \{1, 2\}$ simultaneously choose their instrument vectors (p_i, a_i) . Consumers observe *both* signals at zero search cost, update their prior ρ_i to posteriors, and purchase from the firm that yields the higher

²Throughout the thesis, “fit,” “match,” and “quality realization” are used interchangeably, and so are “prior belief,” “probability of a match,” and “odds of fit.”

expected surplus (ties are broken randomly). Competition thus arises from heterogeneous fit probabilities rather than any spatial or search frictions.

1.7. Solution Concept and Methodological Blueprint

The empirical object of the thesis is theoretical in nature: the closed-form characterization of optimal or equilibrium policies. Accordingly, each essay adopts a two-stage analytical structure that reflects the sequential logic of strategic decision-making and belief formation.

The first stage models consumer behaviour in response to firm actions. Consumers update their beliefs, evaluate their expected utility, and decide whether to purchase, and if dissatisfied, whether to return the product. This demand-side module is grounded in expected-utility theory and produces tractable expressions for purchase and return behaviour under various informational and market settings.

In the second stage, firms anticipate these behavioural responses and optimize their strategic levers subject to economic constraints and competitive pressures. In monopoly environments, this involves solving a constrained optimization problem to identify profit-maximizing policies. In competitive settings, the thesis characterizes equilibrium outcomes under symmetric competition using pure-strategy best responses. When closed-form expressions are intractable, numerical continuation methods are employed.

1.7.1 Notation Summary

Table [1.2](#) summarizes the key symbols and thresholds used throughout the thesis. Study-specific variations are introduced contextually in the relevant chapters.

1.8. Summary of Analytical Contributions and Key Results

Across the baseline monopoly model four recurring patterns are uncovered that reverberate through the later studies. (i) Influence alone: when refunds are unavailable, a persuasion tool

Table 1.2: Glossary of Notations

Notation	Description
$F \in \{G, B\}$	Nature of match between the product and the customer. G : good fit; B : bad fit.
$s \in \{g, b\}$	Signal produced by the influence tool. g : good/positive signal; b : bad/negative signal.
ρ_0	Prior belief about the product fit.
v	Customer valuation if the product has a good fit.
p	Price of the product.
r	Level of refund, $0 \leq r \leq p$; refund fraction $r \in [0, 1]$.
c	Cost of acquiring an influence tool.
γ	Baseline informativeness of the influence tool.
a	Level of influence set by the firm.
k	Threshold of purchase based on prior belief; also corresponds to the prior belief indifference point in duopoly competition.
ρ_g	Threshold of purchase based on positive signal.
ρ_b	Threshold of purchase based on negative signal.
ρ_i	Threshold of prior fit belief defining the protected market region for retailer i .
C	Threshold of influence for the cost of the influence tool.
Γ_f	Threshold of influence for the baseline informativeness of the influence tool under a binary refund.
$\bar{\Delta}$	External cap on price dispersion.
Π	Firm's expected profit.
U	Expected customer utility.
$\tilde{\cdot}$	Thresholds under exogenous refund.
Subscript 0	Scenarios with no refund.
Subscript r	Scenarios with an exogenous refund r .
Subscript i	Index for retailer, where $i = 1, 2$.

is profitable only for mid-range priors – too pessimistic and no one buys, too optimistic and persuasion is redundant – and the preferred influence level rises with both the prior and the tool’s native accuracy. (ii) Influence and refunds: offering even a modest refund shifts the profitable zone for persuasion toward lower priors and raises the cost a retailer is willing to pay for influence, yet refunds and influence remain strategic substitutes rather than complements. (iii) Refund menus matter: if the firm must choose between “no refund” and “full refund”, three clean regimes emerge – full refund with no persuasion for low priors; no refund but heavy persuasion for moderate priors; and neither lever for high priors – whereas the freedom to set a partial refund causes the firm to drop persuasion entirely and fine-tune the refund fraction instead. (iv) Informativeness versus persuasion: boosting baseline informativeness by itself rarely sways highly skeptical consumers, but coupling high accuracy with a tunable bias unlocks the highest profits and allows the retailer to relax refund generosity. (v) Competitive dynamics: extending the model to a spatial duopoly shows that advantage arises from coordinated price–influence bundles: a firm facing an aggressive discounter reallocates budget toward persuasion, whereas a price leader can temper its influence spending – competitive interaction therefore turns price and persuasion into interdependent levers whose joint calibration determines market share and profitability.

Moreover, welfare analysis echoes these patterns: constrained refund menus can yield win-win outcomes (better surplus for both the firm and consumers) when fit uncertainty is high, but a fully flexible refund lets the firm extract almost all surplus through tailored partial refunds, leaving customers worse off than under a strict no-refund policy.

1.9. Thesis Structure

After this introductory chapter, which also sets out the theoretical and methodological groundwork, Chapter 2 presents a comprehensive literature review. Chapters 3, 4, and 5 then deploy that groundwork in three analytical contexts of increasing richness: (i) a monopoly balancing customer influence design and product-return assurances; (ii) an extension that tunes an information medium’s inherent informativeness and calibrates disclosure levels to achieve the

greatest impact; and (iii) A game-theoretic spatial duopoly in which competing retailers simultaneously optimize price and marketing influence, given cross-seller search by consumers and return policy. Chapters 6 and 7 synthesize insights across the studies, draw managerial and policy implications, and outline avenues for future work.

To maintain readability, the body of the thesis states analytical results succinctly; detailed proofs, derivations, auxiliary figures, and supplementary exhibits are collected in Appendices A–C, corresponding to Chapters 3–5.

Chapter 2

Literature Review

2.1. Introduction

This chapter reviews the literature relevant to the interdependent analysis of return policies, influence strategies (electronic word-of-mouth and information provision), and pricing under product fit uncertainty and competition. Five key themes are examined in turn: (i) Return policy and its behavioural and cost implications; (ii) Electronic word-of-mouth (e-WOM), customer influence, and signal credibility; (iii) Baseline informativeness and disclosure design; (iv) Spatial competition and consumer search; and (v) Integrated modelling approaches linking pricing, information design, and post-purchase policies. For each theme, major scholarly developments are outlined, conflicting findings or theoretical viewpoints are highlighted, and gaps pertinent to the research questions are identified. This synthesis provides a foundation for a unified perspective on how marketing influence and operational return policies may be jointly addressed under conditions of fit uncertainty, competitive interaction, or both.

2.2. Return Policy: Behavioural and Cost Implications

2.2.1 Lenient Versus Strict Return Policies

Return policies of retailers, ranging from very lenient (e.g., full money-back guarantees with no questions asked) to very restrictive (e.g., final sale, no returns), have long been recognized as a strategic tool that influences consumer behaviour and firm performance. A lenient return policy serves as a visible signal of product quality and retailer trustworthiness, effectively showing that the seller stands behind its product (Abdulla et al., 2024). Such policies reduce consumers' perceived risk and build purchase confidence by indicating the retailer's confidence in product quality and fulfillment (Jamil and Khan, 2025). Empirical studies affirm that more generous return terms can positively impact purchase intent and sales *ex ante* by lowering the consumer's anticipated regret cost (Abdulla et al., 2022; Rashid et al., 2025). Leniency, however, comes with the well-known downside of higher *ex post* return rates and associated costs; what has been termed the “return policy leniency dilemma” (Duong et al., 2025). Firms must strike a balance between boosting sales through risk reduction and incurring costs of processing returns (inspections, restocking, resale loss, reverse logistics) (Shulman et al., 2011).

Early analytical research in operations and marketing explored this trade-off and derived optimal refund policies for a monopolist. A full refund maximizes consumer assurance but may invite excessive returns or even opportunistic behaviour (e.g., “free renting” of products). On the other hand, a no-refund policy eliminates returns cost but can depress demand from risk-averse consumers. Intermediate policies, partial refunds or restocking fees, often emerge as profit-maximizing because they balance these effects (Shulman et al., 2011). Indeed, models by Hess et al. (1996) and Chu et al. (1998) showed that an optimal restocking fee can outperform both extreme of no returns or full returns, by screening out low-value purchases while preserving enough incentive for consumers to buy. Notably, such return fees can be optimal even without assuming fraudulent or opportunistic consumer behaviour (Shulman et al., 2011). This indicates that the classic rationale for partial refunds holds purely from the perspective of managing fit uncertainty: even honest customers will return mismatched products if allowed, so the firm may

use a partial refund to shift some of that cost back to the consumer and discourage casual trial purchases. Subsequent empirical work has documented that consumers respond to multiple levers of return policy leniency, such as time limits, monetary costs (restocking fees), effort hassle, and scope of returnable products, in different ways. For example, extending the return window or eliminating return shipping fees generally increases initial purchase likelihood, but overly lenient policies may also encourage undesirable behaviours like over-ordering or “wardrobing” (Rashid et al., 2025). Recent comprehensive studies assess five dimensions of leniency and find that not all are equally valued by consumers. Customers particularly value monetary leniency (full refunds with free return shipping) and time leniency, whereas other aspects (such as whether returns require original packaging or receipt) play a lesser role (Abdulla et al., 2022). These nuanced findings underscore that firms can customize return policy levers to strike an optimal balance, for instance, being generous on refund amount but strict on return deadlines might mitigate costs while preserving most demand stimulation.

2.2.2 Signalling and Return Policies

Return policies also carry signalling implications in contexts of quality uncertainty. Signalling theory posits that a high-quality seller will be more willing to offer a generous refund, because fewer customers will exercise it if the product performs well (Association of Accredited Small Business Consultants (AASBC), 2024). A classic result by Moorthy and Srinivasan (1995) formalized that a money-back guarantee can credibly signal high product quality in equilibrium. Interestingly, some later work challenged this intuition: McWilliams (2012) showed scenarios where a low-quality retailer could benefit from offering a generous refund policy, for example, if it allows the low-quality firm to mimic high quality and poach uncertain consumers, especially when high-quality firms find full refunds too costly. This highlights that the signalling value of return policies can be context-dependent and may lead to a pooling equilibrium (both high and low quality offer returns) under certain cost structures, rather than complete separation. Nonetheless, the prevailing view is that lenient return terms generally indicate seller confidence in the product and thus reduce perceived risk. Consistent with that, consumer surveys and

experiments confirm that lenient return policies improve trust and can increase long-run customer loyalty and repurchase intentions ([Abdulla et al., 2022](#); [Jamil and Khan, 2025](#)). Firms often use generous return guarantees as part of a broader marketing strategy to attract customers (especially in online retail where fit uncertainty is high), effectively treating returns as a cost of customer acquisition. For instance, large e-commerce retailers have normalized free returns, accepting higher reverse logistics costs in exchange for greater market share and customer satisfaction. This trend has even strategic implications in competition, as discussed in [2.2.3](#), a firm's return policy choice can affect competitive positioning.

2.2.3 Competition and Return Policy Design

When multiple firms compete, return policy decisions become more complicated as a result of competition. Early models of returns mostly considered a single firm, but [Shulman et al. \(2011\)](#) extended this to a horizontal duopoly where two sellers offer similar products with potential fit mismatch. They found some, perhaps counterintuitive, results. As firms compete to reduce consumers' risk and attract business, one might expect competition to force more lenient policies. However, their findings show that competing firms may charge restocking fees above the cost of returns, and in fact even higher than a monopolist would, in certain equilibrium conditions. The intuition is that with competition, a firm's strict return policy not only saves costs but can also induce strategic switching: a customer who finds the product unsuitable might then switch to the competitor rather than return for a refund; effectively, a strict policy can push the consumer to the rival, which is a loss of sale. To avoid losing such "marginal" consumers, one might think each firm should lower restocking fees. But if both firms simultaneously choose their policies, there are equilibria where each sets a relatively high restocking fee so that once a consumer purchases from one seller, switching to the other after trying the product is not attractive. In other words, refund can be a double-edged sword in competition; it attracts initial purchases, but also makes it easy for customers to "test" one product and then defect to the rival. [Shulman et al. \(2011\)](#) also find that equilibrium restocking fees tend to be higher when consumers have greater uncertainty about fit (less pre-purchase information), even though

higher fees reduce initial willingness to pay. This reflects a competitive discipline: if trying out products is very frictionless for consumers, they might hop between firms more, forcing sellers to use fees to soften this behaviour. Another takeaway is that partial refund policies can arise in competition as a way to balance attracting customers and locking them in, without going to extremes of either no returns (which would scare off new customers) or full refunds (which would encourage cherry-picking among competitors). These theoretical insights align with some real-world observations: in industries with intense competition and product uncertainty (e.g., consumer electronics or apparel online), firms are often seen converging to similar, not completely free, return policies (such as small restocking fees or short return windows) as a form of tacit coordination to avoid excessive cross-firm “free trials”.

2.2.4 Behavioural and Long-Term Considerations

Behaviourally, return policy has further implications that the literature has begun to explore. Some studies note that overly lenient policies may induce moral hazard or changes in consumer psychology. For example, consumers might purchase more items with the intention of returning most (“bracketing” purchases to try options at home) or be less careful with the product ([Balaram et al., 2022](#); [Chang and Yang, 2022](#)). On the flip side, strict policies might erode goodwill and deter future business from dissatisfied customers ([Oghazi et al., 2018](#)). Researchers have also examined long-term effects of return policies: a lenient policy could pay off via increased customer lifetime value, as customers appreciate the flexibility and reciprocate with loyalty, whereas a strict policy might maximize short-term profit per sale but at risk of alienating consumers ([Janakiraman et al., 2016](#); [Narvar Inc., 2022](#)). These considerations imply that the optimal return policy is not only an operations decision but a marketing one, affecting reputation and word-of-mouth ([Powers and Jack, 2013](#)). Thus, the literature increasingly frames return management in a holistic way; part of an overall value proposition to consumers that has both demand-side and cost-side effects ([García-Dastugue et al., 2023](#)). The interplay with word-of-mouth is particularly salient, which leads to the next theme of e-WOM and influence strategies.

2.3. Electronic Word-of-Mouth and Influence Strategies

E-WOM refers to the practice in which customers can publicly share their opinions and experiences about a product, a service, or a company through digital platforms, which may impact future sales of the product or service and the business of the company. E-WOM, encompassing online consumer reviews, ratings, social media posts, and influencer endorsements, has become a critical factor in consumer decision-making, particularly for experience goods and products with fit uncertainty. A vast body of evidence confirms that e-WOM can strongly influence product demand (Rosario et al., 2016). One of the seminal studies by Chevalier and Mayzlin (2006) showed that positive online book reviews significantly boosted relative sales on Amazon compared to Barnes & Noble. Subsequent research has generalized this finding across industries: higher average ratings and favourable customer reviews correlate with increased conversion rates and revenues, while negative reviews can dampen sales (Eslami and Ghasemaghahi, 2018).

2.3.1 Firms' Role in Managing e-WOM

Within this stream of literature, a segment of the studies has investigated the firm's role in managing and controlling e-WOM contents (Godes et al., 2005). Unlike traditional marketer-controlled communications, e-WOM is perceived as user-generated and thus often more credible or diagnostic to consumers. Consumers tend to trust information from peers and third-parties over direct advertisements. Surveys indicate that around 70% of consumers trust online reviews or influencer recommendations as much or more than information coming directly from brands (Deyo, 2023).

The high level of consumer trust in e-WOM has prompted firms to actively manage and leverage it as part of their broader customer influence strategy. This strategy encompasses efforts to shape consumer perceptions and stimulate positive word-of-mouth through both organic initiatives (e.g., encouraging satisfied customers to leave reviews, improving product quality to earn better ratings) and paid tactics (e.g., hiring influencers, offering referral bonuses, or even procuring fake positive reviews). Many studies provide empirical evidence on how such influence is

exerted through specific mechanisms, such as producing or manipulating e-WOM content to stimulate sales, via biased or promotional online reviews (e.g., [Mayzlin et al., 2014](#); [Luca and Zervas, 2016](#)), free samples ([Lin et al., 2019](#)), viral videos (e.g., [Tucker, 2015](#); [Yang et al., 2024](#)), and tweets ([Lee et al., 2018](#)). Accordingly, the literature has examined both firms that passively observe and respond to e-WOM signals, such as those analyzed by [Chen and Xie \(2005\)](#) and [Feng et al. \(2019\)](#), and firms that proactively design and implement structured customer influence programs as in [Chen and Xie \(2008\)](#), [Godes and Mayzlin \(2009\)](#) and [Burtch et al. \(2018\)](#).

2.3.2 Influencer Marketing and Credibility Management

A more targeted stream within this broader literature focuses on influencer marketing, where individuals with social media followings (or reviewers on platforms like YouTube, Instagram) are enlisted to promote products. These studies highlight a key trade-off: while influencer endorsements can personalize marketing and often yield higher engagement and trust than impersonal ads, their credibility must be carefully managed. If consumers suspect that an influencer's review is purely paid or inauthentic, the persuasive power diminishes. Academic findings suggest that influencer recommendations are particularly persuasive when the influencer is perceived as authentic and expert, forging a parasocial relationship with their audience ([Liu and Zheng, 2024](#); [Leite et al., 2024](#)). In fact, influencer endorsements have been found to generate greater trust than traditional celebrity endorsements, precisely because micro-influencers or niche experts are seen as more relatable and credible ([Leite et al., 2024](#)). However, transparency is crucial: disclosing sponsorship can slightly reduce credibility, but failing to disclose can backfire if uncovered. Recent guidelines and research emphasize the importance of authenticity; content should align with the influencer's genuine experiences, and warn that blatant promotional tone or excessive bias will reduce the signal credibility of e-WOM ([Furness, 2023](#)).

The credibility of e-WOM signals has been an active research area. Since online reviews and ratings critically inform consumer beliefs, gaming the system has tempting incentives. [Mayzlin \(2006\)](#) considers a setting where competing firms may use anonymous and promotional messages

to influence consumers' posterior beliefs on the quality of their products, and found that firms may spend more resources promoting low-quality products. [Dellarocas \(2006\)](#) examines a generalized class of settings and identified scenarios under which high-quality firms might have more incentive to inflate their ratings. [Zhuang et al. \(2018\)](#) conduct laboratory experiments to compare the effects of different review manipulation approaches, namely, adding positive reviews and deleting negative ones. [Guan et al. \(2018\)](#) analyze how a firm may strategically disclose its quality information to influence the intertemporal interaction between online reviews and consumers' expectation of the product. More recently, [Pei and Mayzlin \(2022\)](#) explicitly model the affiliation between a firm and an influencer posting reviews for the product and analyzed the impact of such an affiliation on consumer welfare. [Berman et al. \(2022\)](#) investigate strategies for competing firms as well as for the platform in providing biased information to customers. [Hao and Tan \(2019\)](#) study how firms should integrate influencers with livestream selling.

Platforms have responded with algorithmic and community-policing measures to maintain trust, but an arms race persists. The literature on signal quality in e-WOM often frames it in terms of informativeness versus bias. An "informativeness" perspective asks: how accurately do reviews reflect the true quality or fit of a product? Many factors can distort this: self-selection bias in who chooses to leave a review (often extremely satisfied or dissatisfied customers), strategic biases (reviewers who have an incentive might inflate ratings), or even platforms moderating out extreme content. A line of research in information systems and marketing (e.g., [Hu et al., 2009](#)) has studied how the distribution of online ratings itself signals product uncertainty. For example, a product with consistent moderately high reviews might signal more predictable quality than one with very mixed reviews even if the average rating is similar. For experience goods, consumers often interpret a higher variance in reviews as a sign of fit heterogeneity, which is useful information for someone trying to gauge their own likely fit. Firm strategy for managing e-WOM has to consider these dynamics.

2.3.3 Firm Strategy and the Limits of Influence

If a company can “tune” the e-WOM signal, for instance, by selectively amplifying positive content or addressing negative feedback, it can potentially increase sales, but at the risk of reducing the credibility of the signal if done too overtly. Some firms encourage satisfied buyers to post reviews and may even offer small incentives such as future discounts or loyalty points for writing a review.

From a methodological standpoint, quantitative approaches in economic models of information disclosure assess how strategies concerning information accuracy impact consumer behaviour and firm profitability. These models explore the equilibrium outcomes of varying levels of information accuracy, providing insights into optimal marketing strategies ([Gao and Su, 2016](#)). Within signalling theory, which deals with the transmission and perception of information in markets where asymmetrical information is prevalent ([Spence, 1973](#)), much attention has been paid to how these dynamics play out in environments where consumers cannot directly experience products before purchase ([Connelly et al., 2011](#)).

A recent game-theoretic analysis by [Kim and Xu \(2023\)](#) looked at a novel form of influence: offering cashback or rewards to customers for submitting reviews, effectively a refund-for-review policy. They interpret this through signalling theory: a high-quality seller has an incentive to offer a higher reward for reviews (even for negative reviews) because high quality products will naturally earn better word-of-mouth and thus future sales, recouping the cost of rewards. In equilibrium, they find that generous review incentives can signal high product quality, as low-quality competitors would find it too costly to match the reward and suffer the consequences of more numerous mediocre reviews. However, they also warn that such policies might unintentionally encourage fake reviews or unethical behaviour if not properly designed. Consumers or third parties might attempt to exploit the cashback-by-review system by generating reviews solely to collect rewards, undermining the integrity of the system. This reflects a broader theme in the literature: credibility constraints often limit the extent to which firms can engineer e-WOM. If consumers become skeptical – suspecting that reviews are manipulated or biased by company influence – the persuasive effect of e-WOM diminishes. Therefore, firms face an

optimization problem in influence strategy: how strongly to push favourable word-of-mouth without crossing the line that triggers consumer distrust.

Several empirical studies show that heavy-handed manipulation of online chatter can do more harm than good. For instance, [Sparks et al. \(2016\)](#) find that in the hotel industry, firms that engage with customers by responding to complaints and reviews are perceived as more trustworthy and experience greater customer satisfaction. Similarly, [Xie et al. \(2014\)](#) show that restaurants responding publicly to online reviews can enhance perceptions of service quality and increase future patronage. More recent evidence by [Stuhldreier \(2024\)](#) confirms that replying – even selectively – to consumer comments on social-network posts measurably strengthens perceived customer orientation, which in turn raises brand trust, brand image, and repurchase intentions, underscoring the value of transparent interaction in digital word-of-mouth settings. By contrast, covert influence tactics can backfire: when companies are suspected of manipulating reviews, consumer trust deteriorates sharply ([Mayzlin et al., 2014](#); [Luca and Zervas, 2016](#)). These findings underline a credibility threshold; persuasive efforts are effective when they appear authentic and responsive, but they backfire when perceived as deceptive.

In summary, the literature suggests that e-WOM is an strategic lever with both promise and peril: extremely powerful when perceived as authentic but delicate to manage because attempts at influence can erode its very credibility. For firms operating under product fit uncertainty, e-WOM becomes even more crucial; it is effectively a dispersed source of information to help consumers predict if a product will suit them. This connects to the next theme: how firms design and disclose information to inform or persuade consumers about product fit and quality.

2.4. Signal Informativeness and Information Disclosure Design

When consumers face uncertainty about a product – whether it be uncertainty about quality, fit with personal preferences, or performance – any signals or information provided before purchase become critical. The literature on information disclosure and signalling provides a theoretical backbone for understanding how firms might optimally design these signals. Classic economic

theory ([Grossman, 1981](#); [Milgrom, 1981](#)) established the unravelling principle: if sellers can costlessly and credibly disclose quality information, then in equilibrium, even the lowest quality is revealed; high quality sellers fully disclose to distinguish themselves, and then medium quality, and so on, leaving no information hidden ([Celik, 2014](#)). However, in practice, conditions for complete unravelling – such as costless, fully credible disclosure of vertical attributes – rarely hold. Disclosure is often either not fully credible (sellers have incentives to overstate quality) or not costless (acquiring or conveying information may be costly), and consumers may rationally infer that non-disclosed attributes are probably unfavourable. Thus, firms often engage in partial disclosure or signal design, deciding what information to release, what to withhold, and how to frame it in order to influence consumer beliefs optimally.

2.4.1 Bayesian Persuasion and Strategic Signal Design

One useful framework is Bayesian Persuasion ([Kamenica and Gentzkow, 2011](#)), which recasts the problem as the firm (sender) committing to a signalling scheme that balances information and persuasion. Under Bayesian persuasion, the firm can choose a signalling rule (essentially how informative or noisy a signal to send) to maximize its objective, knowing that consumers will update their beliefs rationally upon seeing the signal. This approach can handle cases where the firm might deliberately introduce some noise or bias in the signal to induce more favourable consumer actions, as long as consumers' posterior beliefs remain consistent with Bayes' rule. Influence tools such as online review systems, recommendation algorithms, or influencer campaigns function as mechanisms that generate signals about product fit. Firms can adjust the persuasiveness of these mechanisms by emphasizing favourable aspects, thereby increasing the likelihood of positive signals, or by omitting unfavourable details to shape perception.

The credibility of signals may be enhanced by third-party verification, for example, certification labels, expert reviews, or consumer protection agencies that lend impartial credibility ([Jin and Leslie, 2003](#); [Luca and Zervas, 2016](#)). Some studies (e.g., [Mudambi and Schuff, 2010](#), [Filieri and McLeay, 2014](#)) have looked at mechanisms such as providing “verified purchase” badges to ensure authenticity or increasing visibility of negative reviews to strengthen credibility, as ways

to improve the content of information in the e-WOM signals. Voluntary disclosure through third parties is another strategy: firms might provide samples to independent reviewers, knowing a good product will earn strong endorsements, a strategy akin to facilitating informative word-of-mouth (Zhu and Zhang, 2010; Hu et al., 2017).

The design of rating systems also influences signal informativeness. Binary or coarse rating systems (thumbs-up/thumbs-down, five-star scales) may be simpler but carry less information than detailed multi-dimensional feedback (written comments). However, too much information can overwhelm consumers, leading to attention biases where they focus only on certain cues (Mudambi and Schuff, 2010). A particularly relevant aspect for products with fit uncertainty is how much information to disclose about product specifications and fit predictors. For instance, apparel retailers have introduced tools like virtual fitting rooms, detailed sizing information, or user photos, which effectively provide more precise signals about fit to reduce uncertainty (Pappas, 2016; Gallino and Moreno, 2018).

In summary, research on signal informativeness and disclosure suggests that firms with superior products want to credibly inform consumers as much as possible, but competitive considerations, cost of information, and credibility constraints lead to partial and strategic disclosure. For the thesis context, these insights imply that a firm's control over the "review signal" or any informative cue could be a lever like price or return policy. The firm can decide how accurate or noisy that signal should initially be to maximize profit, keeping in mind that consumers update their beliefs and may become skeptical if the signal lacks credibility. This naturally interacts with competition: if one firm is providing very rosy signals, competitors might try to counteract by providing their own information.

2.5. Spatial Competition and Consumer Search

Competition in markets with differentiated products is often analyzed through spatial models or consumer search models. Spatial price competition refers to the idea that products can be positioned in a characteristic space (literal geography or abstract attributes), and consumers incur

some “distance” cost or disutility if the product does not match their ideal point. The canonical model is Hotelling’s linear city (1929), which provides a simple representation of two firms at locations on a line competing for consumers distributed along that line ([Hotelling, 1929](#)). In Hotelling’s framework (with fixed locations), each firm has a “protected” segment of nearby consumers who will buy from them even if their price is slightly higher, and a “competitive” or overlapping segment in the middle where consumers are almost indifferent and price competition is fierce.

Subsequent models, like Varian’s model of sales (1980), introduced heterogeneity in search behaviour: some consumers are “shoppers” who compare all prices, others are “loyals” who pick one firm without search ([Varian, 1980](#)). In equilibrium, this yields price dispersion: firms occasionally cut price to snag the comparison shoppers, but keep it high otherwise for the loyal segment. In a spatial context, this translates to some consumers effectively not considering the farther store if the cost (distance) is too great, leading to each firm having a partial local monopoly. Search frictions thus allow higher margins, whereas intense competition usually assumes consumers can costlessly pick the best deal.

Modern online markets have reduced some search frictions (price comparison is easier), but other frictions remain (e.g., evaluating fit or quality still takes effort or relies on imperfect information signals). Research on spatial-duopoly with search has looked at scenarios like one-stop versus two-stop search ([Wolinsky, 1986](#); [Stahl, 1989](#); [Armstrong and Vickers, 2010](#)). For example, if consumers can inspect one product for fit, then decide to either keep it or incur the effort to try the competitor, the equilibrium will involve both pricing and returns considerations to manage that sequential search ([Ellison and Ellison, 2009](#)). In such settings, a firm’s first mover advantage might be to get the consumer to try its product first (perhaps via aggressive marketing or lower initial price), hoping inertia or switching costs deter going to the competitor after ([Baye and Morgan, 2001](#); [Johnson et al., 2003](#)). Alternatively, if returns are easy, the consumer might treat the first purchase as a free trial and then switch if unsatisfied, which pushes firms to consider each other’s return policies as part of competitive strategy. This kind of dynamic search behaviour is highlighted in multi-period or two-period models of competition

with learning (Su, 2009; Davis et al., 1998).

Some studies (e.g., Forman et al., 2008; Bell et al., 2015; Gao and Su, 2016; Ratchford, 2001) examine how integrating review information across channels or “showrooming” behaviour (customers see in store, buy online or vice versa) affects competition. Gensler et al. (2017) provide a framework for showrooming, suggesting consumers will incur the cost of seeking additional information if the benefits outweigh costs. Price competition then interacts with these advantages: the firm offering a superior pre-purchase information environment might be able to charge a premium as consumers face less uncertainty, unless the rival compensates with either price cuts or other assurances. This underscores the interplay between search, information, and price in competition. Overall, the literature suggests that where consumers have to actively search or experiment to learn fit, strategies like pricing, location (differentiation), and information provision (including reviews, free trials, or returns) are deeply intertwined.

2.6. Integrated Modelling of Pricing, Information, and Returns

While some of the above themes have often been studied in isolation, a growing body of research seeks to integrate these decisions within unified analytical frameworks. The motivation for such integration is straightforward: in real markets, firms simultaneously set prices, design persuasive communication strategies, and determine post-purchase policies such as refunds, all while anticipating competitor behaviour. However, combining these elements analytically presents substantial challenges, leading much of the earlier literature to vary one strategic dimension at a time.

A particularly underexplored intersection arises when considering customer influence and product returns together. Despite increasing attention to how firms shape consumer beliefs through online reviews or endorsements, most existing models do not account for the possibility of returns post-influence, that is, after consumers have been persuaded but subsequently realize the product is a poor fit. Empirical evidence suggests that “overly positive” or inflated review ratings (Minnema et al. 2016, Sahoo et al. 2017) as well as broader marketing instruments aimed at

boosting sales ([El Kihal and Shehu 2022](#)) can increase return rates, highlighting the importance of incorporating returns into models of e-WOM-driven profitability. Thus, this thesis's focus on jointly optimizing influence and refund policies under fit uncertainty squarely addresses a notable gap in the literature, one that only a few recent studies have started to fill. One such integrated study by [Sun et al. \(2021\)](#) models two competing online retailers that simultaneously choose a price and a return policy in the presence of online reviews. Their analysis yields rich insights such as conditions for win-win outcomes versus prisoner's dilemmas. A key insight is that when the salvage value of returned products is low, generous refund policies already provide substantial assurance to consumers, which limits the additional benefit that positive reviews can offer. In such cases, reviews contribute little to demand while return costs remain significant. In contrast, when no refunds are available, reviews can play a critical role in boosting demand by serving as the primary source of reassurance, although consumers must then bear the full risk of a poor product fit. This reveals a nuanced relationship between the two mechanisms: in some environments, reviews and refunds operate as complements when salvage value is high and return costs are manageable, while in others they act as substitutes in shaping consumer confidence. This kind of interaction illustrates why analyzing these elements jointly can overturn intuitions drawn from simpler, one-dimensional settings: a firm might assume that providing more positive information always improves outcomes. However, the findings of [Sun et al. \(2021\)](#) suggest that when consumer risk is already mitigated through an operational policy such as a full refund, additional information from reviews may intensify price competition or offer limited incremental benefit. On the other hand, when neither firm offers refunds, the introduction of reviews can either soften or intensify competitive dynamics depending on the underlying parameters. This points to a more complex and context-dependent strategic environment.

Another integrated perspective is offered by [Kuksov and Lin \(2010\)](#) who linked information provision strategies with pricing and even returns as an endogenous outcome. They bring together marketing decisions about what information to communicate and operational choices about offering returns, all within a competitive equilibrium setting. Their finding that the low-quality firm may choose to reduce fit uncertainty by offering detailed sizing information or a

generous return policy, while the high-quality firm focuses on highlighting absolute quality, illustrates an integrated equilibrium outcome that would not emerge if the analysis excluded either the information or return decision. It reflects a segmentation strategy; each firm leverages a different lever to maximize profit: one uses info/returns to target suitable customers, the other uses inherent product quality signal. This implies a potential coordination or substitution between marketing and operational tactics among competitors.

2.7. Research Gaps

While key themes such as pricing, product returns, information design, and competition have each been explored in depth, they are often treated in isolation. Much of the existing literature examines firm decisions along a single dimension – e.g., how to price, how to disclose information, how to structure refunds – while holding other levers fixed. Yet, in practice, these choices are deeply interconnected: marketing communications shape expectations that affect return rates, refund generosity alters the credibility of persuasive messages, and all of these interact with competitive pressures. Analytical models that account for such interactions remain relatively sparse. One area that has received limited attention is the coordination between customer influence strategies and post-purchase return policies. Although empirical work has shown that overly positive reviews or promotional messaging can increase both sales and product returns, few analytical models explicitly capture this trade-off. The behavioural and operational consequences of influence, particularly when firms manipulate online reviews, sponsor content, or invest in biased signalling tools, are still not fully integrated into formal models of return policy design.

Additionally, while signalling theory has been extensively used to study information asymmetry in traditional settings, its application to digital influence mechanisms, where firms actively shape the informativeness and credibility of content, is still emerging. The literature remains particularly thin on how signal precision should be optimized. Similarly, the substitution between *ex ante* influence and *ex post* guarantees is conceptually intuitive but rarely formalized.

Finally, despite rising interest in omnichannel and online retail environments, few models explore how firms should simultaneously manage price, online customer influence, and returns in the presence of strategic competition. Competitive settings, in particular, raise new questions about whether aggressive influence tactics or lenient refund policies become more or less attractive when rivals adopt similar or opposing strategies. These underexplored intersections between influence and returns, informativeness and influence, and strategic competition, define a clear gap in the current literature and motivate the integrative approach taken in this thesis.

2.7.1 Positioning

Table 2.1 reveals that prior work rarely treats all three managerial levers – information bias (or signal quality), refund design, and pricing – in a single framework. Of the fifteen flagship studies catalogued, only three endogenize *both* e-WOM manipulation and refund generosity, and none of those allows price to adjust under strategic competition. Likewise, while several papers integrate returns with pricing, they take the information environment as exogenous; conversely, most e-WOM models ignore the operational consequences of increased returns. By embedding influence, baseline informativeness, refund policy, and competitive pricing (Chapter 5) in one unified model of fit uncertainty, this thesis fills the intersection left blank by the existing literature.

Methodologically, this thesis adopts a Bayesian persuasion framework to model how firms may optimally distort review signals to increase sales while maintaining a minimum threshold of credibility. Within this framework, the influence level captures the extent to which the signal is tilted in a positive direction without fully undermining informativeness. This reflects common practices, such as selectively publishing reviews by presenting a mix of genuinely positive and moderate feedback while withholding particularly negative comments, allowing firms to guide consumer perception without provoking disbelief.

Table 2.1: Selected Studies on e–WOM, Refund Design, and Pricing

Study	MS	IL	RP	PR	CU	METH	Main Finding
Davis et al., 1995	M	–	Bin	Exo	Fit	Analytical	MBGs enhance profits and welfare when salvage values are high, mismatch risk is low, return costs are low, and product trial value is high.
Chevalier and Mayzlin, 2006	M	ORG	–	–	Qual.	Empirical	Improved consumer reviews boost relative book sales; one-star reviews often have stronger impact than five-star ones.
Dellarocas, 2006	M	BIAS	–	End	Qual.	Analytical	Forum manipulation can raise or lower information value, but often traps firms in a costly rat race; better filtering helps curb this.
Su, 2009	M	–	Cont	End	Fit	Analytical	Return Policies impacts demand and coordination; they require tailored contracts for optimal supply chain performance.
Anderson et al., 2009	–	–	Bin/ Cont	Fixed	Fit	Empirical	Return option increases demand, with heterogeneity across categories and customer segments.
Kuksov and Lin, 2010	D	–	Bin	End	Fit	Analytical	Low-quality firm may disclose fit; returns serve as information substitute in competition.
Shulman et al., 2011	D	–	Cont	End	Fit	Analytical	Restocking fees above cost can arise; return leniency may backfire.
Mayzlin et al., 2014	D	BIAS	–	–	Qual.	Empirical	Independent hotels post fake reviews to boost self and hurt rivals; manipulation varies by platform.
Janakiraman et al., 2016	–	–	Bin/ Cont	Mix	Fit	Meta	MBG leniency increases purchase more than returns; 11% lift vs 6% rise.
Gensler et al., 2017	D	–	Bin	Fixed	Fit	Empirical	Show-rooming elevates search; cross-channel information mitigates cannibalization.
Li et al., 2019	D	ORG	–	End	Fit	Analytical	Reviews temper price war; offline retailers gain via selective OCR (online customer reviews) integration.
Sun et al., 2021	D	ORG	Bin	End	Fit	Analytical	Reviews and refunds interact: they complement or substitute depending on salvage value and policy.
Pei and Mayzlin, 2022	M	BIAS	–	Exo	Fit	Analytical	Firms balance bias and credibility by affiliating with influencers selectively; optimal affiliation depends on priors, disclosure, and persuasion cost.
Kim and Xu, 2023	M	QUAL	–	End	Qual.	Analytical	Cashback-for-review can credibly signal high quality.
This thesis	M/ D	BIAS/ QUAL	Bin/ Cont	Mix	Fit	Analytical	Price, influence, and refunds interact strategically: influence and returns may complement or substitute depending on policy and disclosure; competitive dynamics reshape these decisions and outcomes.

Key: MS = Market structure (M = monopoly, D = duopoly); IL = Information lever (ORG = organic reviews, BIAS = biased/fake reviews, QUAL = signal-quality choice); RP = Return policy lever (Bin = binary refund, Cont = continuous/partial refund); PR = Pricing decision (Fixed = fixed/exogenous, End = endogenous, Mix = mixed methods); CU = Consumer uncertainty (Fit = product fit uncertainty, Qual. = quality uncertainty); METH = Research method (Analytical = theoretical/model-based, Empirical = econometric or field study, Meta = meta-analysis).

Chapter 3

Influencing Customers and Product Returns

3.1. Overview

Consumers are preponderantly driven by social media and influencers in their discovery of new products and in making purchase decisions ([HubSpot and Brandwatch, 2022](#)). Though the concept of influencers is dynamic, it generally refers to popular social-media users who can influence other consumers' behaviour. They can be Internet celebrities with millions of followers; customers who publish reviews, photos, or videos of their purchases and assist other customers in the buying process; salespeople or authorities with specialized knowledge; or even animals or robots, as the concept continues to evolve ([Avery and Israeli, 2020](#); [McKinsey & Company, 2021](#)). Digital platforms have significantly altered how consumers evaluate products prior to purchase, especially in online or indirect retail environments where products cannot be physically examined. In such contexts, consumers face inherent uncertainty regarding whether a product will suit their individual preferences or usage context; a concept referred to in this study as *fit uncertainty*. Traditional mechanisms to resolve this uncertainty, such as liberal return policies, can impose significant financial and logistical burdens on firms. Meanwhile, firms increasingly turn to influence mechanisms such as curated customer reviews, influencer

endorsements, and algorithmically promoted testimonials to shape consumers' expectations and pre-purchase beliefs. These tools aim to enhance the perceived match between product and buyer, often by selectively amplifying favourable information.

The tide of influencer marketing and the trend in product returns have given rise to interesting questions about how these two elements should integrate with each other. Online orders inspired by influencers still undergo the same customer inspection for fit and match once products are delivered, as they would if bought in brick-and-mortar stores or online without any influence. While viral and influential content may blur consumers' vision of the product match and stimulate more sales initially, they may also contribute to product returns at a later stage. Even though the dynamics of information dissemination and its impact on consumer behaviour have received significant attention from both researchers and practitioners, the interaction between customer influence and product returns remains underexplored. In particular, it is still unclear whether customer influence efforts should account for potential future returns, or whether the design of return policies should be informed by ongoing influence strategies.

This chapter explores how a monopolist can leverage these two strategic levers in an integrated manner to manage fit uncertainty. Specifically, it builds a formal model of the firm's joint decision-making over return policy and influence effort, and analyzes the downstream effects on customer behaviour and welfare. While the general motivations for studying product returns and belief-shaping tools were introduced in Chapter 1, this chapter grounds the discussion in a stylized analytical framework and directly examines the mechanisms by which firm-driven influence and refund strategies interact.

In this context, *influence* refers to the firm's ability to distort the perceived informativeness of pre-purchase signals. Rather than providing neutral or perfectly representative information, the firm strategically adjusts the distribution of a binary signal that customers use to update their expectations about the product's fit. The degree of distortion is modelled as a manipulable parameter, and its effect is evaluated both under fixed and cost-sensitive influence structures. The chapter addresses the following research questions through a structured sequence of firm and consumer decisions:

- **Influence and Refund Strategy:** How should a firm jointly design its influence strategy and refund policy under customer fit uncertainty?
- **Substitutability versus Complementarity:** Do influence tools and return policies act as substitutes or complements in maximizing firm profit?
- **Welfare Implications:** Whether influence mechanisms enhance or erode consumer welfare?
- **Robustness to Influence Cost Structures:** What happens when influence cost increases with the intensity of shaping favourable perception, i.e., when cost is modelled as a function rather than a fixed value?

The firm's tools in this setting are modelled through a multi-stage framework, capturing both pre-purchase and post-purchase decision layers. This enables a clear delineation of how informational tools, which shape expectations, interact with contractual tools, which provide risk-sharing. The model's structure, and its analytical results, offer theoretical insights into the subtle trade-offs firms face when attempting to manage fit risk through design of consumer-facing policies.

3.2. Analytical Setup

In this section, the modelling assumptions, sequence of actions, and roles of the agents in the baseline setting are outlined. The market is characterized by a monopolist offering a product to a mass of consumers whose size is normalized to 1, each facing uncertainty regarding the product's personal fit. The consumer decision process, the firm's available instruments, and their interaction are modelled accordingly.

3.2.1 Product, Preferences, and Uncertainty

Each consumer has a binary realization of product fit $F \in \{G, B\}$, where G indicates a good fit and B a bad fit. A good fit ($F = G$) yields utility $v > 0$, while a bad fit ($F = B$) yields zero

utility. Consumers do not observe their fit prior to purchase, but it is public knowledge that only a fraction $\rho_0 \in [0, 1]$ of the customers would find the product a good match, i.e., $P(G) = \rho_0$ and $P(B) = 1 - \rho_0 = \bar{\rho}_0$.¹ In the absence of any other information, a customer will rely solely on ρ_0 to form its prior belief of the product's fit and to make a purchasing decision.

The monopolist sets a price p for the product and may offer a partial refund $r \in [0, p]$ if the consumer chooses to return the product post-purchase. If the product turns out to be a bad fit and $r > 0$, the consumer returns the item and receives the refund. Consistent with the product returns literature, the price p , the valuation v , and the prior belief of fit ρ_0 are assumed to be exogenously given, while the level of refund $r \in [0, p]$ is determined by the firm.

3.2.2 Sequence of Events

The interaction unfolds in the following sequence:

Stage 1 (Return Policy): The firm announces its return policy $r \in [0, p]$. Then, nature determines whether the product will fit a customer, i.e., $F \in \{G, B\}$ with probability $\{\rho_0, \bar{\rho}_0\}$. The outcome F is not visible to the firm or the customer until Stage 4.

Stage 2 (Influence Configuration): The firm decides whether to obtain an influence tool at cost c , and if so, the influence level $a \in [0, 1]$. If an influence tool is acquired, an informative signal $s \in \{g, b\}$ is generated about the product fit.

Stage 3 (Customer Orders): The customer observes signal s , updates their posterior belief on F , and makes a purchase decision. If no such signal is observed, the customer makes a purchase decision based on the prior belief ρ_0 .

Stage 4 (Product Returns): The customer observes F and decides whether to return the product.

¹Throughout this thesis, the notation $\bar{y}$ is used as a shorthand for $1 - y$.

3.2.3 Modelling the Influence Mechanism

The firm has access to an influence tool, such as selectively amplified reviews or tailored media, which emits a binary signal s about the product's fit and allows the customers to estimate the product fit in a way that could be influenced by the firm. That is, a potential customer may draw a signal from the reviews to infer a match (or mismatch) between herself and the product and then make a purchase decision. The baseline informativeness of this tool (pre-intervention) is γ , where $\gamma > 0.5$, and reflects the signal's correlation with the true product fit. That is:

$$P(s = g \mid F = G) = \gamma, \quad P(s = b \mid F = B) = \gamma. \quad (3.1)$$

In particular, $\gamma \rightarrow 1$ represents a nearly perfect baseline signal that almost always delivers a correct prediction on the product match, whereas $\gamma \rightarrow 0.5$ indicates an unenlightening signal that possesses little information beyond the prior belief of fit. It is assumed that γ is an intrinsic characteristic of the tool and exogenous to the firm, but rather than arising intrinsically from user experience, the delivered signal's informativeness may be actively shaped by the firm. That is, while it can originate entirely from previous users or third party reviewers, the firm may instead exert influence to skew content in favour of purchase, through tactics such as incentivizing favourable reviews, eliciting testimonials from loyal customers, or suppressing negative feedback.² As a result, a customer's belief about product fit may be shaped by both organically generated cues and the firm-imposed influence. Regardless, the firm may manipulate the precision of the signal by choosing an influence level $a \in [0, 1]$, representing the degree of distortion. The resulting signal distributions are then given by:

$$P(g \mid G, a) = \gamma + a(1 - \gamma), \quad P(g \mid B, a) = a\gamma + (1 - \gamma). \quad (3.2)$$

When no influence is imposed by the firm, i.e., $a = 0$, the signal is reduced to its baseline form, namely, $P(g|G, 0) = P(g|G) = \gamma$ and $P(g|B, 0) = P(g|B) = \bar{\gamma}$. Under full influence,

²Examples include offering rewards for positive reviews, flagging unfavourable comments for moderation, or selectively displaying content.

i.e., $a = 1$, the signal is always positive, such that $P(g|G, 1) = P(g|B, 1) = 1$. In essence, a higher level of influence results in an increased $P(g|G, a)$, which corresponds to the true positive rate (TPR, or recall) of the influence tool, that is, the probability that a matched product is accompanied by a positive signal. This feature is assumed to be statistically traceable and contractible. In line with existing literature (e.g., [Kamenica and Gentzkow, 2011](#), [Pei and Mayzlin, 2022](#)), the signal distribution (the baseline signal quality γ and influence level a) is considered common knowledge among customers and they incorporate it when updating their belief $P(s|a)$ about the probability that the product is a good fit after observing the signal. For example, signal quality may be inferred from users' personal review experiences in the absence of compensation. The level of influence may also be disclosed directly by the firm or inferred via third-party investigation. For instance, influencers are often required to disclose any "material connection" with a brand ([Federal Trade Commission, 2023](#)). Online retailers such as HomeDepot and Amazon label sponsored reviews with statements like "This review was collected as part of a promotion" or "Vine Customer Review of Free Product" and browser extensions such as Fakespot and ReviewMeta assist consumers in detecting subtler forms of influence. These expressions imply that as a increases, the signal becomes more favourable (i.e., more likely to emit g) regardless of true fit, but at the expense of both true and perceived informativeness. This modelling choice captures the strategic use of information distortion by the firm to enhance the likelihood of purchase without changing the underlying product quality or characteristics.

Observing signal s and the influence level a (or merely the signal distribution), a customer can update its posterior belief of the product's fit using Bayes' rule and make a purchasing decision accordingly. In particular, the probabilities of a good fit with signal g and b under influence

level a are, respectively,

$$\begin{aligned} P(G | g, a) &= \frac{P(g | G, a) \cdot P(G)}{P(g | a)} = \frac{P(g | G, a) \cdot P(G)}{P(g | G, a) \cdot P(G) + P(g | B, a) \cdot P(B)} \\ &= \frac{(\gamma + a\bar{\gamma})\rho_0}{(\gamma + a\bar{\gamma})\rho_0 + (\bar{\gamma} + a\gamma)\bar{\rho}_0}, \end{aligned} \quad (3.3)$$

$$\begin{aligned} P(G | b, a) &= \frac{P(b | G, a) \cdot P(G)}{P(b | a)} = \frac{P(b | G, a) \cdot P(G)}{P(b | G, a) \cdot P(G) + P(b | B, a) \cdot P(B)} \\ &= \frac{\bar{\gamma}\rho_0}{\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0}. \end{aligned} \quad (3.4)$$

and $P(B|g, a) = 1 - P(G|g, a)$ and $P(B|b, a) = 1 - P(G|b, a)$. In particular, $P(G|g, a)$, namely, the *precision* of the signal, decreases with the level of influence a (see Lemma A.2 in Appendix A).

3.3. Benchmark: No Refund ($r = 0$)

The benchmark case in which no refund is offered by the firm, i.e., $r = 0$, is first considered. This case also corresponds to a reduced instance consistent with the setup in Pei and Mayzlin (2022). By applying backward induction, it is noted that the customer's decision at Stage 4 becomes trivial, as no product return is possible.

At Stage 3, in the absence of any signal (i.e., when no influence tool is acquired by the firm at Stage 2), a purchase is made if and only if

$$\rho_0 v - p \geq 0. \quad (3.5)$$

Let the price-to-value ratio be denoted by $k = \frac{p}{v}$. The purchasing condition in this case simplifies to $\rho_0 \geq k$.

In the presence of a positive signal, i.e., $s = g$, the purchase is made if and only if

$$P(G|g, a)v - p \geq 0, \quad (3.6)$$

which is equivalent to

$$\rho_0 \geq \rho_g(a) = \frac{k(1 - \gamma\bar{a})}{\bar{k}(\gamma + a\bar{\gamma}) + k(1 - \gamma\bar{a})}. \quad (3.7)$$

In contrast, when a negative signal is received, i.e., $s = b$, the product is purchased if only if

$$P(G | b, a)v - p \geq 0, \quad (3.8)$$

which implies

$$\rho_0 \geq \rho_b = \frac{k\gamma}{\bar{k}\bar{\gamma} + k\gamma}. \quad (3.9)$$

In summary, the threshold conditions $\rho_g(a)$, k , and ρ_b characterize the minimum belief level required for purchase under the cases of a positive signal, no signal, and a negative signal, respectively.

Further analytical implications of this benchmark is developed in Lemma 3.1.

Lemma 3.1. *The purchase thresholds without refund satisfy $\rho_g(a) \leq k \leq \rho_b$.*

Then, the customer's purchasing decision follows:

- if $\rho_0 < \rho_g(a)$, a customer will never purchase;
- if $\rho_g(a) \leq \rho_0 < k$, a customer will purchase only if a positive signal is observed ($s = g$);
- if $k \leq \rho_0 < \rho_b$, a customer will purchase if there is no signal or a positive signal ($s = g$), but will not purchase if it observes a negative signal ($s = b$);
- if $\rho_0 \geq \rho_b$, the customer will always purchase, with or without the signal, and whether the signal is positive or negative.

Back in Stage 2, the firm needs to decide whether or not to acquire the influence tool, and if so, the level of influence. If the firm chooses not to acquire an influence tool, it will earn a solid p if the prior $\rho_0 \geq k$ and 0 otherwise. The expected profit function can then be written as

$$\Pi_0(N) = \mathbb{1}\{\rho_0 \geq k\}p, \quad (3.10)$$

where the subscript 0 represents the scenario with no refund, the N represents the choice of no influence tool, and $\mathbb{1}\{\cdot\}$ is an indicator function that yields 1 if a statement is true and 0 otherwise. If the firm acquires an influence tool with an influence level a , the expected profit follows

$$\Pi_0(a) = P(g|a)\mathbb{1}\{\rho_0 \geq \rho_g(a)\}p + P(b|a)\mathbb{1}\{\rho_0 \geq \rho_b\}p - c. \quad (3.11)$$

The first component of the above captures the chance of a positive signal under influence level a and the respective profit. The second component concerns the scenario of a negative signal. The last term represents the cost for acquiring the influence tool. Let $a_0^* = \arg \max_{0 \leq a \leq 1} \Pi_0(a)$. The following proposition characterizes the optimal solution for the firm at this stage.

Proposition 3.1. Optimal Influence Configuration with No Refund

Suppose the firm offers customers no refund. Then,

- (i) *when $\rho_0 < \rho_g(0)$, the firm will not acquire an influence tool and the customers will not make a purchase; the firm's expected profit is $\Pi_0 = 0$;*
- (ii) *when $\rho_g(0) \leq \rho_0 < k$, there exists $C > 0$ such that*
 - (a) *if $c \leq C$, the firm will acquire an influence tool and set the influence level*

$$a_0^* = \left(\frac{\gamma}{\bar{\gamma}} - \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} \right) / \left(\frac{\gamma}{\bar{\gamma}} \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} - 1 \right)$$

and a customer will purchase if a positive signal is observed; the firm's expected profit is $\Pi_0 = \Pi_0(a_0^*) = C - c$;

- (b) *otherwise, the firm will not acquire an influence tool and the customers will not make a purchase; the firm's expected profit is $\Pi_0 = 0$;*
- (iii) *when $\rho_0 \geq k$, the firm will not acquire an influence tool and the customers will always purchase; the firm's expected profit is $\Pi_0 = p$.*

The proposition reveals several important principles for influencing customers. First, customer influencing activities should focus on the zone where the purchase is endangered. For instance, the firm should never intervene and influence the customers when they will purchase based on

prior belief, i.e., $\rho_0 \geq k$. Second, within the “endangered purchase” zone, i.e., $\rho_0 < k$, the firm should only invest in the influence tool when the prior belief is not too low and the cost of the tool is not too high, i.e., $\rho_g(0) \leq \rho_0 < k$ and $c \leq C$. Otherwise, influencing customers is not worthwhile, and the firm should just let that business go. Lastly, the optimal influence level a_0^* , whenever applicable, should generally increase with the prior belief ρ_0 , the baseline signal quality of the influence tool γ , and the valuation in case of a match v ; but decrease with the price p . These relationships are depicted in Figures 3.1, 3.2, and 3.3, respectively, with parameters set as follows: optimal influence level is evaluated with $p = 1$, $v = 2$, $c = 0$, $\gamma = 0.7$, and $\rho_0 = 0.45$, wherever applicable.




Figure 3.1: Optimal Influence Level varying with the prior belief

Lastly, it is emphasized that influence should, in principle, only be exerted in the region where a purchase would not occur based on the prior belief of fit, i.e., when $\rho_0 < k$. However, the application of influence is found not to affect the customer’s expected utility. Specifically, the expected utility of the customer remains unchanged regardless of whether an influence tool is employed, as formalized in Corollary 3.1. This finding implies that the influence mechanism primarily serves to enhance the firm’s expected profit rather than to benefit the customer directly.




Figure 3.2: Optimal Influence Level varying with the signal quality




Figure 3.3: Optimal Influence Level varying with price/valuation

Corollary 3.1. Customer Welfare under No Refund

When no refund is allowed, the expected customer utility $U_0 = 0$ when $\rho_0 < k$ and $U_0 = \rho_0 v - p \geq 0$ otherwise.

3.4. Influencing Customers under an Exogenous Refund Rate

This section investigates the setting in which the firm accommodates partial refunds and determines its level of influence based on an exogenously specified refund rate $r \in (0, p]$. Such refund rates typically arise when firms must conform to prevailing industry standards or uphold commitments to customer service expectations. Moreover, in practice, adjusting the level of customer influence is generally more feasible than altering the return policy.

At Stage 4, the customer observes the product fit outcome F and decides whether to retain or return the item. If the product is a good fit ($F = G$), it is kept, yielding a net utility of $v - p$. If the product is a poor fit ($F = B$), it is returned, resulting in a net utility of $r - p$.

At Stage 3, the purchase decision is made based on the available signal. In the absence of a signal, the customer relies on the prior belief ρ_0 to evaluate the expected utility at Stage 4. If a purchase is made, the expected utility is given by

$$U_r = \rho_0(v - p) + \bar{\rho}_0(r - p) = \rho_0 v + \bar{\rho}_0 r - p. \quad (3.12)$$

Let $\tilde{k} = \frac{p-r}{v-r}$. A purchase occurs if and only if $U_r \geq 0$, or equivalently, $\rho_0 \geq \tilde{k}$. If a positive signal is observed, the customer updates beliefs and purchases if

$$U_r = P(G|g, a)v + P(B|g, a)r - p \geq 0. \quad (3.13)$$

This condition simplifies to $\rho \geq \tilde{\rho}_g(a)$, where

$$\tilde{\rho}_g(a) = \frac{\tilde{k}(1 - \gamma\bar{a})}{\tilde{k}(\gamma + a\bar{\gamma}) + \tilde{k}(1 - \gamma\bar{a})}. \quad (3.14)$$

Similarly, with a negative signal, the customer proceeds with the purchase if

$$U_r = P(G|b, a)v + P(B|b, a)r - p \geq 0, \quad (3.15)$$

which holds if and only if $\rho \geq \tilde{\rho}_b$, where

$$\tilde{\rho}_b = \frac{\tilde{k}\gamma}{\tilde{k}\tilde{\gamma} + \tilde{k}\gamma}. \quad (3.16)$$

These thresholds, $\tilde{k}$, $\tilde{\rho}_g(a)$, and $\tilde{\rho}_b$, thus characterize the customer's purchase decision under no signal, a positive signal, and a negative signal, respectively.

Lemma 3.2. *Suppose the firm will refund $r \in (0, p]$ for any returns. Then, at Stage 3,*

- *if $\rho_0 < \tilde{\rho}_g(a)$, a customer will never purchase;*
- *if $\tilde{\rho}_g(a) \leq \rho_0 < \tilde{k}$, a customer will purchase only if a positive signal is observed;*
- *if $\tilde{k} \leq \rho_0 < \tilde{\rho}_b$, a customer will purchase if there is no signal or a positive signal, but will not purchase if a negative signal is observed;*
- *if $\rho_0 \geq \tilde{\rho}_b$, the customer will always purchase.*

In summary, customer purchase decisions follow the same structure as those without a refund (Lemma 3.1). However, a refund will generally lower the thresholds for purchase (Corollary 3.2), and the influence tool will be used when the market possesses a lower prior belief about the fit (ρ_0).

Corollary 3.2. Impact of Refund on Customer Choice

The purchasing thresholds $\tilde{k}$, $\tilde{\rho}_g(a)$, and $\tilde{\rho}_b$ decrease in the level of refund r . In particular, $\tilde{k} \leq k$, $\tilde{\rho}_g(a) \leq \rho_g(a)$, and $\tilde{\rho}_b \leq \rho_b$.

Consistent with Corollary 3.2, Figure 3.4 shows that the minimum positive-signal purchase threshold decreases with the refund rate across a range of price–valuation ratios. As refunds rise, the return policy insures against misfit and induces purchases from consumers who are less optimistic about product fit even after a positive review. Holding the refund fixed, lower price–valuation ratios likewise reduce the threshold, so less-optimistic consumers can still be persuaded by positive review signals. All curves meet at the full-refund case, where the threshold




Figure 3.4: Minimum Return-affected Purchase Threshold given a Positive Signal varying with Refund r ($p = 1, \gamma = 0.7$)




Figure 3.5: Return-affected price/valuation for Purchase varying with Refund r ($p = 1, \gamma = 0.7$)

collapses to zero; under a full-refund policy, the thresholds characterized in Lemma 3.2 are nonbinding because consumers would always purchase under an MBG.

Figure 3.5 presents the complementary shift toward more pessimistic consumers: the return-affected price-valuation ratio of purchase with no signal is decreasing in the refund amount, so

the threshold moves left as refunds rise. This mirrors the effect observed in Figure 3.4 for positive signals, that is, higher refunds leads to effective conversion in less-optimistic consumers without a signal. However, the decline here is noticeably steeper. Consequently, as the threshold shifts left, the region in which the refund meaningfully affects behaviour also contracts, tightening the application zone at higher refunds.

Back in Stage 2, the firm compares the expected profit across all scenarios to determine if an influence tool will be acquired, and if so, the level of influence to use. If the firm decides to go without an influence tool, by Lemma 3.2, a customer will only purchase when $\rho_0 \geq \tilde{k}$, and, with a probability $\bar{\rho}_0$, the customer will come back for a refund r at Stage 4. The firm's expected profit without an influence tool is then

$$\Pi_r(N) = \mathbb{1}\{\rho_0 \geq \tilde{k}\}(p - \bar{\rho}_0 r). \quad (3.17)$$

On the other hand, if an influence tool is acquired, it will produce a positive signal ($s = g$) with probability $P(g|a)$. By Lemma 3.2 (ii)-(iv), a customer will order as long as $\rho_0 \geq \tilde{\rho}_g(a)$, and the odds they will claim a refund in Stage 4 can be represented by $P(B|g, a)$. This gives rise to the first component in Equation 3.18. The influence tool may also produce a negative signal ($s = b$) with probability $P(b|a)$, and the fraction of the expected profit in this scenario is provided in the second component of Equation 3.18. Overall, the firm's expected profit with an influence tool can be written as follows:

$$\begin{aligned} \Pi_r(a) = & P(g|a) \mathbb{1}\{\rho_0 \geq \tilde{\rho}_g(a)\} [p - P(B|g, a)r] \\ & + P(b|a) \mathbb{1}\{\rho_0 \geq \tilde{\rho}_b\} [p - P(B|b, a)r] - c. \end{aligned} \quad (3.18)$$

Denote $a_r^* = \arg \max_{0 \leq a \leq 1} \Pi_r(a)$ as the optimal influence level should an influence tool be acquired. The following proposition characterizes the firm's optimal decision in Stage 2:

Proposition 3.2. Optimal Influence Configuration with Refund

Suppose the firm will issue a refund $r \in (0, p]$ for each product return. Then, at Stage 2,

- (i) when $\rho_0 < \tilde{\rho}_g(0)$, the firm will not acquire an influence tool and there will be no purchase;
the firm's expected profit is $\Pi_r = 0$;
- (ii) when $\tilde{\rho}_g(0) \leq \rho_0 < \tilde{k}$, there exists $\tilde{C} > 0$ such that
 - (a) if $c \leq \tilde{C}$, the firm will acquire an influence tool and set the influence level as

$$a_r^* = \left(\frac{\gamma}{\gamma} - \frac{\bar{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} \right) / \left(\frac{\gamma}{\gamma} \frac{\bar{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} - 1 \right)$$

and the customers will purchase if a positive signal is observed; the firm's expected profit is $\Pi_r = \Pi_r(a_r^*) = \tilde{C} - c$;

- (b) otherwise, the firm will not acquire an influence tool, and there will be no purchase;
the firm's expected profit is $\Pi_r = 0$;
- (iii) when $\rho_0 \geq \tilde{k}$, the firm will not acquire an influence tool and, the customers will always purchase; the firm's expected profit is $\Pi_r = p - \bar{\rho}_0 r$.

Proposition 3.2 establishes that, when a refund is offered, the structure of the firm's influence decision is preserved, though the purchasing thresholds become more lenient, as analyzed in Corollary 3.2. Specifically, given the prior belief of fit, increased refund generosity may compromise the use of an influence tool by shifting its applicable domain toward more pessimistic priors. Nevertheless, as long as the tool remains in use, the presence of a refund leads to more aggressive customer influencing behaviour. This effect is explicitly characterized below:

Corollary 3.3. Impact of Refund on Influencing Customers

As the level of refund r increases, a firm may discontinue the application of an influence tool if the prior belief of fit is relatively high; otherwise, the firm should continue the use of the influence tool and

- (i) the cost threshold $\tilde{C}$ increases in the level of refund r . In particular, $\tilde{C} \geq C$.

(ii) the optimal level of influence a_r^* increases in the level of refund r . In particular, $a_r^* \geq a_0^*$.

Indeed, with more refunds, the firm should either cease using the influence tool so to reduce the amount of returns, or, keep using the tool but more aggressively in order to benefit from a higher level of *true positive rate*. Corollary 3.3 finds that it is more beneficial to take the first (cease) approach when the prior is already optimistic, as the influence cost is unlikely to deliver more revenue from the matched customers than the loss in additional refunds to the *falsely* matched ones. If the prior is less optimistic, the firm should instead take the second (continue) approach and impose more influence, hence there will be sufficient *truly* matched customers persuaded to make the purchase. For the latter case, Corollary 3.3 (i) finds that as the amount of refund increases, the upper bound for $\tilde{C}$ will improve in favour of the acquisition of an influence tool. Corollary 3.3 (ii) suggests that the firm will also engage in a higher level of influence a_r^* , so to increase the *true positive rate* and mitigate the loss from more refunds. In summary, depending on the joint effect of contrary forces, the level of refund may have a mixed impact on the influence decision.

Figures 3.6 and 3.7 illustrate the aggregate impact of the refund level on the use of the influence tool, the cost threshold $\tilde{C}$, and the optimal influence level a_r^* , respectively. Throughout both figures, the price is set to $p = 1$, the valuation under a good fit is fixed at $v = 2$, and the baseline signal quality is taken as $\gamma = 0.7$. To isolate the effect of refunds, the cost of acquiring the influence tool is set to $c = 0$.

In accordance with Proposition 3.2 and Corollary 3.3, the influence tool becomes relevant only when the prior belief ρ_0 falls below the price-to-value ratio, i.e., $\rho_0 < k = \frac{p}{v} = 0.5$. In the figures, darker curves correspond to more restrictive refund levels, while lighter curves represent more generous refund policies. Each curve begins at $\tilde{\rho}_g(0)$ and ends at $\tilde{k}$, depending on the specified level of refund r .

As depicted in the graphs, an increase in the refund level shifts the application zone of the influence tool, $\rho_0 \in [\tilde{\rho}_g(0), \tilde{k})$, to the left, toward more pessimistic prior beliefs. Although the cost threshold $\tilde{C}$ generally increases with r , a higher refund may exclude a given prior belief from the tool's application zone, effectively lowering the usable cost threshold at higher refund

levels. For example, at $\rho_0 = 0.32$, a no-refund policy enables the use of the influence tool for any cost below 0.46. With a 50% refund policy, the allowable cost threshold increases to 0.61. However, under 75% or 90% refund policies, the tool is not utilized at $\rho_0 = 0.3$.

More broadly, under a no-refund policy, the cost threshold ranges from 0.42 to 1 across all levels of ρ_0 . For a 50% refund policy, the threshold spans from 0.25 to 0.67; for a 75% refund policy, from 0.14 to 0.40; and for a 90% refund policy, from 0.06 to 0.18.




Figure 3.6: Threshold for the Cost of the Influence Tool $\tilde{C}$ as a Function of the prior belief ρ_0 ($p = 1, v = 2, \gamma = 0.7$)

The optimal level of influence a_r^* is subject to similar effects as a result of the shift in the application zone of the influence tool: at $\rho_0 = 0.32$, a no-refund policy would set an influence level of 0.05, while a 50% refund policy would raise the influence level to 0.86. This pattern however does not hold for the 75% or 90% refund policy, as no influence tool will be adopted at those levels of refund. Despite these variations, the optimal level of influence consistently ranges from 0 to 1 across all kinds of prior beliefs and for any given refund policy. In particular, the optimal level of influence is more sensitive to changes in prior belief under a more generous refund.

Figure 3.8 plots the optimal influence level over its application zone. Within this region, the optimal influence increases with the refund amount: higher refunds allow the firm to




Figure 3.7: Optimal Level of Influence a_r^* as a Function of Prior Belief ρ_0
($p = 1, v = 2, \gamma = 0.7$)




Figure 3.8: Optimal Level of Influence a_r^* as a Function of
the Refund r ($p = 1, \gamma = 0.7, \rho_0 = 0.5$)

deploy more persuasion and raise the likelihood of purchase. Taken with Corollary 3.2 and Figures 3.4 and 3.5, this pattern underscores the mediating role of the positive-review-driven purchase threshold: refunds push this threshold downward, relaxing the effective upper bound on influence and enabling the firm to use the influence tool more freely.

The impact of the refund on the optimal level of influence can be more conspicuously illustrated

for a given prior belief as in Figures 3.9–3.10. In both graphs, a more generous refund r (lighter curves) generally yields a higher level of influence a_r^* as long as the influence tool is employed. For instance, when the price-to-valuation ratio is high, as depicted in the right panel, the influence tool will be adopted for all refund levels when the baseline informativeness γ is greater than 0.68, and a 90% refund stimulates the most aggressive influence level (the lightest curve) versus no refund with the minimum amount of influence (the darkest curve). When the price-to-valuation ratio is low, as in the left panel, the more generous refunds of 75% and 90% will not be accompanied with any influence tool, as the low price and generous refund themselves are sufficient to induce the purchase. As the refund becomes more frugal, the firm is better off expanding the purchase zone by picking up the influence tool and improve the *true positive rate*. At the same time, the influence action will also reduce the signal *precision*. Even though a less precise signal will lead to more returns *ex post*, a lower refund can effectively limit the loss the firm would suffer in this regard. In this scenario, interestingly, the level of influence may increase with the refund rate, as can be seen from the two dark curves in Figure 3.9. This suggests that when the influence tool stays in place, a more generous refund may invite more aggressive influence actions.




Figure 3.9: Optimal Level of Influence a_r^* as a Function of the Baseline Informativeness γ for low price/valuation ($p = 1$, $v = 2$, $\rho_0 = 0.3$)




Figure 3.10: Optimal Level of Influence a_r^* as a Function of the Baseline Informativeness γ for high price/valuation ($p = 1.8, v = 2, \rho_0 = 0.6$)

Corollary 3.4. Impact of Refund on Welfare

- (i) The firm's expected profit Π_r increases in the level of refund r when $\rho_0 \leq \tilde{k}$ and decreases in the level of refund otherwise.
- (ii) Customer utility $U_r = 0$ when $\rho_0 \leq \tilde{k}$ and $U_r = \rho_0 v + \bar{\rho}_0 r - p \geq 0$, increasing in the level of refund otherwise.

Corollary 3.4 summarizes the impacts of refunds on the parties' welfare. When the prior belief of fit is low ($\rho_0 < \tilde{k}$), customer utility stays at 0 regardless of the refund amount r , but the firm's expected profit increases with r . In particular, the increase in the firm's expected profit is strict when an influence tool is employed, i.e., conditions in Proposition 3.2 (ii)(a) are satisfied. This implies that allowing product returns on top of influencing the customers has non-trivial value to the firm as well as to the whole supply chain when the chance of a match between the product and the customers is pessimistic.

On the other hand, when the prior belief of fit is high ($\rho_0 \geq \tilde{k}$), allowing product returns will hurt the firm's expected profit but enhance the customer utility. This impact is most significant among customers with moderate prior beliefs, e.g., $\rho_0 \in [\tilde{k}, k]$. Without a refund, these customers will

be under configured influence ($a_0^* \geq 0$), resulting in only a fraction of them making the purchase (those observing a positive signal) and possibly enjoying the product. However, with refunds r , all customers in this range will have the chance to learn the true fit through purchase and to return the product if it does not match their preferences.

Thus far, it has been established that the level of refund shifts the application zone of the influence tool and induces a non-monotonic change in the firm's expected profit. However, determining the appropriate refund level for the firm to adopt from the outset remains unresolved. In the following section, the Stage 1 problem is addressed by examining the firm's optimal selection among potential return policies, accounting for the corresponding optimal customer influence configuration associated with each policy.

3.5. Joint Refund Design and Influence Configuration

In this section, the optimal joint design of the return policy and the influence tool configuration is examined. Specifically, the refund level r is endogenously selected by the firm, while the influence tool is employed and configured according to the structure identified in Proposition 3.2.

Section 3.5.1 begins with a binary decision space for the refund policy, i.e., $r \in \{0, p\}$, wherein the firm chooses between a no-refund and a full-refund option. This setup captures a wide range of practical contexts in which firms prefer simple, easy-to-communicate return rules to reduce administrative burden.

In Section 3.5.2, the more general setting is analyzed, where the refund level is selected from a continuous interval, $r \in [0, p]$, representing a benchmark scenario in which the firm has full flexibility to offer no refund, a partial refund, or a full refund at its discretion.

3.5.1 Binary Refund

When $r \in \{0, p\}$, the problem is reduced to comparing the no-refund scenario as solved in Proposition 3.1 against a full refund ($r = p$) as solved in Proposition 3.2.

Proposition 3.3. Optimal Binary Refund

Suppose the firm can choose between no refund and a full refund in Stage 1. Then,

- (i) when $\rho_0 < \rho_g(0)$, the firm will offer a full refund $r^* = p$ without acquiring any influence tool; the customers will always purchase; and the firm's expected profit is $\Pi = \rho_0 p$;
- (ii) when $\rho_g(0) \leq \rho_0 \leq k$, there exist $C_f \in [0, C]$ and $\Gamma_f \in [\frac{1}{2}, 1]$ such that
 - (a) if $c \leq C_f$ and $\gamma \geq \Gamma_f$, the firm will adopt a no-refund policy, $r^* = 0$, acquire an influence tool, and set the level of influence $a^* = a_0^*$; the customers will only purchase if a positive signal is observed; and the firm's expected profit is $\Pi = C - c$.
 - (b) otherwise, the firm will offer a full refund $r^* = p$ without acquiring any influence tool; the customers will always purchase; and the firm's expected profit is $\Pi = \rho_0 p$;
- (iii) when $\rho_0 > k$, the firm will adopt a no-refund policy $r^* = 0$ without any influence tool; the customers will always purchase; and the firm's expected profit is $\Pi = p$.

As an example, Figure 3.11 illustrates the optimal joint binary refund and influence level when the price-to-value ratio is $k = \frac{p}{v} = 0.75$. In regions I and II where $\rho_0 \leq 0.63$, there will be




Figure 3.11: Joint Binary Refund and Influence Configuration ($k = 0.75$, $\gamma = 0.7$, $c = 0$):
I/II – Full Refund, No Influence; **III** – No Refund, Moderate to High Influence;
IV – No Refund, No Influence.

no influence tool but a full refund, and the customers will always purchase. In region III, as ρ_0

increases from 0.63 to 0.75, the firm will acquire an influence tool and configure the influence level from 0.18 to 1 in accordance with the ρ_0 ; the customer will purchase if the influence tool yields a positive signal, keeping in mind that no refund will be offered if the signal is a false positive. In region IV, the firm will not offer any refunds or acquire any influence tools; the customers will always purchase.

As can be observed, endogenizing the refund introduces structural changes to the purchase decision. A joint configuration of binary refund and influence level will allow customers with all kinds of prior beliefs to purchase and learn the true fit of the product. In particular, a refund will encroach on the application zone of the influence tool (marked as region II in Figure 3.11), where the level of influence is marginal. This implies that, at times, it is more beneficial to let the customers try the product and discover its true fit *ex post* than to impose influence and stimulate nonrefundable sales *ex ante*.

These effects are rooted in Proposition 3.3, which suggests that a firm will not acquire an influence tool when the customers hold a rather optimistic or pessimistic prior belief about product fit; in cases where the prior is moderate, if the influencing cost is high, or if the baseline informativeness of the influence tool is low. In these scenarios, the firm will offer a full refund if the prior belief is moderate or pessimistic, and no refund if the prior belief is optimistic. In other words, the firm will only adopt an influence tool if the prior belief is moderate, the influencing cost is below some threshold, and the signal is reasonably informative. When adopted, the influence tool will be accompanied by a no-refund policy.

Compared to the no-refund scenario, the optimal binary refund restricted the adoption of the influence tool through a tighter cost threshold $C_f \leq C$ and a minimum requirement Γ_f for the degree of informativeness of the influence tool. Figures 3.12 and 3.13 demonstrate how these two thresholds vary with the prior belief ρ_0 and the price-to-value ratio k . In all curves, the price is fixed at $p = 1$ and the baseline level of informativeness of the influence tool is $\gamma = 0.7$. The solid dark curves plot the thresholds with $k = p/v = 0.5$, or equivalently, $v = 2$. At this price level, a firm offering the optimal binary refund would only employ an influence tool when the prior $\rho_0 \in [0.3, 0.5]$. Within this range, $\gamma = 0.7$ is consistently above Γ_f , thus meeting the




Figure 3.12: Threshold for the Cost of the Influence Tool
under the Optimal Binary Refund ($p = 1, \gamma = 0.7$)




Figure 3.13: Threshold for the Level of Informativeness of the Influence Tool
under the Optimal Binary Refund ($p = 1, \gamma = 0.7$)

minimum requirement of baseline informativeness. In addition, the acquisition of an influence tool is conditioned upon $c \leq C_f$, which varies from 0.12 to 0.5 for $\rho_0 \in [0.3, 0.5]$. Compared

to Figure 3.6, where the cost threshold C under a no-refund policy for the same range of ρ_0 varies from 0.42 to 1, it can be observed that a binary refund substantially limits the adoption of the influence tool. As the price-to-value ratio k increases to 0.75, the application zone for an influence tool is shifted to be within $\rho_0 \in [0.56, 0.75]$. In this area, however, $\gamma = 0.7$ only exceeds Γ_f when $\rho_0 \geq 0.63$. That is, the minimum requirement of the baseline informativeness of the influence tool further shrinks its application zone to $\rho_0 \in [0.63, 0.75]$, within which an even lower but still positive cost threshold C_f can be observed. When k goes up to 0.9, the actual application zone of the influence tool $\rho_0 \in [0.88, 0.9]$ is even narrower due to the minimum requirement on informativeness, along with a more stringent cost threshold. This provides numerical illustration for the following result:

Corollary 3.5. Full Refund vs. Costless Influence Tool

Under the optimal binary refund, the firm may adopt a full refund policy without acquiring the influence tool, even if the latter is costless.




Figure 3.14: Expected Profit (Π) and Customer Utility (U) under the Optimal Binary Refund ($p = 1, v = 2, \gamma = 0.7, c = 0$)

In this context, Figures 3.14 and 3.15 illustrate numerical examples that further examine the effect described in Corollary 3.5. A comparison between high and low price-to-value ratios indicates that the effect is substantially more pronounced for high price-to-value products.

Specifically, when $k = 0.9$, the range of priors $\rho_0 \in [0.79, 0.88]$ over which the influence tool would be optimally applied is entirely eliminated by the availability of a full refund. By contrast,




Figure 3.15: Expected Profit (Π) and Customer Utility (U) under an Optimal Binary Refund ($p = 1.8, v = 2, \gamma = 0.7, c = 0$)

when $k = 0.5$, the zone of influence remains unchanged. This comparison highlights that the joint design of refund and influence mechanisms is particularly important for high price-to-value products.

Lastly, the impact of a binary refund policy on the welfare of both parties is examined. It is evident that an optimally chosen binary refund structure enhances the firm's expected profit relative to offering either no refund or a full refund in isolation. This outcome is illustrated in the left panel of Figure 3.14 and Figure 3.15, corresponding to a low-price setting ($p = 1, k = \frac{p}{v} = 0.5$) and a high-price setting ($p = 1.8, k = \frac{p}{v} = 0.9$), respectively. Specifically, the optimal binary refund policy yields a higher expected profit than a no-refund policy when the prior belief is pessimistic, and it surpasses the full-refund policy in performance when the prior belief is optimistic. Within this same region, the binary refund scheme also strictly enhances customer welfare compared to the no-refund alternative. This result is formalized in Corollary 3.6.

Overall, a binary refund enhances the expected profit and customer utility when the prior belief of a match is low. Further, there are scenarios where all parties are better off under a full refund than no refund but influences on customers, e.g., in Figure 3.15 when $\rho_0 \in [0.79, 0.88]$.

Corollary 3.6. Customer Welfare under Optimal Binary Refund

(i) Under the optimal binary refund, when $\rho_g(0) \leq \rho_0 < k$, $c < C_f$, and $\gamma \geq \Gamma_f$, the

customer's expected utility is $U = 0$; otherwise, $U = \rho_0(v - p) \geq 0$ if $\rho_0 < k$ and $U = \rho_0 v - p \geq 0$ if $\rho_0 \geq k$.

- (ii) Compared to no refund, endogenizing a binary refund enhances the customer utility. In particular, the enhancement is strict when $\rho_0 < \rho_g(0)$ or when $\rho_g(0) \leq \rho_0 < k$, $c \in [C_f, C]$ and $\gamma \in [0, \Gamma_f]$.

3.5.2 Continuous Refund

The analysis proceeds by considering a setting in which the refund level is selected from a continuous range, $r \in [0, p]$, allowing for partial refunds to be offered at the firm's discretion.

Based on Proposition 3.2, for any given prior belief ρ_0 , the firm may choose a refund level that results in one of three outcomes: no purchase occurs; purchase is conditioned on receiving a positive signal, with the influence tool employed and configured accordingly; or purchase occurs unconditionally.

The optimal continuous refund policy is characterized in the following proposition:

Proposition 3.4. Optimal Continuous Refund

If a firm may freely choose its refund level r within $[0, p]$, it will never acquire any influence tools. Instead,

- (i) *when $\rho_0 < k$, the firm would offer a partial refund $r^* = \frac{p - \rho_0 v}{\rho_0}$, and the expected profit is $\Pi = \rho_0 v$;*
- (ii) *otherwise, the firm would offer no refund $r^* = 0$, and the expected profit is $\Pi = p$.*

Interestingly, the result suggests that the firm can solely rely on partial refunds, without using an influence tool to achieve the maximum expected profit. In particular, the refund should be next to full for the most pessimistic prior belief ($r^* \rightarrow 1$ when $\rho_0 \rightarrow 0$) and gradually reduce as the prior improves, until it hits the point ($\rho_0 = k$) where customers would always make a purchase, after which no refund would apply ($r^* = 0$).

This aligns with the existing literature stating that partial refunds can be used to screen and discriminate against customers with heterogeneous valuation uncertainties (e.g., [Courty and Li, 2000](#), [Akan et al., 2015](#), albeit in a more dynamic setting). In a generic setting, as considered in this paper, the result also implies the robustness of partial refunds in the presence of customer influence.




Figure 3.16: Expected Profit (Π) and Customer Utility (U):
Exogenous Refund vs. Optimal Continuous Refund ($p = 1, v = 2, \gamma = 0.7, c = 0$)

The expected profit under the optimal continuous refund policy is plotted against the expected profits under various exogenous refund levels (darker curves representing lower refunds and lighter curves representing higher refunds) in the left panel of Figure 3.16. As illustrated, the optimal continuous refund achieves the maximum attainable expected profit across all feasible refund levels, thereby rendering the use of an influence tool unnecessary.

Simultaneously, the optimal continuous refund policy also extracts customer surplus to its fullest extent. As shown in the right panel of Figure 3.16 and summarized in Corollary 3.7, the expected customer utility increases with the level of exogenous refund, while the optimal continuous refund delivers the same level of customer utility as a no-refund policy.

Corollary 3.7. Customer Welfare under Optimal Continuous Refund

Under the optimal continuous refund, the customer's expected utility is $U = 0$ when $\rho_0 \leq k$ and $U = \rho_0 v - p \geq 0$ otherwise.

3.6. Robustness Analysis with Performance–Linked Influence Cost

To examine the robustness of the baseline model, an extension is introduced by incorporating performance-based costs associated with influence tools, particularly when third-party reviewers are involved. In this setting, resources are allocated by the firm to external reviewers to generate reviews that deliver a desired level of distorted fit signals. This variation corresponds to a special case of the general framework, wherein the influence mechanism is deployed in a broader sense to shape customer purchase decisions. The analysis is conducted using the benchmark scenario for tractability, while noting that the insights can be extended to more general cases. The primary difference from the conditions specified in Proposition 3.1 concerns the characterization of an inactive player following Stage 2 (the influence decision), which is described as follows:

Stage 2 (Influence Configuration): The firm decides whether to affiliate with a third-party reviewer at a cost of $\tau(a)$ and, if so, determines the influence level a .

Stage 3 (Reviewer Decision): The reviewer learns the fit of the product by investing c and reports the finding $s \in \{g, b\}$ to the followers, based on the influence level a .

Here, to incentivize the reviewer to influence customers at level a , the firm allocates resources according to an increasing concave function $\tau(a)$. These resources are used to establish the necessary conditions for achieving the desired level of engagement, which may include offering exclusive samples or other incentives to the reviewer. In this way, the allocated resources act as both facilitator and a form of compensation in exchange for the reviewer’s participation.

The reviewer will accept to incur a cost c to examine the product and obtain a fit signal and publish a review if $\tau(a) \geq c$; otherwise, there will be no review and the consumer makes a purchase decision at Stage 4 based on their prior.

Let a^* be a maximizer of $\Pi(a)$, and

$$C = (\rho_0(\gamma + a^*(1 - \gamma)) + (1 - \rho_0)(1 - \gamma + a^*\gamma))p \quad (3.19)$$

and

$$(\gamma + \rho_0 - 2\gamma\rho_0)p > \tau'(0) \quad (3.20)$$

or equivalently, the firm's profit is strictly increasing in affiliation level.

Proposition 3.5. Optimal Influence with Performance-based Cost Function

In the benchmark scenario, the optimal solution is as follows:

- (i) *when $\rho_0 < \rho_g(0)$, there will be no reviews ($a^* = 0$) and also no purchase; the firm's expected profit is $\Pi = 0$;*
- (ii) *when $\rho_g(0) \leq \rho_0 < k$, there exists $C > 0$ such that*
 - (a) *If $C \geq \tau(a^*) \geq c$, then the firm will pay the reviewer $\tau(a^*)$ and set the influence level*

$$a_0^* = \left(\frac{\gamma}{\bar{\gamma}} - \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} \right) / \left(\frac{\gamma}{\bar{\gamma}} \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} - 1 \right)$$

and a customer will purchase if a positive signal is observed; the firm's expected profit is $\Pi_0 = \Pi_0(a_0^*) = C - \tau(a^*)$;

- (b) *Otherwise, there are no efforts to generate reviews with distorted signal from the firm, no review and therefore, no purchase; the firm's expected profit is $\Pi_0 = 0$;*
- (iii) *when $\rho_0 \geq k$, then there is no effort for distortion in the review signal ($a^* = 0$) and the consumer always buys the product; the firm's expected profit is $\Pi_0 = p$.*

In comparison to Proposition 3.1, the objective remains the maximization of profit through influencing the probability of purchase. The introduction of the performance-based cost $\tau(a)$ and the reviewer's examination cost c reflects practical considerations, representing real-world situations in which influencing a reviewer or third-party agent entails a trade-off between the level of influence exerted and the associated costs.

Proposition 3.5 demonstrates that, despite these additional considerations, the fundamental structure of the firm's decision-making process remains intact. The optimal influence level a^* continues to be determined by balancing the expected profit against cost considerations,

regardless of whether such costs are direct, as in the original modelling, or mediated through a third-party reviewer. Consequently, the inclusion of a third-party reviewer and performance-based costs does not alter the solution's structure. This outcome supports the simplifications adopted in the main model and indicates that these simplifications do not compromise the model's validity or applicability to more intricate, real-world contexts. As such, the robustness of the main model is reinforced, confirming that the overall analytical approach remains appropriate even under extended cost configurations.

3.7. Summary

In this chapter, a baseline monopoly setting with fit uncertainty was developed in which influence and refunds were treated as policy levers. The purchase thresholds under no refund were characterized and the demand regions were identified. It was established that influence is relevant only in the endangered–purchase region where the match prior is not too pessimistic and only when the acquisition cost is sufficiently low; consumer utility remains unaffected in this case, whereas expected profit may rise.

When refunds were introduced exogenously, the purchase thresholds were shown to shift downward, moving the application domain of influence toward more pessimistic priors. Conditional on use, the optimal influence level increases with the refund, and the cost threshold that justifies adopting influence also rises. Under a binary refund choice, the structure of the influence decision was preserved but the admissible conditions narrowed; full refunds were optimal without influence over a sizable range, while influence – when used – was paired with no refund. Under a continuous refund, the optimal policy was shown to rely exclusively on partial refunds, rendering influence unnecessary and extracting surplus up to the profit bound. Welfare implications follow directly; notably revealing zones in which both the firm and consumers are made better off by refund availability.

A robustness extension with performance–linked influence costs and a third–party reviewer preserved the qualitative structure of the solution, reinforcing the modelling simplifications

adopted. The analysis provides a baseline against which subsequent chapters deepen the information design: Chapter 4 endogenizes baseline signal accuracy and studies optimal strategic information disclosure and its upper bound under favourable priors. Chapter 5 embeds pricing, competition and spatial search on a relative-match probability spectrum.

Chapter 4

The Impact of Baseline Informativeness in Customer Influence and Returns

4.1. Overview

In markets for experience goods, where consumers cannot fully assess product fit or quality prior to purchase, firms rely heavily on external information mechanisms to facilitate consumer decision-making. Digital platforms have expanded the range of such mechanisms, including customer reviews, influencer endorsements, and expert evaluations, which collectively serve as pre-purchase signals. In parallel, return policies function as mechanisms that allow consumers to resolve fit uncertainty, particularly in settings where pre-purchase sampling or trials are impractical.

The credibility and informativeness of these pre-purchase signals play a critical role in shaping both initial purchase decisions and subsequent return behaviour. While a growing body of literature has examined the effects of e-WOM and influencer marketing, limited attention has been devoted to the systematic management of baseline signal precision: the intrinsic accuracy

of the information environment prior to the application of active influence strategies. In practice, firms often retain some ability to shape this baseline precision through the selection of endorsers, disclosure of product details, or targeted reviewer engagement. Managing this baseline informativeness constitutes a foundational element of the firm's broader customer influence strategy.

This phase builds upon the analytical foundation established in Chapter 3, further integrating signalling theory and utility theory to examine how firms should coordinate signal informativeness, influence intensity, and refund policy design. The investigation centres on two primary research questions:

- **Signal Precision in Digital Persuasion:** How does the pre-influence signal precision affect the overall effectiveness of digital marketing efforts in shaping purchase and return decisions?
- **Signal Accuracy and Policy Optimization:** In what ways does the level of baseline signal accuracy alter the optimal configuration of influence actions and refund terms?

4.2. Analytical Setup

The analysis presented in Chapter 3 treated the baseline informativeness of the information provision mechanism as an intrinsic attribute of the influence tool, exogenous to the firm's control. In this phase, attention is directed toward examining the effects of optimizing the inherent informativeness of the information medium, denoted by γ , through strategic management of pre-influence information disclosure. The information medium, such as a review platform, is assumed to be established and continuously generating signals. However, the informativeness of these signals, which reflects their inherent accuracy and relevance, may be configured by the firm through deliberate information facilitation and precision enhancement efforts. In this framework, the firm is afforded the ability to adjust the precision of the information medium while also applying additional influence to the autonomously generated signals. For analytical tractability and to isolate the strategic interactions of interest, the costs associated with both

information provision and customer influence are normalized to zero throughout the analysis.

Section 4.3 examines a setting in which the influence level is externally imposed, restricting the firm's role to optimizing the informativeness of the information medium. Section 4.4 analyzes the joint configuration of baseline informativeness and influence level, investigating how these factors can be strategically coordinated to improve effectiveness under both no-refund and exogenous-refund scenarios. Lastly, Section 4.5 considers the fully endogenous case, where all relevant decision variables are jointly optimized.

4.3. Exogenous Influence Level

Externally imposed influence levels may arise due to regulatory requirements, platform policies, market conditions, consumer feedback, contractual partnerships, or technological limitations. Such external constraints require firms to reconcile their internal marketing objectives with imposed boundaries. In this analysis, it is assumed that, given a predetermined influence level $0 < a \leq 1$, the firm retains discretion over the baseline signal quality $\gamma \in [0.5, 1]$, thereby determining the informativeness of the initial signal prior to any influence distortion. Consequently, the signal distributions are adjusted to $P(g|G, \gamma)$ and $P(g|B, \gamma)$, following the structure defined in Section 3.2.3. The corresponding posterior probabilities of product fit upon observing either a favourable or unfavourable signal are updated to $P(G|g, \gamma)$ and $P(G|b, \gamma)$, as previously established in equations 3.3 and 3.4. The no-refund scenario ($r = 0$) serves as the benchmark setting for the subsequent analysis of exogenously determined influence levels.

In terms of the sequence of events, Stage 2 will be updated to:

Stage 2 (Signal Configuration): Aware of the ordained influence level a , the firm decides the baseline signal quality γ . A signal $s \in \{g, b\}$ is produced based on the influenced signal distributions.

Using backward induction, since there will be no product return, if there is a positive signal, i.e.,

$s = g$, the customer will make a purchase at Stage 3, if and only if $P(G|g, \gamma)v - p \geq 0$, i.e.,

$$\rho_0 \geq \rho_g(\gamma) = \frac{k(1 - \gamma\bar{a})}{\bar{k}(\gamma + a\bar{\gamma}) + k(1 - \gamma\bar{a})}. \quad (4.1)$$

For a negative signal, i.e., $s = b$, the customer will only make a purchase if $P(G|b, \gamma)v - p \geq 0$, or equivalently,

$$\rho \geq \rho_b(\gamma) = \frac{k\gamma}{\bar{k}\bar{\gamma} + k\gamma}. \quad (4.2)$$

Overall, $\rho_g(\gamma)$ and $\rho_b(\gamma)$ represent the threshold for purchase when there is a positive and a negative signal, respectively. Further, it can be shown that

Lemma 4.1. *The customer's purchasing decision follows immediately as:*

- if $\rho_0 < \rho_g(\gamma)$, a customer will never purchase;
- if $\rho_g(\gamma) \leq \rho_0 < \rho_b(\gamma)$, a customer will purchase only if a positive signal is observed ($s = g$);
- if $\rho_0 \geq \rho_b(\gamma)$, the customer will always purchase whether the signal is positive or negative.

Back in Stage 2, the firm needs to set the baseline signal quality of the influence tool, and the expected profit follows

$$\Pi_0(\gamma) = P(g|\gamma)\mathbb{1}\{\rho_0 \geq \rho_g(\gamma)\}p + P(b|\gamma)\mathbb{1}\{\rho_0 \geq \rho_b(\gamma)\}p. \quad (4.3)$$

Let $\gamma^* = \arg \max_{0.5 \leq \gamma \leq 1} \Pi_0(\gamma)$. The following proposition characterizes the optimal solution for the firm at this stage.

Proposition 4.1. Optimal Baseline Signal Quality Configuration with No Refund

Suppose the firm offers customers no refund. Then,

- (i) when $\rho_0 < \rho_g(1)$, the firm can set $\gamma_0^* \in [0.5, 1]$ and regardless, the customers will not make a purchase; the firm's expected profit is $\Pi_0^* = 0$;
- (ii) when $\rho_g(1) \leq \rho_0 < k$,
 - (a) if $\rho_0 \geq 0.5$, the firm will set the signal quality level at $\gamma_0^* = 1$ and the customers will purchase if a positive signal is observed; the firm's expected profit is $\Pi_0^* = (\rho_0 + a\bar{\rho}_0)p$;
 - (b) otherwise, the firm will set the signal quality level at

$$\gamma_0^* = \frac{k\bar{\rho}_0 - a\bar{k}\rho_0}{\bar{a}(k\bar{\rho}_0 + \bar{k}\rho_0)},$$

and the customers will make a purchase when there is a good signal; the firm's expected profit is $\Pi_0(\gamma^*) = P(g|\gamma^*)p$;

- (iii) when $\rho_0 \geq k$, the firm may set any $\gamma_0^* \in [0.5, \frac{\rho_0\bar{k}}{\rho_0 k + \bar{\rho}_0 k}]$ and regardless, the customers will always make a purchase; the firm's expected profit is $\Pi_0^* = p$.




Figure 4.1: Optimal Baseline Signal Quality varying with customer prior
($k = 0.7$, $a = 0.3$)




Figure 4.2: Optimal Baseline Signal Quality varying with the Influence Level
($k = 0.7, \rho_0 = 0.4$)




Figure 4.3: Optimal Baseline Signal Quality varying with price/valuation
($a = 0.3, \rho_0 = 0.4$)

The results of Proposition 4.1 are exemplified in Figure 4.1, which illustrates the behaviour of the optimal baseline signal quality level in relation to customers' prior belief in fit. The grey panel on the left side, corresponding to Proposition 4.1(i), represents a prior belief where no effort in adjusting the signal's informativeness is likely to be fruitful. In other words, any level of baseline signal quality in this region fails to lead to customer's purchase decision. Similarly, the shaded area on the right-hand side panel corresponds to any baseline signal quality level that does not

alter the customer's decision, given that the decision not to purchase is already determined by the prior belief in this region. The corresponding analysis – Part (iii) of Proposition 4.1 – yields the perhaps counter-intuitive result that, for optimistic priors ($\rho_0 \geq k$), the firm *must not* raise baseline accuracy above $\gamma^{\max} = \rho_0 \bar{k} / (\rho_0 \bar{k} + \bar{\rho}_0 k) < 1$. Intuitively, when consumers already believe the product is very likely to fit, any signal – positive or negative – adds little persuasive value; the purchase will occur even if the information channel is noisy. Increasing γ beyond $\gamma^{\max}$ would sharpen the signal to the point where a *negative* draw could overturn the favourable prior and deter the sale, while a *positive* draw confers no incremental benefit. Hence excessive baseline informativeness reduces expected profit by introducing downside risk without offsetting upside, so the firm optimally caps γ rather than pursuing perfect credibility in situations where confidence is already high.

For customers who believe there is more than a 50 % chance of a match (third panel from the left), it is optimal for the firm to maximize its effort in improving signal quality to generate the most accurate fit signal. This is because when a positive signal is necessary for a purchase and has a higher likelihood of occurring than a negative signal, it benefits the firm to enhance its effectiveness to the fullest extent.

Conversely, when a positive signal is needed for a purchase but has a lower likelihood of occurring (second panel), the optimal baseline signal quality level decreases as prior belief increases but rises with an increase in the exogenous influence level. This trend suggests that the firm may aim for less accuracy in this scenario, possibly to exploit distortions in the updating process to secure a purchase. However, as shown in Figures 4.2 and 4.3, within this zone, the optimal baseline signal quality increases with both the exogenous influence level and the specific price-to-valuation ratio that ensures non-negative utility for the customer upon making a purchase.

4.4. Joint Configuration of Baseline Informativeness and Influence

In this setting, control over both the baseline signal quality and the level of influence is assumed to reside with the firm. The second stage of the analysis sequence is therefore characterized as follows:

Stage 2 (Signal Configuration): The firm decides the signal quality γ and whether to obtain an influence tool at no cost, and if so, the influence level a . A signal $s \in \{g, b\}$ is produced based on nature's choice of F , the baseline signal quality γ , and the influence level a .

4.4.1 No Refund

The benchmark case in which no refund is offered, i.e., $r = 0$, is first considered. By backward induction, the decision of the customer at Stage 4 becomes irrelevant, as product returns do not occur.

At Stage 3, in the absence of any influence exerted by the firm (i.e., when no influence tool is acquired at Stage 2), the customer receives an unbiased signal. If a positive signal is observed, $s = g$, a purchase is made if and only if $P(G|g, a, \gamma)v - p \geq 0$, which is equivalent to the condition $\rho_0 \geq \rho_g(a, \gamma)$, where

$$\rho_g(a, \gamma) = \frac{k(1 - \gamma\bar{a})}{\bar{k}(\gamma + a\bar{\gamma}) + k(1 - \gamma\bar{a})}. \quad (4.4)$$

For a negative signal, $s = b$, the customer proceeds with the purchase if $P(G|b, a, \gamma)v - p \geq 0$, or equivalently, if $\rho_0 \geq \rho_b(\gamma)$, where

$$\rho_b(\gamma) = \frac{k\gamma}{\bar{k}\bar{\gamma} + k\gamma}. \quad (4.5)$$

At Stage 2, both the baseline signal quality and the level of influence are jointly determined by the firm. The expected profit follows

$$\Pi_0(a, \gamma) = P(g|a, \gamma)\mathbb{1}\{\rho_0 \geq \rho_g(a, \gamma)\}p + P(b|a, \gamma)\mathbb{1}\{\rho_0 \geq \rho_b(\gamma)\}p. \quad (4.6)$$

The first component of the above captures the chance of a positive signal under baseline signal quality level γ and influence level a , as well as the respective profit. The second component concerns the scenario of a negative signal. Let $a_0^* = \arg \max_{0.5 \leq a \leq 1} \Pi_0(a, \gamma)$ and $\gamma_0^* = \arg \max_{0.5 \leq \gamma \leq 1} \Pi_0(a, \gamma)$. The following proposition characterizes the optimal solution for the firm at this stage.

Proposition 4.2. Optimal Influence and Signal Quality Configuration with No Refund

Suppose the firm offers customers no refund. Then,

(i) when $\rho_0 < k$,

the firm will set the influence and signal quality level at

$$a_0^* = \frac{\rho_0 \bar{k}}{k \bar{\rho}_0},$$

and

$$\gamma_0^* = 1,$$

the customers will purchase if a positive signal is observed; the firm's expected profit is

$$\Pi_0^* = \rho_0 v;$$

(ii) when $\rho_0 \geq k$, the firm may set any $\gamma_0^* \in [0.5, \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}]$ and use any influence level $a_0^* \in [0, 1]$ and regardless, the customers will always make a purchase; the firm's expected profit is $\Pi_0^* = p$.

The analysis of the firm's optimal strategies, as outlined in Proposition 4.2, is illustrated in Figure 4.4, which summarizes the behaviour of both the optimal baseline signal quality and the optimal influence level relative to consumers' prior belief. Below the price-to-valuation ratio, represented in Panel I, the firm benefits from maximizing the baseline signal quality to make the increasing influence (which rises with prior belief) as effective as possible. However, beyond this threshold, as shown in the blue shaded area (II), and in the absence of associated costs, the firm can apply any influence level. Yet, as indicated by the grey shaded area (III), the baseline signal quality is constrained by an upper limit to prevent discouraging customers with certain




Figure 4.4: Optimal Baseline Signal Quality, γ^* , and Influence Level, a^* varying with ρ_0 ($k = 0.7$)

price-to-valuation ratios from making a purchase. Overall, compared to the results in Section 4.3, the firm's control over the influence mechanism ensures that even the most pessimistic customers can be influenced to make a purchase, leaving no hope lost.

4.4.2 Exogenous Refund Rate

The scenario in which the firm offers partial refunds and simultaneously determines the baseline signal quality and the level of influence under an exogenous refund rate $r \in (0, p]$ is now examined. Following the same analytical approach as previously employed in Section 3.4, the purchase threshold corresponding to a positive signal can be derived as follows:

$$\tilde{\rho}_g(a, \gamma) = \frac{\tilde{k}(1 - \gamma\bar{a})}{\tilde{k}(\gamma + a\bar{\gamma}) + \tilde{k}(1 - \gamma\bar{a})}, \quad (4.7)$$

where the customer will make the purchase if $\rho \geq \tilde{\rho}_g(a, \gamma)$ given a positive signal. And with a negative signal, the customer will purchase if $\rho \geq \tilde{\rho}_b(\gamma)$ where

$$\tilde{\rho}_b(\gamma) = \frac{\tilde{k}\gamma}{\tilde{k}\bar{\gamma} + \tilde{k}\gamma}. \quad (4.8)$$

The customer's purchasing choice can then be summarized as follows:

Lemma 4.2. *Suppose the firm will refund $r \in (0, p]$ for any returns. Then, at Stage 3,*

- *if $\rho_0 < \tilde{\rho}_g(a, \gamma)$, a customer will never purchase;*
- *if $\tilde{\rho}_g(a, \gamma) \leq \rho_0 < \tilde{\rho}_b(\gamma)$, a customer will purchase only if a positive signal is observed ($s = g$);*
- *if $\rho_0 \geq \tilde{\rho}_b(\gamma)$, the customer will always purchase.*

At Stage 2, the firm needs to simultaneously set the signal quality and influence level in the presence of a return policy. The influence tool will produce a positive signal ($s = g$) with probability $P(g|a, \gamma)$. A customer will order as long as $\rho_0 \geq \tilde{\rho}_g(a, \gamma)$, and the odds they will claim a refund in Stage 4 can be represented by $P(B|g, a, \gamma)$. This gives rise to the first component in Equation 4.9. The influence tool may also produce a negative signal ($s = b$) with probability $P(b|a, \gamma)$, and the fraction of the expected profit in this scenario is provided in the second component of Equation 4.9. Overall, the firm's expected profit with an influence tool can be written as follows:

$$\begin{aligned} \Pi_r(a, \gamma) = & P(g|a, \gamma) \mathbb{1}\{\rho_0 \geq \tilde{\rho}_g(a, \gamma)\} [p - P(B|g, a, \gamma)r] \\ & + P(b|a, \gamma) \mathbb{1}\{\rho_0 \geq \tilde{\rho}_b(\gamma)\} [p - P(B|b, a, \gamma)r]. \end{aligned} \tag{4.9}$$

The following proposition characterizes the firm's optimal decision in Stage 2:

Proposition 4.3. Optimal Influence and Signal Quality Configuration with Refund

Suppose the firm will issue a refund $r \in (0, p]$ for each product return. Then, at Stage 2,

(i) when $\rho_0 < \tilde{k}$,

the firm will set the influence and signal quality level at

$$a_r^* = \frac{\rho_0 \tilde{k}}{\tilde{k} \bar{\rho}_0},$$

and

$$\gamma_r^* = 1,$$

the customers will purchase if a positive signal is observed; the firm's expected profit is

$$\Pi_r^* = \rho_0 v;$$

(ii) when $\rho_0 \geq \tilde{k}$, the firm may set any $\gamma_r^* \in [0.5, \frac{\rho_0 \tilde{k}}{\rho_0 \tilde{k} + \bar{\rho}_0 \tilde{k}}]$ and use any influence level $a_r^* \in [0, 1]$ and regardless, the customers will always make a purchase; the firm's expected profit is $\Pi^* = p - \bar{\rho}_0 r$.

Under an exogenous refund policy, the firm's decision structure mirrors that of the no-refund scenario. However, as established in Corollary 4.1, a higher exogenous refund amount leads to an increase in the optimal influence level.

Corollary 4.1. Impact of Refund on Optimal Influence Level

The optimal influence level a_r^* , is increasing in r within its application zone.

4.5. Joint Design of Return Policy, Baseline Signal Informativeness, and Influence Configuration

In this section, the optimal joint design of the return policy and information influence configuration is examined. Specifically, the firm is assumed to have full control over the key strategic levers: the refund level r , the influence intensity a , and the baseline signal quality γ . These

instruments are jointly optimized to maximize the firm's expected profit by balancing pre-purchase persuasion, post-purchase flexibility, and the initial credibility of information signals. This comprehensive approach enables a unified strategy that integrates both marketing influence and operational policy under conditions of fit uncertainty.

The following proposition characterizes the firm's optimal strategy:

Proposition 4.4. Optimal Influence, Baseline Signal Quality, and Continuous Refund

Suppose that the firm may freely choose its refund level r within $[0, p]$,

(i) when $\rho_0 < k$, the firm has two options that are equivalent in terms of expected profit:

(a) it can either offer a partial refund

$$r^* = \frac{p - \rho_0 v}{\bar{\rho}_0},$$

and use any influence and signal quality level;

(b) or, offer no refund $r^* = 0$, and set

$$a^* = \frac{\rho_0 \bar{k}}{k \bar{\rho}_0},$$

and

$$\gamma^* = 1,$$

both of which will earn the same expected profit of $\Pi^* = \rho_0 v$.

(ii) when $\rho_0 \geq k$,

$$r^* = 0,$$

and the firm may set any $\gamma^* \in [0.5, \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}]$ and use any influence level $a^* \in [0, 1]$, and regardless, the customers will always make a purchase; the firm's expected profit is $\Pi^* = p$.

Proposition 4.4 shows that, in a world with zero implementation frictions, a seller can achieve the same expected profit in two diametrically opposed ways: (i) promise a customized partial refund

that perfectly offsets consumers' downside risk, or (ii) offer *no* refund but upgrade the information environment to full accuracy and then influence it just enough to trigger purchase. Both routes leave buyers with identical expected utility and therefore generate the same willingness to pay, so the firm is indifferent. This indifference is not a weakness of the model; it is a deliberate cost-free benchmark. Essentially, this comparison between the two scenarios boils down to analyzing the costs associated with different tools. However, once even modest real-world frictions – marketing outlays, reverse-logistics costs, or credibility penalties – are introduced, the tie between pre-purchase persuasion and post-purchase reassurance disappears and one lever becomes strictly superior to the other for a given context. The proposition therefore serves as a “knife-edge” reference point against which more realistic cost structures can be compared.

4.5.1 Binary Refund Case

In the specific scenario where the firm is restricted to a binary refund choice, i.e., either offering no refund ($r = 0$) or a full refund ($r = p$), the structure of the optimal strategy simplifies considerably. In this setting, offering a full refund removes any risk borne by the customer, making any investment in information disclosure or customer influence unnecessary since the refund fully guarantees customer satisfaction regardless of fit realization. Therefore, the firm's profit-maximizing approach is to forgo offering refunds altogether and instead rely entirely on its ability to configure the informativeness of the signal and apply influence to stimulate purchases while maintaining customer risk.

When a full refund ($r = p$) is offered, the customer bears no risk regardless of product fit. In such a case, customers always purchase whenever $v \geq p$, rendering any effort to influence customer perception or improve information informativeness, redundant. Consequently, offering a full refund provides no incremental gain in profit beyond guaranteeing purchases that would already occur under favourable valuations.

Conversely, under no refund ($r = 0$), customer purchasing decisions depend on their updated beliefs about product fit, which the firm can actively shape through configuring γ and a . Given $\rho_0 < k$, the firm can optimally calibrate a and γ to induce purchases only when the posterior fit

belief exceeds the price threshold, while preserving profit margins by not refunding unsatisfied customers. The optimal influence level derived from the indifference condition is $a^* = \frac{\rho_0 \bar{k}}{k \rho_0}$. Since increasing γ improves the informativeness of the signal and enhances purchase conversion, it is optimal to set $\gamma^* = 1$. The resulting expected profit equals $\Pi^* = \rho_0 v$. Consistent with Proposition 4.1, for $\rho_0 \geq k$, no intervention is necessary as customers will purchase even without additional influence.

The resulting optimal strategies and profit implications under binary refund options are formally summarized in Corollary 4.2.

Corollary 4.2. Binary Refund

Under binary refund choice, $r \in \{0, p\}$, the optimal return policy is to offer no refund, $r^* = 0$, and the optimal influence configuration follows the solution characterized in Proposition 4.4 (i)(b) and (ii).

Similar to Section 4.5, the binary-refund result should be interpreted as a deliberately *friction-free* benchmark. No restocking fees, reverse-logistics expenses, or hassle costs for either party have been modelled. In practice, those frictions would govern the choice, though the qualitative insight that pre-purchase influence and post-purchase refunds act as alternative risk-mitigation levers remains intact once realistic costs are added; only the precise cut-off between the two instruments moves.

4.6. Summary

In this chapter, baseline informativeness was treated as a design variable and its interaction with influence and refunds was analyzed. With an exogenous influence level and no refunds, purchase thresholds were shown to depend on the information disclosure accuracy, and an exogenous upper bound on baseline informativeness was identified under favourable priors. That is, when $\rho_0 \geq k$, default information accuracy must not exceed $\gamma^{\max}$, since excess precision can induce sales-detering negative signals without any offsetting gains. Under joint control of baseline

informativeness and influence intensity, and in the absence of refunds, baseline informativeness was optimally set to its maximum for beliefs below the price–value threshold, and the accuracy upper bound for optimistic prior beliefs continued to apply.

Under exogenous refunds, the solution structure was preserved, while the optimal influence intensity was shown to increase with the refund rate. In the fully endogenous setting, however, two profit-equivalent routes emerged for relatively low prior beliefs: either a tailored partial refund or no refund combined with maximum baseline accuracy prior to influence, indicating strategic flexibility. Under a binary refund restriction, relying on information configuration without refunds was favoured over offering a full refund. Taken together, these results formalize when signal credibility should be capped or maximized *ex ante*, when persuasion should be intensified, and how refunds substitute for information tools. The transition to a framework that permits analysis of pricing and competition is presented in [Chapter 5](#).

Chapter 5

Strategic Influence and Pricing

Decisions with Post-Purchase

Considerations in Duopoly Markets

5.1. Overview

Competition for experience goods rarely hinges on price alone. When customers are uncertain about *fit*, any attribute that nudges their posterior beliefs, brand reputation, online reviews, influencer endorsements, or a money-back guarantee becomes strategically significant. This chapter extends the thesis-wide framework to a Hotelling-style duopoly in which two retailers simultaneously configure *prices* and *influence strategies*. The interaction of these instruments is examined in a setting where each customer observes both firms' signals, updates match probabilities accordingly, and may later exploit the posted refund.

Competition involving differentiated products is often shaped by spatial dynamics, where each firm's position in an attribute space determines its relative advantage over competitors. One of the key insights from spatial competition models, as first conceptualized by [Hotelling \(1929\)](#), is that differentiation can soften price competition. Firms positioned closer to their customers'

preferences enjoy a degree of loyalty that allows for modest price increases without substantial loss of demand. However, too little differentiation intensifies price undercutting as firms aggressively compete over nearly identical customer bases, while too much differentiation results in segmented markets where competition weakens as firms dominate distinct niches. In other words, Hotelling's linear city illustrates how mild differentiation softens price rivalry, whereas excessive similarity triggers predatory pricing.

Typical location models, like the Hotelling framework, traditionally describe consumer preferences based on physical proximity to retailers. However, these models have been extended to consider various factors beyond mere location, such as brand preference and service quality, with "distance" reflecting a form of differentiation rather than just geographic separation (e.g., [Li et al., 2019](#), [Bar-Gill, 2020](#)). Re-interpreting "location" as *mismatch probability* translates that insight to *fit uncertainty*. This leads to the intuition that similarly, each retailer would enjoy a protected region of consumers for whom its product is the likely match, and competes only for the *border segment*. This emphasizes the centrality of customer influence strategies in shaping competitive outcomes. In omnichannel information environments, retailers curate entire information ecosystems – sponsored unboxing videos, algorithm-driven recommendations, loyalty emails – long before the price tag is inspected by the customer. Differences in framing, visibility, or content quality cause heterogeneity in perceived fit across customer segments. Competition between differentiated products often demonstrates how informational asymmetries can sustain distinct consumer segments despite functional similarities across offerings. These realities motivate a model in which influence intensity is endogenous, costly, and potentially asymmetric. Building on Chapters [3](#) and [4](#), a duopoly market is examined. Four main questions guide the analysis:

- **Cross-Seller Information and Market Sorting:** How does the availability of cross-seller information shape consumer sorting and the equilibrium mix of prices and persuasion in a duopoly?
- **Competitive Strategy Design:** What are the optimal pricing and influence strategies for retailers in competitive markets?

- **Prior Beliefs and Strategic Interaction:** How do the interactions of influence and pricing strategies vary under different consumer search and belief distributions?
- **Return Policies in Competitive Contexts:** What is the impact of return policy on influence design and pricing decisions in shaping competitive outcomes?

By modelling these interdependencies, this phase contributes to a comprehensive understanding of how firms can strategically align influence efforts, information precision, and operational policies in increasingly complex digital markets.

5.2. Analytical Setup

This section presents an extension of spatial competition models, advancing the thesis’s progression from monopoly benchmarks to competitive duopoly settings. The analysis focuses on simultaneous influence competition, where both retailers strategically determine their influence intensities and prices to maximize profits within a Nash equilibrium framework.

In this framework, the model adopts a probabilistic representation of product fit, wherein consumers’ positions along the unit line reflect their subjective likelihood of perceiving a product as well-matched to their preferences, both in absolute terms and relative to competing alternatives. This perspective aligns with competitive markets where consumers evaluate offerings based on perceived fit rather than purely objective attributes or physical proximity. For instance, two retailers may emphasize distinct value propositions or product features, leading consumers to self-position along a “fit expectation spectrum” according to their personal preferences. Such modelling captures the reality that purchasing decisions are shaped by subjective perceptions influenced by marketing efforts and information exposure, rather than by inherent product characteristics alone. This approach facilitates tractable analysis of complex competitive environments by linking consumer heterogeneity to strategic seller behaviour.

5.2.1 Sellers

The market consists of two competing sellers, Seller 1 and Seller 2, offering Product 1 and Product 2, respectively. The two products are assumed to be vertically identical, implying equivalence in terms of objective quality and functionality. However, consumers may form heterogeneous expectations regarding product fit based on each seller's marketing activities and informational presence. The selling prices are denoted by p_1 for Product 1 and p_2 for Product 2. It is assumed that the marketing efforts of the two sellers operate independently, with no cross-seller spillover effects. The actual product fit from either seller, denoted by $F_i \in \{G, B\}$, remains unknown to all players, and is only revealed after the purchase.

5.2.2 Customers

If the product ultimately proves to be a good match ($F_i = G$), the customer assigns it a valuation of $v > 0$; otherwise ($F_i = B$), the valuation is zero. Consumers are distributed along a unit interval, representing the spectrum of their optimism regarding the suitability of each product. The market is assumed to be fully covered. It is common knowledge that a fraction $\rho_0 \in [0, 1]$ of consumers associates Product 2 with a favourable fit, such that $P(F_2 = G) = \rho_0$ and $P(F_2 = B) = \bar{\rho}_0 = 1 - \rho_0$, while the remaining $1 - \rho_0$ proportion of consumers considers Product 1 a match. Alternatively, the location of a consumer on this line reflects their subjective probability of fit for the competing products: a customer positioned at $\rho_0 \in [0, 1]$ assigns a probability of $\bar{\rho}_0$ to Product 1 being a fit, and ρ_0 to Product 2. In the absence of additional information, customers base their prior beliefs solely on this intra-seller fit optimism spectrum to form expectations and make purchasing decisions. Each customer selects the retailer that maximizes their expected utility.

5.2.3 Customer Influence

In modern competitive environments, retailers no longer compete solely on intrinsic product attributes or prices but actively shape the informational environments that drive consumer beliefs.

Marketing efforts increasingly rely on diverse information provision mechanisms, including online reviews, social media engagement, influencer campaigns, targeted advertisements, and recommendation systems. These elements shape consumers' initial perceptions of fit and influence competitive dynamics even before actual purchase decisions are made.

Thus, the firms are assumed to have access to tools that assist customers in estimating the true fit of the products. These tools represent various forms of information intermediaries that generate informative, yet imperfect, signals about product fit. It is assumed that a certain degree of distortion may be introduced by the firms to influence the signals – signals that would not have emerged autonomously in the absence of such intervention. The signal produced by the influence tool is modelled as a binary variable $S_i \in \{g, b\}$, where $S_i = g$ and $S_i = b$ indicate a favourable and an unfavourable fit with Seller i , respectively.

The baseline quality of this signal is denoted by $\gamma_i \in (0.5, 1)$, reflecting its precision in the absence of any firm-imposed influence, such that

$$P(S_i = g|F_i = G) = P(S_i = b|F_i = B) = \gamma_i. \quad (5.1)$$

In particular, $\gamma_i \rightarrow 1$ represents a nearly perfect baseline signal that almost always yields an accurate prediction of product fit, whereas $\gamma_i \rightarrow 0.5$ reflects a non-informative signal that conveys minimal information beyond the prior fit belief. It is further assumed that the inherent signal precision is identical for both sellers, i.e., $\gamma_1 = \gamma_2$. This reflects situations where both firms operate on the same information provision platform or mechanism – for example, a shared online review platform, a common third-party evaluator, or a standardized influencer network – thereby subjecting both to the same baseline signal quality. The influence cost c is normalized to zero for both firms.

Given the baseline signal quality γ , each seller may set an influence level $a_i \in [0, 1]$, which will distort the signal distribution to

$$P(S_i = g|F_i = G, a_i) = \gamma + a_i\bar{\gamma} \quad \text{and} \quad P(S_i = g|F_i = B, a_i) = a_i\gamma + \bar{\gamma}. \quad (5.2)$$

While consumer search frictions have been extensively studied and shown to significantly affect pricing dynamics – sometimes even collapsing competition entirely under minimal search costs (Diamond, 1971) – this dissertation adopts a framework that abstracts away from such search costs. This modelling choice is motivated by the increasing transparency and accessibility of information in modern digital marketplaces. In today’s highly integrated retail environments, consumers typically face minimal barriers to information access due to the ubiquity of on-line reviews, influencer content, recommendation algorithms, and platform-mediated product comparisons. Unlike traditional markets where consumers may have encountered substantial effort or cost to discover alternative offerings, modern digital platforms effectively aggregate and disseminate product information, allowing consumers to effortlessly observe key product attributes, pricing, and fit-related signals across multiple sellers. The presence of these sophisticated information provision mechanisms – often managed directly or indirectly by firms through strategic influence – serves to substantially mitigate or eliminate the very search frictions that would otherwise generate the pricing power described in classical search theory.

Therefore, while acknowledging that search costs have important implications in some contexts, for the class of digital influence settings characterized by proactive firm-driven information provision and pervasive consumer access to cross-seller signals, assuming negligible search costs offers a reasonable approximation. Accordingly, it is assumed that consumers observe both the signal S_i and the influence level a_i for each retailer, allowing them to update their subjective beliefs about product fit prior to making purchasing decisions.

Employing Bayes’ rule, as is standard in signalling literature, consumers update their posterior beliefs after observing the signals. Specifically, the posterior probabilities of Product 1 being a good fit, conditional on receiving signals $S_1 = g$ and $S_1 = b$ under influence level a_1 , are given by

$$\begin{aligned}
P(F_1 = G|S_1 = g, a_1) &= \frac{P(S_1 = g|F_1 = G, a_1)P(F_1 = G)}{P(S_1 = g|a_1)} \\
&= \frac{P(S_1 = g|F_1 = G, a_1)P(F_1 = G)}{P(S_1 = g|F_1 = G, a_1)P(F_1 = G) + P(S_1 = g|F_1 = B, a_1)P(F_1 = B)} = \\
&\quad \frac{(\gamma + a_1\bar{\gamma})\bar{\rho}_0}{(\gamma + a_1\bar{\gamma})\bar{\rho}_0 + (\bar{\gamma} + a_1\gamma)\rho_0}, \tag{5.3}
\end{aligned}$$

and

$$\begin{aligned}
P(F_1 = G|S_1 = b) &= \frac{P(S_1 = b|F_1 = G, a_1)P(F_1 = G)}{P(S_1 = b)} \\
&= \frac{P(S_1 = b|F_1 = G, a_1)P(F_1 = G)}{P(S_1 = b|F_1 = G, a_1)P(F_1 = G) + P(S_1 = b|F_1 = B, a_1)P(F_1 = B)} = \\
&\quad \frac{\bar{\gamma}\bar{\rho}_0}{\bar{\gamma}\bar{\rho}_0 + \gamma\rho_0}; \tag{5.4}
\end{aligned}$$

also, $P(F_1 = B|S_1 = g, a_1) = 1 - P(F_1 = G|S_1 = g, a_1)$, and $P(F_1 = B|S_1 = b) = 1 - P(F_1 = G|S_1 = b)$.

Similarly, for Seller 2,

$$P(F_2 = G|S_2 = g, a_2) = \frac{(\gamma + a_2\bar{\gamma})\rho_0}{(\gamma + a_2\bar{\gamma})\rho_0 + (\bar{\gamma} + a_2\gamma)\bar{\rho}_0}, \tag{5.5}$$

and

$$P(F_2 = G|S_2 = b) = \frac{\bar{\gamma}\rho_0}{\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0}; \tag{5.6}$$

also, $P(F_2 = B|S_2 = g, a_2) = 1 - P(F_2 = G|S_2 = g, a_2)$, and $P(F_2 = B|S_2 = b) = 1 - P(F_2 = G|S_2 = b)$.

5.3. Benchmark Case: Exogenous Prices

In this section, the analysis begins with the case of exogenous prices, which serves as a benchmark for subsequent developments. Under this setting, both firms determine their investment levels in reviewer affiliation or customer influence, while prices remain exogenously given.

5.3.1 Timing of the Game

The sequence of the game can be described as follows:

Stage 1 (Prior Formation): Nature determines if a product will fit a customer, i.e., $F_i = \{G, B\}$ with probability $\{\rho_0, \bar{\rho}_0\}$ for Product 2 and $\{\bar{\rho}_0, \rho_0\}$ for Product 1, respectively. The outcome F_i is not visible to the sellers or the customer until Stage 3.

Stage 2 (Influence Configuration): Seller 1 and 2 set the influence level a_i with selling price p_i given. Signals $S_i \in \{g, b\}$ are produced based on Equation 5.2.

Stage 3 (Purchase Decision): The customer observes both signals S_i , forms posterior belief on F_i , and chooses the seller to buy from, by comparing the expected utility of purchasing each product.

5.3.2 Equilibrium Outcomes

By backward induction, at Stage 3, the customer will buy from Seller i if the expected utility of Product i is greater than that of the other Product.

Even when products provide comparable end-use value, heterogeneity in information framing, content visibility, and perceived credibility across different channels results in variation in consumer beliefs. For instance, certain consumer segments may perceive a product as highly compatible based on positive influencer endorsements or online reviews, while others – exposed to different informational cues – may remain skeptical. This variability alters the competitive balance between firms, introducing an additional layer of strategic complexity that extends beyond simple price competition.

Several real-world product categories exemplify these dynamics. In the high-end smartphone market, competition between Apple’s iPhones (iOS) and flagship Android devices from manufacturers like Samsung and Google reflects differentiated value propositions targeting distinct consumer segments. Apple’s iOS ecosystem appeals to users who value integration, brand prestige, security, and seamless experience, while Android offers customization, hardware variety, and innovative features. While both platforms satisfy similar core needs, the probability that a consumer finds a good fit with one platform typically complements their likelihood of fitting with the other. Similarly, in the gaming console market, Microsoft’s Xbox and Sony’s PlayStation offer distinct ecosystems and exclusive features, targeting overlapping but differentiated consumer bases. In other words, despite their unique offerings, these products ultimately serve the same valuation goal: meeting consumers’ gaming and entertainment needs. In both markets, differentiation generates inherent regions of competitive strength and border regions of fierce

competition.

Thus, in the absence of any signal, the customer will be indifferent between the two sellers when $(1 - \rho_0)v - p_1 = \rho_0v - p_2$. Denote the indifference price-to-value ratio $k = \frac{1}{2} - \frac{p_1 - p_2}{2v}$. When $\rho_0 < k$, the customer will make a purchase from seller 2 if and only if there is an asymmetrical positive signal inclination towards Product 2, indicating a stronger tendency for favourable perceptions and expectations about product fit compared to Product 1. Specifically, when $S_1 = b$ and $S_2 = g$, the customer will make a purchase from Seller 2 if

$$P(F_2 = G \mid S_2 = g, a_2)v - p_2 - P(F_1 = G \mid S_1 = b)v + p_1 \geq 0. \quad (5.7)$$

It is assumed that there exists $0 \leq \rho_1 \leq k$ such that the inequality holds whenever $\rho_0 \geq \rho_1$, where ρ_1 is defined as

$$\rho_1(p_1, p_2, a_2) = \frac{-\sqrt{A_1} - B_1}{C_1}. \quad (5.8)$$

The explicit forms of A_1 , B_1 , and C_1 are provided in Appendix C (Eqs. C.1 – C.3) for completeness.

When $\rho_0 > k$, the customer will make a purchase from Seller 1 if and only if there is an asymmetrical positive signal inclination towards Product 1, indicating a stronger tendency for favourable perceptions and expectations about product fit compared to Product 2. Specifically, when $S_1 = g$ and $S_2 = b$, the customer will make a purchase from Seller 1 if and only if

$$P(F_1 = G \mid S_1 = g, a_1)v - p_1 - P(F_2 = G \mid S_2 = b)v + p_2 \geq 0.$$

It is assumed that there exists $k \leq \rho_2 \leq 1$ such that this inequality holds when $\rho_0 \leq \rho_2$, where ρ_2 is defined as:

$$\rho_2(p_1, p_2, a_1) = \frac{-\sqrt{A_2} + B_2}{C_2}. \quad (5.9)$$

The explicit forms of A_2 , B_2 , and C_2 are provided in Appendix C (Eqs. C.4 – C.6) for completeness.

In sum, ρ_1 and ρ_2 represent the threshold for choosing Product 2 and 1 over the other respectively, while price-to-valuation ratio is in favour of the other seller (the stronger competitor). Further, it can be shown that

Lemma 5.1. *Suppose that the consumer valuation is normalized to $v = 1$. Then the product choice thresholds satisfy $\rho_1 \leq k \leq \rho_2$.*

The customer's purchasing decision follows:

- if $\rho_0 < \rho_1$, a customer will always purchase Product 1;
- if $\rho_1 \leq \rho_0 < k$,
 - (a) a customer will purchase Product 2 only if a positive signal from Seller 2 and a negative signal from Seller 1 is observed ($S_1 = b$ and $S_2 = g$);
 - (b) otherwise, the customer buys Product 1.
- if $k \leq \rho_0 \leq \rho_2$,
 - (a) a customer will purchase Product 1 only if a positive signal from Seller 1 and a negative signal from Seller 2 is observed ($S_1 = g$ and $S_2 = b$);
 - (b) otherwise, the customer buys Product 2.
- if $\rho_0 > \rho_2$, a customer will always purchase Product 2.

As observed in equations 5.8 and 5.9, the secure market boundaries for each seller – denoted by ρ_1 and ρ_2 – are not functions of their own influence levels, but instead depend solely on the influence exerted by their competitor. This asymmetry emerges from the structure of the influence mechanism itself, which is designed to increase the probability of generating a favourable signal but does not actively suppress the generation of unfavourable signals. In other words, a seller cannot directly reduce the odds of a customer receiving a bad signal; they can only increase the likelihood that their own product is viewed more positively by the influence tool.

This feature reflects a common characteristic of real-world persuasive systems, such as digital marketing or influencer endorsements: marketing efforts are more effective at boosting positive

sentiment or awareness than at preventing exposure to negative impressions. For example, a brand may invest heavily in targeted campaigns or influencer collaborations to ensure more customers encounter enthusiastic reviews. However, this does not prevent the audience from encountering criticisms or forming skepticism. Accordingly, a seller’s persuasive actions end up shaping the perception of the *competitor’s* purchase certainty boundary, not their own.

This strategic externality – wherein one seller’s influence alters the protected market size of the rival – underscores a key tension in competitive persuasion: actions intended to increase one’s signal effectiveness may paradoxically erode the informativeness of signals market-wide, influencing who benefits from default (certainty-based) customer decisions. This insight also lays the groundwork for the interpretation of Figure 5.1 and 5.2, which visualizes how these interdependent thresholds evolve with varying influence intensities.




Figure 5.1: Seller 1’s secure market boundary ρ_1 vs. Seller 2’s Influence Level a_2
 $(p_1 = 0.7, p_2 = 0.6, v = 1, \gamma = 0.7)$

Figure 5.1 and 5.2 illustrate how the boundaries of the protected market, where consumers opt to purchase regardless of the realization of the signal, are influenced by the competitor’s level of influence effort. Taken together, these thresholds, to which customer choices have been mapped in Lemma 5.1, are shown to vary systematically with changes in the opposing seller’s strategy.

Somewhat counterintuitively, the protected market expands as the competing seller’s influence level increases. This effect arises from the public observability of influence effort: as the




Figure 5.2: Seller 2's secure market boundary ρ_2 vs. Seller 1's Influence Level a_1
 $(p_1 = 0.7, p_2 = 0.6, v = 1, \gamma = 0.7)$

influence level of the competitor rises, the chance of receiving a favourable signal (e.g., a persuasive review) improves. However, this comes at the cost of reduced credibility: consumers become more skeptical, and the posterior probability of a match conditional on a good signal declines. As a result, even though good signals become more frequent, they are less trusted, leading to an expansion of the region on the low-fit end of the unit line where consumers decline to purchase despite receiving a favourable signal. Consequently, the certainty threshold shifts outward, enlarging the protected market in favour of the focal seller.




Figure 5.3: Seller 1's secure market boundary vs Baseline Signal Quality
 $(p_1 = 0.7, p_2 = 0.6, v = 1, a_2 = 0.5)$




Figure 5.4: Seller 2's secure market boundary vs Baseline Signal Quality
 $(p_1 = 0.7, p_2 = 0.6, v = 1, a_2 = 0.5)$

Figures 5.3 and 5.4 illustrate how the certainty purchase thresholds for the two firms respond to changes in the baseline signal quality parameter, γ . Specifically, the certainty threshold for Firm 1 is shown to decrease with higher γ , whereas that of Firm 2 increases. This implies that the protected market region, defined as the set of consumer types who purchase with certainty, regardless of signal realization, contracts for both sellers as γ rises.

This shift indicates that as the baseline informativeness increases, the overall signal environment becomes more reliable, and customers rely more heavily on the signals themselves rather than defaulting to their baseline valuation. Consequently, fewer consumers are located in the certainty regions where signal realization is irrelevant, and more are moved into the interior of the figurative unit line, where decisions depend on the signal outcome. In this way, higher baseline signal quality weakens the gravitational pull of the certainty regions and heightens the role of persuasion.

From a strategic perspective, this transition amplifies the power and stakes of influence efforts. As these zones shrink, firms are no longer protected by loyalist-type consumers whose actions are unaffected by competitor persuasion. Instead, a larger share of the market becomes contestable through signal shaping, intensifying the competitive pressure to invest in influence. Practically, this mirrors the dynamics observed in digital marketplaces: as consumer reviews, ratings, or

third-party information become more accurate or trusted (that is, γ increases), the importance of swaying the undecided middle grows, and firms must compete more aggressively to capture it. Thus, improvements in signal quality diminish both the static advantage and expand the domain where dynamic persuasion can be leveraged.

To complement the analytical results, representative numerical examples are presented in Figures 5.5, 5.6, and 5.7, which illustrate the behaviour of the purchase thresholds ρ_1 , ρ_2 , and the intermediate threshold k under varying model parameters. These plots serve to reinforce the structural properties established in Lemma 5.1, and confirm that the condition $0 \leq \rho_1 \leq k \leq \rho_2 \leq 1$ is preserved across the tested parameter configurations.




Figure 5.5: Purchase thresholds ρ_1 , ρ_2 , and k as a function of a_1
($a_2 = 0.5$, $\Delta p = 0.1$)

In Figure 5.5, it can be observed that the threshold ρ_2 decreases monotonically with increasing a_1 , while ρ_1 and k remain unchanged. In contrast, Figure 5.6 shows that ρ_1 increases with a_2 , whereas ρ_2 and k exhibit no dependence on the focal seller's own influence level. These trends are consistent with the model's analytical structure, where each seller's secure market boundary is determined solely by the competitor's influence effort. As a result, the protected market size for each seller is expanded as the competitor intensifies their persuasion effort, while remaining unaffected by their own actions. This asymmetry reflects the core feature of the




Figure 5.6: Purchase thresholds ρ_1 , ρ_2 , and k as a function of a_2
($a_1 = 0.4$, $\Delta p = 0.1$)

influence mechanism, wherein persuasion affects the informativeness of signals but does not suppress negative outcomes directly.




Figure 5.7: Purchase thresholds ρ_1 , ρ_2 , and k as a function of $\Delta p = p_1 - p_2$
($a_1 = 0.4$, $a_2 = 0.5$)

In Figure 5.7, the dependence of the thresholds on the price difference Δp is exemplified. As the price gap increases, all three thresholds decline smoothly, maintaining the inequality

$0 \leq \rho_1 \leq k \leq \rho_2 \leq 1$ throughout the range. At the extremes where $|\Delta p| = 1$, all thresholds converge toward the boundary of the unit interval. This behaviour is consistent with the economic interpretation that extreme pricing renders persuasion ineffective, as price signals dominate consumer decisions and marginal influence becomes unresponsive (Linnemer, 2002).

In stage 2, the sellers decide their level of influence. Their expected profit will be as follows:

$$\begin{aligned}\Pi_1(a_1, a_2) &= \mathbb{1}\{\rho_0 < \rho_1\}p_1 \\ &\quad + (1 - P(S_1 = b)P(S_2 = g | a_2)) \mathbb{1}\{\rho_1 \leq \rho_0 < k\}p_1 \\ &\quad + P(S_1 = g | a_1)P(S_2 = b) \mathbb{1}\{k \leq \rho_0 < \rho_2\}p_1 - c_1\end{aligned}\tag{5.10}$$

$$\begin{aligned}\Pi_2(a_1, a_2) &= \mathbb{1}\{\rho_0 \geq \rho_2\}p_2 \\ &\quad + (1 - P(S_1 = g | a_1)P(S_2 = b)) \mathbb{1}\{k \leq \rho_0 < \rho_2\}p_2 \\ &\quad + P(S_1 = b)P(S_2 = g | a_2) \mathbb{1}\{\rho_1 \leq \rho_0 < k\}p_2 - c_2\end{aligned}\tag{5.11}$$

In this formulation, the first component is associated with the payoff arising from the protected market regions of each seller, where consumer choices are dominated by the respective firm. The second component is attributed to the insecure regions, in which customers may be lost due to competitive signal profiles. The third component corresponds to the regions primarily influenced by the opposing firm, where customers may still be captured as a result of a favorable signal imbalance. The cost parameters c_1 and c_2 are used to denote the respective customer influence expenditures incurred by sellers 1 and 2. For the purpose of simplifying the analysis, the cost terms c_1 and c_2 are normalized to zero within the payoff functions. This normalization does not influence the strategic decisions or the relative profitability of alternative actions, as the cost components would introduce only additive constants to the profit levels. In particular, their inclusion would affect only the feasibility conditions or shift the thresholds, without altering the fundamental structure of the decision-making process. As such, the core dynamics of competition and signalling are isolated, with any cost-related adjustments deferred as a post-optimization step.

The following proposition characterizes the equilibrium outcomes of the firms at this stage.

Proposition 5.1. Optimal Influence Strategies (No Refund)

Suppose that the sellers offer customers no refund and the consumer valuation is normalized to $v = 1$. Then,

- (i) when $\rho_0 < \rho_1(a_2 = 0)$, neither firm has an incentive to exert influence,

$$a_1^* = a_2^* = 0,$$

and the expected profits are $\Pi_1 = p_1$ and $\Pi_2 = 0$.

- (ii) when $\rho_1(a_2 = 0) \leq \rho_0 < k$, the optimal strategies are $a_1^* = 0$ and a_2^* as defined in closed form in Appendix C (see C.34). The corresponding expected payoffs are $\Pi_1(a_2^*)$ and $\Pi_2(a_2^*)$ (see C.35 and C.36).

- (iii) when $k \leq \rho_0 \leq \rho_2(a_1 = 0)$, the optimal strategies $a_2^* = 0$ and a_1^* as defined in closed form in Appendix C (see C.40). The corresponding expected payoffs are $\Pi_1(a_1^*)$ and $\Pi_2(a_1^*)$ (see C.41 and C.42).

- (iv) when $\rho_2(a_1 = 0) > \rho_0$, neither firm has an incentive to exert influence,

$$a_1^* = a_2^* = 0,$$

and the expected profits are $\Pi_1 = 0$ and $\Pi_2 = p_2$.

Proposition 5.1 provides several key insights into how sellers strategically deploy influence in a competitive market based on customer-product fit. These results highlight several structural features of strategic influence in competitive settings, with two particularly prominent insights emerging. First, the presence of kinks and non-differentiable points in the influence and profit functions reflects underlying threshold effects: sellers should anticipate abrupt shifts in strategic value as consumers approach key segment boundaries. It is not only the level of customer fit that matters, but also its volatility; influence becomes especially valuable in regions where small changes in perceived fit can lead to major shifts in market control. Second, the interdependence

of influence decisions is not fixed: in some regions, firms' investments in influence act as strategic substitutes, while in others, they may be mutually reinforcing.

Figures 5.8 and 5.9 illustrate these insights from both ends of the misfit spectrum: Figure 5.8 focuses on regions where Seller 1 dominates (i.e., lower values of ρ_0), while Figure 5.9 captures the opposite end, where Seller 2 gains a relative advantage as ρ_0 increases.




Figure 5.8: Optimal Influence Level of Seller 2, a_2^* , and expected profits vs. misfit level, ρ_0 ($\gamma = 0.7$, $v = 1$, $p_1 = 0.7$, and $p_2 = 0.6$).

- **Influence in Protected Regions:** When consumers fall within a seller's protected region, defined by spatial advantage in product fit (i.e., $\rho_0 < \rho_1$ for Seller 1 and $\rho_0 > \rho_2$ for Seller 2), influence investment is unnecessary. In these zones, customers are effectively “loyal”, and both sellers refrain from exerting influence, as it yields no marginal benefit.
- **Influence in Endangered or Competitive Regions:** In the intermediate region, where neither seller has a clear fit advantage and consumers are susceptible to switching, the non-dominant seller incrementally increases their influence level. When the misfit is extreme (i.e., far from their preferred segment), influence starts at zero and rises with the increasing probability of a favourable match. This continues until influence reaches its maximum




Figure 5.9: Optimal Influence Level of Seller 1, a_1^* , and expected profits vs. misfit level, ρ_0 ($\gamma = 0.7$, $v = 1$, $p_1 = 0.7$, and $p_2 = 0.6$)

allowable level. Once the non-dominant seller becomes dominant (i.e., as fit and price advantages align), further influence becomes unnecessary and ineffective. Meanwhile, the previously dominant seller's influence gains strategic importance.

- **Influence Follows Fit Potential:** For each seller, the optimal influence level a_i^* tends to increase with their relative fit advantage. As the probability of matching with customers declines, influence levels fall, eventually reaching zero at the boundary of the competitor's protected region. This reflects the alignment between influence incentives and market opportunity.

In summary, the role of influence is highly dynamic and contingent on spatial positioning. Sellers adjust their influence levels in response to competitive conditions and customer fit probabilities. Influence is most effective in marginal regions and declines when a seller either fully dominates or has no competitive leverage.

5.4. Customer Influence and Pricing

In this section, the benchmark model is extended by allowing prices to be determined endogenously by the firms. Whereas prices were previously held fixed to isolate the role of customer influence, they are now treated as strategic variables to examine how pricing decisions interact with spatial competition and influence efforts. By endogenizing prices, the joint effects of influence and pricing on market outcomes are captured, offering a more complete understanding of competitive dynamics when both levers are under firm control. Accordingly, the timing of the game is modified in Stage 2 as follows, relative to Section 5.3.1:

Stage 2 (Pricing and Influence Configuration): Sellers 1 and 2 simultaneously choose their influence levels a_i and selling prices p_i ; Signals $S_i \in \{g, b\}$ are produced based on Equation 5.2.

5.4.1 Equilibrium Outcomes

By backward induction, the purchase thresholds, customer decision, and profit functions of the firms follow the structure established in Section 5.3.2. Specifically, ρ_1 (Equation 5.8) and ρ_2 (Equation 5.9) denote the secure market boundaries for Products 1 and 2, respectively.

In stage 2, the sellers decide their selling prices and level of influence.

The following proposition characterizes the equilibrium outcomes at this stage.

Proposition 5.2. Optimal Influence and Pricing Strategies

Suppose that the sellers offer customers no refund and the consumer valuation is normalized to $v = 1$. Then, the optimal prices are $p_1^* = p_2^* = 1$, while the optimal influence level for Seller 1 and 2 are given by a_1^* and a_2^* , respectively (see C.61 and C.52 for the closed-form solutions).

The corresponding expected profits are $\Pi_1(p_1^*, a_2^*)$ and $\Pi_2(p_2^*, a_2^*)$ (see C.62 and C.63).

Proposition 5.2 reveals several strategic insights into seller behaviour under endogenous pricing and influence decisions. It shows that influence and pricing interact strategically: dominance in fit allows price rigidity, while influence offers a lever for underdogs to remain relevant.

Compared to Proposition 5.1, the edge cases (parts (i) and (iv)) – where the dominant firm effectively enjoys a local monopoly, enabling it to charge the highest price and earn the maximum profit, while the competitor’s price becomes irrelevant – no longer emerge under endogenous pricing. This outcome aligns with standard intuitions from spatial or perceptual differentiation models. Many real-world markets impose explicit or implicit caps on price dispersion. For example, zone-pricing statutes in retail gasoline markets restrict the maximum price gap that a station may charge relative to nearby competitors (Verlinda, 2008; Barron et al., 2004). Likewise, “fare-parity” (or most-favoured-nation) clauses in airline distribution contracts effectively bound the difference between the highest and lowest publicly available fares on a route (Borenstein and Rose, 1994; European Commission, 2018).

Even when such regulations do not explicitly prohibit price differences, they indirectly constrain firms’ ability to use price as a manipulative signal, such as signalling product quality, exclusivity, or urgency. By mandating greater transparency in digital markets, for instance, the EU’s Digital Markets Act (DMA) and Digital Services Act (DSA), these frameworks reduce the persuasive power of extreme or hidden pricing strategies, particularly those that rely on consumer naivety or incomplete information (European Parliament and Council, 2022). Furthermore, these regulations discourage segmentation practices in which certain groups of consumers are systematically exposed to biased information or discriminatory pricing. In doing so, they aim to foster fairer competition and protect consumer autonomy in markets increasingly shaped by algorithmic targeting and personalized influence. Such regulations carve out secure markets for each firm, preventing extreme undercutting that would otherwise collapse the protected regions to measure zero.

Introducing a symmetric bound $|\Delta p| < \bar{\Delta}$ into Proposition 5.2 therefore yields a practically relevant special case, formalized in Corollary 5.1.

Corollary 5.1. Bounded Price Differences and Minimum Protected Regions

Impose, in addition to the assumptions of Proposition 5.2, a regulatory cap on price dispersion,

$$-\bar{\Delta} < \Delta p < \bar{\Delta},$$

then:

- (a) *the secure-market thresholds satisfy $\rho_1^{\min} = \rho_1(a_2 = 0, \Delta p = \bar{\Delta})$ and $\rho_2^{\max} = \rho_2(a_1 = 0, \Delta p = -\bar{\Delta})$*
- (b) *when $\rho_0 < \rho_1(a_2 = 0, \Delta p = \bar{\Delta})$, neither firm has an incentive to exert influence, $a_1^* = a_2^* = 0$, and the optimal prices*

$$p_1^* = 1, \text{ and } p_2^* \in [0, 1],$$

yield the expected profits $\Pi_1^* = p_1^* = 1$ and $\Pi_2^* = 0$.

- (c) *when $\rho_0 > \rho_2(a_1 = 0, \Delta p = -\bar{\Delta})$, neither firm has an incentive to exert influence, $a_1^* = a_2^* = 0$, and the optimal prices*

$$p_1^* \in [0, 1], \text{ and } p_2^* = 1,$$

yield the expected profits $\Pi_1^* = 0$ and $\Pi_2^* = 1$.

- (d) *the optimal prices, influence levels, and profits in Proposition 5.2 remain valid for the competitive region, i.e., $\rho_0 \in [\rho_1(a_2 = 0, \Delta p = \bar{\Delta}), \rho_2(a_1 = 0, \Delta p = -\bar{\Delta})]$.*

The results in Corollary 5.1 further reinforce the consistency with the exogenous pricing case, particularly in terms of the structural features of the equilibrium. In the competitive region, where potential purchases are endangered by the competitor's influence actions, the dominant seller's influence level becomes irrelevant. However, by endogenizing prices, firms are able to fully exploit opportunities in the protected regions and extract maximum profits.

Compared to Proposition 5.2, the introduction of the external constraint $|\Delta p| < \bar{\Delta}$ does not bind the equilibrium pricing decisions, as the resulting optimal prices lie within the allowable range. Nonetheless, it affects the secure market thresholds, since ρ_1 and ρ_2 depend on the *feasible* – rather than the realized – range of Δp . This effect is illustrated in Figure 5.10, which illustrates




Figure 5.10: Contour Plot of the competitive region size as a function of Price Difference and Influence Level ($\gamma = 0.7$).

the variation in the width of the competitive region, measured by $\rho_2 - \rho_1$, with respect to the price difference Δp and the associated influence level a_i . Darker regions correspond to a wider competitive zone, indicating a larger range of customer beliefs where both firms compete for demand. The competitive region is widest when the allowable $\Delta p \approx 0$, reflecting moderate pricing strategies. In contrast, as pricing becomes extreme ($|\Delta p| \rightarrow 1$) or influence becomes stronger, the overlap between firms' appeal narrows, reducing competition. Thus, imposing a tighter cap on the allowable price difference leads to a visible expansion of the competitive region, aligning with findings in the prior literature and practice: excessive price undercutting can lead to market segmentation and reduced consumer choice. Consequently, as formally stated in Corollary 5.1, the effective domain over which customer influence strategies apply is widened under bounded price dispersion.

Additionally, imposing an upper bound on sellers' influence levels – whether through platform

design, regulatory oversight, or advertising constraints – can have the similar effect: it potentially expands the competitive region by reducing the extent to which firms can secure captive market segments. As influence becomes constrained, neither firm can dominate the consumer belief space through persuasion, thereby increasing the overlap in appeal and intensifying head-to-head competition. On a side note, it is revealed that while extreme pricing alone can potentially eliminate competition entirely, influence by itself cannot achieve the same effect. This highlights that, in this framework, pricing is a stronger lever than influence in swaying customer decisions.

Another key insight emerging from Proposition 5.2 and Corollary 5.1 (part (b) and (c)) is the role of influence in enabling price rigidity. In the competitive region, the non-dominant seller leverages influence to sustain a high price, despite having a lower baseline probability of fit. This illustrates that influence is not only a tool for demand capture but also a strategic mechanism to avoid price cuts.

Moreover, the dominant seller does not engage in aggressive price responses to counter the competitor’s growing influence. Instead, its pricing remains inflexible, perhaps counterintuitively, despite increased persuasion from the rival. This lack of price responsiveness may be explained by the dual effect of influence: although the competitor improves their conversion rates, their influence also shrinks the dominant seller’s competitive region, thereby enlarging the latter’s secure market. As a result, the dominant seller benefits from a more committed, non-convertible, customer base that is willing to pay the maximum price, mitigating the need for price concessions.

Furthermore, a *co-escalation of competition* is observed as the sellers’ perceived fit becomes more balanced, i.e., when ρ_0 approaches the centre of the belief spectrum. In this highly competitive regime, both firms respond by simultaneously increasing their prices and influence levels, reflecting the intensity of strategic rivalry. However, consistent with the findings in Proposition 5.1, the *dominant seller continues to earn higher profits* than the non-dominant one. While pricing enables the weaker seller to compete on dimensions beyond perceptions, it does not allow them to surpass the dominant firm’s profitability.

5.5. The Impact of Product Returns

In this section, the analysis is extended to consider the case where both sellers offer product return options, allowing us to examine the impact of post-purchase factors. This introduces an additional stage in the game compared to Section 5.4, in which the consumer decides whether to keep or return the product after observing the actual fit, F . The refund policy is treated as exogenous – either mandated by the platform or shaped by consumer expectations – but it may differ across firms depending on the operational factor.

Accordingly, the backward induction now begins from the final stage of this updated game sequence:

Stage 4 (Return Decision): The customer evaluates the purchased product, realizes the true fit, and decides whether to keep or make a return based on the actual fit.

5.5.1 Equilibrium Outcomes

In the scenario where the firm allows some refunds based on an exogenous refund factor $r_i \in (0, 1]$. The consumer will receive $r_i p_i$ in exchange for the returned product.

At Stage 4, the customer must make a product return decision upon observing F_i . If the product matches the customer ($F_i = G$), the customer will keep the item with a positive net utility $v - p_i$; otherwise, if the fit is bad ($F_i = B$), the customer will return it to at least partially recover the price paid and end up with a non-positive net utility $r_i p_i - p_i$.

By backward induction, at Stage 3, the customer will buy from Seller i if the expected utility of Product i is greater than that of the other. In the absence of any signal, the customer will be indifferent between the two sellers when $\bar{\rho}_0 v - p_1 + \rho_0 r_1 p_1 = \rho_0 v - p_2 + \bar{\rho}_0 r_2 p_2$. Denote the influence-independent indifference point $\tilde{k} = \frac{v - p_1 + (1 - r_2)p_2}{2v - (r_1 p_1 + r_2 p_2)}$.

When $\rho_0 < \tilde{k}$, the customer will make a purchase from seller 2 if and only if there is an asymmetrical positive signal inclination towards Product 2, indicating a stronger tendency for favourable perceptions and expectations about product fit compared to Product 1, that is, when $S_1 = b$ and $S_2 = g$, the customer will make a purchase from seller 2 if and only if

$$P(F_2 = G|S_2 = g, a_2)v - p_2 + P(F_2 = B|S_2 = g, a_2)r_2p_2 - P(F_1 = G|S_1 = b)v + p_1 - P(F_1 = B|S_1 = b)r_1p_1 \geq 0, \text{ i.e.,}$$

$$\rho_0 \geq \tilde{\rho}_1 = \frac{\sqrt{\tilde{A}_1^2 + \tilde{B}_1 \cdot \tilde{C}_1 + \tilde{D}_1}}{\tilde{E}_1}. \quad (5.12)$$

The explicit forms of $\tilde{A}_1$, $\tilde{B}_1$, and $\tilde{C}_1$, $\tilde{D}_1$, and $\tilde{E}_1$ are provided in Appendix C (Eqs. C.16 – C.20) for completeness.

When $\rho_0 > \tilde{k}$, the customer will make a purchase from seller 1 if and only if there is an asymmetrical positive signal inclination towards Product 1, indicating a stronger tendency for favourable perceptions and expectations about product fit compared to Product 2, i.e., when $S_1 = g$ and $S_2 = b$, the customer will make a purchase from seller 1 if and only if $P(F_1 = G|S_1 = g, a_1)v - p_1 + P(F_1 = B|S_1 = g, a_1)r_1p_1 - P(F_2 = G|S_2 = b)v + p_2 - P(F_2 = B|S_2 = b)r_2p_2 \geq 0$, i.e.,

$$\rho_0 \geq \tilde{\rho}_2 = -\frac{\tilde{A}_2 + \sqrt{\tilde{B}_2^2 + \tilde{C}_2 \cdot \tilde{D}_2}}{\tilde{E}_2}. \quad (5.13)$$

The explicit forms of $\tilde{A}_2$, $\tilde{B}_2$, and $\tilde{C}_2$, $\tilde{D}_2$, and $\tilde{E}_2$ are provided in Appendix C (Eqs. C.21 – C.25) for completeness.

In essence, the thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ represent the minimum prior belief in product fit, ρ_0 , required for a seller to win the customer's purchase decision under asymmetric signals. Specifically, $\tilde{\rho}_1$ is the threshold above which Seller 2 can win the sale despite Seller 1 having the fit advantage, while $\tilde{\rho}_2$ defines the point below which Seller 1 can overcome a fit prior disadvantage. These thresholds thus capture the influence of persuasion on the consumer's decision while accounting for product return policy effect, therefore, serving as refund-adjusted purchase thresholds that characterize the customer's susceptibility to being swayed away from the prior-favoured seller.

To further explore the behaviour of the belief thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ under varying return policies, a graphical exploration is conducted by varying the refund factors r_1 and r_2 while fixing the other parameters to provide insight into the firms' strategic positioning in terms of persuasion

and the effect of refund attractiveness.




Figure 5.11: Effect of Seller 1's refund rate r_1 on choice thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ ($p_1 = 0.7, p_2 = 0.6, \gamma = 0.7, v = 1, a_i = 0.4, r_2 = 0.9$)




Figure 5.12: Effect of Seller 2's refund rate r_2 on choice thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ ($p_1 = 0.7, p_2 = 0.6, \gamma = 0.7, v = 1, a_i = 0.4, r_1 = 0.9$)

The impact of Seller 1's refund factor, r_1 , on the purchase thresholds $\tilde{\rho}_1$ and $\tilde{\rho}_2$ is illustrated in Figure 5.11. Both thresholds increase as r_1 rises, indicating that Seller 1's protected market expands under a more lenient return policy. At the same time, Seller 2's protected market shrinks,

making it more demanding for Seller 2 to attract customers. This suggests that generous return policies by Seller 1 not only improve its defensive position but also raise the bar for competitive incursions.

A similar pattern is observed in Figure 5.12, which illustrates the influence of Seller 2's refund factor on the belief thresholds. As r_2 increases, both $\tilde{\rho}_1$ and $\tilde{\rho}_2$ decrease, indicating that a more generous return policy by Seller 2 increases its own protected market and weakens Seller 1's ability to retain customers. This reinforces the strategic role of return leniency in expanding consumer retention and shaping competitive boundaries.

By definition, for $0 \leq \tilde{\rho}_1 \leq \tilde{\rho}_2 \leq 1$ consumer's purchase choice in Stage 3 follows the same structure in Lemma 5.1 with the refund-adjusted thresholds.

In stage 2, the sellers decide their selling price and level of influence. Their expected profits are as follows:

$$\begin{aligned}
\Pi_1(a_1, a_2, p_1, p_2) &= \mathbb{1}\{\rho_0 < \tilde{\rho}_1\}p_1 - \rho_0 r_1 p_1 \\
&+ (1 - P(S_1 = b)P(S_2 = g | a_2)) \mathbb{1}\{\tilde{\rho}_1 \leq \rho_0 < k_r\}p_1 \\
&- (1 - P(F_1 = B | S_1 = b)P(F_1 = B | S_2 = g | a_2)) \mathbb{1}\{\tilde{\rho}_1 \leq \rho_0 < k_r\}r_1 p_1 \\
&+ P(S_1 = g | a_1)P(S_2 = b) \mathbb{1}\{k_r \leq \rho_0 < \tilde{\rho}_2\}p_1 \\
&- P(F_1 = B | S_1 = g, a_1)P(F_1 = B | S_2 = b) \mathbb{1}\{k_r \leq \rho_0 < \tilde{\rho}_2\}r_1 p_1 - c_1, \tag{5.14}
\end{aligned}$$

and

$$\begin{aligned}
\Pi_2(a_1, a_2, p_1, p_2) &= \mathbb{1}\{\rho_0 \geq \tilde{\rho}_2\}p_2 - \bar{\rho}_0 r_2 p_2 \\
&+ (1 - P(S_1 = g | a_1)P(S_2 = b)) \mathbb{1}\{k_r \leq \rho_0 < \tilde{\rho}_2\}p_2 \\
&- (1 - P(F_2 = B | S_1 = g, a_1)P(F_2 = B | S_2 = b)) \mathbb{1}\{k_r \leq \rho_0 < \tilde{\rho}_2\}r_2 p_2 \\
&+ P(S_1 = b)P(S_2 = g | a_2) \mathbb{1}\{\tilde{\rho}_1 \leq \rho_0 < k_r\}p_2 \\
&- P(F_2 = B | S_1 = b)P(F_2 = B | S_2 = g, a_2) \mathbb{1}\{\tilde{\rho}_1 \leq \rho_0 < k_r\}r_2 p_2 - c_2. \tag{5.15}
\end{aligned}$$

In the return-adjusted model, the expected profit of each seller is determined not only by the revenue from successful sales but also by the expected cost associated with product returns. The prior belief space is partitioned into three regions and in each region, the revenue term is accompanied by a corresponding refund liability that depends on the firm's refund factor r_i and the likelihood that a customer who initially accepts the product ultimately seeks a return. In the protected region $\rho_0 < \tilde{\rho}_1$, Seller 1 is chosen regardless of the signal outcome, yielding revenue p_1 per buyer. However, a refund is issued with probability ρ_0 , resulting in an expected cost of $\rho_0 r_1 p_1$. A symmetric logic applies for Seller 2 in the region $\rho_0 > \tilde{\rho}_2$, where the refund liability is governed by $\bar{\rho}_0 r_2 p_2$.

Within the insecure regions, the final purchase decision is influenced by the signal and the probability that sellers are chosen in terms of realized signals; while the expected refund liability cost in this region is in terms of the probability that bad fits are ultimately realized, given the same signal profiles. Thus, the influence variables a_i affect the profit functions through the probability of generating favourable signals, and through the likelihood of receiving negative feedback despite favourable and purchase-enabling signals. By construction, increasing a_i raises the probability of securing marginal customers, while also increasing exposure to returns and the magnitude of r_i governs the weight of this trade-off.

At this stage, it would be desirable to characterize the equilibrium prices and influence levels in closed form to enable explicit comparative-static predictions of the sellers' strategic responses. However, the first-order conditions and subsequent isolations yield a system of coupled high-degree polynomials whose coefficients include nested radicals. As a result, standard symbolic methods do not produce a tractable solution, and a closed-form analytical characterization is precluded. To investigate the model's strategic implications nonetheless, equilibrium prices, influence levels, and profits are computed numerically over a grid of refund parameters and prior beliefs. The outcomes are reported in Table 5.1, which illustrates how each seller responds to changes in refund policy while all other primitives are held constant. The results are obtained for $\gamma = 0.7$, with the rival firm's refund fixed at $r_i = 0.7$, and correspond to three representative prior beliefs: $\rho_0 = 0.2$ to reflect consumers initially inclined toward Product 1, $\rho_0 = 0.8$ for

those inclined toward Product 2, and $\rho_0 = 0.5$ for consumers with more balanced or uncertain initial preferences. Valuation, $v = 1$, and cost parameters are set to zero ($c_i = 0$).¹

Table 5.1: Numerical Equilibrium Results with Different Refund Values

ρ_0	r_1	a_1^*	a_2^*	p_1^*	p_2^*	Π_1^*	Π_2^*	r_2	a_1^*	a_2^*	p_1^*	p_2^*	Π_1^*	Π_2^*
0.2	0.1	0	1	1	1	0.53	0.27	0.1	0	1	0	1	0	0.36
	0.2	0	1	1	1	0.44	0.27	0.2	0	1	0	1	0	0.35
	0.3	0	1	1	1	0.34	0.27	0.3	0	1	0	1	0	0.33
	0.4	0	1	1	1	0.25	0.27	0.4	0	1	0	1	0	0.32
	0.5	0	1	1	1	0.15	0.27	0.5	0	1	0	1	0	0.30
	0.6	0	1	1	1	0.06	0.27	0.6	0	1	0	1	0	0.28
	0.7	0	0	1	1	0	0.27	0.7	0	1	0	1	0	0.27
	0.8	0	0	1	1	0	0.27	0.8	0	1	0	1	0	0.25
	0.9	0	0	1	1	0	0.27	0.9	1	1	1	0.92	0.59	0.12
	1.0	0	0	1	1	0	0.27	1.0	1	1	1	1	0.59	0.10
0.5	0.1	1	0	1	1	0.48	0.05	0.1	1	1	1	1	0.05	0.48
	0.2	1	0	1	1	0.45	0.05	0.2	1	1	1	1	0.05	0.45
	0.3	1	0	1	1	0.43	0.05	0.3	1	1	1	1	0.05	0.43
	0.4	1	0	1	1	0.40	0.05	0.4	1	1	1	1	0.05	0.40
	0.5	1	1	0.87	1	0.15	0.33	0.5	1	1	1	0.87	0.15	0.33
	0.6	1	1	1	1	0.35	0.05	0.6	1	1	1	0.93	0.33	0.10
	0.7	1	1	1	1	0.33	0.05	0.7	1	1	1	1	0.33	0.05
	0.8	0	1	0	1	0	0.33	0.8	1	0	1	0	0.33	0
	0.9	0	1	0	1	0	0.33	0.9	1	0	1	0	0.33	0
	1.0	0	1	0	1	0	0.33	1.0	1	0	1	0	0.33	0
0.8	0.1	1	0	1	0	0.32	0	0.1	1	0	0.08	0.65	0	0.34
	0.2	1	0	1	0	0.25	0	0.2	1	0	0.06	0.65	0	0.28
	0.3	1	0	1	0	0.19	0	0.3	1	0	0	0.65	0	0.22
	0.4	1	0	1	0	0.12	0	0.4	1	0	0	0.65	0	0.16
	0.5	1	0	1	0	0.06	0	0.5	1	0	0	0.67	0	0.14
	0.6	0	0	0	0	0	0	0.6	1	0	0	0.67	0	0.10
	0.7	0	0	0	0	0	0	0.7	1	0	0	0	0	0
	0.8	0	0	0	0	0	0	0.8	1	0	0	0	0	0
	0.9	0	0	0	0.70	0	0.01	0.9	1	0	0	0	0	0
	1.0	0	0	0	0.70	0	0.03	1.0	1	0	0	0	0	0

Table 5.1 shows that refunds systematically redistribute profit toward the seller that is *ex ante* disadvantaged by the prior, but only when the disadvantaged seller itself raises its refund generosity. When consumers are predisposed to favour Seller 1 ($\rho_0 = 0.2$) for example, Seller 1's profit falls sharply as its own refund r_1 increases, declining from 0.53 at $r_1 = 0.1$ to zero once $r_1 \geq 0.7$. Seller 2's profit, by contrast, remains fixed at 0.27 because its refund is held constant; consequently, the non-dominant seller overtakes the dominant one for high values of r_1 . A mirror effect appears at $\rho_0 = 0.8$, where Seller 2 is dominant: as r_2 rises on

¹Equilibria were computed using a grid-search routine implemented in Python (package `scipy.optimize`), and the admissible root within $[0, 1]$ was retained for each parameter configuration. The solver uses `scipy.optimize.root` with a hybrid Newton-bisection algorithm, and an absolute tolerance of 10^{-8} .

the right-hand panel, Seller 2's profit collapses to zero, while Seller 1 (now the disadvantaged seller) cannot secure any profits, but still reflects the symmetry in terms of remaining unchanged. This collapse is not unique to a specific market position or seller identity. Rather, the results suggest a broader structural constraint: for most types of customers, there exists a threshold level of refund generosity beyond which no positive price is sustainable. Once this limit is surpassed, the equilibrium price and resulting profit both converge to zero. This pattern reflects the compounded effect of rising refund liabilities and eroded pricing power, effectively nullifying the seller's ability to extract surplus from any consumer segment, regardless of prior disposition.

The balanced-prior case, $\rho_0 = 0.5$, reveals the clearest evidence that the refund policy has a threshold effect which is indicative of its active role on the entire strategic package of influence, pricing, and profitability. This segment also exhibits markedly stronger competitive pressure than the two polar segments. Because neither seller enjoys a clear *ex ante* asymmetry, both are compelled to deploy their full strategic arsenal until refund levels rise above an impairing threshold, where decisions begin to sharply deviate from the normalized benchmarks.

Notably, the seller facing an unfavourable prior (i.e., the non-dominant seller) earns the higher profit at times. This is a counterintuitive but consistent pattern observed across all customer segments. The reason lies in the refund mechanism itself: refund liabilities scale with the size of secured purchases. The dominant seller, who enjoys a higher baseline demand, bears a proportionally higher refund cost, particularly when it sets a high r_i , which erodes its margin. The disadvantaged seller, operating with a smaller protected market base, can afford to offer a less generous refund, thereby limiting excessive return-related costs while leveraging influence to extract whatever conversions are possible. Meanwhile, prices generally adjust in response to changes in the strategic value of influence, reinforcing the idea that refund policies – rather than price cuts – serve as the primary lever through which the disadvantaged firm can capture additional demand.

5.6. Summary

This chapter extended the thesis-wide framework to a duopoly in which prices and persuasion were chosen simultaneously and customers observed both sellers' signals before purchasing and potentially returning. Competition was shown to pivot on belief thresholds that partition the market into protected regions and a contested interior. A strategic externality emerged: raising one seller's influence expanded the rival's protected region by making positive signals more frequent yet less credible, while greater baseline informativeness shrank both protected regions and shifted demand into the marginal zone where persuasion is most effective.

With prices endogenized, influence and pricing interacted in predictable but asymmetric ways. Influence proved redundant inside protected regions and valuable only where purchases were endangered. Imposing bounds on price dispersion widened the marginal zone and increased the domain over which influence mattered, mirroring the effect of higher baseline signal quality.

Allowing product returns re-shaped the belief thresholds: more generous refunds expanded a seller's protected region but also raised refund liabilities, so that beyond a generosity threshold prices and profits collapsed toward zero. Refunds therefore can secure uncertain demand, but excessive reliance on refunds is profit-eroding. Taken together, these mechanisms delineate where persuasion is worth deploying and how pricing power co-moves with persuasion under duopolistic fit uncertainty.

Chapter 6

Discussion and Implications

6.1. Chapter Overview and Purpose

Omnichannel retail – defined as the seamless integration of physical and digital touchpoints (Verhoef et al., 2015) – has reshaped the way consumers *search for, evaluate, and purchase* products. Shoppers now form expectations of product fit long before they walk into a store or click the *Buy* button. For managers, strategic levers operate in tandem: the informativeness of third-party signals (baseline accuracy), the degree of influence a firm chooses to exert on those signals, price and the generosity of post-purchase refunds. Academic research has tended to examine these levers in isolation, yet retailers must set them *jointly*. This dissertation is aimed at closing that gap by developing a modular analytical framework and applying it across monopoly and duopoly settings. This framework reveals the underlying interdependencies: substitution, complementarity, and the emergence of pivot points and role reversals.

The discussion proceeds in two steps. First, an integrative synthesis is provided in which the principal findings of the three essays are consolidated and the circumstances under which the interactions among refunds, influence, and pricing confirm (or overturn) conventional intuitions, are highlighted. Second, practical and managerial insights are distilled: the analytical results are translated into observable decision rules, such as whether budget should be directed towards influencer marketing or towards more generous refunds as well as the associated implications.

6.2. Coalescing Findings Across Studies

6.2.1 Revisiting the Core Questions

This section revisits the core research questions introduced at the outset and reflects on the unified insights that have emerged across the three analytical studies. The central aim has been to understand how firms in uncertain, online retail environments can coordinate the three key levers to manage information provision and competition. Building on the objectives set in Chapter 1 and the phase-specific inquiries introduced throughout the thesis, this section synthesizes the findings across the various themes. Accordingly, particular attention is given to those questions that support integrative reflection, offer a broader perspective, and inform the interpretation of results.

The Principal Interaction: Though a central theme throughout the thesis, the interaction is most notably illustrated in the first essay, which focuses on a monopoly setting and demonstrates how the simultaneous availability of both an influence tool and a refund policy alters the firm's optimal behaviour compared to having access to only one of these levers.

To begin with, without refunds, there is a certain range of intermediate prior-fit probability where acquiring the influence tool and exerting a positive influence level is optimal (to induce uncertain consumers to buy). When a refund policy is introduced, the range of priors for which influence is used shifts – the firm is able to use influence at more pessimistic prior beliefs because the refund option will catch the mismatches, reducing the downside for customers. In other words, returns enable the encroachment of the application zone of customer influence tools by allowing the firm to be more aggressive in marketing relying on the presence of a safety net for customers. This represents a case of strategic complementarity between returns and influence in achieving the desired goals.

The insights on optimal firm behaviour are further extended by endogenizing the refund decision and examining the impact of varying return policies. Under a binary refund setting, the integrated model of dual instruments yields coordinated strategies. For instance, when the prior consumer

belief is very low, the optimal policy may involve offering a full refund while refraining from exerting influence, reflecting the firm's preference to reduce perceived risk and encourage trial without incurring additional marketing costs. When the prior belief is moderate, a strategy of no refund coupled with high influence becomes optimal, as persuasion is necessary to drive purchases despite an increased risk of mismatch, making refund commitments less desirable. Conversely, when the prior is high – indicating a strong likelihood of product fit – neither refunds nor influence are required, as consumers are expected to buy and retain the product regardless. These distinct regimes underscore the strategic interplay between influence and refund decisions: the optimal deployment of one instrument is contingent upon the policy applied to the other. This pattern reflects a substitution effect at the micro level where influence and refund tools are used in place of one another depending on the consumer belief. However, from a macro perspective, the instruments act as complements, jointly shaping the firm's strategy across different market segments. That is, while one lever may be emphasized over the other in any given regime, their combined availability enhances the firm's ability to tailor responses to varying customer profiles and belief structures.

Welfare Implications: While exploring different scenarios within the core tradeoff, it is observed that well-designed refund policies can create strategic value beyond their immediate role in shaping behaviour. This additional value can also benefit customers, particularly when the firm's freedom in configuring refund terms is limited.

The welfare effects of product returns in conjunction with influence strategies exhibit variation across different scenarios. When the refund rate is exogenously fixed, allowing product returns alongside customer influence yields non-trivial gains for the firm – especially when the prior belief of product fit is low. In such cases, all surplus is captured by the firm, with no spillover to the customer. By contrast, when the prior belief of fit is high, exogenous refunds tend to enhance customer welfare but reduce the firm's expected profit.

When the firm can configure the rate itself through a binary refund, both the firm and the customers will be strictly better off when the prior is relatively low, and experience no change in expected profit or welfare when it is high. However, when full discretion is granted to configure

the level of refunds, the firm optimally offers partial refunds scaled to the likelihood of a fit, allowing it to extract all remaining surplus. In this setting, customers may be strictly worse off even compared to a no-refund scenario, particularly when the prior fit probability is low.

Taken together, these findings highlight that the welfare implications of returns and influence are mainly shaped by the structure of the firm's refund decision space. A win-win outcome is more likely to emerge when product misfit is likely and refund terms are externally capped or otherwise restricted.

Strategic Informativeness: When the possibility of strategically configuring the informativeness of customer-facing signals – used as the basis for exerting influence – emerges as an option, it introduces a new and powerful lever for managing consumer uncertainty. This thesis advances existing approaches by integrating the configuration of baseline informativeness *ex ante* with the application of influence and return policies *ex post*. In doing so, it demonstrates how tuning the signal's credibility upstream enhances the effectiveness and flexibility of downstream persuasion tools and post-purchase policies.

The analysis reveals that while strategically tweaking baseline informativeness is not sufficient to achieve full market coverage – that is, effectively delivering persuasive messaging to the broad base of initially uncertain or outside consumers – it significantly enhances the impact of subsequent influence efforts if applied properly. In particular, when influence levels are externally constrained, strategically managing signal precision *ex ante* primarily helps with increasing the conversion rate of consumers with relatively optimistic prior beliefs. By contrast, when influence is endogenously determined, the ability to configure baseline informativeness becomes critical for expanding persuasive reach across a broader spectrum of consumer segments.

Interestingly, the firm does not always benefit from improving the informativeness of the information mechanism. In cases where *ex post* influence is constrained and consumers perceive the likelihood of a product misfit to be higher than that of a good fit – but not so pessimistic as to render persuasion futile – the firm may actually benefit from *reducing* baseline informativeness. In such instances, blurring signal precision can dilute negative priors and allow the influence

tool to sway purchase decisions more effectively, thereby improving conversion and profitability. This observation also rationalizes the emergence of informativeness caps under highly optimistic prior beliefs, which function to prevent over-precision from undermining the purpose by dissuasion. This underscores the importance of coordinating both the passive credibility of information channels and the active persuasion strategies that build upon them.

In the extended setting – where all three instruments (refund policy, influence level, and baseline informativeness) are under the firm’s control – a central insight emerges: a partial refund is sufficient to achieve the maximum expected profit in all cases. This flexibility allows the firm to substitute between influence and refunds, using signal quality as a complementary lever to fine-tune outcomes. For instance, when refund options are limited to binary choices (e.g., full or none), this adaptability enables the firm to avoid offering full refunds while still capturing optimal profit levels.

Competition and Instrument Interplay: While many of the earlier insights extend naturally to competitive environments, the presence of strategic interaction introduces additional complexity – and opportunity. The third essay investigates these dynamics, showing how two firms simultaneously choose prices and influence levels in a horizontally differentiated market. A key insight is the asymmetric deployment of influence across market segments. The weaker competitor in any given segment typically uses minimal influence for targeting consumers whose purchase decisions remain uncertain. As the organic likelihood of fit improves, persuasive efforts are intensified, enabling the firm to contest demand more aggressively – until the prior becomes sufficiently favourable to render the firm dominant. This shifting pattern highlights the role of influence as a tactical equalizer for the underdog in contested or marginal regions, and aligns with the theoretical expectation that the cost-effectiveness of persuasion decreases as the distance between the consumer’s prior and the firm’s ideal point widens. Influence is thus most productively directed at marginal customers, whose preferences are not yet firmly aligned with either competitor and remain malleable.

A targeted influence, wherein the intensity of persuasion is contingent on both strategic positioning and perceived demand elasticity induces heterogeneity in firms’ strategic behaviour.

However, this heterogeneity does not extend to pricing: in general, both firms tend to choose identical prices across different prior settings. That is, in some cases, influence acts as a substitute for price cuts depending on the fit conditions, but not vice versa. This asymmetry suggests that influence induces price rigidity, elevating its strategic value as a differentiator in competitive contexts.

Mandated return policies likewise, become strategically significant considerations in competitive settings, as they expand the firm's reach to risk-averse consumers. More importantly, refund is shown to serve as a strategic pivot under specific conditions – particularly in balanced-prior scenarios where competition is most intense. In such cases, offering refunds reduces the pressure on rivals to escalate their own strategic levers; a pattern that contrasts sharply with no-refund scenarios. In this regard, numerical examples show that refund generosity can mitigate the classic prisoner's dilemma: while mutual restraint would be beneficial, heightened competitive pressure drives both firms toward maximal instrument activation (e.g., aggressive influence and pricing). Refund generosity, however, enables more moderate instrument use, softening this arms race dynamic.

When one firm is disadvantaged, a generous refund policy offered by the competitor can allow the weaker firm to remain profitable by avoiding excessive return burdens. In this sense, even firms with weaker baseline appeal can sustain competitiveness. This reallocation of profitability demonstrates how externally enforced safety-net mechanisms such as refunds can distort market power and overturn expected outcomes under asymmetrical situations.

These dynamics underscore the importance of flexibility in strategic design and highlight the complex interdependencies among the instruments. Crucially, the interaction of instruments is not merely complementary or substitutive in static terms. Their strategic interplay often induces regime shifts. Especially in the “borderline” region where consumer preferences lie between the two firms' offerings: small adjustments to price, refund policies, or persuasive signals can decisively sway purchasing behaviour. Influence becomes most potent in these marginal areas, while refund generosity governs which firm secures marginal advantage. The broader takeaway, therefore, is that competitive advantage emerges less from reliance on any single lever than from

well-calibrated combinations of pricing, influence, and return policies.

6.3. Practical Insights and Strategic Relevance

In general, the findings are not intended to be interpreted in a literal or prescriptive manner, but rather as conceptual tools that offer broader strategic understandings and serve to define possibilities. While derived from stylized models and idealized conditions, these insights can shed light on real-world dynamics that firms frequently encounter. Importantly, the goal is not to claim exact behavioural replication in all business settings, but to offer theoretically grounded explanations that align with, or help anticipate, certain managerial decisions observed in practice.

Before mapping these insights directly onto observations, however, it is essential to outline several key considerations that affect this translational process. These include contextual differences, implementation feasibility, consumer behaviour variability, and institutional constraints – all of which shape how and whether the theoretical mechanisms can be meaningfully applied in practice.

6.3.1 Practical Relevance and Translational Considerations

Navigating the Trade-Off: Generous refunds and aggressive marketing can, to some extent, be used as substitutes in encouraging purchase, though they also amplify each other's costs. Persuasive efforts tend to increase the number of buyers – including mismatches – while higher influence levels and greater tolerance for associated costs become feasible under more lenient refund policies. Nevertheless, a net benefit can be achieved if these tools are managed strategically. This may help explain why high return rates are often observed in online markets that also invest heavily in marketing. In such cases, returns are not treated as a failure, but rather accepted as a cost of broadening the customer base through persuasive messaging.

In practice, a firm must determine *which* reassurance lever to activate once implementation frictions are reintroduced. The analysis suggests that the choice between refunds and influence

depends on their relative costs, which can be assessed along four dimensions:

- (i) Cash versus Fulfillment Outlay. Refunds tie up working capital and incur salvage or restocking costs (Anderson et al., 2009), whereas influence spending draws directly from the marketing budget.
- (ii) Timing of Reassurance. Influence mitigates *entry* risk by expanding the consideration set, while refunds mitigate *exit* risk by absorbing mismatches (Davis et al., 1995).
- (iii) Operational Bandwidth. Lenient returns place demands on inspection and refurbishment capacity (Shulman et al., 2011), while aggressive influence may trigger platform sanctions or erode credibility if perceived as manipulative (Dellarocas, 2006; Mayzlin et al., 2014).
- (iv) Credibility and Salvage Environment. High baseline review accuracy may reduce the value of additional persuasive effort (Hu et al., 2009), while low salvage value diminishes the appeal of refund generosity (Sun et al., 2021).

Taken together, these trade-offs suggest that influence and refunds function as equivalent forms of insurance against fit uncertainty – yet they are priced and constrained differently. Decisions regarding their use should therefore be informed by firm-specific factors such as logistics costs, marketing budgets, and credibility constraints. Whether reassurance is best provided *before* purchase through information design, or *after* purchase through refund policies, depends on how these constraints are balanced.

Benchmark Value of Continuous Refund: The analysis of different return policy configurations suggests that a continuous refund mechanism is more efficient in addressing match uncertainty than any combination of binary refund and influence tools. This implies that, when feasible, firms may benefit from implementing instruments such as restocking fees, non-refundable surcharges (e.g., shipping), deposits, or store credits to retain an appropriate level of sales revenue.

In practice, however, implementing the theoretically optimal partial refund levels may prove challenging. First, customers may resist or find it difficult to accept a complex refund policy

that appears designed solely to serve the firm’s interests. Moreover, the proposed partial refund scheme imposes additional administrative burdens on the firm, particularly when consumer priors vary across time or product categories. As a result, it may be more practical to adopt a standardized refund policy and rely on configurable, lower-cost customer influence mechanisms to accommodate such variations. In this context, the identified optimal refunds can serve as a useful benchmark for assessing the effectiveness of simplified policy designs and influence configurations. Nonetheless, the literature has recognized several rationales for implementing partial refunds, reinforcing their practical relevance despite implementation challenges. Prior studies have attributed the use of partial refunds to factors such as consumer opportunism (Chu et al., 1998), heterogeneous valuation uncertainty (Courty and Li, 2000), capacity constraints (Xie and Gerstner, 2007), aggregate demand uncertainty (Su, 2009), market competition (Guo, 2009; Shulman et al., 2011), return logistics and channel considerations (Shulman et al., 2010), product line management (Huang and Zhang, 2020), and salvage value opportunities in omnichannel environments (Nageswaran et al., 2020). These findings support the broader insight that, while continuous refunds offer theoretical efficiency, partial refunds remain a valuable and context-dependent tool in real-world retail practice.

Temporal and Multidimensional Considerations in Return Policies: While full or free refunds are a common feature in retail, they are often subject to additional dimensions that constrain their applicability in practice. One such dimension is *time*, where refund eligibility is restricted to a limited return window. In many cases, returns made after this period result in only a percentage of the purchase price being reimbursed – or may be outright denied (e.g., Amazon’s policy on restocking fees¹). From the consumer’s perspective, this temporal boundary transforms what is nominally a “full refund” into a probabilistic benefit, contingent on their ability to make a timely return decision.

This perspective becomes especially relevant in settings where product evaluation requires time and use – for example, a consumer purchasing a Virtual Reality (VR) gaming accessory from GameStop Canada, which imposes a return window of only 7 or 15 days.² When the likelihood

¹Amazon.ca (2025)

²EB Games Canada (2025)

of evaluating the product and initiating a return within this constrained period is uncertain, the consumer's expected value of the refund effectively declines. In such cases, the perceived utility of a full refund policy may align more closely with that of a partial refund.

These practical constraints emphasize the multidimensional nature of real-world return policies. Beyond the nominal refund amount, factors such as timing restrictions, restocking fees, return channels, and conditional eligibility shape both firm strategy and consumer perception. Consequently, the binary refund structures modelled in this thesis can also be viewed as benchmarks or abstractions that capture the core strategic trade-offs, while acknowledging that implementation in practice involves richer design space and administrative complexity.

Search Frictions and the Price–Influence Nexus: Online retail and information channels lower monetary search cost but leaves substantial *cognitive* cost in assessing fit; ease of returning a mismatched item effectively lowers that cognitive burden. Hence price, persuasive messaging, and refund generosity form a three-way trade-off: a steep discount or a lenient return may substitute for heavy influence spending, whereas a credible, informative signal may allow the firm to tighten its refund policy without harming demand. While this framework abstracts from explicit search frictions (e.g., Diamond's Paradox), these considerations suggest that its core insights remain robust to such real-world complexities – at least in terms of identifying the presence of strategic interactions, if not their exact direction.

6.4. From Model to Managerial Implications

Operationalizing the Refund–Influence Trade-off: The interplay between return policy design and persuasive influence actions reflects a critical strategic dimension in managing consumer fit uncertainty. While several intuitive aspects of this relationship are affirmed through the analysis,³ less intuitive patterns are also revealed. In particular, the notable substitution effect,⁴ underscore the importance of coordinating decisions to achieve an integrated reassurance

³E.g., the increased use of influence tools when consumers hold moderate priors and influence costs are relatively low.

⁴Generous refund policies can reduce the need for costly influence efforts, while strategic deployment of influence can mitigate the need for lenient return policies.

strategy.

Observable industry practices reflect both exogenous and endogenous applications of this trade-off. In some contexts, return policies are externally constrained by product type, regulatory norms, or operational feasibility. For instance, major appliance retailers such as Structube typically do not offer free returns on oversized items. Similarly, healthcare and personal medical devices are generally non-refundable due to safety and efficacy concerns.⁵ By contrast, the fashion and clothing sectors routinely offer generous return policies and rank among the highest categories for return frequency.⁶ This leniency is mirrored by intensive use of influencer marketing in these same categories, suggesting that influence and refunds are not merely coexisting but strategically co-deployed.⁷

Endogenous coordination is also evident in firms that control both return policy and marketing investment. L.L.Bean provides a notable case: once known for an unlimited return and satisfaction guarantee, the company revised its policy in February 2018, replacing the unlimited return window with a one-year return period requiring proof of purchase in response to customer misuse.⁸ Coinciding with this shift, L.L.Bean's marketing team intensified its digital and influencer-based campaigns to reinforce product quality and brand trust.⁹

Similarly, airlines such as Emirates and Turkish Airlines have implemented tiered refund structures, offering full refunds for business and first-class bookings and only partial refunds for economy fares. Influence efforts are also concentrated on premium offerings, as seen in their partnerships with lifestyle influencers and viral social media campaigns featuring first-class experiences.¹⁰ This practice reflects the insight that products associated with generous refunds may warrant higher influence spending, particularly when both levers are within the firm's control.

The electronics giant, Samsung, offers a further illustration of this alignment. Its extended return

⁵In accordance with Standard 19 of the Standards of Practice for Pharmacists and Pharmacy Technicians, pharmacy professionals are advised against accepting returned drugs, healthcare products, aids, or devices for reuse. This caution is based on the potential risk of product tampering or incorrect storage, which may compromise their safety and efficacy (DaSilva, 2022).

⁶PowerReviews Team (2021); Statista (2025)

⁷SocialPubli (2025); Lionize (2025)

⁸Austen Hufford (2018)

⁹Marketing Dive (2020)

¹⁰New York Post (2023b, a)

periods for flagship products (such as supersized QLED TVs) are paired with targeted influence campaigns and affiliate marketing content.¹¹ This focused promotional effort contrasts with the more limited refund options and muted marketing associated with other models in the same category, underscoring a deliberate linkage between return leniency and influence intensity.

Collectively, these industry examples support the proposition that influence and refund policies function as reassurance instruments that can be strategically adjusted in tandem. As highlighted in Section 3.4, influence should be directed toward different product segments depending on the refund environment. Under more generous policies, influence may be most effective for niche or unfamiliar items; under stricter policies, it should target well-matched, widely accepted products. When managed jointly, these tools allow firms to mitigate risk, expand reach, and align persuasion with post-purchase outcomes.

Configuring Credibility for Persuasive Impact: A key insight emerging from the analysis is that the effectiveness of downstream influence and refund strategies depends critically on the baseline informativeness of the signal presented to consumers. One practical lever firms possess to modulate this informativeness is the selection of influencers used in their marketing campaigns. Specifically, the degree of expertise and credibility associated with an influencer can shape consumers' interpretation of product signals, thereby acting as a form of *signal calibration* prior to persuasion.

In practice, this choice often manifests as a trade-off between *reach* and *credibility*. Brands seeking to maximize exposure may partner with high-profile, general-audience influencers who possess broad appeal but limited category expertise. Such associations may dilute signal precision, particularly when the product in question involves a non-trivial risk of mismatch. By contrast, firms confident in product fit or targeting consumers with relatively optimistic priors tend to prefer more credible or technically informed channels – for instance, through expert reviewers, niche content creators, or domain-specific Key Opinion Leaders (KOLs).

This dynamic is observable in multiple industries. In the skincare sector, for example, brands like La Roche-Posay and Paula's Choice routinely collaborate with dermatologists and licensed

¹¹ John Archer (2023); Samsung Electronics (2025)

estheticians to deliver high-credibility product messaging.¹² Conversely, fashion or lifestyle brands often employ mega-influencers with no clear subject-matter expertise to blur specific fit expectations and instead appeal to broader lifestyle narratives.

Empirical research supports this strategic differentiation. A recent meta-analysis finds that while both celebrities and social media influencers can drive engagement, the *perceived credibility of the messenger* strongly moderates persuasive effectiveness (Lee et al., 2024). Another survey notes that over 60% of consumers report reduced trust in influencer endorsements when the influencer lacks direct relevance or expertise in the product domain (Agility PR Solutions, 2023).

These practical trends reinforce the thesis findings: firms benefit from configuring baseline informativeness based on the likely prior beliefs of their consumer base. When prior beliefs are highly optimistic, increased precision through credible channels is recommended. When priors are slightly negative – but not fully pessimistic – deliberately blurring the signal by selecting less authoritative influencers may enhance downstream persuasion by softening consumer skepticism. Thus, the type of influencer chosen is not merely a branding decision, but a tactical instrument for upstream signal design that shapes the firm’s broader strategic toolkit.

Strategic Pivot under Price Dispersion Constraints: While pricing has traditionally served as a central lever in competitive markets, emerging regulatory frameworks and digital commerce dynamics are increasingly limiting firms’ ability to engage in unrestricted price differentiation. These regulations reduce the persuasive power of extreme or opaque pricing tactics, particularly those relying on consumer naivety or informational asymmetries.

Moreover, in online environments where consumers can easily compare prices, excessive dispersion risks triggering price arbitrage or regulatory scrutiny. These pressures are further reinforced by platform policies and contractual practices. Digital platforms can impose *de facto* minimum-advertised-price (MAP) caps that suppress price dispersion.¹³ For instance, MAP agreements can enforce uniformity in pricing across distribution channels, as seen in the U.S. iPhone market.

¹²Vogue Business (2023)

¹³42Signals (2023)

Apple’s MAP clauses prevent major carriers such as AT&T, Verizon, and T-Mobile (the “Big 3”) from advertising prices below the manufacturer’s suggested retail price (MSRP), thereby suppressing within-product price dispersion.

This strategic shift is clearly observable in real-world campaigns. Launch promotions for the iPhone 15, for example, featured identical handset prices across all major U.S. carriers, yet varied significantly in (i) 30–45-day “no-risk” return windows, (ii) inflated trade-in credits, and (iii) influencer-driven narratives on network reliability disseminated across TikTok and YouTube. These differences, while not reflected in pricing, allowed carriers to signal product value, reduce perceived risk, and capture consumers who might otherwise be indifferent among identical price offerings.¹⁴

This pattern aligns closely with the thesis insights: when price signalling is muted, firms shift along the strategic frontier by emphasizing persuasive content or rely on generous guarantees in regulated or uniform-price environments – particularly in digital ecosystems where price flexibility is structurally diminished.

Refund Generosity as a Competitive Safety Net: The direct-to-consumer (DTC) mattress sector demonstrates how generous refunds can moderate a classic prisoner’s-dilemma in influence spending. In 2016, both Casper and Purple escalated acquisition tactics – heavy promo codes, podcast sponsorships, and social-media influencers – driving marketing expenses above 35% of revenue. After each firm introduced a 100-night, no-questions-asked refund with free pickup in 2017, the competitive equilibrium shifted: Casper’s market share fell to 27% by 2019, average selling prices stabilized around \$1000, and broad-spectrum promo codes became holiday-only events. This mirrors the balanced-prior numerical simulations: a high refund-generosity parameter mitigates the arms-race outcome predicted under no-refund conditions.¹⁵

Moreover, the same policy enabled smaller entrant Tuft & Needle to stay profitable despite lower brand equity – supporting the asymmetry result that refund “safety nets” can reallocate market

¹⁴The three major U.S. carriers currently list the iPhone 15 at identical base prices, as verified through their official websites. Despite uniform pricing, they differentiate through other levers: AT&T emphasizes trade-in credits, Verizon promotes generous return windows, and T-Mobile leverages influencer-driven messaging and bundled perks.

¹⁵Casper Sleep Inc. (2020); VentureTwins (2019); Forbes (2020); Seeking Alpha Contributor (2019); Purple Innovation Inc. (2023)

power and keep weaker players viable.¹⁶

¹⁶Tuft & Needle (2025)

Chapter 7

Conclusion and Future Research

7.1. Key Insights and Contributions to the Knowledge

7.1.1 Integrated Retail Strategy

This thesis advances the literature by analyzing seemingly disparate strategies and policy designs *jointly*, to offer an integrated view of influence, refunds, and competition. Where prior studies tended to isolate these levers, the present framework demonstrates that they form a strategic system in which a constraint on one lever (e.g., a regulatory consumer protection obligation) induces compensatory shifts in the others. It is shown that refunds and influence can act either as *substitutes* – expanding consumer reach while avoiding costly influence, time-consuming configuration of the information environment, or promotional arms races – or as *complements* when competitive intensity or consumer uncertainty demands both reassurance and persuasion. By formalizing these interactions, the work extends classical signalling models into digital settings in which post-purchase guarantees matter.

7.1.2 Signal Design Before Persuasion

A second contribution of this thesis is the endogenization of *baseline signal credibility* as an upstream decision, marking the intersection where crafted informativeness meets the subtle

influence of perception. Firms are often capable of configuring the initial informativeness of the information channel – through, for example, the choice of influencer expertise, review curation strategies, or the granularity of product descriptions – *prior* to deploying influence or setting refund terms. While signalling theory has long been used to study strategic communication under information asymmetry, and its applications have extended to digital environments where direct product experience is lacking, limited attention has been paid to how firms control the default reliability and precision of the signals transmitted through digital channels.

This thesis advances a more nuanced understanding of how such supervisory control over baseline signal precision shapes the effectiveness of influence and refund levers in electronic word-of-mouth and digital marketing contexts. The analysis uncovers a pivotal non-monotonicity: enhancing baseline precision is profitable when consumers hold optimistic priors, whereas deliberate *blurring* of the signal can improve profitability when priors are moderately pessimistic and downstream persuasion tools are constrained. This finding rationalizes observed practices in which firms alternate between using domain experts and broad-reach influencers depending on product complexity and audience skepticism.

7.1.3 Practical Implications

This thesis demonstrates that price, influence, and refund policies form an interdependent toolkit for managing consumer fit uncertainty. When one lever is constrained, costly, or ineffective, firms pivot toward the others – substituting or complementing as internal or external pressures dictate. Thus, the choice is not only spatially, but also strategically transitional – a firm’s strategy must be adaptable, not static, even in seemingly stable market positions. If calibrated appropriately, this interdependence fosters greater resilience, competitive robustness, and long-term strategic flexibility. Accordingly, managerial implications, offer actionable guidance on how to coordinate influence strategies with refund policy design in digital retail platforms. Viewed from the perspective of implementation, three principles emerge. First, when feasible, generous yet *partial* refunds are typically sufficient to maximize profits, enabling firms to avoid the higher cost of full money-back guarantees. Second, refund generosity should be

directed toward products that carry the greatest risk of misfit, whereas influence efforts are best targeted at consumer segments characterized by intermediate uncertainty. Third, in markets where price competition is externally restricted, persuasive content and post-purchase guarantees become critical differentiators. Together, these insights offer a strategic framework for balancing reassurance and persuasion in environments marked by fit uncertainty and digital transparency. While the model is intentionally stylized, it represents a first attempt to examine these issues and provides a foundation for richer dynamic and behavioural extensions that may better guide managerial execution in increasingly transparent and regulated digital markets.

7.2. Limitations and Boundary Conditions

Behavioural Assumptions and Consumer Heterogeneity: The analysis presumes Bayesian, risk-neutral consumers who update beliefs rationally and incur no psychological costs from returns. Behavioural evidence suggests several departures from these premises. First, *loss aversion* can make refund options loom larger than equivalent gains, thereby altering willingness to try a product (Kahneman and Tversky, 1979). Second, *trust heuristics* determine how much weight consumers place on influencer claims; endorsements by unfamiliar or low-expertise sources can sharply reduce perceived diagnosticity (Metzger and Flanagin, 2013), while highly technical signals may overwhelm cognitively constrained buyers. Third, experimental work on remote purchases indicates that the marginal persuasion benefit of increasingly generous returns declines once buyers begin to suspect hidden strings or fraud (Wood, 2001); overly generous guarantees may trigger suspicion (“too good to be true”). Finally, real-world data show that hassle costs are heterogeneous: technology-savvy shoppers who can generate return labels and schedule courier pick-ups online, incur much lower time and effort costs than consumers lacking digital proficiency, leading to higher return propensities in the former group (Hsiao and Chen, 2014; UPS, 2021). Accounting for such behavioural biases and cost heterogeneity would provide a richer test of the substitution patterns identified in this thesis and could shift the optimal mix of influence and refund instruments.

Static Framework vs. Dynamic Markets: The model is static (or, at most, two-period), abstracting from inventory carry-over, word-of-mouth diffusion across cohorts, and repeat purchase. In practice, reviews accumulate, refund costs reveal product quality, and prices adjust over time. Dynamic learning models, where early refunds influence later demand or where firms cycle between penetration and skimming strategies, remain largely unexplored when refunds, influence, and price interact simultaneously. Extending the framework to a multi-period horizon with endogenous review generation and inventory decisions is a promising research avenue.

Return-Logistics: Throughout the analysis, consumer returns is treated as partial or full refunds, with no additional cost incurred by the seller beyond the refunded amount. That is, no separate write-off or recovery cost is considered. In practice, reverse-logistics expenses are a complex reality: inspection, refurbishment, repackaging, secondary-market liquidation, and disposal, each carry variable costs that depend on volume, product condition, and channel capacity. Salvage values can also fluctuate with seasonality or aftermarket demand, creating feedback between return rates and realized recovery. By abstracting these operational nuances, the framework may understate the marginal cost of generous refunds, especially when high influence levels drive large, unpredictable return flows. Enriching the model to include stochastic salvage value, capacity-dependent processing costs, and potential resale revenues would yield more precise insights on when refunds and influence remain substitutes, and when escalating reverse-logistics costs make influence the only preferred lever.

7.3. Directions for Future Research

Behavioural Extensions. Future work should relax fully rational, risk-neutral consumers. Incorporating loss aversion and reference dependence could generate asymmetric responses to refunds and price cuts, amplifying the value of money-back guarantees when losses loom larger than gains (Kahneman and Tversky, 1979; Köszegi and Rabin, 2006, 2007). Endowment and status-quo effects suggest “refund aversion” – reluctance to return once possession shifts reference points – potentially dampening realized returns even when *ex ante* purchase is stimulated

([Thaler, 1980](#); [Kahneman et al., 1990](#)). Heuristic persuasion and social proof (e.g., credibility cues, authority, and liking) may also interact with baseline informativeness and disclosure to shape posterior beliefs in systematically biased ways ([Cialdini, 2007](#)). Embedding such frictions could overturn some comparative statics of the baseline model (e.g., the monotone relation between refund generosity and optimal influence).

Dynamic Learning and Repeat Purchase. A multi-period setting would allow reviews, prices, and refund terms to co-evolve as reputations form and decay. Online reputation dynamics and word-of-mouth propagation provide microfoundations for belief updating and strategic disclosure ([Dellarocas, 2003](#); [Chevalier and Mayzlin, 2006](#); [Moe and Trusov, 2011](#)). Richer reputation models where future demand depends on current refund resolution or responsiveness, could link post-purchase policies to long-run market power ([Resnick and Zeckhauser, 2002](#); [Cabral and Hortaçsu, 2010](#)). On the supply side, learning–pricing interactions can be analyzed in a dynamic optimization setting where influence and baseline information precision serve as controls and belief updates govern the evolution of optimal policies.

Influencer Contracts and Commission Schemes. An explicit contracting layer (e.g., pay-per-conversion, tiered commissions, clawbacks for returned sales) would align creator incentives with firm profitability and refund exposure. This connects to measurement and attribution in digital advertising (incrementality vs. selection) and to mechanism design for performance-contingent payments (e.g., [Lewis and Rao, 2015](#); [Edelman et al., 2007](#)). Empirical influencer-marketing work on credibility, disclosure, and follower composition suggests contract terms should also respond to heterogeneous persuasion elasticities across audiences ([De Veirman et al., 2017](#); [Federal Trade Commission, 2023](#)).

Operational Integration. Linking information tools to inventory, fulfillment timing, and cross-channel returns would unify marketing “momentum” with reverse-logistics capacity and costs. The returns/MBG literature and reverse-logistics research provide natural primitives for embedding capacity limits, processing frictions, and salvage values (e.g., [Davis et al., 1995](#); [Rogers and Tibben-Lembke, 2001](#); [Guide and Van Wassenhove, 2009](#)). Omnichannel operations show

that store-based pickup/return and channel integration reshape both demand and the incidence of returns, with implications for where and when to deploy influence ([Gallino and Moreno, 2014](#)).

Technology and Policy Shifts. Roughly 70% of returns trace to fit/style mismatches, indicating that better discovery and sizing tools are first-order levers for prevention ([McKinsey & Company, 2023](#)). AI-driven recommendation and predictive fit may compress uncertainty and reduce the need for generous refunds, but could also amplify persuasion asymmetries if credibility cues are unevenly distributed. On the policy side, updated endorsement and disclosure rules along with platform obligations restrict manipulative practices and increase the salience of baseline credibility ([Federal Trade Commission, 2023](#)). Modelling these shocks as exogenous shifts in feasible influence space or disclosure costs would help firms adapt reassurance strategies to changing regulatory and technological frontiers.

Bibliography

John Archer. Samsung q80c review: mid-range qled marvel, and a great tv value. TechRadar web article, June 2023. URL <https://www.techradar.com/televisions/samsung-q80c-review>. Published June 2023; accessed July 11 2025.

42Signals. What is a map policy and why should brands care? E-commerce blog, July 2023. URL <https://www.42signals.com/blog/what-is-map-policy/>. Accessed July 2025.

Huseyn Abdulla, James D. Abbey, and Michael Ketzenberg. How consumers value retailer's return policy leniency levers: An empirical investigation. *Production and Operations Management*, 31(4):1719–1733, 2022. ISSN 1059-1478. doi: 10.1111/poms.13640.

Huseyn Abdulla, James D. Abbey, and Michael Ketzenberg. The point of no return? restrictive changes to lenient return policies and consumer reactions to them. *Journal of Operations Management*, 79(1):1–15, 2024. ISSN 0272-6963. doi: 10.1002/joom.1346.

Agility PR Solutions. Do consumers trust celebrity endorsements anymore? new research takes a look, 2023. URL <https://www.agilitypr.com/pr-news/social-media-influencer-marketing/do-consumers-trust-celebrity-endorsements-anymore-new-research-takes-a-look>. Accessed July 11, 2025.

Melih Akan, Baris Ata, and James D. Dana. Revenue management by sequential screening. *Journal of Economic Theory*, 159:728–774, 2015. doi: 10.1016/j.jet.2015.07.004.

M. Serkan Altug. The impact of online product reviews on retailer's pricing and return policy

- decisions. *SSRN Electronic Journal*, 2023. doi: 10.2139/ssrn.3914886. Working paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3914886.
- Amazon.ca. Refund timelines – amazon customer service. Web page, September 2025. URL <https://www.amazon.ca/gp/help/customer/display.html?nodeId=GKQNFKFK5CF3C54B>. Accessed July 2025.
- Eric T. Anderson, Karsten Hansen, and Duncan I. Simester. The option value of returns: Theory and empirical evidence. *Marketing Science*, 28(3):405–423, 2009. doi: 10.1287/mksc.1080.0432.
- Mark Armstrong and John Vickers. A model of delegated project choice. *Econometrica*, 78(1): 213–244, 2010. doi: 10.3982/ECTA7965.
- Association of Accredited Small Business Consultants (AASBC). Product or service guarantees boost business, May 2024. URL <https://www.aasbc.com/media/post/product-or-service-guarantees-boost-business>.
- Austen Hufford. L.L.Bean shifts marketing focus from print to digital & tv, February 2018. URL <https://www.wsj.com/articles/l-l-bean-to-end-unlimited-returns-citing-customer-misuse-1518194780/>.
- Jill Avery and Ayelet Israeli. Influencer marketing. Technical note 520-075, Harvard Business School, March 2020. URL <https://www.hbs.edu/faculty/Pages/item.aspx?num=57842>.
- Aditya Balaram, Olga Perdikaki, and Michael R. Galbreth. Bracketing of purchases to manage size uncertainty: Should online retailers be worried? *Naval Research Logistics*, 69(1-2): 3–24, 2022. doi: 10.1002/nav.22048. Issue Online: July 3, 2022; Version of Record Online: January 22, 2022.
- Oren Bar-Gill. Consumer misperception in a hotelling model: With and without price discrimination. *Journal of Institutional and Theoretical Economics (JITE)*, 176(24):180–203, 2020. doi: 10.1628/jite-2020-0021.

- John M. Barron, Brian A. Taylor, and John R. Umbeck. Number of sellers and price dispersion: Empirical evidence from the retail gasoline market. *Journal of Political Economy*, 112(2): 322–352, 2004.
- Michael R. Baye and John Morgan. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 91(3):454–474, 2001. doi: 10.1257/aer.91.3.454.
- David Bell, Santiago Gallino, and Antonio Moreno. Showrooms and information provision in omni-channel retail. *Production and Operations Management*, 24(3):360–362, 2015. doi: 10.1111/poms.12258_2. Reprinted 2024.
- Ron Berman, Hangcheng Zhao, and Yi Zhu. Strategic recommendation algorithms: Overselling and demarketing information designs. *SSRN Electronic Journal*, 2022. doi: 10.2139/ssrn.4301489.
- Severin Borenstein and Nancy L. Rose. Competition and price dispersion in the u.s. airline industry. *Journal of Political Economy*, 102(4):653–683, 1994.
- Gordon Burtch, Yili Hong, Ravi Bapna, and Gaurav Vishwanath. Stimulating online reviews by combining financial incentives and social norms. *Management Science*, 64(5):2065–2082, 2018. doi: 10.1287/mnsc.2016.2715.
- Luís Cabral and Ali Hortaçsu. The dynamics of seller reputation: Evidence from eBay. *The Journal of Industrial Economics*, 58(1):54–78, 2010. doi: 10.1111/j.1467-6451.2010.00405.x.
- Capital One Shopping. Average retail return rate (2025 data): ecommerce vs in-store. <https://capitaloneshopping.com/research/average-retail-return-rate/>, 2025. Accessed: 2025-06-01.
- Casper Sleep Inc. Form s-1 registration statement. U.S. Securities and Exchange Commission filing, January 2020. URL <https://www.sec.gov/Archives/edgar/data/1598674/000104746920000166/a2240404zs-1.htm>. Reports sales/marketing expense 35% of revenue; accessed 11 July 2025.

- Levent Celik. Information unraveling revisited: Disclosure of horizontal attributes. *The Journal of Industrial Economics*, 62(1):113–136, 2014. doi: 10.1111/joie.12039.
- Hsiu-Hua Chang and Ting-Shan Yang. Consumer rights or unethical behaviors: Exploring the impacts of retailer return policies. *Journal of Retailing and Consumer Services*, 64:102779, 2022. ISSN 0969-6989. doi: 10.1016/j.jretconser.2021.102779.
- Yeon-Koo Che. Customer return policies for experience goods. *Journal of Industrial Economics*, 44(1):17–24, 1996.
- Yubo Chen and Jinhong Xie. Third-party product review and firm marketing strategy. *Marketing Science*, 24(2):218–240, 2005. ISSN 1526-548X.
- Yubo Chen and Jinhong Xie. Online consumer review: Word-of-mouth as a new element of marketing communication mix. *Management Science*, 54(3):477–491, 2008. doi: 10.1287/mnsc.1070.0810.
- Judith A. Chevalier and Dina Mayzlin. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*, 43(3):345–354, 2006. ISSN 0022-2437. doi: 10.1509/jmkr.43.3.345.
- Wujin Chu, Eitan Gerstner, and James D. Hess. Managing dissatisfaction: How to decrease customer opportunism by partial refunds. *Journal of Service Research*, 1(2):140–155, 1998. ISSN 1094-6705. doi: 10.1177/109467059811008.
- Robert B. Cialdini. *Influence: The Psychology of Persuasion*. Harper Business, New York, revised ed. edition, 2007.
- Brian L. Connelly, S. Trevis Certo, R. Duane Ireland, and Christopher R. Reutzel. Signaling theory: A review and assessment. *Journal of Management*, 37(1):39–67, 2011. doi: 10.1177/0149206310388419.
- Pascal Courty and Hao Li. Sequential screening. *The Review of Economic Studies*, 67(4): 697–717, 2000. doi: 10.1111/1467-937X.00152.

- Tomasia DaSilva. Calgary consumer questions pharmacy return policy after buying a ‘faulty’ medical device. Global News (Calgary), May 2022. URL <https://globalnews.ca/news/8883212/calgary-faulty-medical-device-return-issues/>. Accessed July 11, 2025.
- Scott Davis, Eitan Gerstner, and Michael Hagerty. Money back guarantees in retailing: Matching products to consumer tastes. *Journal of Retailing*, 71(1):7–22, 1995. doi: 10.1016/0022-4359(95)90010-1.
- Scott Davis, Michael Hagerty, and Eitan Gerstner. Return policies and the optimal level of “hassle”. *Journal of Economics and Business*, 50(5):445–460, 1998. doi: 10.1016/S0148-6195(98)00013-7.
- Marijke De Veirman, Veroline Cauberghe, and Liselot Hudders. Marketing through Instagram influencers: The impact of number of followers and product divergence on brand attitude. *International Journal of Advertising*, 36(5):798–828, 2017. doi: 10.1080/02650487.2017.1348035.
- Chrysanthos Dellarocas. The digitization of word-of-mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424, 2003. doi: 10.1287/mnsc.49.10.1407.17308.
- Chrysanthos Dellarocas. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management Science*, 52(10):1577–1593, 2006. doi: 10.1287/mnsc.1060.0567.
- Jessica Deyo. 81% of consumers embraced influencer marketing in the past year, study finds, February 2023. URL <https://www.marketingdive.com/news/influencer-marketing-success-matter-study-2023/643310/>. News article.
- Peter A. Diamond. A model of price adjustment. *Journal of Economic Theory*, 3(2):156–168, 1971. doi: 10.1016/0022-0531(71)90013-5.
- Quang Huy Duong, Li Zhou, Meng Meng, Le Thuy An Dang, and Tiep Duy Nguyen. Striking the right balance: Customising return policy leniency for managing customer online return

- proclivity and satisfaction. *Journal of Retailing and Consumer Services*, 85:104315, 2025. ISSN 0969-6989. doi: 10.1016/j.jretconser.2025.104315.
- EB Games Canada. Online returns and exchanges. Help Center article, July 2025. URL <https://ebgamescanada.zendesk.com/hc/en-us/articles/21002678426132-Online-Returns-and-Exchanges>. Published approximately 2 months ago.
- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007. doi: 10.1257/aer.97.1.242.
- Souad El Kihal and Edlira Shehu. It’s not only what they buy, it’s also what they keep: Linking marketing instruments to product returns. *Journal of Retailing*, 98(3):558–571, 2022. doi: 10.1016/j.jretai.2022.01.002.
- Glenn Ellison and Sara F. Ellison. Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452, 2009. doi: 10.3982/ECTA5708.
- Seyed Pouyan Eslami and Maryam Ghasemaghahi. Effects of online review positiveness and review score inconsistency on sales: A comparison by product involvement. *Journal of Retailing and Consumer Services*, 45:74–80, 2018. ISSN 0969-6989. doi: 10.1016/j.jretconser.2018.08.003.
- European Commission. Commission regulation on fare-parity clauses in air transport distribution. Regulation (EU) 2018/1117, 2018.
- European Parliament and Council. Regulation (EU) 2022/1925 on contestable and fair markets in the digital sector (Digital Markets Act). *Official Journal of the European Union*, L 265, 21 September 2022, pp. 1–66, 2022. URL <https://eur-lex.europa.eu/eli/reg/2022/1925/oj>. Entered into force 1 November 2022.
- Federal Trade Commission. Disclosures 101 for social media influencers, 2023. URL https://www.ftc.gov/system/files/ftc_gov/pdf/1001A_Influencer%20Guide_508.pdf.

- Ju Feng, Xiaoquan (Michael) Li, and Xiao Zhang. Online product reviews-triggered dynamic pricing: Theory and evidence. *Information Systems Research*, 30(4):1107–1123, 2019. doi: 10.1287/isre.2019.0852.
- Raffaele Filieri and Fraser McLeay. E-wom and accommodation: An analysis of the factors that influence travelers’ adoption of information from online reviews. *Journal of Travel Research*, 53(1):44–57, 2014. doi: 10.1177/0047287513481274.
- Financial Times. Us retail’s multibillion-dollar returns problem. <https://www.ft.com/content/0fc7e452-3664-435b-b5d4-68529bce1df3/>, 2024. Accessed: 2025-06-01.
- Forbes. Don’t buy casper, the latest money-losing ipo. *Forbes*, February 2020. URL <https://www.forbes.com/sites/greatspeculations/2020/02/05/dont-buy-casper-the-latest-money-losing-ipo/>. Accessed July 2025.
- Chris Forman, Anindya Ghose, and Batia Wiesenfeld. Examining the relationship between reviews and sales: The role of reviewer identity disclosure in electronic markets. *Information Systems Research*, 19(3):291–313, 2008. doi: 10.1287/isre.1080.0193.
- Adam Furness. Trust, credibility and authenticity: The influencer marketing blueprint, September 2023. URL <https://www.warc.com/newsandopinion/opinion/trust-credibility-and-authenticity-the-influencer-marketing-blueprint/en-gb/6347>. Opinion article.
- Santiago Gallino and Antonio Moreno. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science*, 60(6): 1434–1451, 2014. doi: 10.1287/mnsc.2014.1950.
- Santiago Gallino and Antonio Moreno. The value of fit information in online retail: Evidence from a randomized field experiment. *Manufacturing & Service Operations Management*, 20(4):767–787, 2018. doi: 10.1287/msom.2017.0686.

- Fei Gao and Xuanming Su. Online and offline information for omnichannel retailing. *Manufacturing & Service Operations Management*, 19(1):84–98, 2016. doi: 10.1287/msom.2016.0593.
- Sebastián J. García-Dastugue, Rahul Nilakantan, Carl Marcus Wallenburg, and Shashank Rao. On merchandise return policy, entrepreneurial internet retail, and customer reviews – insights from an observational study. *Journal of Business Logistics*, 2023. doi: 10.1111/jbl.12342.
- Sonja Gensler, Scott A. Neslin, and Peter C. Verhoef. The showrooming phenomenon: It’s more than just about price. *Journal of Interactive Marketing*, 38(1):29–43, 2017. doi: 10.1016/j.intmar.2017.01.003.
- David Godes and Dina Mayzlin. Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science*, 28(4):721–739, 2009. doi: 10.1287/mksc.1080.0444.
- David Godes, Dina Mayzlin, Yubo Chen, Sanjog Das, Chrysanthos Dellarocas, Bernd Pfeiffer, Barak Libai, Sharan Sen, Mengze Shi, and Peeter Verlegh. The firm’s management of social interactions. *Marketing Letters*, 16(3–4):415–428, 2005. doi: 10.1007/s11002-005-5902-4.
- Sanford J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law & Economics*, 24(3):461–483, 1981. doi: 10.1086/466104.
- Xu Guan, Yulan Wang, Zelong Yi, and Ying-Ju Chen. Inducing consumer online reviews via disclosure. Technical report, SSRN Electronic Journal, April 2018. Forthcoming in *Production and Operations Management*. Available at SSRN: <https://ssrn.com/abstract=3571985> or <http://dx.doi.org/10.2139/ssrn.3571985>.
- V. Daniel R. Guide and Luk N. Van Wassenhove. The evolution of closed-loop supply chain research. *Operations Research*, 57(1):10–18, 2009. doi: 10.1287/opre.1080.0628.
- Liang Guo. Service cancellation and competitive refund policy. *Marketing Science*, 28(5): 901–917, 2009.
- Jamie Gutfreund. How is ltk’s advanced technology driving \$5b in affiliate sales? <https://www.forbes.com/sites/jamiegutfreund/2024/12/12/>

- [how-is-ltks-advanced-technology-driving-5b-in-affiliate--sales/](#), 2024. Accessed: 2025-06-01.
- Lin Hao and Yong Tan. Who wants consumers to be informed? facilitating information disclosure in a distribution channel. *Information Systems Research*, 30(1):34–49, 2019. doi: 10.1287/isre.2017.0770.
- James D. Hess, Wujin Chu, and Eitan Gerstner. Controlling product returns in direct marketing. *Marketing Letters*, 7(4):307–317, 1996. ISSN 0923-0645. doi: 10.1007/BF00435538.
- Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929. doi: 10.2307/2224214.
- Lu Hsiao and Ying-Ju Chen. Return policy: Hassle-free or your money-back guarantee? *Naval Research Logistics*, 61(5):403–417, 2014.
- Nan Hu, Paul A. Pavlou, and Jie (Jennifer) Zhang. Overcoming the j-shaped distribution of product reviews. *Communications of the ACM*, 52(10):144–147, 2009. doi: 10.1145/1562764.1562800.
- Nan Hu, Paul A. Pavlou, and Jie Zhang. On self-selection biases in online product reviews. *MIS Quarterly*, 41(2):449–475, 2017.
- Xiao Huang and Dan Zhang. Service product design and consumer refund policies. *Marketing Science*, 39(2):366–381, 2020.
- HubSpot and Brandwatch. 2022 state of u.s. consumer trends report. <https://offers.hubspot.com/2022-consumer-trends-report-download>, 2022. Retrieved July 5, 2023.
- Influencer Marketing Hub. Influencer marketing hub: Influencer marketing benchmark report 2024. <https://influencermarketinghub.com/influencer-marketing-benchmark-report/>, 2024. Accessed: 2025-06-01.
- Raja Ahmed Jamil and Tariq Iqbal Khan. Lenient return policies and religiosity: enhancing

- consumer confidence, well-being and intentions. *Journal of Islamic Marketing*, ahead-of-print (ahead-of-print), 2025. doi: 10.1108/JIMA-03-2024-0143.
- Narayan Janakiraman, Holly A. Syrdal, and Ryan Freling. The effect of return policy leniency on consumer purchase and return decisions: A meta-analytic review. *Journal of Retailing*, 92(2):226–235, 2016. ISSN 0022-4359. doi: 10.1016/j.jretai.2015.11.002.
- Ginger Zhe Jin and Phillip Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451, 2003. doi: 10.1162/003355303321675428.
- Eric J. Johnson, Steven Bellman, and Gerald Lohse. Cognitive lock-in and the power law of practice. *Journal of Marketing*, 67(2):62–75, 2003. Available at SSRN: <https://ssrn.com/abstract=1324766>.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. doi: 10.2307/1914185.
- Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6):1325–1348, 1990. doi: 10.1086/261737.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011. doi: 10.1257/aer.101.6.2590.
- Suzanne Kapner. Shoppers fight back as retailers crack down on returns. *The Wall Street Journal*, 2023. URL <https://www.wsj.com/business/retail/online-retail-return-policy-fees-8444b5fe>. Published online November 2023; accessed 2025-06-01.
- Jeong-Yoo Kim and Wei Xu. A signaling theory of the online consumer review policy. *Review of Network Economics*, 22(2):53–84, 2023. doi: 10.1515/rne-2023-0041.
- Botond Köszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006. doi: 10.1093/qje/121.4.1133.

- Botond Kőszegi and Matthew Rabin. Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073, 2007. doi: 10.1257/aer.97.4.1047.
- KPMG International. Making returns pay: The hidden cost and carbon footprint of product returns. Technical report, KPMG International Cooperative, 2024. Industry white paper.
- Dmitri Kuksov and Ying Lin. Information provision in a vertically differentiated competitive marketplace. *Marketing Science*, 29(1):122–138, 2010. URL <http://www.jstor.org/stable/40608094>.
- Jungmi Lee, Nathan Walter, James L. Hayes, and Guy J. Golan. Do influencers influence? a meta-analytic comparison of celebrities and social media influencers’ effects. *Social Media + Society*, 10(3), 2024. doi: 10.1177/20563051241269269. URL <https://doi.org/10.1177/20563051241269269>. Original work published 2024.
- Sangwoo Lee, Liangfei Qiu, and Andrew B. Whinston. Sentiment manipulation in online platforms: An analysis of movie tweets. *Production and Operations Management*, 27(3): 393–416, 2018. doi: 10.1111/poms.12805.
- Filipe P. Leite, Nayara Pontes, and Bruno Schivinski. Influencer marketing effectiveness: giving competence, receiving credibility. *Journal of Travel & Tourism Marketing*, 41(3):307–321, 2024. doi: 10.1080/10548408.2024.2317748.
- Randall A. Lewis and Justin M. Rao. The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, 130(4):1941–1973, 2015. doi: 10.1093/qje/qjv023.
- Yiming Li, Gang Li, Giri Kumar Tayi, and T. C. E. Cheng. Omni-channel retailing: Do offline retailers benefit from online reviews? *International Journal of Production Economics*, 218 (1):43–61, 2019. doi: 10.1016/j.ijpe.2019.05.002.
- Zhi Lin, Yiting Zhang, and Yong Tan. An empirical study of free product sampling and rating bias. *Information Systems Research*, 30(1):260–275, 2019. doi: 10.1287/isre.2018.0801.

- Laurent Linnemer. Price and advertising as signals of quality when some consumers are informed. *International Journal of Industrial Organization*, 20(7):931–947, 2002. ISSN 0167-7187. doi: [https://doi.org/10.1016/S0167-7187\(01\)00081-9](https://doi.org/10.1016/S0167-7187(01)00081-9).
- Lionize. What top brand categories use influencer marketing. Blog post, January 2025. URL <https://lionize.ai/post/what-top-brand-categories-use-influencer-marketing/>.
- Xiaoyu Liu and Xiaohui Zheng. The persuasive power of social media influencers in brand credibility and purchase intention. *Humanities and Social Sciences Communications*, 11:15, 2024. doi: 10.1057/s41599-023-02512-1.
- Michael Luca and Georgios Zervas. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science*, 62(12):3412–3427, 2016. doi: 10.1287/mnsc.2015.2304.
- Marketing Dive. L.l.beam shifts marketing focus from print to digital & tv, November 2020. URL <https://www.marketingdive.com/news/ll-beam-shifts-marketing-focus-from-print-to-digital-tv/447507/>.
- Dina Mayzlin. Promotional chat on the internet. *Marketing Science*, 25(2):155–163, 2006. doi: 10.1287/mksc.1050.0137.
- Dina Mayzlin, Yaniv Dover, and Judith Chevalier. Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review*, 104(8):2421–2455, 2014. doi: 10.1257/aer.104.8.2421.
- McKinsey & Company. Returning to order: Improving returns management for apparel companies, 2021. URL <https://www.mckinsey.com/industries/retail/our-insights/returning-to-order-improving-returns-management-for-apparel-companies>. Industry white paper.
- McKinsey & Company. What is influencer marketing? April 10, 2023. <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-influencer-marketing#/> Retrieved on July 5, 2023., 2023.

- Bruce McWilliams. Money-back guarantees: Helping the low-quality retailer. *Management Science*, 58(8):1521–1524, 2012. doi: 10.1287/mnsc.1110.1497.
- Miriam J. Metzger and Andrew J. Flanagin. Credibility and trust of information in online environments. In S. Shyam Sundar, editor, *The Handbook of the Psychology of Communication Technology*, pages 445–466. Wiley, 2013.
- Nathaniel Meyersohn. Free returns are going away, 2023. URL <https://www.cnn.com/2023/12/26/business/returns-shopping-gifts-charge>. 2025-06-04.
- Paul R. Milgrom. Rational expectations, information acquisition, and competitive bidding. *Econometrica*, 49(4):921–943, 1981. doi: 10.2307/1912511.
- Annemiek Minnema, Tammo H. A. Bijmolt, Sonja Gensler, and Thorsten Wiesel. To keep or not to keep: Effects of online customer reviews on product returns. *Journal of Retailing*, 92(3):253–267, 2016. doi: 10.1016/j.jretai.2016.03.001.
- Wendy W. Moe and Michael Trusov. The value of social dynamics in online product ratings forums. *Journal of Marketing Research*, 48(3):444–456, 2011. doi: 10.1509/jmkr.48.3.444.
- Sridhar Moorthy and Kannan Srinivasan. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Science*, 14(4):442–466, 1995. ISSN 0732-2399. URL <http://www.jstor.org/stable/184139>.
- Susan M. Mudambi and David Schuff. What makes a helpful review? a study of customer reviews on amazon.com. *MIS Quarterly*, 34(1):185–200, 2010. doi: 10.2307/20721420.
- Leela Nageswaran, Soo-Haeng Cho, and Alan Scheller-Wolf. Consumer return policies in omnichannel operations. *Management Science*, 66(12):5558–5575, 2020.
- Narvar Inc. Returns benchmark report. Industry white paper, Narvar Inc., 2022. URL <https://corp.narvar.com/returns-benchmark-report>.
- National Retail Federation and Happy Returns. 2024 consumer returns in the retail industry.

- <https://nrf.com/research/2024-consumer-returns-retail-industry>, 2024. Accessed: 2025-12-18.
- New York Post. What turbulence is like flying first class: ‘turnt drinking free flow dom perignon’. October 3, 2023. <https://nypost.com/2023/10/03/what-turbulence-is-like-flying-first-class-turnt-drinking-free-flow-dom-perignon/>. Retrieved on November 18, 2023., 2023a.
- New York Post. First-class canine is flying in style. Video article on New York Post, July 2023b. URL <https://nypost.com/video/first-class-canine-is-flying-in-style/>. Accessed July 11, 2025.
- Pejvak Oghazi, Stefan Karlsson, Daniel Hellström, and Klas Hjort. Online purchase return policy leniency and purchase decision: Mediating role of consumer trust. *Journal of Retailing and Consumer Services*, 41:190–200, 2018. ISSN 0969-6989. doi: 10.1016/j.jretconser.2017.12.007.
- Myle Ott, Claire Cardie, and Jeff Hancock. Estimating the prevalence of deception in online review communities. *arXiv preprint arXiv:1204.2804*, 2012. doi: 10.48550/arXiv.1204.2804.
- Nikolaos Pappas. Marketing strategies, perceived risks, and consumer trust in online buying behaviour. *Journal of Retailing and Consumer Services*, 29:92–103, 2016. doi: 10.1016/j.jretconser.2015.11.007.
- Anran Pei and Dina Mayzlin. Influencing social media influencers through affiliation. *Marketing Science*, 41(3):593–615, 2022. doi: 10.1287/mksc.2021.1322.
- PowerReviews. 2023 apparel shopping trends: Merchandise returns. <https://www.powerreviews.com/research/apparel-shopping-trends-2023/merchandise-returns/>, 2023. Accessed: 2025-06-01.
- PowerReviews Team. Consumer survey: Returns in retail in 2021. PowerReviews Web page, June 2021. URL <https://www.powerreviews.com/consumer-survey-retail-returns-2021/>. Based on insights from 7,688 U.S. consumers.

- Thomas L. Powers and Eric P. Jack. The influence of cognitive dissonance on retail product returns. *Psychology & Marketing*, 30(8):724–735, 2013. doi: 10.1002/mar.20640.
- Purple Innovation Inc. Purple innovation reports fourth quarter and full year 2022 results. Press release, March 2023. URL <https://investors.purple.com/news/news-details/2023/Purple-Innovation-Reports-Fourth-Quarter-and-Full-Year-2022-Results/default.aspx>. Advertising spend down 77.7 % YoY; accessed 11 July 2025.
- Radial. Online fashion retailers’ guide to reducing returns in 2024. <https://www.radial.com/insights/online-fashion-retailers-guide-to-reducing-returns-in-2024>, 2024. Accessed: 2025-06-01.
- Neha Rashid, Ali Imtiaz, Jamshid Ali, and Shahryar Sorooshian. Mitigating the impact of purchase uncertainty on reverse supply chain through impulsive consumer buying with moderation of lenient return policies. *Cleaner Logistics and Supply Chain*, 14:100207, 2025. ISSN 2772-3909. doi: 10.1016/j.clscn.2025.100207.
- Brian T. Ratchford. The economics of consumer knowledge. *Journal of Consumer Research*, 27(4):397–411, 2001. ISSN 0093-5301. doi: 10.1086/319617.
- Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010. doi: 10.1086/657922.
- Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay’s reputation system. In Michael R. Baye, editor, *The Economics of the Internet and E-Commerce*, volume 11 of *Advances in Applied Microeconomics*, pages 127–157. JAI Press/Elsevier, 2002.
- Dale S. Rogers and Ronald S. Tibben-Lembke. An examination of reverse logistics practices. *Journal of Business Logistics*, 22(2):129–148, 2001. doi: 10.1002/j.2158-1592.2001.tb00007.x.
- Andrea B. Rosario, Francesca Sotgiu, Katja De Valck, and Tammo H. A. Bijmolt. The effect of

- electronic word of mouth on sales: A meta-analytic review of platform, product, and metric factors. *Journal of Marketing Research*, 53(3):297–318, 2016. doi: 10.1509/jmr.14.0380.
- Nachiketa Sahoo, Chrysanthos N. Dellarocas, and Shuba Srinivasan. The impact of online product reviews on product returns. Research Paper No. 2491276, Boston University School of Management, March 2017. Available at SSRN: <https://ssrn.com/abstract=2491276> or <http://dx.doi.org/10.2139/ssrn.2491276>.
- Samsung Electronics. What is samsung’s return policy? Samsung Customer Support, May 2025. URL <https://order-help.us.samsung.com/hc/en-us/articles/360057028614-What-is-Samsung-s-return-policy>. Accessed July 11, 2025.
- Seeking Alpha Contributor. Purple and casper: A tale of two mattresses. *Seeking Alpha*, March 2019. URL <https://seekingalpha.com/article/4375416-purple-and-casper-a-tale-of-two-mattresses>. Accessed July 2025.
- Jeffrey D. Shulman, Anne T. Coughlan, and R. Canan Savaskan. Optimal reverse channel structure for consumer product returns. *Marketing Science*, 29(6):1071–1085, 2010.
- Jeffrey D. Shulman, Anne T. Coughlan, and R. Canan Savaskan. Managing consumer returns in a competitive environment. *Management Science*, 57(2):347–362, 2011. ISSN 0025-1909. doi: 10.1287/mnsc.1100.1274.
- Paul Skeldon. Under the influence: What influences consumer purchases uncovered, June 2023. URL <https://internetretailing.net/under-the-influence-understanding-what-influences-consumer-purchases-uncovered/>.
- SocialPubli. Socialpubli — influencer marketing platform. Website, 2025. URL <https://socialpubli.com/>. Accessed July 11, 2025.
- Beverley A. Sparks, Kevin Kam Fung So, and Graham L. Bradley. Responding to negative online reviews: The effects of hotel responses on customer inferences of trust and concern. *Tourism Management*, 53:74–85, 2016. ISSN 0261-5177. doi: 10.1016/j.tourman.2015.09.011.

- Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973. doi: 10.2307/1882010.
- II Stahl, Dale O. Oligopolistic Pricing with Sequential Consumer Search. *American Economic Review*, 79(4):700–712, 1989.
- Statista. Global influencer marketing market size 2017–2024. <https://www.statista.com/statistics/1092819/global-influencer-market-size/>, 2024. Accessed: 2025-06-01.
- Statista. Clothing & shoes are the most returned online purchases. Statista Web Chart, April 2025. URL <https://www.statista.com/forecasts/997235/most-returned-online-purchases-by-category-in-the-us>. 25% of U.S. online shoppers returned clothing in the past year, 17% returned shoes.
- Sanja Maria Stuhldreier. Unlocking (re)purchase potential through corporate responsiveness on social networks: The role of perceived customer orientation. *Journal of Retailing and Consumer Services*, 81:104041, 2024. doi: 10.1016/j.jretconser.2024.104041.
- Xuanming Su. Consumer return policies and supply chain performance, 2009.
- Miao Sun, Jing Chen, Ye Tian, and Yufei Yan. The impact of online reviews in the presence of customer returns. *International Journal of Production Economics*, 232:107929, 2021. doi: 10.1016/j.ijpe.2020.107929.
- Richard H. Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, 1(1):39–60, 1980. doi: 10.1016/0167-2681(80)90051-7.
- Catherine E. Tucker. The reach and persuasiveness of viral video ads. *Marketing Science*, 34(2):281–296, 2015. doi: 10.1287/mksc.2014.0874.
- Tuft & Needle. What is the 100-night sleep trial. Customer Help Article, 2025. Accessed July 2025.

- UK Competition and Markets Authority. Fake online reviews research. Technical report, UK Government, 2023. URL <https://assets.publishing.service.gov.uk/media/6447c00c529eda000c3b03c5/fake-online-reviews-research.pdf>.
- UPS. Pulse of the online shopper 2021: Returns experience matters, 2021. URL <https://about.ups.com/content/dam/upsstories/assets/documents/pulse-of-the-online-shopper-2021.pdf>. Consumer survey showing tech-savvy shoppers incur lower hassle costs when returning online purchases.
- Hal R. Varian. A model of sales. *The American Economic Review*, 70(4):651–659, 1980. doi: 10.2307/1803562.
- VentureTwins. Four things to learn about d2c economics from casper's s-1. *Medium*, 2019. URL <https://medium.com/@venturetwins/four-things-to-learn-about-d2c-economics-from-caspers-s-1-9e117e446cc1>. Accessed July 2025.
- Peter C. Verhoef, P. K. Kannan, and Jeffrey J. Inman. From multi-channel retailing to omni-channel retailing: Introduction to the special issue on multi-channel retailing. *Journal of Retailing*, 91(2):174–181, 2015. doi: 10.1016/j.jretai.2015.02.005.
- Jeremy A. Verlinda. Do rockets rise faster and feathers fall slower in an atmosphere of local market power? *Journal of Industrial Economics*, 56(3):581–612, 2008.
- Vogue Business. The business of influence: Why the era of celebrity influencers is over, 2023. URL <https://www.voguebusiness.com/companies/the-business-of-influence-why-the-era-of-celebrity-influencers-is-over>. Accessed July 11, 2025.
- Asher Wolinsky. True Monopolistic Competition as a Result of Imperfect Information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986.
- Stacey L. Wood. Remote purchase environments: The influence of return policy leniency on

- two-stage decision processes. *Journal of Marketing Research*, 38(2):157–169, 2001. doi: 10.1509/jmkr.38.2.157.18848.
- Jinhong Xie and Eitan Gerstner. Service escape: Profiting from customer cancellations. *Marketing Science*, 26(1):18–30, 2007.
- Karen L. Xie, Zili Zhang, and Ziqiong Zhang. The business value of online consumer reviews and management response to hotel performance. *International Journal of Hospitality Management*, 43:1–12, 2014. ISSN 0278-4319. doi: 10.1016/j.ijhm.2014.07.007.
- Manjit S. Yadav and Paul A. Pavlou. Marketing in computer-mediated environments: Research synthesis and new directions. *Journal of Marketing*, 78(1):20–40, 2014. doi: 10.1509/jm.12.0020.
- Jeremy Yang, Juanjuan Zhang, and Yuhan Zhang. Engagement that sells: Influencer video advertising on tiktok. *Marketing Science*, 44(2):247–267, 2024. doi: 10.1287/mksc.2021.0107.
- Feng Zhu and Xiaoquan M. Zhang. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of Marketing*, 74(2):133–148, 2010. doi: 10.1509/jm.74.2.133.
- Mengzhou Zhuang, Geng Cui, and Ling Peng. Manufactured opinions: The effect of manipulating online product reviews. *Journal of Business Research*, 87:24–35, 2018. doi: 10.1016/j.jbusres.2018.02.016.

Appendix A

Chapter 3

Lemma A.1. (i) $\frac{\partial \rho_g(a)}{\partial a} > 0$, $\frac{\partial \rho_g(a)}{\partial k} > 0$ and $\rho_g(1) = k$. (ii) $\frac{\partial \tilde{\rho}_g(a)}{\partial a} > 0$, $\frac{\partial \tilde{\rho}_g(a)}{\partial \tilde{k}} > 0$ and $\tilde{\rho}_g(1) = \tilde{k}$.

Proof. (i) Since $\gamma \geq 1/2 \geq \bar{\gamma}$, there is $\frac{\partial \rho_g(a)}{\partial a} = \frac{k\bar{k}(\gamma^2 - \bar{\gamma}^2)}{[k(\gamma + a\bar{\gamma}) + k(1 - \gamma\bar{a})]^2} \geq 0$. Also, $\rho_g(a) = \frac{k(1 - \gamma\bar{a})}{\gamma + a\bar{\gamma} - k\bar{a}(\gamma - \bar{\gamma})}$ clearly increases in k . With $a = 1$ and $\bar{a} = 0$, it can easily be verified that $\rho_g(1) = \frac{k}{k+k} = k$. (ii) The proof is similar to (i) but replacing k with $\tilde{k}$. ■

Lemma A.2. $P(G|g, a)$ decreases in a .

Proof. Note that $P(G|g, a) = \frac{(\gamma + a\bar{\gamma})\rho_0}{(\gamma + a\bar{\gamma})\rho_0 + (\bar{\gamma} + a\gamma)\bar{\rho}_0} = \frac{\gamma\rho_0 + a\bar{\gamma}\rho_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0 + a(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0)}$. Since $\gamma > 1/2$, there is $\frac{\gamma\rho_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0} \geq \bar{\rho}_0 \geq \frac{\bar{\gamma}\bar{\rho}_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0}$. Hence $\gamma\rho_0(\bar{\gamma}\bar{\rho}_0 + \gamma\bar{\rho}_0) \geq \bar{\gamma}\bar{\rho}_0(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0)$. Then, $\frac{\partial P(G|g, a)}{\partial a} = \frac{\bar{\gamma}\bar{\rho}_0[\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0 + a(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0)] - (\bar{\gamma}\bar{\rho}_0 + \gamma\bar{\rho}_0)(\gamma\rho_0 + a\bar{\gamma}\rho_0)}{[\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0 + a(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0)]^2} = \frac{\bar{\gamma}\bar{\rho}_0(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0) - (\bar{\gamma}\bar{\rho}_0 + \gamma\bar{\rho}_0)\gamma\rho_0}{[\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0 + a(\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0)]^2} \leq 0$. ■

Lemma A.3. $P(G|g, a) \in [\rho_0, \frac{\gamma\rho_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0}]$, $P(B|g, a) \in [\frac{\bar{\gamma}\bar{\rho}_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0}, \bar{\rho}_0]$.

Proof. By Lemma A.2, $P(G|g, a)$ decreases in a and $P(B|g, a) = 1 - P(G|g, a)$ increases in a . Therefore, $\rho_0 = P(G|g, 1) \leq P(G|g, a) \leq P(G|g, 0) = \frac{\gamma\rho_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0}$, and $\frac{\bar{\gamma}\bar{\rho}_0}{\gamma\rho_0 + \bar{\gamma}\bar{\rho}_0} = 1 - P(G|g, 0) = P(B|g, 0) \leq P(B|g, a) \leq P(B|g, 1) = 1 - P(G|g, 1) = \bar{\rho}_0$. ■

Lemma 3.1

Proof. Since $\gamma \geq 1/2 \geq \bar{\gamma}$, there is $1 - \gamma\bar{a} \leq \gamma + a\bar{\gamma}$. Thus $\rho_g(a) = \frac{k(1-\gamma\bar{a})}{k(\gamma+a\bar{\gamma})+k(1-\gamma\bar{a})} \leq \frac{k}{k+k} = k$ and $\rho_b = \frac{k\gamma}{k\bar{\gamma}+k\gamma} \geq \frac{k}{k+k} = k$. ■

Proposition 3.1

Proof. Based on Lemma 3.1, the firm's optimal influencing decision can be analyzed as follows:

(i). If $\rho_0 \in [0, \rho_g(0))$, the customer will never buy the product; hence no influence tool will be acquired and $\Pi_0^* = 0$.

(ii). If $\rho_0 \in [\rho_g(0), k)$, the firm can go without an influence tool and earns an expected profit $\Pi_0(N) = 0$, or it can acquire the influence tool and impose a low level of influence such that $\rho_0 < \rho_g(a)$ and earns $\Pi_0(0) = 0 - c$, or impose a high level of influence satisfying $\rho_g(a) \leq \rho_0$ such that the customer only purchases when the signal $s = g$ with expected profit $\Pi_0(a) = P(g|a)p - c$ where

$$P(g|a) = P(g|G, a)P(G) + P(g|B, a)P(B) = (\gamma + a\bar{\gamma})\rho_0 + (a\gamma + \bar{\gamma})\bar{\rho}_0. \quad (\text{A.1})$$

In the latter case, the firm's optimal influencing level can be solved by

$$\begin{aligned} \Pi_0(a_0^*) &= \max_{0 \leq a \leq 1} P(g|a)p - c \\ \text{s.t.} \quad &\rho_g(a) \leq \rho_0. \end{aligned} \quad (\text{A.2})$$

Since $P(g|a) = (\gamma + \rho_0 - 2\gamma\rho_0)a + 1 - (\gamma + \rho_0 - 2\gamma\rho_0)$ increases in a , and $\rho_g(a)$ increases in a , the constraint should be binding. Solving $\rho_g(a_0^*) = \rho_0$ yields

$$a_0^* = \frac{\frac{\gamma}{\rho_0 k} - \frac{\bar{\rho}_0 k}{\rho_0 k}}{\frac{\gamma}{\rho_0 k} - 1}, \quad (\text{A.3})$$

which falls within $[0, 1]$ as $\rho_g(0) \leq \rho_0 < k$.

With $\rho_g(a_0^*) = \rho_0$, there is $P(G|g, a_0^*) = k$. Thus $P(g|a_0^*) = \frac{P(g|G, a_0^*)P(G)}{P(G|g, a_0^*)} = \frac{(\gamma + a_0^* \gamma) \rho_0}{k} = \frac{(\frac{\gamma}{\gamma} - 1) \bar{\rho}_0}{\frac{\gamma}{\gamma} \frac{\rho_0 k}{\gamma} - 1} \frac{\bar{\rho}_0}{k}$, and $\Pi_0(a_0^*) = P(g|a_0^*)p - c = \frac{(\frac{\gamma}{\gamma} - 1) \bar{\rho}_0}{\frac{\gamma}{\gamma} \frac{\rho_0 k}{\gamma} - 1} \frac{\bar{\rho}_0}{k} p - c = \frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \rho_0 v - c$. The firm will acquire the influence tool if and only if $\gamma \geq \Gamma$ and $\Pi_0(a_0^*) \geq \Pi_0(N) = 0$, or equivalently, when $c \leq C$ where

$$C = \frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \rho_0 v. \quad (\text{A.4})$$

Overall, in this scenario, the firm will acquire the influence tool and set influence level a_0^* if and only if $c \leq C$. Otherwise, there will be no influence tool acquisition and the customer will not purchase.

(iii). If $\rho_0 \in [k, \rho_b)$, the customer buys when $s = g$ or there is no signal, and will not buy if $s = b$. Then $\Pi_0(a) \leq p - c \leq p = \Pi_0(N)$. If $\rho_0 \in [\rho_b, 1]$, the consumer will always buy in all scenarios. Clearly, there is no need to acquire an influence tool either. Therefore, when $\rho_0 \geq k$, the firm will not acquire an influence tool and the customer will always buy. ■

Proof. Corollary 3.1. According to Proposition 3.1, when $\rho_0 < \rho_g(0)$, there will be no purchase. The customer's expected utility $U_0 = 0$. When $\rho_g(0) \leq \rho_0 < k$, without an influence tool, the customer's expected utility is 0 as before. With the influence tool and the level of influence set to a_0^* , the expected utility of the customer is $U_0 = P(g|a_0^*)[P(G|g, a_0^*)v - p]$. Since $\rho_g(a_0^*) = \rho_0$, there is $P(G|g, a_0^*)v - p = 0$. Thus the customer's expected utility is again $U_0 = 0$. When $\rho_0 \geq k$, the customer will always purchase based on prior belief. The expected utility is then $U_0 = \rho_0 v - p \geq 0$. ■

Lemma 3.2

Proof. Similarly to the proof of Lemma 3.1, it can be shown that $\tilde{\rho}_g(a) \leq \tilde{k} \leq \tilde{\rho}_b$ by replacing k with $\tilde{k}$. ■

Proof. Corollary 3.2. $\tilde{k} = \frac{p-r}{v-r} = 1 - \frac{v-p}{v-r}$ decreases in r . $\tilde{\rho}_g(a) = \frac{\tilde{k}(1-\gamma\bar{a})}{k(\gamma+a\bar{\gamma})+\tilde{k}(1-\gamma\bar{a})} = \frac{1-\gamma\bar{a}}{\bar{a}\bar{\gamma}-\bar{a}\gamma+(\gamma+a\bar{\gamma})/\tilde{k}}$ increases in $\tilde{k}$, and thus decreases in r . $\tilde{\rho}_b = \frac{\tilde{k}\gamma}{k\bar{\gamma}+\tilde{k}\gamma} = \frac{\gamma}{\gamma-\bar{\gamma}+\bar{\gamma}/\tilde{k}}$ increases in $\tilde{k}$, and thus decreases in r . It can be verified that $k = \tilde{k}|_{r=0}$, $\rho_g(a) = \tilde{\rho}_g(a)|_{r=0}$, and $\rho_b = \tilde{\rho}_b|_{r=0}$. Thus the inequalities hold $\forall r \in (0, p]$. ■

Proposition 3.2

Proof. Based on Lemma 3.2,

- (i). If $\rho_0 \in [0, \tilde{\rho}_g(0))$, the customer will never purchase; hence the firm will not take any action and $\Pi_r^* = 0$.
- (ii). If $\rho_0 \in [\tilde{\rho}_g(0), \tilde{k})$, the firm may choose not to get the influence tool and live with $\Pi_r(N) = 0$, or acquire the tool and set a low level of influence such that $\rho_0 < \tilde{\rho}_g(a)$ and earns $\Pi_r(a) = 0 - c < 0$, or a high level of influence satisfying $\tilde{\rho}_g(a) \leq \rho_0$ such that the customer purchases only when the signal produces $s = g$, which in this case yields an expected profit $\Pi_r(a) = P(g|a) [p - P(B|g, a)r] - c$ where $P(g|a)$ follows Equation (A.1). Solving the following optimization problem gives rise to the optimal influence level a_r^* in the last option:

$$\begin{aligned} \Pi_r(a_r^*) = \max_{0 \leq a \leq 1} & P(g|a) [p - P(B|g, a)r] - c \\ \text{s.t.} & \tilde{\rho}_g(a) \leq \rho_0. \end{aligned} \quad (\text{A.5})$$

By Lemma A.2, $[p - P(B|g, a)r] = p - r + P(G|g, a)r$ increases in a . Also, both $P(g|a)$ and $\tilde{\rho}_g(a)$ increase in a . There should be $\tilde{\rho}_g(a_r^*) = \rho_0$, i.e.,

$$a_r^* = \left(\frac{\gamma}{\bar{\gamma}} - \frac{\bar{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} \right) / \left(\frac{\gamma}{\bar{\gamma}} \frac{\bar{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} - 1 \right), \quad (\text{A.6})$$

which falls within $[0, 1]$ as $\tilde{\rho}_g(0) \leq \rho_0 < \tilde{k}$. Further, there are $P(G|g, a_r^*) = \tilde{k}$, $P(g|a_r^*) = \frac{P(g|G, a_r^*)P(G)}{P(G|g, a_r^*)} = \frac{(\gamma+a_r^*\bar{\gamma})\rho_0}{k} = \frac{(\frac{\gamma}{\bar{\gamma}}-1)\bar{\rho}_0}{\frac{\gamma}{\bar{\gamma}}\frac{\bar{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}}-1} \frac{\bar{\rho}_0}{k}$, and the expected profit associated with acquiring the tool and setting a high level of influence is $\Pi_r(a_r^*) = P(g|a_r^*)[p - P(B|g, a_r^*)r] - c =$

$\frac{\frac{\gamma}{\gamma}-1}{\frac{\gamma}{\gamma}-\frac{\rho_0 k}{\rho_0 k}-1} \frac{\bar{\rho}_0}{k} (p - \bar{k}r) - c = \frac{\frac{\gamma}{\gamma}-1}{\frac{\gamma}{\gamma}-\frac{\rho_0 k}{\rho_0 k}-1} \frac{\bar{\rho}_0}{k} v \tilde{k} - c = \frac{\frac{\gamma}{\gamma}-1}{\frac{\gamma}{\gamma}-\frac{\rho_0 k}{\rho_0 k}} \rho_0 v - c$. The firm will choose this option if $\Pi_r(a_r^*) \geq \Pi_r(N) = 0$, or equivalently, $c \leq \tilde{C}$ where

$$\tilde{C} = \frac{(\frac{\gamma}{\gamma}-1)}{\frac{\gamma}{\gamma}-\frac{\rho_0 k}{\rho_0 k}} \rho_0 v. \quad (\text{A.7})$$

Overall, if $c \leq \tilde{C}$, the firm will acquire the influence tool and set the influence level to a_r^* ; otherwise, it will not acquire the tool and the customer will not purchase.

(iii). If $\rho_0 \in [\tilde{k}, \tilde{\rho}_b)$, the customer purchases when $s = g$ or there is no signal, and will not buy if $s = b$. The firm may either not acquire a tool and end up with expected profit $\Pi_r(N) = p - \bar{\rho}_0 r$, or acquire a tool and set the influence level to some a with expected profit $\Pi_r(a) = P(g|a) [p - P(B|g, a)r] - c$. Note that $\Pi_r(a) = P(g|a) [p - P(B|g, a)r] - c \leq P(g|a) [p - P(B|g, a)r] + P(b|a) [p - P(B|b, a)r] - c = p - \bar{\rho}_0 r - c \leq p - \bar{\rho}_0 r = \Pi_r(N)$. Therefore, the firm would not acquire a tool and the customer will always purchase.

If $\rho_0 \in [\tilde{\rho}_b, 1]$, the consumer will always buy, regardless of the signals. Similarly, there is then no need to acquire an influence tool. ■

Proof. Corollary 3.3. Consider $r' > r$. By Corollary 3.2, $\tilde{k}' < \tilde{k}$. By Proposition 3.2 (iii), if the prior $\rho_0 \in (\max\{\tilde{k}', \tilde{\rho}_g(0)\}, \tilde{k})$, the firm should discontinue the use of influence tool; otherwise, for $\rho_0 \in (\tilde{\rho}_g(0), \max\{\tilde{k}', \tilde{\rho}_g(0)\})$, the influence should continue. In the latter case:

(i). $\tilde{C} = \frac{(\frac{\gamma}{\gamma}-1)}{\frac{\gamma}{\gamma}-\frac{\rho_0 k}{\rho_0 k}} \rho_0 v$ decreases in $\tilde{k}$, and thus increases in r . $\tilde{\Gamma} = \frac{\bar{\rho}_0 \tilde{k}}{\bar{\rho}_0 \tilde{k} + \rho_0 k}$ increases in $\tilde{k}$, and thus decreases in r . Given the above, $C = \tilde{C}|_{r=0} \leq \tilde{C}$ and $\Gamma = \tilde{\gamma}|_{r=0} \geq \tilde{\Gamma}$.

(ii). By Equation A.6, a_r^* decreases in $\tilde{k}$, and hence increases in r . Thus $a_0^* \leq a_r^*$. ■

Proof. Corollary 3.4. (i). Following Proposition 3.2, the firm's expected profit $\Pi_r = \tilde{C} - c$ or 0 when $\rho_0 \leq \tilde{k}$, and $\Pi_r = p - \bar{\rho}_0 r$ or $\Pi_r = 0$ otherwise. By (i), $\Pi_r = \tilde{C} - c$ increases in r , and $\Pi_r = p - \bar{\rho}_0 r$ decreases in r . This proves the statement.

(ii). The proof is similar to Corollary 3.1. When $\rho_0 \geq \tilde{k}$, $U_r = \rho_0 v + \bar{\rho}_0 r - p \geq 0$ and increases in r . ■

Proposition 3.3

Proof. Following Proposition 3.2, with full refund $r = p$, there will be $\tilde{\rho}_g(0) = \tilde{k} = 0$. Comparing that with $r = 0$ (Proposition 3.1), we have:

(i). When $\rho_0 \in [0, \rho_g(0))$, with no refund, the firm's expected profit is $\Pi_0 = 0$. With full refund, no influence tool will be acquired and $\Pi_1 = \rho_0 p$. Clearly, the seller will go with the full refund $r^* = p$ and without an influence tool in this case.

(ii). When $\rho_0 \in [\rho_g(0), k)$, with no refund, the seller will use the influence tool at level $a_0^* \geq 0$ with expected profit $\Pi(a_0^*) = C - c$ when $c \leq C$ and $\gamma \geq \Gamma$. With full refund, the firm will not acquire an influence tool and earn an expected profit $\Pi = \rho_0 p$. Comparing the two, the firm will choose no refund if and only if $c \leq C - \rho_0 p = \rho_0 p \left[\frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \frac{1}{k} - 1 \right]$. To ensure $\rho_0 p \left[\frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \frac{1}{k} - 1 \right] \geq 0$, we need $\frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \geq k$, which is equivalent to $\gamma \geq \Gamma_f = \frac{1 - \frac{\rho_0 k}{\rho_0}}{2 - k - \frac{\rho_0 k}{\rho_0}}$. Since $\rho_0 < k$, it can be verified that $\Gamma_f \in [\frac{1}{2}, 1]$. Denote $C_f = \max\{0, \rho_0 p \left[\frac{(\frac{\gamma}{\gamma} - 1)}{\frac{\gamma}{\gamma} - \frac{\rho_0 k}{\rho_0 k}} \frac{1}{k} - 1 \right]\} \leq C$. The firm will then choose $r^* = 0$ and $a^* = a_0^*$ if $c \leq C_f$ and $\gamma \geq \Gamma_f$; otherwise, it will set $r^* = p$ with no influence tool, and the expected profit will be $\Pi = \rho_0 p$.

(iii). When $\rho_0 \in [k, 1]$, no influence tool will be employed under either refund scheme and the customer will always purchase. Thus the firm should offer no refund $r^* = 0$ with expected profit $\Pi = p$. ■

Proof. Corollary 3.5. The statement follows Proposition 3.3 (ii)-(b) when $\gamma < \Gamma_f$ and $c = 0$. ■

Proof. Corollary 3.6. (i). According to Proposition 3.3, when 1) $\rho_0 < \rho_g(0)$, or 2) $\rho_g(0) \leq \rho_0 < k$, and $c \geq C_f$ or $\gamma \leq \Gamma_f$, there will be a full refund without an influence tool, and customers will always purchase. The customer's expected utility is then $U = \rho_0(v - p) \geq 0$. When $\rho_g(0) \leq \rho_0 < k$ and $c \leq C_f$ and $\gamma \geq \Gamma_f$, the expected utility of the customer is $U = P(g|a_0^*)[P(G|g, a_0^*)v - p] = 0$. When $\rho_0 \geq k$, there will be no refund nor an influence tool, and the customer will always purchase. The expected utility is then $U = \rho_0 v - p \geq 0$.

(ii). Following (i), customer utility $U \geq 0$ when $\rho_0 < \rho_g(0)$ or when $\rho_g(0) \leq \rho_0 < k$, $c \in [C_f, C]$ and $\gamma \in [\Gamma, \Gamma_f]$. By Corollary 3.1, customer utility $U_0 = 0$ under the said conditions and $U_0 = U$ otherwise. ■

Proposition 3.4

Proof. By Lemma A.1, $\tilde{\rho}_g(0)$ increases in $\tilde{k}$, and hence decreases in r , and $\rho_g(0) = \tilde{\rho}_g(0)$ when $r = 0$.

(i). When $\rho_0 \in [0, \rho_g(0)]$, there exists an r_0 such that $\tilde{\rho}_g(0) = \rho_0$ when $r = r_0$, $\tilde{\rho}_g(0) < \rho_0$ when $r > r_0$, and $\rho_0 < \tilde{\rho}_g(0)$ when $r < r_0$. There also exists an $r_k \geq r_0$ such that $\tilde{k} = \rho_0$ when $r = r_k$, $\tilde{k} < \rho_0$ when $r > r_k$, and $\rho_0 < \tilde{k}$ when $r < r_k$. Following Proposition 4.3, the firm can then set a refund level $r \leq r_0$ with no influence tool and no customer will purchase, or $r_0 \leq r \leq r_k$ with influence level $a = a_r^* > 0$ where customers will purchase upon a positive signal, or $r \geq r_k$ with no influence tool where the customer will always purchase. For $r \leq r_0$, the expected profit $\Pi = 0$. For $r_0 \leq r \leq r_k$, $\Pi_r(a_r^*) = \max\{\tilde{C} - c, 0\}$ increases in r according to Corollary 3.3 (i). Thus $\Pi = \max_{r_0 \leq r \leq r_k} \Pi_r(a_r^*) = \Pi(a_r^*)|_{r=r_k} = \max\{\rho_0 v - c, 0\}$. For $r \geq r_k$, the expected profit $\Pi = \max_{r \geq r_k} \{p - \bar{\rho}_0 r\} = p - \bar{\rho}_0 r_k = \rho_0 v$. Comparing the three options, in this scenario, the firm should set $r^* = r_k = \frac{p - \rho_0 v}{\rho_0}$ without any influence tool with $\Pi = \rho_0 v$.

When $\rho_0 \in (\rho_g(0), k)$, there is $\rho \geq \tilde{\rho}_g(0)$ for any r . The firm can then set a refund level $r \leq r_k$ with an influence level $a = a_r^* > 0$ and where customers will purchase upon a positive signal; or $r \geq r_k$ with no influence tool and where the customer will always purchase. Similar to the reasoning above, the firm should set $r^* = r_k = \frac{p - \rho_0 v}{\rho_0}$ without any influence tool and $\Pi = \rho_0 v$.

(ii). When $\rho_0 \in [k, 1]$, there is $\rho_0 \geq \tilde{k}$ for any r , and customers always purchase. Therefore, the firm should offer no refund $r^* = 0$ nor use any influence tool, and the expected profit $\Pi = p$.

■

Proof. **Corollary 3.7.** According to Proposition 3.4, when $\rho_0 < k$, there will be no signal, and the customer's expected utility under partial refund $r^* = \frac{p - \rho_0 v}{\rho_0}$ will be $U = \rho_0 v + \bar{\rho}_0 r^* - p = 0$. When $\rho_0 \geq k$, there will be no refund and the customer will always purchase. The expected utility is then $U = \rho_0 v - p \geq 0$.

■

Proposition 3.5

Proof. Based on Lemma A.1, the firm's optimal affiliation decision can be analyzed as follows:

- (i). if $\rho_0 \in [0, \rho_g(0))$, the customer will never buy the product hence $a^* = 0$ and $\Pi = 0$.
- (ii). if $\rho_0 \in [\rho_g(0), k)$, the firm can either use a low level of influence such that $\rho_0 < \rho_g(a)$ and earn $\Pi(a) = 0 - \tau(a)$, or a high influence level satisfying $\rho_g(a) \leq \rho_0$ such that the consumer buys only when influencer's signal $s = g$ with expected profit $\Pi = P(g|a) \cdot p - \tau(a)$, where $P(g|a) = P(g|G, a)P(G) + P(g|B, a)P(B) = [\gamma + a(1 - \gamma)]\rho_0 + (a\gamma + 1 - \gamma)(1 - \rho_0) = (\gamma + \rho_0 - 2\gamma\rho_0)a + 1 - (\gamma + \rho_0 - 2\gamma\rho_0)$.

In the latter case, the firm's optimal affiliation can be solved by

$$\Pi(a^*) = \max_a P(g|a) \cdot p - \tau(a) \quad (\text{A.8})$$

s.t.

$$\tau(a) - c \geq 0$$

$$\rho_g(a) \leq \rho_0.$$

Since $\tau(a)$ is concave and increasing in a and $\frac{\partial(P(g|a) \cdot p)}{\partial a} = (\gamma + \rho_0 - 2\gamma\rho_0) \cdot p > \tau'(0)$, the firm's profit is increasing in the influence level. If the first constraint is temporarily relaxed (to be revisited later), and given that $\rho_g(a)$ increases in a , it follows that $\rho_g(a^*) = \rho_0$. Solving this yields

$$a^* = \frac{\gamma \bar{k} \rho_0 - k \bar{\rho}_0 \bar{\gamma}}{\gamma k \bar{\rho}_0 - \bar{\gamma} k \rho_0}. \quad (\text{A.9})$$

It can be shown that $0 \leq a^* \leq 1$ under the condition $\gamma(\rho_0 + k) > k(1 - \rho_0)$. Upon revisiting the first constraint, it is established that the firm offers $\tau(a^*)$ to the reviewer in order to induce influence at level a^* if $\tau(a^*) > c$, thereby ensuring that the requested review is generated. If this condition is not satisfied, no offer is made, no review is produced, and the resulting profit is $\Pi^*(0) = 0$.

By comparing both cases, it can be determined that influence will be employed at level a^* when the expected profit satisfies

$$\Pi(a^*) = [\rho_0(\gamma + a^*(1 - \gamma)) + (1 - \rho_0)(1 - \gamma + a^*\gamma)] p - \tau(a^*) \geq 0. \quad (\text{A.10})$$

Letting

$$C = [\rho_0(\gamma + a^*(1 - \gamma)) + (1 - \rho_0)(1 - \gamma + a^*\gamma)] p, \quad (\text{A.11})$$

the condition for using influence can be expressed as

$$c \leq \tau(a^*) \leq C. \quad (\text{A.12})$$

(iii). if $\rho_0 \in [k, \rho_b)$, a customer buys when influencer's signal $s = g$ or there is no review, and

will not buy if $s = b$. Accordingly, the influence scenario's problem is similar to the Proposition 3.1: since the customer always buys, no review is observed, thus the firm does not affiliate with reviewers for customer influence ($a^* = 0$). If $\rho_0 \in [\rho_b, 1]$, the customer will always buy. Therefore, there is no need to put any effort for producing a review signal and ($a^* = 0$). ■

Appendix B

Chapter 4

Lemma B.1. (i) $\frac{\partial \rho_g(\gamma)}{\partial \gamma} < 0$, $\frac{\partial \rho_g(\gamma)}{\partial k} > 0$ and $\rho_g(0.5) = k$. (ii) $\frac{\partial \tilde{\rho}_g(\gamma)}{\partial \gamma} < 0$, $\frac{\partial \tilde{\rho}_g(\gamma)}{\partial \tilde{k}} > 0$ and $\tilde{\rho}_g(0.5) = \tilde{k}$.

Proof. (i) There is $\frac{\partial \rho_g(\gamma)}{\partial \gamma} = -\frac{k \bar{a} \bar{k} (1+a)}{[k + \bar{k} a + \gamma \bar{a} (1-2k)]^2} \leq 0$. Also, since $\gamma \geq 1/2 \geq \bar{\gamma}$, $\rho_g(\gamma) = \frac{k(1-\gamma \bar{a})}{\gamma + a\gamma - k\bar{a}(\gamma - \bar{\gamma})}$ increases in k . With $\gamma = 0.5$ and $\bar{\gamma} = 0.5$, it can easily be verified that $\rho_g(0.5) = k$. (ii) The proof is similar to (i) but replacing k with $\tilde{k}$. ■

Lemma B.2. (i) $\frac{\partial \rho_b(\gamma)}{\partial \gamma} > 0$, $\frac{\partial \rho_b(\gamma)}{\partial k} > 0$ and $\rho_b(0.5) = k$. (ii) $\frac{\partial \tilde{\rho}_b(\gamma)}{\partial \gamma} > 0$, $\frac{\partial \tilde{\rho}_b(\gamma)}{\partial \tilde{k}} > 0$ and $\tilde{\rho}_b(0.5) = \tilde{k}$.

Proof. (i) There is $\frac{\partial \rho_b(\gamma)}{\partial \gamma} = \frac{k \bar{k}}{[k + \gamma(2k-1)]^2} \geq 0$. Also, since $\frac{\partial \rho_b(\gamma)}{\partial k} = \frac{\gamma \bar{\gamma}}{[\bar{\gamma} - k(\bar{\gamma} - \gamma)]^2} \geq 0$, $\rho_b(\gamma)$ increases in k . With $\gamma = 0.5$ and $\bar{\gamma} = 0.5$, it can easily be verified that $\rho_b(0.5) = k$. (ii) The proof is similar to (i) but replacing k with $\tilde{k}$. ■

Lemma 4.1

Proof. It suffices to show that $\rho_g(\gamma) \leq \rho_b(\gamma)$. Lemma B.1 (i) shows that $\rho_g(\gamma)$ is strictly decreasing in $\gamma \in [0.5, 1]$, whereas Lemma B.2 (i) shows that $\rho_b(\gamma)$ is strictly increasing on the same interval. Since $\rho_g(0.5) = \rho_b(0.5) = k$, it follows that $\rho_b(\gamma) \geq \rho_g(\gamma), \forall \gamma \in [0.5, 1]$. ■

Proposition 4.1

Proof. Based on Lemma 4.1, the firm's optimal influencing decision can be analyzed as follows:

- (i). If $\rho_0 \in [0, \rho_g(1))$, the customer will never buy the product; hence, with no cost, the firm can set the baseline signal quality at any level in the domain. The firm will be earning $\Pi^* = 0$.
- (ii). If $\rho_0 \in [\rho_g(1), \rho_b(0.5) = k)$, the customer only purchases when the signal $s = g$ with expected profit $\Pi(\gamma) = P(g|\gamma)p$ where

$$P(g|\gamma) = P(g|G, \gamma)P(G) + P(g|B, \gamma)P(B) = (\gamma + a\bar{\gamma})\rho_0 + (a\gamma + \bar{\gamma})\bar{\rho}_0. \quad (\text{B.1})$$

The firm's optimal signal quality level can be solved by

$$\begin{aligned} \Pi(\gamma^*) &= \max_{0.5 \leq \gamma \leq 1} P(g|\gamma)p \\ \text{s.t.} \quad &\rho_g(\gamma) \leq \rho_0 \leq \rho_b(\gamma) \end{aligned} \quad (\text{B.2})$$

Since $\frac{\partial P(g|\gamma)}{\partial \gamma} = (1 - a)(2\rho_0 - 1)$, the objective function decreases in γ for $\rho_0 < 0.5$, and increases in γ for $\rho_0 \geq 0.5$. Moreover, we know that $\rho_g(\gamma)$ decreases in γ , and $\rho_b(\gamma)$ increases in γ . Thus, when $\rho_0 \geq 0.5$, $\gamma^* = 1$.

The optimal profit in this case will be $\Pi(1) = P(g|1)p = (\rho_0 + a\bar{\rho}_0)p$.

When $\rho_0 < 0.5$, either $\rho_g(\gamma^*) = \rho_0$ or $\rho_b(\gamma^*) = \rho_0$. Solving the latter yields $\gamma^* = \frac{\rho_0 - k}{\rho_0 \bar{k}}$ which is negative for $\rho_0 \in [\rho_g(1), \rho_b(0.5) = k)$. While, the former yields

$$\gamma^* = \frac{k\bar{\rho}_0 - a\bar{k}\rho_0}{\bar{a}(k\bar{\rho}_0 + \bar{k}\rho_0)}, \quad (\text{B.3})$$

which falls within $[0.5, 1]$ for $\frac{ak}{ak+k} \leq \rho_0 \leq k$. Checking $\rho_0 \leq \rho_b(\gamma^*) = \frac{k^2 \bar{\rho}_0}{k\rho_0 + k\bar{\rho}_0 - akk}$ reveals that the constraint is always satisfied in the above mentioned range for ρ_0 . The optimal profit will be

$$\Pi(\gamma^*) = P(g|\gamma^*)p = \frac{k(\bar{\rho}_0^2 - \rho_0^2) + a\rho_0\bar{\rho}_0}{k\bar{\rho}_0 + \rho_0}, \quad (\text{B.4})$$

which is positive for $\frac{a\rho_0\bar{\rho}_0}{2\rho_0-1} > k$.

(iii). If $\rho_0 \in [\rho_b(0.5) = k, 1)$, the customers will always buy the product; hence, with no cost, the firm can set the baseline signal quality at any level in the domain as long as $\rho_0 > \rho_b(\gamma^*)$ or $\rho_0 \geq \frac{k\gamma^*}{k\bar{\gamma}^* + k\gamma^*}$ which is equal to $\gamma^* \leq \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}$. The firm will be earning $\Pi^* = p$. ■

Proposition 4.2

Proof. There is $\frac{\partial \rho_g(a, \gamma)}{\partial \gamma} < 0$, $\frac{\partial \rho_g(a, \gamma)}{\partial k} > 0$, $\rho_g(a, 0.5) = k$, $\frac{\partial \rho_g(a, \gamma)}{\partial a} > 0$, $\frac{\partial \rho_g(a, \gamma)}{\partial k} > 0$, and $\rho_g(1, \gamma) = k$. And since $\gamma \geq 1/2 \geq \bar{\gamma}$, there is $\gamma + a\bar{\gamma} \geq 1 - \gamma\bar{a}$. Thus $\rho_g(a, \gamma) = \frac{k(1-\gamma\bar{a})}{k(\gamma+a\bar{\gamma})+k(1-\gamma\bar{a})} \leq \frac{k}{k+k} = k$ and $\rho_b(\gamma) = \frac{k\gamma}{k\bar{\gamma}+k\gamma} \geq \frac{k}{k+k} = k$.

Based on Lemma 4.1, the firm's optimal influencing decision can be analyzed as follows:

(i). If $\rho_0 \in [0, \rho_b(0.5) = k)$, the customer only purchases when the signal $s = g$ with expected profit $\Pi(a, \gamma) = P(g|a, \gamma)p$ where

$$P(g|a, \gamma) = P(g|G, a, \gamma)P(G) + P(g|B, a, \gamma)P(B) = (\gamma + a\bar{\gamma})\rho_0 + (a\gamma + \bar{\gamma})\bar{\rho}_0. \quad (\text{B.5})$$

The firm's optimal signal quality and influencing level can be solved by

$$\begin{aligned} \Pi(a^*, \gamma^*) &= \max_{0 \leq a \leq 1 \text{ and } 0.5 \leq \gamma \leq 1} P(g|a, \gamma)p \\ \text{s.t.} \quad &\rho_g(a, \gamma) \leq \rho_0 \leq \rho_b(\gamma). \end{aligned} \quad (\text{B.6})$$

Since $\frac{\partial P(g|a, \gamma)}{\partial \gamma} = (1-a)(2\rho_0-1)$, the objective function clearly increases in a but decreases in γ for $\rho_0 < 0.5$ and increases in γ when $\rho_0 \geq 0.5$. Moreover, $\rho_g(a, \gamma)$ decreases in γ , thus for $\rho_0 \geq 0.5$ there is $\gamma^* = 1$ and $\rho_g(a^*, 1) = \rho_0$ or $\rho_0 = \frac{ka^*}{k+ka^*}$ which yields $a^* = \frac{\rho_0 \bar{k}}{k\rho_0}$. It can be shown that $0 \leq a^* = \frac{\rho_0 \bar{k}}{k\rho_0} \leq 1$ for $\rho_0 < k$ and $\rho_b(1) = 1$.

Otherwise (i.e., $\rho_0 < 0.5$), checking $\rho_0 = \rho_b(\gamma^*)$ leads to a negative γ^* since $\rho_b(\gamma^*) = \rho_0$ yields $\gamma^* = \frac{\rho_0 \bar{k} - \bar{\rho}_0 k}{\rho_0 k}$. Moving on to $\rho_g(a^*, \gamma^*) = \rho_0$ leads to the following

$$a^* = \left(\frac{\gamma^*}{\bar{\gamma}} - \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} \right) / \left(\frac{\gamma^*}{\bar{\gamma}} \frac{\bar{\rho}_0 k}{\rho_0 \bar{k}} - 1 \right).$$

Replacing $a^*(\gamma)$ in the objective function it can be rewritten as below and maximized over γ :

$$\Pi^*(\gamma) = \left(\bar{\rho}_0 + \frac{(\gamma - [\frac{\bar{\rho}_0 k}{\rho_0 k}] \bar{\gamma})(\frac{\rho_0}{\bar{\rho}_0} + \gamma)}{\gamma [\frac{\bar{\rho}_0 k}{\rho_0 k}] - \bar{\gamma}} \right) p.$$

Since, $\rho_0 < 0.5$ and $k > \rho_0$, it can be shown that $[\frac{\bar{\rho}_0 k}{\rho_0 k}] > 1$, and consequently, that the above expression is increasing in γ . Thus $\gamma^* = 1$, and similarly $\rho_g(a^*, 1) = \rho_0$ or $\rho_0 = \frac{k a^*}{k + k a^*}$ which yields $a^* = \frac{\rho_0 \bar{k}}{k \rho_0}$.

The optimal profit in both cases will be $\Pi(a^*, 1) = P(g|a^*, 1)p = \rho_0 v$ which is positive.

(ii). If $\rho_0 \in [\rho_b(0.5) = k, 1)$, the customer will always buy the product; hence, with no cost, the firm can set the baseline signal quality at any level in the domain as long as $\rho_0 > \rho_b(\gamma^*)$ or $\rho_0 \geq \frac{k \gamma^*}{k \bar{\gamma}^* + k \gamma^*}$ which is equal to $\gamma^* \leq \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}$. Also, as there is no cost of influence tool, the optimal influence level can be set at any $a^* \in [0, 1]$. The firm will be earning $\Pi^* = p$. ■

Lemma 4.2

Proof. The argument follows the same steps as in Lemma 4.1 based on Lemma B.1 (ii) and Lemma B.2 (ii); the details are therefore omitted to avoid repetition. ■

Proposition 4.3

Proof. There is $\frac{\partial \tilde{\rho}_g(a, \gamma)}{\partial \gamma} < 0$, $\frac{\partial \tilde{\rho}_g(a, \gamma)}{\partial k} > 0$, $\tilde{\rho}_g(a, 0.5) = \tilde{k}$, $\frac{\partial \tilde{\rho}_g(a, \gamma)}{\partial a} > 0$, $\frac{\partial \tilde{\rho}_g(a, \gamma)}{\partial k} > 0$, and $\tilde{\rho}_g(1, \gamma) = \tilde{k}$. And since $\gamma \geq 1/2 \geq \bar{\gamma}$, there is $\gamma + a \bar{\gamma} \geq 1 - \gamma \bar{a}$. Thus $\tilde{\rho}_g(a, \gamma) = \frac{\tilde{k}(1 - \gamma \bar{a})}{k(\gamma + a \bar{\gamma}) + k(1 - \gamma \bar{a})} \leq \frac{\tilde{k}}{k + k} = \tilde{k}$ and $\rho_b(\gamma) = \frac{\tilde{k} \gamma}{k \bar{\gamma} + k \gamma} \geq \frac{\tilde{k}}{k + k} = \tilde{k}$.

Based on Lemma 4.2, the firm's optimal influencing decision can be analyzed as follows:

(i). If $\rho_0 \in [0, \tilde{\rho}_b(0.5) = \tilde{k})$, the customer only purchases when the signal $s = g$ with expected profit $\Pi(a, \gamma) = P(g|a, \gamma)[p - P(B|g, a, \gamma)r]$. The firm's optimal signal quality and influencing level can be solved by

$$\begin{aligned} \Pi(a_r^*, \gamma_r^*) = \max_{0 \leq a \leq 1 \text{ and } 0.5 \leq \gamma \leq 1} & P(g|a, \gamma)[p - P(B|g, a, \gamma)r] \\ \text{s.t.} & \tilde{\rho}_g(a, \gamma) \leq \rho_0 \leq \tilde{\rho}_b(\gamma). \end{aligned} \quad (\text{B.7})$$

Since $\frac{\partial P(B|g, a, \gamma)}{\partial \gamma} < 0$, $P(B|g, a, \gamma)$ is decreasing in γ . Also, $\frac{\partial P(g|a, \gamma)}{\partial \gamma} = (1 - a)(2\rho_0 - 1)$. Subsequently, it can be shown that the objective function increases in a but decreases in γ for $\rho_0 < 0.5$ and increases in γ when $\rho_0 \geq 0.5$. Moreover, $\tilde{\rho}_g(a, \gamma)$ decreases in γ , thus for $\rho_0 \geq 0.5$ there is $\gamma_r^* = 1$ and $\tilde{\rho}_g(a_r^*, 1) = \rho_0$ or $\rho_0 = \frac{\tilde{k}a_r^*}{k + \tilde{k}a_r^*}$ which yields $a_r^* = \frac{\rho_0 \tilde{k}}{k \rho_0}$. It can be shown that $0 \leq a_r^* = \frac{\rho_0 \tilde{k}}{k \rho_0} \leq 1$ for $\rho_0 < \tilde{k}$ and $\tilde{\rho}_b(1) = 1$.

Otherwise (i.e., $\rho_0 < 0.5$), checking $\rho_0 = \tilde{\rho}_b(\gamma_r^*)$ leads to a negative γ_r^* since $\tilde{\rho}_b(\gamma_r^*) = \rho_0$ yields $\gamma_r^* = \frac{\rho_0 \tilde{k} - \tilde{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}}$. Moving on to $\tilde{\rho}_g(a_r^*, \gamma_r^*) = \rho_0$ leads to the following

$$a_r^* = \left(\frac{\gamma_r^*}{\gamma_r^*} - \frac{\tilde{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} \right) / \left(\frac{\gamma_r^*}{\gamma_r^*} \frac{\tilde{\rho}_0 \tilde{k}}{\rho_0 \tilde{k}} - 1 \right). \quad (\text{B.8})$$

Replacing $a^*(\gamma_r)$ in the objective function and simplifying, it can be rewritten as below and maximized over γ_r :

$$\Pi(\gamma_r) = \frac{2\gamma - 1}{\gamma - \frac{\tilde{\gamma} \rho_0 \tilde{k}}{\tilde{\rho}_0 \tilde{k}}} \rho_0 v. \quad (\text{B.9})$$

The above expression is clearly increasing in γ_r . Thus $\gamma_r^* = 1$, and similarly $\tilde{\rho}_g(a_r^*, 1) = \rho_0$ or $\rho_0 = \frac{\tilde{k}a_r^*}{k + \tilde{k}a_r^*}$ which yields $a_r^* = \frac{\rho_0 \tilde{k}}{k \rho_0}$.

The optimal profit in both cases will be $\Pi(a_r^*, 1) = P(g|a_r^*, 1)[p - P(B|g, a_r^*, 1)r] = \rho_0 v$.

(ii). If $\rho_0 \in [\tilde{\rho}_b(0.5) = \tilde{k}, 1)$, the customer will always buy the product. The firm can set the baseline signal quality at any level in the domain as long as $\rho_0 > \rho_b(\gamma_r^*)$ or $\rho_0 \geq \frac{\tilde{k}\gamma_r^*}{k\gamma_r^* + k\gamma_r^*}$ which

is equal to $\gamma_r^* \leq \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}$. Also, as there is no cost of influence tool, the optimal influence level can be set at any $a^* \in [0, 1]$. The firm will be earning $\Pi_r^* = p - \bar{\rho}_0 r$. ■

Proof. Corollary 4.1. It can be shown that $\frac{\partial a_r^*}{\partial r} = \frac{\rho_0(v-p)}{(1-\rho_0)^2(p-r)^2} \geq 0$. Hence, the optimal level of influence, where it is applicable, increases in refund. ■

Proposition 4.4

Proof. It can be shown that $\tilde{\rho}_g(r)$ increases in $\tilde{k}$, and hence decreases in r , and $\rho_g(r=0) = \tilde{\rho}_g(r=0)$.

(i). If $\rho_0 \in [0, k)$, the firm has two options in order to look forward to a purchase from customers. It may choose to use the optimal influence, signal quality, and offer the optimal partial refund and gain $\Pi^*(a^*, \gamma^*, r^*)$. Or offer a partial refund based on the consumer's expected utility that ensures a non-negative payoff with a purchase, that is $r^* = \frac{p-\rho_0 v}{\rho_0}$, with any influence and signal quality (since they cost zero).

For the former case, and based on Proposition 4.2 and 4.3, the firm may earn $\rho_0 v$ by setting $a^* = \frac{\rho_0 \bar{k}}{k \rho_0}$, $\gamma^* = 1$, and $r^* = 0$ which is equal to the payoff of the latter case under $r^* = \frac{p-\rho_0 v}{\rho_0}$ and any influence and signal quality level.

(ii). If $\rho_0 \in [k, 1)$, there is $\rho_0 \geq \tilde{k}$ for any r , and customers always buy. Therefore, the firm should offer no refund $r^* = 0$. But it can use any signal quality level in the domain as long as $\rho_0 > \rho_b(\gamma^*)$ or $\rho_0 \geq \frac{k \gamma^*}{k \bar{\gamma}^* + k \gamma^*}$ which is equal to $\gamma^* \leq \frac{\rho_0 \bar{k}}{\rho_0 k + \bar{\rho}_0 k}$. Also, as there is no cost of influence tool, the optimal influence level can be set at any $a^* \in [0, 1]$. The firm will be earning $\Pi^* = p$. ■

Appendix C

Chapter 5

C.1. Explicit Analytical Forms

The parameters A_i , B_i , C_i , N_i , and D_i in ρ_i and Π_i where $\Delta p = p_1 - p_2$, are given by:

$$A_1 = 4\gamma^2 [a_2^2(\Delta p)^2 + ((a_2 - 3)a_2 + 1)v^2] + a_2^2(\Delta p)^2 + 4a_2\gamma(v^2 - a_2(\Delta p)^2) + 4\bar{a}_2^2\gamma^4v^2 - 8\bar{a}_2^2\gamma^3v^2 \quad (\text{C.1})$$

$$B_1 = (1 - 2\gamma) [-2\bar{a}_2\gamma - a_2 + 2] \Delta p - 2\bar{\gamma}v [-\bar{a}_2\gamma + 1] \quad (\text{C.2})$$

$$C_1 = 2(2\gamma - 1) [-\bar{a}_2(1 - 2\gamma)\Delta p + v] \quad (\text{C.3})$$

$$A_2 = 4\gamma^2 [a_1^2(\Delta p)^2 + ((a_1 - 3)a_1 + 1)v^2] + a_1^2(\Delta p)^2 + 4a_1\gamma(v^2 - a_1(\Delta p)^2) + 4\bar{a}_1^2\gamma^4v^2 - 8\bar{a}_1^2\gamma^3v^2 \quad (\text{C.4})$$

$$B_2 = 2a_1\gamma^2(v - 2\Delta p) + 2\gamma((1 - 2a_1) + a_1v) + a_1\Delta p \quad (\text{C.5})$$

$$C_2 = 2(2\gamma - 1) [\bar{a}_1(1 - 2\gamma)\Delta p + v] \quad (\text{C.6})$$

$$\begin{aligned} N_1 = & 4\Delta p\gamma^2\rho_0^2 - 4\Delta p\gamma^2\rho_0 + \Delta p\gamma^2 - 4\Delta p\gamma\rho_0^2 + 4\Delta p\gamma\rho_0 - \Delta p\gamma \\ & + \Delta p\rho_0^2 - \Delta p\rho_0 + 8\gamma^3\rho_0^4 - 16\gamma^3\rho_0^3 + 10\gamma^3\rho_0^2 - 2\gamma^3\rho_0 \\ & - 12\gamma^2\rho_0^4 + 28\gamma^2\rho_0^3 - 21\gamma^2\rho_0^2 + 7\gamma^2\rho_0 - \gamma^2 + 6\gamma\rho_0^4 \\ & - 16\gamma\rho_0^3 + 14\gamma\rho_0^2 - 6\gamma\rho_0 + \gamma - \rho_0^4 + 3\rho_0^3 - 3\rho_0^2 + \rho_0 \end{aligned} \quad (\text{C.7})$$

$$\begin{aligned} D_1 = & 4\Delta p\gamma^2\rho_0^2 - 4\Delta p\gamma^2\rho_0 + \Delta p\gamma^2 - 4\Delta p\gamma\rho_0^2 + 4\Delta p\gamma\rho_0 - \Delta p\gamma \\ & + \Delta p\rho_0^2 - \Delta p\rho_0 + 2\gamma^2\rho_0 - \gamma^2 - 2\gamma\rho_0 + \gamma \end{aligned} \quad (\text{C.8})$$

$$\begin{aligned} N_2 = & -8\gamma^3\rho_0^4 + 16\gamma^3\rho_0^3 - 10\gamma^3\rho_0^2 + 2\gamma^3\rho_0 + 12\gamma^2\rho_0^4 - 28\gamma^2\rho_0^3 \\ & + 21\gamma^2\rho_0^2 - 5\gamma^2\rho_0 - 6\gamma\rho_0^4 + 16\gamma\rho_0^3 - 14\gamma\rho_0^2 + 4\gamma\rho_0 + \rho_0^4 \\ & - 3\rho_0^3 + 3\rho_0^2 - \rho_0 \end{aligned} \quad (\text{C.9})$$

$$\begin{aligned} D_2 = & 4\Delta p\gamma^2\rho_0^2 - 4\Delta p\gamma^2\rho_0 + \Delta p\gamma^2 - 4\Delta p\gamma\rho_0^2 + 4\Delta p\gamma\rho_0 - \Delta p\gamma \\ & + \Delta p\rho_0^2 - \Delta p\rho_0 + 2\gamma^2\rho_0 - \gamma^2 - 2\gamma\rho_0 + \gamma \end{aligned} \quad (\text{C.10})$$

$$\begin{aligned} N_3 = & \rho_0(8\gamma^3\rho_0^3 - 16\gamma^3\rho_0^2 + 10\gamma^3\rho_0 - 2\gamma^3 - 12\gamma^2\rho_0^3 + 20\gamma^2\rho_0^2 - 9\gamma^2\rho_0 + \gamma^2 \\ & + 6\gamma\rho_0^3 - 8\gamma\rho_0^2 + 2\gamma\rho_0 - \rho_0^3 + \rho_0^2) \end{aligned} \quad (\text{C.11})$$

$$\begin{aligned} D_3 = & 4\Delta p\gamma^2\rho_0^2 - 4\Delta p\gamma^2\rho_0 + \Delta p\gamma^2 - 4\Delta p\gamma\rho_0^2 + 4\Delta p\gamma\rho_0 - \Delta p\gamma \\ & + \Delta p\rho_0^2 - \Delta p\rho_0 + 2\gamma^2\rho_0 - \gamma^2 - 2\gamma\rho_0 + \gamma \end{aligned} \quad (\text{C.12})$$

$$\begin{aligned}
N_4 = & 4\Delta p \gamma^2 \rho_0^2 - 4\Delta p \gamma^2 \rho_0 + \Delta p \gamma^2 - 4\Delta p \gamma \rho_0^2 + 4\Delta p \gamma \rho_0 - \Delta p \gamma + \Delta p \rho_0^2 - \Delta p \rho_0 \\
& - 8\gamma^3 \rho_0^4 + 16\gamma^3 \rho_0^3 - 10\gamma^3 \rho_0^2 + 2\gamma^3 \rho_0 + 12\gamma^2 \rho_0^4 - 20\gamma^2 \rho_0^3 + 9\gamma^2 \rho_0^2 + \gamma^2 \rho_0 - \gamma^2 \quad (\text{C.13}) \\
& - 6\gamma \rho_0^4 + 8\gamma \rho_0^3 - 2\gamma \rho_0^2 - 2\gamma \rho_0 + \gamma + \rho_0^4 - \rho_0^3
\end{aligned}$$

$$\begin{aligned}
D_4 = & 4\Delta p \gamma^2 \rho_0^2 - 4\Delta p \gamma^2 \rho_0 + \Delta p \gamma^2 - 4\Delta p \gamma \rho_0^2 + 4\Delta p \gamma \rho_0 - \Delta p \gamma \\
& + \Delta p \rho_0^2 - \Delta p \rho_0 + 2\gamma^2 \rho_0 - \gamma^2 - 2\gamma \rho_0 + \gamma \quad (\text{C.14})
\end{aligned}$$

$$\begin{aligned}
N_5 = & -8\gamma^3 \rho_0^4 + 16\gamma^3 \rho_0^3 - 10\gamma^3 \rho_0^2 + 2\gamma^3 \rho_0 + 12\gamma^2 \rho_0^4 - 20\gamma^2 \rho_0^3 + 9\gamma^2 \rho_0^2 + \gamma^2 \rho_0 - \gamma^2 \quad (\text{C.15}) \\
& - 6\gamma \rho_0^4 + 8\gamma \rho_0^3 - 2\gamma \rho_0^2 - 2\gamma \rho_0 + \gamma + \rho_0^4 - \rho_0^3
\end{aligned}$$

The parameters $\tilde{A}_i$, $\tilde{B}_i$, $\tilde{C}_i$, $\tilde{D}_i$, and $\tilde{E}_i$ in $\tilde{\rho}_i$ are given by:

$$\begin{aligned}
\tilde{A}_1 = & \bar{a}_2 \gamma^2 (p_1(4 - r_1) - p_2(4 - 3r_2) - 2v) + \gamma(p_1(4a_2 + r_1 - 6) + p_2(2a_2 \\
& (r_2 - 2) - 5r_2 + 6) - 2(a_2 - 2)v) - a_2 p_1 + a_2 p_2 + 2p_1 + 2p_2 r_2 - 2(p_2 + v) \quad (\text{C.16})
\end{aligned}$$

$$\tilde{B}_1 = 4\bar{\gamma}(2\gamma - 1)(\bar{a}_2\gamma - 1)(p_1 - p_2\bar{r}_2 - v) \quad (\text{C.17})$$

$$\tilde{C}_1 = -\bar{a}_2 p_1(\gamma(r_1 - 2) + 1) + \bar{a}_2 \gamma p_2(r_2 - 2) - p_2(a_2 + r_2) + p_2 + v \quad (\text{C.18})$$

$$\begin{aligned}
\tilde{D}_1 = & p_1(\bar{a}_2 \gamma^2(4 - r_1) + \gamma(4a_2 + r_1 - 6) - a_2 + 2) \\
& + a_2 p_2(-4\gamma\bar{\gamma} - 3\gamma^2 r_2 + 2\gamma r_2 + 1) - 2a_2 v \gamma \bar{\gamma} + p_2(-4\gamma^2 + 6\gamma - 2) \quad (\text{C.19}) \\
& + p_2 r_2(3\gamma^2 - 5\gamma + 2) - 2v\bar{\gamma}^2
\end{aligned}$$

$$\tilde{E}_1 = 2(2\gamma - 1)(\bar{a}_2 p_1(\gamma(1 + \bar{r}_1) - 1) - \bar{a}_2 \gamma p_2(1 + r_2) - p_2(a_2 + r_2) + p_2 + v) \quad (\text{C.20})$$

$$\begin{aligned}\tilde{A}_2 = & -\bar{a}_1\gamma^2(p_1(r_1 - 4) + p_2(4 - 3r_2) + 2v) + \gamma(p_1(4a_1 + r_1 - 2) \\ & + (2a_1 - 1)p_2(r_2 - 2) - 2a_1v) + a_1(p_2 - p_1)\end{aligned}\quad (\text{C.21})$$

$$\begin{aligned}\tilde{B}_2 = & -(\bar{a}_1\gamma^2(p_1(4 - r_1) + p_2(4 - 3r_2) + 2v)) + \gamma(-p_1(4a_1 + r_1 - 2) \\ & + p_2(-2a_1r_2 + 4a_1 + r_2 - 2) + 2a_1v) + a_1(p_1 - p_2)\end{aligned}\quad (\text{C.22})$$

$$\tilde{C}_2 = 4\gamma(1 - 2\gamma)(a_1\bar{\gamma} + \gamma)(p_1 - p_2\bar{r}_2 - v) \quad (\text{C.23})$$

$$\tilde{D}_2 = p_1(\bar{a}_1\gamma(1 + \bar{r}_1) + a_1 - \bar{r}_1) + \bar{a}_1p_2(-\gamma(1 + \bar{r}_2) + 1) - v \quad (\text{C.24})$$

$$\tilde{E}_2 = 2(1 - 2\gamma)(-p_1(\bar{a}_1\gamma(1 + \bar{r}_1) + a_1 - \bar{r}_1) + \bar{a}_1p_2(\gamma(1 + \bar{r}_2) - 1) + v) \quad (\text{C.25})$$

C.2. Analytical Proofs and Numerical Continuations

Lemma C.1. (i) $k \in [0, 1]$, (ii) For v normalized to 1, $\min \rho_1(a_2) = \rho_1(0)$, and $\max \rho_2(a_1) = \rho_2(0)$, (iii) For v normalized to 1, $\rho_1(\Delta p, a_2)$ and $\rho_2(\Delta p, a_1)$ are strictly decreasing in Δp .

Proof. (i) Since k is clearly increasing in p_2 and decreasing in p_1 , and $0 \leq p_1, p_2 \leq v$, $\min(k) = 0$ for $p_1 = v$ and $p_2 = 0$. Similarly, $\max(k) = 1$ for $p_1 = 0$ and $p_2 = v$.

(ii) The derivative $\frac{\partial \rho_1(a_2)}{\partial a_2}$ is given by a highly non-linear expression whose real roots cannot be expressed in elementary closed form.¹ As a result, the identification of any extrema must be carried out numerically. For this purpose, the function $\rho_1(a_2)$, Figure C.1, and its derivative were evaluated over the interval $0 < a_2 < 1$, using the parameter values $p_1 = 0.7$, $p_2 = 0.6$, $v = 1$, and $\gamma = 0.7$.

The derivative $\frac{d\rho_1(a_2)}{da_2}$, as illustrated in Figure C.2, remains strictly positive throughout the

¹Symbolic algebra systems such as MATHEMATICA return irreducible quartic polynomials that are not solvable in radicals.

domain, indicating that $\rho_1(a_2)$ is monotonically increasing and attains its minimum at the lower boundary $a_2 \rightarrow 0^+$.




Figure C.1: Function $\rho_1(a_2)$ over the interval $a_2 \in [0, 1]$




Figure C.2: First derivative of function $\rho_1(a_2)$ with respect to a_2 over the interval $a_2 \in [0, 1]$

For the numerical analysis of the function $\rho_2(a_1)$, Figure C.4, the objective is to examine its behaviour and identify the maximum by substituting the parameter values $p_1 = 0.7$, $p_2 = 0.6$, $v = 1$, and $\gamma = 0.7$. The function $\rho_2(a_1)$ and its derivative were evaluated numerically over the interval $0 < a_1 < 1$. The presence of critical points was investigated by examining whether the derivative $\frac{d\rho_2(a_1)}{da_1}$ vanishes within the domain.

As shown in Figure C.4, the derivative $\frac{d\rho_2(a_1)}{da_1}$ remains strictly positive over the interval, indicating the absence of interior critical points. Therefore, $\rho_2(a_1)$ is strictly increasing throughout, and its maximum is attained at the boundary value $a_1 \rightarrow 0^+$.




Figure C.3: Function $\rho_2(a_1)$ over the interval $a_1 \in [0, 1]$




Figure C.4: First derivative of function $\rho_2(a_1)$ with respect to a_1 over the interval $a_1 \in [0, 1]$

Grid Analysis of $\rho_1(a_2)$ and $\rho_2(a_1)$: To ensure the robustness of the numerical analysis, it is extended by evaluating $\rho_1(a_2)$ and $\rho_2(a_1)$ over their entire effective parameter ranges at $v = 1$. Specifically:

- $\Delta p = p_1 - p_2$ ranging from -1 to 1 ,
- γ ranging from 0.5 to 1 ,
- a_2 and a_1 ranging from 0 to 1 .

These evaluations ensure that the analysis accounts for all plausible variations of these parameters.

Observations:

- $\rho_1(a_2)$ **Grid Analysis** (Figure C.5 and C.6): The grid confirms that $\rho_1(a_2)$ remains strictly increasing for the entire parameter space with respect to a_2 , and the function's minimum occurs at the lower boundary $a_2 \rightarrow 0^+$.
- $\rho_2(a_1)$ **Grid Analysis** (Figure C.7 and C.8): Similarly, $\rho_2(a_1)$ is strictly decreasing with respect to a_1 , with its maximum at $a_1 \rightarrow 0^+$, supporting the conclusion of the boundary behaviour.




Figure C.5: 4D Grid Analysis of $\rho_1(a_2)$ for parameters Δp , γ , and a_2




Figure C.6: 4D Grid Analysis of $\frac{\partial \rho_1(a_2)}{\partial a_2}$ for parameters Δp , γ , and a_2

(iii) **Grid Analysis of $\rho_1(\Delta p)$ and $\rho_2(\Delta p)$:** Similarly, the numerical continuation is based on a grid analysis as follows:

- $\frac{\partial \rho_1(\Delta p)}{\partial \Delta p}$ **Grid Analysis** (Figure C.9): For $v = 1$, the grid confirms that $\rho_1(\Delta p)$ remains strictly increasing for the entire parameter space with respect to Δp , and the function's




Figure C.7: 4D Grid Analysis of $\rho_2(a_1)$ for parameters Δp , γ , and a_1




Figure C.8: 4D Grid Analysis of $\frac{\partial \rho_2(a_1)}{\partial a_1}$ for parameters Δp , γ , and a_1

minimum occurs at the lower boundary $\Delta p \rightarrow 1^-$.

- $\frac{\partial \rho_2(\Delta p)}{\partial \Delta p}$ **Grid Analysis** (Figure C.10): Similarly, $\rho_2(\Delta p)$ is strictly decreasing with respect to Δp , with its maximum at $\Delta p \rightarrow -1^+$, supporting the boundary behaviour conclusion.

Lemma 5.1 (Numerical Search)




Figure C.9: 4D Grid Analysis of $\rho_1(\Delta p)$ for parameters Δp , γ , and a_2




Figure C.10: 4D Grid Analysis of $\rho_2(\Delta p)$ for parameters Δp , γ , and a_1

Proof. Fixing $v = 1$ and letting the parameters range over the compact box

$$\mathcal{D} = \{(\Delta p, \gamma, a) \mid \Delta p \in [-1, 1], \gamma \in [0.50, 1.00], a \in [0, 1]\}, \quad (\text{C.26})$$

the following functions are defined

$$g_1(\Delta p, \gamma, a) = k - \rho_1(\Delta p, \gamma, a), \quad (\text{C.27})$$

and

$$g_2(\Delta p, \gamma, a) = \rho_2(\Delta p, \gamma, a) - k. \quad (\text{C.28})$$

Their full distributions are visualized in Figures C.11 and C.12. No node attains a negative value, hence on the sampled grid $g_1 \geq 0$ and $g_2 \geq 0$. Thus

$$g_1(\Delta p, \gamma, a) \geq 0 \quad \text{and} \quad g_2(\Delta p, \gamma, a) \geq 0 \quad \forall (\Delta p, \gamma, a) \in \mathcal{D}. \quad (\text{C.29})$$

Interpretation. Because $g_1 \geq 0$ there is $\rho_1(\Delta p, \gamma, a) \leq k$ on $\mathcal{D}$; likewise $g_2 \geq 0$ implies $\rho_2(\Delta p, \gamma, a) \geq k$. Integrating the two yields the ordering of the product choice thresholds

$$\rho_1(\Delta p, \gamma, 1) \leq k \leq \rho_2(\Delta p, \gamma, 1). \quad (\text{C.30})$$




Figure C.11: 4D Grid Analysis of ρ_1 for parameters Δp , γ , and a_2




Figure C.12: 4D Grid Analysis of ρ_2 for parameters Δp , γ , and a_1

■

Proposition 5.1

Proof. Based on Lemma 5.1 and Lemma C.1, the firms' optimal influencing decision can be analyzed as follows:

(i). If $\rho_0 \in [0, \rho_1(a_2 = 0))$, the customer will never buy from seller 2; therefore, there is no need for influence, and in the absence of influence cost, Seller 1 can set the influence level at $a_1 \in [0, 1]$.

Similarly, there is no chance of selling for Seller 2 at any influence level, therefore $a_2 \in [0, 1]$. This implies that in realistic settings where influence is costly, both sellers optimally choose no influence, that is, $a_1^* = a_2^* = 0$.

Substituting this yields optimal profits: $\Pi_1^* = p_1$ and $\Pi_2^* = 0$.

(ii). If $\rho_0 \in [\rho_1(a_2 = 0), k)$, Seller 1's payoff function Π_1 is independent of its decision variable a_1 . Consequently, any choice $a_1^* \in [0, 1]$ is optimal for Seller 1, i.e., not using influence $a_1^* = 0$ is optimal; its payoff can be determined by substituting Seller 2's decision into the payoff function.

For Seller 2, the optimal influence decision is obtained by maximizing the associated payoff function within this range of priors:

$$\begin{aligned} \Pi_2(a_2^*) &= \max_{0 \leq a_2 \leq 1} P(S_1 = b) \cdot P(S_2 = g|a_2) \cdot p_2 \\ \text{s.t.} \quad &\rho_1 \leq \rho_0 \leq k \end{aligned} \quad (\text{C.31})$$

Since $\Pi_2 = (\bar{\gamma}\bar{\rho}_0 + \gamma\rho_0) [(\gamma + a_2\bar{\gamma})\rho_0 + (\bar{\gamma} + a_2\gamma)\bar{\rho}_0] p_2$, Seller 2's payoff function is strictly increasing in a_2 . By Lemma C.1, ρ_1 is also strictly increasing in a_2 , establishing a one-to-one correspondence between ρ_1 and a_2 . Consequently, the optimal choice a_2^* can be determined by solving

$$\rho_1(a_2^*) = \rho_0 \quad (\text{C.32})$$

for a_2^* , which yields:

$$a_2^* = \frac{(\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0)) [\Delta p \bar{\rho}_0 + (\gamma - \bar{\rho}_0) + \gamma \Delta p (\rho_0 - \bar{\rho}_0)]}{\bar{\gamma} \gamma (\bar{\rho}_0 - \rho_0) + \Delta p (-\rho_0 + \gamma(\rho_0 - \bar{\rho}_0)) (\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0))} \quad (\text{C.33})$$

It can be shown that $a_2^* \in [0, 1]$ for

$$\rho_0 \in \left[\frac{\bar{\gamma}(1-\Delta p)}{\Delta p(\gamma-\bar{\gamma})+1}, \frac{\Delta p(\bar{\gamma}-\gamma)-2\bar{\gamma}+\sqrt{\Delta p^2(\gamma-\bar{\gamma})^2+4\gamma\bar{\gamma}}}{2(\gamma-\bar{\gamma})} \right].$$

Therefore, since a_2^* increases monotonically in ρ_0 , it follows that $a_2^* = 0$ for ρ_0 below the lower bound, and $a_2^* = 1$ for ρ_0 above the upper bound. The optimal a_2^* can thus be expressed as a piecewise function of ρ_0 :

$$a_2^* = \begin{cases} 0, & \text{if } \rho_0 < \rho_0^{\min} \\ \frac{(\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0)) [\Delta p \bar{\rho}_0 + (\gamma - \bar{\rho}_0) + \gamma \Delta p (\rho_0 - \bar{\rho}_0)]}{\bar{\gamma} \gamma (\bar{\rho}_0 - \rho_0) + \Delta p (-\rho_0 + \gamma(\rho_0 - \bar{\rho}_0)) (\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0))}, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ 1, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.34})$$

The optimal profit functions can be expressed as:

$$\Pi_1^* = \begin{cases} (-4\gamma^2\rho_0^2 + 4\gamma^2\rho_0 - \gamma^2 + 4\gamma\rho_0^2 - 6\gamma\rho_0 + 2\gamma - \rho_0^2 + 2\rho_0) p_1, & \text{if } \rho_0 < \rho_0^{\min} \\ \frac{N_1}{D_1} p_1, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ (-2\gamma\rho_0 + \gamma + \rho_0) p_1, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.35})$$

since $\Pi_1^* =$

$$[1 - P(S_1 = b).P(S_2 = g \mid a_2^*)] p_1 = [1 - (\bar{\gamma}\bar{\rho}_0 + \gamma\rho_0) ((\gamma + a_2^*\bar{\gamma})\rho_0 + (\bar{\gamma} + a_2^*\gamma)\bar{\rho}_0)] p_1; \text{ and}$$

$$\Pi_2^* = \begin{cases} (4\gamma^2\rho_0^2 - 4\gamma^2\rho_0 + \gamma^2 - 4\gamma\rho_0^2 + 6\gamma\rho_0 - 2\gamma + \rho_0^2 - 2\rho_0 + 1) p_2, & \text{if } \rho_0 < \rho_0^{\min} \\ \frac{N_2}{D_2} p_2, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ (1 - \bar{\gamma}\rho_0 - \gamma\bar{\rho}_0) p_2, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.36})$$

where N_i and D_i are defined in equations C.7 – C.10 and $\rho_0^{\min}, \rho_0^{\max}$ are given by $\rho_0^{\min} = \frac{\bar{\gamma}(1-\Delta p)}{\Delta p(\gamma-\bar{\gamma})+1}$, and $\rho_0^{\max} = \frac{\Delta p(\bar{\gamma}-\gamma)-2\bar{\gamma}+\sqrt{\Delta p^2(\gamma-\bar{\gamma})^2+4\gamma\bar{\gamma}}}{2(\gamma-\bar{\gamma})}$.

(iii). If $\rho_0 \in [k, \rho_2(a_1 = 0)]$, Seller 2's payoff function Π_2 becomes independent of its decision variable a_2 . Consequently, any choice $a_2^* \in [0, 1]$ yields the same payoff, and the optimal profit can be evaluated by substituting Seller 1's decision into the payoff function. For Seller 1, the optimal influence strategy within this range of priors is determined by maximizing the corresponding payoff function:

$$\begin{aligned} \Pi_1(a_1^*) &= \max_{0 \leq a_1 \leq 1} P(S_1 = g|a_1)].P(S_2 = b).p_1 \\ \text{s.t.} \quad &k \leq \rho_0 \leq \rho_2. \end{aligned} \quad (\text{C.37})$$

Since $\Pi_1 = (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0) [(\gamma + a_1\bar{\gamma})\bar{\rho}_0 + (\bar{\gamma} + a_1\gamma)\rho_0] p_1$, Seller 1's payoff function is strictly

increasing in a_1 . According to Lemma C.1, ρ_2 is strictly decreasing in a_1 , establishing a one-to-one correspondence between ρ_2 and a_1 . Therefore, the optimal value a_1^* can be obtained by solving the equation

$$\rho_2(a_1^*) = \rho_0 \quad (\text{C.38})$$

for a_1^* , which yields:

$$a_1^* = \frac{[\Delta p(-2\gamma\rho_0 + \gamma + \rho_0) - \gamma + \rho_0](-2\gamma\rho_0 + \gamma + \rho_0)}{\Delta p(4\gamma^2\rho_0^2 - 4\gamma^2\rho_0 + \gamma^2 - 4\gamma\rho_0^2 + 4\gamma\rho_0 - \gamma + \rho_0^2 - \rho_0) + 2\gamma^2\rho_0 - \gamma^2 - 2\gamma\rho_0 + \gamma} \quad (\text{C.39})$$

It can be shown that $a_1^* \in [0, 1]$ for

$$\rho_0 \in \left[\frac{\Delta p(\bar{\gamma} - \gamma) + 2\gamma + \sqrt{\Delta p^2(\gamma - \bar{\gamma})^2 + 4\gamma\bar{\gamma}}}{2(\gamma - \bar{\gamma})}, \frac{\gamma(1 - \Delta p)}{\Delta p(\bar{\gamma} - \gamma) + 1} \right].$$

Therefore, since a_1^* decreases monotonically in ρ_0 , it follows that $a_1^* = 1$ for ρ_0 below the lower bound, and $a_1^* = 0$ for ρ_0 above the upper bound. The optimal a_1^* can thus be expressed as a piecewise function of ρ_0 :

$$a_1^* = \begin{cases} 1, & \text{if } \rho_0 < \rho_0^{\min} \\ \text{expression given in C.39}, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ 0, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.40})$$

The optimal profit functions can be expressed as:

$$\Pi_1^* = \begin{cases} (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)p_1, & \text{if } \rho_0 < \rho_0^{\min} \\ \frac{N_3}{D_3}p_1, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)^2p_1, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.41})$$

and since $\Pi_2^* = [1 - P(S_1 = g|a_1^*) \cdot P(S_2 = b)] \cdot p_2 = [1 - ((\gamma + a_1^*\bar{\gamma})\bar{\rho}_0 + (\bar{\gamma} + a_1^*\gamma)\rho_0)(\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)]p_2$,

$$\Pi_2^* = \begin{cases} [1 - (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)]p_2, & \text{if } \rho_0 < \rho_0^{\min} \\ \frac{N_4}{D_4}p_2, & \text{if } \rho_0 \in [\rho_0^{\min}, \rho_0^{\max}] \\ [1 - (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)^2]p_2, & \text{if } \rho_0 > \rho_0^{\max} \end{cases} \quad (\text{C.42})$$

where N_i and D_i are defined in equations C.11 – C.14 and $\rho_0^{\min}, \rho_0^{\max}$ are given by $\rho_0^{\min} = \frac{\Delta p(\bar{\gamma}-\gamma)+2\gamma+\sqrt{\Delta p^2(\gamma-\bar{\gamma})^2+4\gamma\bar{\gamma}}}{2(\gamma-\bar{\gamma})}$, and $\rho_0^{\max} = \frac{\gamma(1-\Delta p)}{\Delta p(\bar{\gamma}-\gamma)+1}$.

(iv) If $\rho_0 \in (\rho_2(a_1 = 0), 1]$, the customer will never buy from seller 1; therefore, there is no need for influence, and in the absence of influence cost, Seller 2 can set the influence level at $a_2 \in [0, 1]$.

Similarly, there is no chance of selling for Seller 1 at any influence level, therefore $a_1 \in [0, 1]$. This implies that in realistic settings where influence is costly, both sellers optimally choose not to have influence, that is, $a_1^* = a_2^* = 0$.

Substituting this yields optimal profits: $\Pi_1^* = 0$ and $\Pi_2^* = p_2$. ■

Proposition 5.2

Proof. Relying on Lemmas 5.1 and C.1, the firms' optimal influence and pricing choices are obtained by simultaneously maximizing their respective profit functions and solving for the Nash Equilibrium. Because $\rho_1^{\min}(a_2, \Delta p) = 0$ and $\rho_2^{\max}(a_1, \Delta p) = 1$, the first indicator terms in Π_1 and Π_2 (equations 5.10 and 5.11) evaluate to zero; thus, the profit functions reduce to

$$\begin{aligned} \Pi_1(p_1^*, a_1^*, a_2, p_2) = \\ \max_{0 \leq a_1 \leq 1, 0 \leq p_1 \leq 1} \left[(1 - P(S_1 = b) \cdot P(S_2 = g \mid a_2)) \cdot \mathbb{1}\{\rho_1 \leq \rho_0 < k\} \cdot p_1 \right. \\ \left. + P(S_1 = g \mid a_1) \cdot P(S_2 = b) \cdot \mathbb{1}\{k \leq \rho_0 < \rho_2\} \cdot p_1 \right] \end{aligned} \quad (\text{C.43})$$

and

$$\begin{aligned} \Pi_2(p_1, a_1, a_2^*, p_2^*) = \\ \max_{0 \leq a_2 \leq 1, 0 \leq p_2 \leq 1} \left[(1 - P(S_1 = g \mid a_1) \cdot P(S_2 = b)) \cdot \mathbb{1}\{k \leq \rho_0 < \rho_2\} \cdot p_2 \right. \\ \left. + P(S_1 = b) \cdot P(S_2 = g \mid a_2) \cdot \mathbb{1}\{\rho_1 \leq \rho_0 < k\} \cdot p_2 \right] \end{aligned} \quad (\text{C.44})$$

When $\rho_1 \leq \rho_0 \leq k$, equations C.43 and C.44 become

$$\begin{aligned} \Pi_1(p_1^*, a_2, p_2) = \max_{0 \leq p_1 \leq 1} [1 - P(S_1 = b)] \cdot P(S_2 = g \mid a_2) \cdot p_1 \\ \text{s.t.} \quad \rho_1 \leq \rho_0 \leq k \end{aligned} \quad (\text{C.45})$$

and

$$\begin{aligned} \Pi_2(p_1, a_2^*, p_2^*) = \max_{0 \leq a_2 \leq 1, 0 \leq p_2 \leq 1} P(S_1 = b) \cdot P(S_2 = g \mid a_2) \cdot p_2 \\ \text{s.t.} \quad \rho_1 \leq \rho_0 \leq k \end{aligned} \quad (\text{C.46})$$

Since $\Pi_1 = [1 - P(S_1 = b)] \cdot P(S_2 = g \mid a_2) \cdot p_1$ is strictly increasing in p_1 , and $k(\Delta p)$ is strictly decreasing in p_1 , it follows that the threshold condition $k(p_1^*, p_2) = \rho_0$ must hold. This yields the optimal price

$$p_1^* = 1 + p_2 - 2\rho_0. \quad (\text{C.47})$$

According to Lemma C.1, the threshold ρ_1 is strictly increasing in both a_2 and p_2 , due to its dependence on the price gap $\Delta p = p_1 - p_2$. Since Π_2 is also strictly increasing in a_2 and p_2 , the equilibrium condition $\rho_1(a_2^*, p_2^*) = \rho_0$ must be satisfied. Solving for a_2 , the expression $a_2^* = \rho_1^{-1}(p_1, p_2)$ can be substituted into Seller 2's payoff function:

$$\Pi_2(p_1, p_2^*) = \max_{0 \leq p_2 \leq 1} (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0) [(\gamma + \rho_1^{-1}(p_1, p_2) \cdot \bar{\gamma})\rho_0 + (\bar{\gamma} + \rho_1^{-1}(p_1, p_2) \cdot \gamma)\bar{\rho}_0] p_2. \quad (\text{C.48})$$

By substituting Seller 1's best response into the objective function and noting that ρ_1 depends on p_1 and p_2 only through Δp , it follows that $\Delta p^* = p_1^* - p_2 = 1 - 2\rho_0$ is constant. Therefore, $a_2^* = \rho_1^{-1}$ is constant with respect to p_2 , which implies that $\Pi_2(p_2)$ is strictly increasing in p_2 .

As a result, the optimal price for Seller 2 is attained at the upper bound, i.e.,

$$p_2^* = 1 \quad \text{for } \rho_1 \leq \rho_0 \leq k. \quad (\text{C.49})$$

Substituting into the expression for p_1^* , yields

$$p_1^* = 1 + p_2 - 2\rho_0 = 2\bar{\rho}_0 \quad \text{for } \rho_1 \leq \rho_0 \leq k. \quad (\text{C.50})$$

In particular, for $\rho_0 \leq 0.5$, it follows that $p_1^* = 1$ (the upper bound for p_1).

As for the optimal a_2^* , since $\Pi_2 = (\bar{\gamma}\bar{\rho}_0 + \gamma\rho_0)[(\gamma + a_2\bar{\gamma})\rho_0 + (\bar{\gamma} + a_2\gamma)\bar{\rho}_0]p_2$ and ρ_1 are strictly increasing in a_2 , solving $\rho_1(p_1 = 1, p_2 = 1, a_2^*) - \rho_0 = 0$ yields

$$a_2^* = \frac{(\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0))(\gamma - \bar{\rho}_0)}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}. \quad (\text{C.51})$$

Since $0 \leq a_2 \leq 1$ and considering its monotonicity with respect to ρ_0 , it can be shown that the optimal influence level for $\rho_1 \leq \rho_0 \leq k$ is determined as follows:

$$a_2^* = \begin{cases} 0, & \text{if } \rho_0 < \bar{\gamma} \\ 1, & \text{if } \rho_0 > \frac{\gamma + \sqrt{\gamma\bar{\gamma}} - 1}{\gamma - \bar{\gamma}} \\ \frac{(\bar{\rho}_0 + \gamma(\rho_0 - \bar{\rho}_0))(\gamma - \bar{\rho}_0)}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}, & \text{otherwise} \end{cases} \quad \text{for } \rho_1 \leq \rho_0 \leq k \quad (\text{C.52})$$

When $k \leq \rho_0 \leq \rho_2$, equations C.43 and C.44 become

$$\Pi_1(p_1^*, a_1^*, p_2) = \max_{0 \leq a_1 \leq 1 \text{ and } 0 \leq p_1 \leq 1} P(S_1 = g|a_1).P(S_2 = b).p_1 \quad (\text{C.53})$$

s.t. $k \leq \rho_0 \leq \rho_2$

and

$$\Pi_2(a_1, p_1, p_2^*) = \max_{0 \leq p_2 \leq 1} [1 - P(S_1 = g|a_1).P(S_2 = b)].p_2 \quad (\text{C.54})$$

s.t. $k \leq \rho_0 \leq \rho_2$

Since $\Pi_2 = [1 - P(S_1 = g | a_1) \cdot P(S_2 = b)] \cdot p_2$ and $k(\Delta p)$ are strictly increasing in p_2 , the

threshold condition $k(p_1, p_2^*) = \rho_0$ implies

$$p_2^* = p_1 + 2\rho_0 - 1. \quad (\text{C.55})$$

According to Lemma C.1, the threshold ρ_2 is strictly decreasing in both a_1 and p_1 , due to its dependence on the price gap $\Delta p = p_1 - p_2$. Since Π_1 is also strictly increasing in a_1 and p_1 , it follows that the equilibrium condition $\rho_2(a_1^*, p_1^*) = \rho_0$ must hold. Solving for a_1 , the expression $a_1^* = \rho_2^{-1}(p_1^*, p_2)$ can be substituted into the profit function to yield

$$\Pi_1(p_1^*, p_2) = \max_{0 \leq p_1 \leq 1} [(\gamma + \rho_2^{-1}(p_1, p_2) \cdot \bar{\gamma}) \cdot \bar{\rho}_0 + (\bar{\gamma} + \rho_2^{-1}(p_1, p_2) \cdot \gamma) \cdot \rho_0] (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0) \cdot p_1. \quad (\text{C.56})$$

By substituting Seller 2's best response into the objective function, and noting that ρ_2 depends only on p_1 and p_2 through Δp , it follows that $\Delta p^* = p_1 - p_2^* = 2\rho_0 - 1$. This implies that $a_1^* = \rho_2^{-1}$ is constant with respect to p_1 , and thus $\Pi_1(p_1)$ is strictly increasing in p_1 . Consequently, the optimal price is achieved at the upper bound

$$p_1^* = 1 \quad \text{for } k \leq \rho_0 \leq \rho_2. \quad (\text{C.57})$$

Substituting into the best response function yields

$$p_2^* = p_1 + 2\rho_0 - 1 = 2\rho_0 \quad \text{for } k \leq \rho_0 \leq \rho_2, \quad (\text{C.58})$$

and in particular, $p_2^* = 1$ for $\rho_0 \geq 0.5$.

As for the optimal influence level a_1^* , observe that the profit function

$$\Pi_1 = (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0) [(\gamma + a_1\bar{\gamma})\bar{\rho}_0 + (\bar{\gamma} + a_1\gamma)\rho_0] p_1, \quad (\text{C.59})$$

is strictly increasing in a_1 , while the threshold ρ_2 is strictly decreasing in a_1 . Therefore, the equilibrium value a_1^* is determined by solving the condition $\rho_2(p_1 = 1, p_2 = 1, a_1^*) = \rho_0$,

yielding:

$$a_1^* = \frac{(\rho_0 - \gamma)(-2\gamma\rho_0 + \gamma + \rho_0)}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}. \quad (\text{C.60})$$

Since $0 \leq a_1 \leq 1$ and considering its monotonicity with respect to ρ_0 , it can be shown that the optimal influence level is determined as follows:

$$a_1^* = \begin{cases} 0, & \text{if } \gamma < \rho_0 \\ 1, & \text{if } \rho_0 < \frac{\gamma - \sqrt{\gamma\bar{\gamma}}}{\gamma - \bar{\gamma}} \\ \frac{((\rho_0 - \gamma)(-2\gamma\rho_0 + \gamma + \rho_0))}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}, & \text{otherwise} \end{cases} \quad \text{for } k \leq \rho_0 \leq \rho_2 \quad (\text{C.61})$$

Integrating the analysis for $\rho_1 \leq \rho_0 \leq \rho_2$, the optimal prices are $p_1^* = p_2^* = 1$. Since the first component of C.43 is independent of a_1 , the optimal influence levels of the whole market can be expressed as $a_1^* = \text{C.61}$. Similarly, the first component of C.44 is independent of a_2 . Thus, $a_2^* = \text{C.52}$ holds for the whole market.

The optimal profit functions can be expressed as:

$$\Pi_1^* = \begin{cases} (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)^2, & \text{if } \gamma < \rho_0 \\ \bar{\gamma}\rho_0 + \gamma\bar{\rho}_0, & \text{if } \rho_0 < \frac{\gamma - \sqrt{\gamma\bar{\gamma}}}{\gamma - \bar{\gamma}} \\ \frac{N_3}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}, & \text{otherwise} \end{cases} \quad (\text{C.62})$$

$$\Pi_2^* = \begin{cases} 1 - (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0), & \text{if } \rho_0 < \bar{\gamma} \\ 1 - (\bar{\gamma}\rho_0 + \gamma\bar{\rho}_0)^2, & \text{if } \rho_0 > \frac{\gamma + \sqrt{\gamma\bar{\gamma}} - 1}{\gamma - \bar{\gamma}} \\ \frac{N_5}{\bar{\gamma} \cdot \gamma(\bar{\rho}_0 - \rho_0)}, & \text{otherwise} \end{cases} \quad (\text{C.63})$$

where N_3 and N_5 are defined in equations C.11 and C.15. ■

Proof. Corollary 5.1.

- (a) Immediate from Lemma C.1 by restricting the feasible set for Δp and re-evaluating ρ_1, ρ_2 under the bound $\bar{\Delta}$.
- (b) If $\rho_0 \in [0, \rho_1(a_2 = 0, \Delta p = v = 1))$, the customer will never buy from Seller 2, eliminating any incentive for Seller 2 to exert influence. Moreover, in the absence of influence costs, Seller 1 is free to choose any influence level $a_1 \in [0, 1]$, but has no incentive to do so either. This implies that in realistic settings where influence is costly, both sellers optimally choose no influence, that is, $a_1^* = a_2^* = 0$.
Given this, Seller 1 can set the highest possible price $p_1^* = 1$ equal to zero yields the optimal price for Seller 1, $p_1^* = v = 1$.
Since Seller 2 does not sell, its price can be any value within the range $p_2 \in [0, v]$. The corresponding expected profits are $\Pi_1^* = p_1^* = 1$ and $\Pi_2^* = 0$.
- (c) If $\rho_0 \in (\rho_2(a_1 = 0, \Delta p = -v = -1), 1]$, by symmetry, all results follow those in (b) with seller labels reversed, that is, $a_1^* = a_2^* = 0$, $p_2^* = 1$, and $\Pi_1^* = 0$. The corresponding expected profits are: $\Pi_1^* = 0$ and $\Pi_2^* = 1$.
- (d) Proposition 5.2 shows that, in equilibrium, both sellers charge the monopoly price,

$$p_1^* = p_2^* = 1 \implies \Delta p^* = p_1^* - p_2^* = 0.$$

Hence the regulatory cap $-\bar{\Delta} < \Delta p < \bar{\Delta}$ is *automatically* satisfied, i.e. it is non-binding in the pricing stage.

However, the cap $|\Delta p| < \bar{\Delta}$ *does* affect the secure-market thresholds because ρ_1 and ρ_2 depend on the *feasible*—not the realised—range of Δp . ■